

CONTENTS

LIARS AND OUTLIERS

Chapter 1: Overview	1
Part I: The Science of Trust	15
Chapter 2: A Natural History of Security	17
Chapter 3: The Evolution of Cooperation	27
Chapter 4: A Social History of Trust	41
Chapter 5: Societal Dilemmas	51
Part II: A Model of Trust	61
Chapter 6: Societal Pressures	63
Chapter 7: Moral Pressures	75
Chapter 8: Reputational Pressures	87
Chapter 9: Institutional Pressures	103
Chapter 10: Security Systems	123
Part III: The Real World	137
Chapter 11: Competing Interests	139
Chapter 12: Organizations	155
Chapter 13: Corporations	173
Chapter 14: Institutions	195
Part IV: Conclusions	205
Chapter 15: How Societal Pressures Fail	207
Chapter 16: Technological Advances	225
Chapter 17: The Future	243

CARRY ON

Chapter 1: The Business and Economics of Security	1
Chapter 2: Crime, Terrorism, Spying, and War	41
Chapter 3: Human Aspects of Security	117
Chapter 4: Privacy and Surveillance	155
Chapter 5: Psychology of Security	189
Chapter 6: Security and Technology	209
Chapter 7: Travel and Security	259
Chapter 8: Security, Policy, Liberty, and Law	279

"Without trust, nothing can be achieved. *Liars and Outliers* is a brilliant analysis of the role of trust in society and business." —**Klaus Schwab**, *Founder and Executive Chairman, World Economic Forum*

LIARS & OUTLIERS

ENABLING THE TRUST THAT SOCIETY NEEDS TO THRIVE



Bruce Schneier

Bestselling Author of *Secrets and Lies* and *Beyond Fear*

Advance Praise for **Liars and Outliers**

“A rich, insightfully fresh take on what security really means!” —DAVID ROPEIK
author of *How Risky is it, Really?*

“Schneier has accomplished a spectacular tour de force: an enthralling ride through history, economics, and psychology, searching for the meanings of trust and security. A must read.” —ALESSANDRO ACQUISTI

Associate Professor of Information Systems and Public Policy
at the Heinz College, Carnegie Mellon University

“*Liars and Outliers* offers a major contribution to the understandability of these issues, and has the potential to help readers cope with the ever-increasing risks to which we are being exposed. It is well written and delightful to read.”

—PETER G. NEUMANN
Principal Scientist in the SRI International
Computer Science Laboratory

“Whether it’s banks versus robbers, Hollywood versus downloaders, or even the Iranian secret police against democracy activists, security is often a dynamic struggle between a majority who want to impose their will, and a minority who want to push the boundaries. *Liars and Outliers* will change how you think about conflict, our security, and even who we are.” —ROSS ANDERSON

Professor of Security Engineering at Cambridge University
and author of *Security Engineering*

“Readers of Bruce Schneier’s *Liars and Outliers* will better understand technology and its consequences and become more mature practitioners.” —PABLO G. MOLINA
Professor of Technology Management
Georgetown University

“*Liars and Outliers* is not just a book about security—it is the book about it. Schneier shows that the power of humour can be harnessed to explore even a serious subject such as security. A great read!” —FRANK FUREDI

Professor Emeritus, School of Social Policy, Sociology and Social Research
The University of Kent at Canterbury and author of
On Tolerance: A Defence of Moral Independence

“This fascinating book gives an insightful and convincing framework for understanding security and trust.”

—JEFF YAN

Founding Research Director, Center for Cybercrime and Computer Security
Newcastle University

“By analyzing the moving parts and interrelationships among security, trust, and society, Schneier has identified critical patterns, pressures, levers, and security holes within society. Clearly written, thoroughly interdisciplinary, and always smart, *Liars and Outliers* provides great insight into resolving society’s various dilemmas.”

—JERRY KANG

Professor of Law, UCLA

“By keeping the social dimension of trust and security in the center of his analysis, Schneier breaks new ground with an approach that’s both theoretically grounded and practically applicable.”

—JONATHAN ZITTRAIN

Professor of Law and Computer Science, Harvard University
and author of *The Future of the Internet—And How to Stop It*

“Eye opening. Bruce Schneier provides a perspective you need to understand today’s world.”

—STEVEN A. LEBLANC

Director of Collections, Harvard University
and author of *Constant Battles: Why We Fight*

“An outstanding investigation of the importance of trust in holding society together and promoting progress. *Liars and Outliers* provides valuable new insights into security and economics.”

—ANDREW ODLYZKO

Professor, School of Mathematics, University of Minnesota

“What Schneier has to say about trust—and betrayal—lays a groundwork for greater understanding of human institutions. This is an essential exploration as society grows in size and complexity.”

—JIM HARPER

Director of Information Policy Studies, CATO Institute
and author of *Identity Crisis: How Identification Is Overused and Misunderstood*

“Society runs on trust. *Liars and Outliers* explains the trust gaps we must fill to help society run even better.”

—M. ERIC JOHNSON

Director, Glassmeyer/McNamee Center for Digital Strategies
Tuck School of Business at Dartmouth College

“An intellectually exhilarating and compulsively readable analysis of the subtle dialectic between cooperation and defection in human society. Intellectually rigorous and yet written in a lively, conversational style, *Liars and Outliers* will change the way you see the world.”

—DAVID LIVINGSTONE SMITH

Associate Professor of Philosophy, University of New England
and author of *Less Than Human: Why We Demean, Enslave, and Exterminate Others*

“Schneier tackles trust head on, bringing all his intellect and a huge amount of research to bear. The best thing about this book, though, is that it’s great fun to read.”

—ANDREW MCAFEE

Principal Research Scientist, MIT Center for Digital Business
and co-author of *Race Against the Machine*

“Bruce Schneier is our leading expert in security. But his book is about much more than reducing risk. It is a fascinating, thought-provoking treatise about humanity and society and how we interact in the game called life.”

—JEFF JARVIS

author of *Public Parts: How Sharing in the Digital Age
Improves the Way We Work and Live*

“Both accessible and thought provoking, *Liars and Outliers* invites readers to move beyond fears and anxieties about security in modern life to understand the role of everyday people in creating a healthy society. This is a must-read!”

—DANAH BOYD

Research Assistant Professor in Media, Culture, and Communication
New York University

“Trust is the *sine qua non* of the networked age and trust is predicated on security. Bruce Schneier’s expansive and readable work is rich with insights that can help us make our shrinking world a better one.”

—DON TAPSCOTT

co-author of *Macrowikinomics: Rebooting Business and the World*

“An engaging and wide-ranging rumination on what makes society click. Highly recommended.”

—JOHN MUELLER

Senior Research Scientist, Mershon Center, Ohio State University
and author of *Overblown: How Politicians and the Terrorism Industry Inflate
National Security Threats, and Why We Believe Them*



Liars and Outliers

Liars and Outliers

Enabling the Trust That
Society Needs to Thrive

Bruce Schneier



John Wiley & Sons, Inc.

Liars and Outliers: Enabling the Trust That Society Needs to Thrive

Published by
John Wiley & Sons, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2012 by Bruce Schneier
Published by John Wiley & Sons, Inc., Indianapolis, Indiana
Published simultaneously in Canada

ISBN: 978-1-118-14330-8
ISBN: 978-1-118-22556-1 (ebk)
ISBN: 978-1-118-23901-8 (ebk)
ISBN: 978-1-118-26362-4 (ebk)

Manufactured in the United States of America
10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or website may provide or recommendations it may make. Further, readers should be aware that Internet websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Not all content that is available in standard print versions of this book may appear or be packaged in all book formats. If you have purchased a version of this book that did not include media that is referenced by or accompanies a standard print version, you may request this media by visiting <http://booksupport.wiley.com>. For more information about Wiley products, visit us at www.wiley.com.

Library of Congress Control Number: 2011944879

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.



Credits

Executive Editor

Carol Long

Project Editor

Tom Dinse

Senior Production Editor

Debra Banninger

Copy Editor

Kim Cofer

Editorial Manager

Mary Beth Wakefield

Freelancer Editorial Manager

Rosemarie Graham

Marketing Manager

Ashley Zurcher

Business Manager

Amy Knies

Production Manager

Tim Tate

**Vice President and Executive Group
Publisher**

Richard Swadley

Vice President and Executive Publisher

Neil Edde

Associate Publisher

Jim Minatel

Project Coordinator, Cover

Katie Crocker

Proofreader

Nancy Carrasco

Indexer


Johnna Dinse

Cover Designer

Ryan Sneed

Cover Concept

Luke Fretwell



Contents

<i>A Note for Readers</i>	<i>xiii</i>
1 Overview	1
PART I THE SCIENCE OF TRUST	15
2 A Natural History of Security	17
3 The Evolution of Cooperation	27
4 A Social History of Trust	41
5 Societal Dilemmas	51
PART II A MODEL OF TRUST	61
6 Societal Pressures	63
7 Moral Pressures	75
8 Reputational Pressures	87
9 Institutional Pressures	103
10 Security Systems	123

PART III THE REAL WORLD	137
11 Competing Interests	139
12 Organizations	155
13 Corporations	173
14 Institutions	195
PART IV CONCLUSIONS	205
15 How Societal Pressures Fail	207
16 Technological Advances	225
17 The Future	243
<i>Acknowledgments</i>	249
<i>Notes</i>	251
<i>References</i>	289
<i>About the Author</i>	349
<i>Index</i>	351



A Note for Readers

This book contains both notes and references. The notes are explanatory bits that didn't make it into the main text. These are indicated by superscript numbers in both the paper and e-book formats. The references are indicated by links in the main text.

High-resolution versions of the figures can be found at www.schneier.com/10.



1

Overview

Just today, a stranger came to my door claiming he was here to unclog a bathroom drain. I let him into my house without verifying his identity, and not only did he repair the drain, he also took off his shoes so he wouldn't track mud on my floors. When he was done, I gave him a piece of paper that asked my bank to give him some money. He accepted it without a second glance. At no point did he attempt to take my possessions, and at no point did I attempt the same of him. In fact, neither of us worried that the other would. My wife was also home, but it never occurred to me that he was a sexual rival and I should therefore kill him.

Also today, I passed several strangers on the street without any of them attacking me. I bought food from a grocery store, not at all concerned that it might be unfit for human consumption. I locked my front door, but didn't spare a moment's worry at how easy it would be for someone to smash my window in. Even people driving cars, large murderous instruments that could crush me like a bug, didn't scare me.

Most amazingly, this worked without much overt security. I don't carry a gun for self-defense, nor do I wear body armor. I don't use a home burglar alarm. I don't test my food for poison. I don't even engage in conspicuous displays of physical prowess to intimidate other people I encounter.

It's what we call "trust." Actually, it's what we call "civilization."

All complex ecosystems, whether they are biological ecosystems like the human body, natural ecosystems like a rain forest, social ecosystems like an open-air market, or socio-technical ecosystems like the global financial system or the Internet, are deeply interlinked. Individual units within those ecosystems are interdependent, each doing its part and relying on the other units to do their parts as well. This is neither rare nor difficult, and complex ecosystems abound.

At the same time, all complex ecosystems **contain parasites**. Within every interdependent system, there are individuals who try to subvert the system to their own ends. These could be tapeworms in our digestive tracts, thieves in a bazaar, robbers disguised as plumbers, spammers on the Internet, or companies that move their profits offshore to evade taxes.

Within complex systems, there is a fundamental tension between what I'm going to call cooperating, or acting in the group interest; and what I'll call defecting, or acting against the group interest and instead in one's own self-interest. Political philosophers have recognized this antinomy since Plato. We might individually want each other's stuff, but we're collectively better off if everyone respects property rights and no one steals. We might individually want to reap the benefits of government without having to pay for them, but we're collectively better off if everyone pays taxes. Every country might want to be able to do whatever it wants, but the world is better off with international agreements, treaties, and organizations. In general, we're collectively better off if society limits individual behavior, and we'd each be better off if those limits didn't apply to us individually. That doesn't work, of course, and most of us recognize this. Most of the time, we realize that it is in our self-interest to act in the group interest. But because parasites will always exist—because some of us steal, don't pay our taxes, ignore international agreements, or ignore limits on our behavior—we also need security.

Society runs on trust. We all need to trust that the random people we interact with will cooperate. Not trust completely, not trust blindly, but be reasonably sure (whatever that means) that our trust is well-founded and they will be trustworthy in return (whatever *that* means). This is vital. If the number of parasites gets too large, if too many people steal or too many people don't pay their taxes, society no longer works. It doesn't work both because there is so much theft that people can't be secure in their property, and because even the honest become suspicious of everyone else. More importantly, it doesn't work because the social contract breaks down: society is no longer seen as providing the required benefits. Trust is largely habit, and when there's not enough trust to be had, people stop trusting each other.

The devil is in the details. In all societies, for example, there are instances where property is legitimately taken from one person and given to another: taxes, fines, fees, confiscation of contraband, theft by a legitimate but despised ruler, etc. And a societal norm like "everyone pays his or her taxes" is distinct from any discussion about what sort of tax code is fair. But while we might disagree

about the extent of the norms we subject ourselves to—that’s what politics is all about—we’re collectively better off if we all follow them.

Of course, it’s actually more complicated than that. A person might decide to break the norms, not for selfish parasitical reasons, but because his moral compass tells him to. He might help escaped slaves flee into Canada because slavery is wrong. He might refuse to pay taxes because he disagrees with what his government is spending his money on. He might help laboratory animals escape because he believes animal testing is wrong. He might shoot a doctor who performs abortions because he believes abortion is wrong. And so on.

Sometimes we decide a norm breaker did the right thing. Sometimes we decide that he did the wrong thing. Sometimes there’s consensus, and sometimes we disagree. And sometimes those who dare to defy the group norm become catalysts for social change. Norm breakers rioted against the police raids of the Stonewall Inn in New York in 1969, at the beginning of the gay rights movement. Norm breakers hid and saved the lives of Jews in World War II Europe, organized the Civil Rights bus protests in the American South, and assembled in unlawful protest at Tiananmen Square. When the group norm is later deemed immoral, history may call those who refused to follow it heroes.

In 2008, the U.S. real estate industry collapsed, almost taking the global economy with it. The causes of the disaster are complex, but were in a large part caused by financial institutions and their employees subverting financial systems to their own ends. They wrote mortgages to homeowners who couldn’t afford them, and then repackaged and resold those mortgages in ways that intentionally hid real risk. Financial analysts, who made money rating these bonds, gave them high ratings to ensure repeat rating business.

This is an example of a failure of trust: a limited number of people were able to use the global financial system for their own personal gain. That sort of thing isn’t supposed to happen. But it did happen. And it will happen again if society doesn’t get better at both trust and security.

Failures in trust have become global problems:

- The Internet brings amazing benefits to those who have access to it, but it also brings with it new forms of fraud. Impersonation fraud—now called identity theft—is both easier and more profitable than it was pre-Internet. Spam continues to undermine the usability of e-mail. Social networking sites deliberately make it hard for people to effectively manage [their own privacy](#). And antagonistic behavior threatens almost every Internet community.

- Globalization has improved the lives of people in many countries, but with it came an increased threat of global terrorism. The terrorist attacks of 9/11 were a failure of trust, and so were the government overreactions in the decade following.
- The financial network allows anyone to do business with anyone else around the world; but easily hacked financial accounts mean there is enormous profit in fraudulent transactions, and easily hacked computer databases mean there is also a global market in (terrifyingly cheap) stolen credit card numbers and personal dossiers to enable those fraudulent transactions.
- Goods and services are now supplied worldwide at much lower cost, but with this change comes tainted foods, unsafe children's toys, and the outsourcing of data processing to countries with different laws.
- Global production also means more production, but with it comes environmental pollution. If a company discharges lead into the atmosphere—or chlorofluorocarbons, or nitrogen oxides, or carbon dioxide—that company gets all the benefit of cheaper production costs, but the environmental cost falls on everybody else on the planet.

And it's not just global problems, of course. Narrower failures in trust are so numerous as to defy listing. Here are just a few examples:

- In 2009–2010, officials of Bell, California, [effectively looted](#) the city's treasury, awarding themselves unusually high salaries, often for part-time work.
- Some early online games, such as Star Wars Galaxy Quest, collapsed due to [internal cheating](#).
- The senior executives at companies such as WorldCom, Enron, and Adelphia inflated their companies' stock prices through fraudulent accounting practices, awarding themselves huge bonuses but destroying the companies in the process.

What ties all these examples together is that the interest of society was in conflict with the interests of certain individuals within society. Society had some normative behaviors, but failed to ensure that enough people cooperated and

followed those behaviors. Instead, the defectors within the group became too large or too powerful or too successful, and ruined it for everyone.



This book is about trust. Specifically, it's about trust within a group. It's important that defectors not take advantage of the group, but it's also important for everyone in the group to trust that defectors won't take advantage.

“Trust” is a complex concept, and has a lot of flavors of meaning. Sociologist Piotr Sztompka wrote that “[trust is a bet](#) about the future contingent actions of others.” Political science professor Russell Hardin wrote: “[Trust involves](#) giving discretion to another to affect one's interests.” These definitions focus on trust between individuals and, by extension, their trustworthiness.¹

When we trust people, we can either trust their intentions or their actions. The first is more intimate. When we say we trust a friend, that trust isn't tied to any particular thing he's doing. It's a general reliance that, whatever the situation, he'll do the right thing: that he's trustworthy. We trust the friend's intentions, and know that his actions will be informed by those intentions.²

The second is less intimate, what sociologist Susan Shapiro calls [impersonal trust](#). When we don't know someone, we don't know enough about her, or her underlying motivations, to trust her based on character alone. But we can trust her future actions.³ We can trust that she won't run red lights, or steal from us, or cheat on tests. We don't know if she has a secret desire to run red lights or take our money, and we really don't care if she does. Rather, we know that she is likely to follow most social norms of acceptable behavior because the consequences of breaking these norms are high. You can think of this kind of trust—that people will behave in a trustworthy manner even if they are not inherently trustworthy—more as confidence, and the corresponding trustworthiness as compliance.⁴

In another sense, we're reducing trust to consistency or predictability. Of course, someone who is consistent isn't necessarily trustworthy. If someone is a habitual thief, I don't trust him. But I do believe (and, in another sense of the word, trust) that he will try to steal from me. I'm less interested in that aspect of trust, and more in the positive aspects. In *The Naked Corporation*, business strategist Don Tapscott [described trust](#), at least in business, as the expectation that the other party will be honest, considerate, accountable, and transparent. When two people are consistent in this way, we call them cooperative.

In today's complex society, we often trust systems more than people. It's not so much that I trusted the plumber at my door as that I trusted the systems that produced him and protect me. I trusted the recommendation from my insurance company, the legal system that would protect me if he did rob my house, whatever the educational system is that produces and whatever insurance system bonds skilled plumbers, and—most of all—the general societal systems that inform how we all treat each other in society. Similarly, I trusted the banking system, the corporate system, the system of police, the system of traffic laws, and the system of social norms that govern most behaviors.⁵

This book is about trust more in terms of groups than individuals. I'm not really concerned about how specific people come to trust other specific people. I don't care if my plumber trusts me enough to take my check, or if I trust that driver over there enough to cross the street at the stop sign. I'm concerned with the general level of impersonal trust in society. Francis Fukuyama's definition nicely captures the term as I want to use it: "[Trust is the expectation](#) that arises within a community of regular, honest, and cooperative behavior, based on commonly shared norms, on the part of other members of that community."

Sociologist Barbara Misztal identified [three critical functions](#) performed by trust: 1) it makes social life more predictable, 2) it creates a sense of community, and 3) it makes it easier for people to work together. In some ways, trust in society works like oxygen in the atmosphere. The more customers trust merchants, the easier commerce is. The more drivers trust other drivers, the smoother traffic flows. Trust gives people the confidence to deal with strangers: because they know that the strangers are likely to behave honestly, cooperatively, fairly, and sometimes even altruistically. The more trust is in the air, the healthier society is and the more it can thrive. Conversely, the less trust is in the air, the sicker society is and the more it has to contract. And if the amount of trust gets too low, society withers and dies. A [recent example](#) of a systemic breakdown in trust occurred in the Soviet Union under Stalin.

I'm necessarily simplifying here. Trust is relative, fluid, and multidimensional. I trust Alice to return a \$10 loan but not a \$10,000 loan, Bob to return a \$10,000 loan but not to babysit an infant, Carol to babysit but not with my house key, Dave with my house key but not my intimate secrets, and Ellen with my intimate secrets but not to return a \$10 loan. I trust Frank if a friend vouches for him, a taxi driver as long as he's displaying his license, and Gail as long as she hasn't been drinking. I don't trust anyone at all with my computer password. I trust my brakes to stop the car, ATM machines to dispense money from my account, and Angie's List to recommend a qualified plumber—even though

I have no idea who designed, built, or maintained those systems. Or even who Angie is. In the language of this book, we all need to trust each other to follow the behavioral norms of our group.

Many other books talk about the [value of trust](#) to society. This book explains how society establishes and maintains that trust.⁶ Specifically, it explains how society enforces, evokes, elicits, compels, encourages—I'll use the term *induces*—trustworthiness, or at least compliance, through systems of what I call *societal pressures*, similar to sociology's social controls: coercive mechanisms that induce people to cooperate, act in the group interest, and follow group norms. Like physical pressures, they don't work in all cases on all people. But again, whether the pressures work against a particular person is less important than whether they keep the scope of defection to a manageable level across society as a whole.

A manageable level, but not too low a level. Compliance isn't always good, and defection isn't always bad. Sometimes the group norm doesn't deserve to be followed, and certain kinds of progress and innovation require violating trust. In a police state, everybody is compliant but no one trusts anybody. A too-compliant society is a stagnant society, and defection contains the seeds of social change.

This book is also about security. Security is a type of a societal pressure in that it induces cooperation, but it's different from the others. It is the only pressure that can act as a physical constraint on behavior regardless of how trustworthy people are. And it is the only pressure that individuals can implement by themselves. In many ways, it obviates the need for intimate trust. In another way, it is how we ultimately induce compliance and, by extension, trust.

It is essential that we learn to think smartly about trust. Philosopher Sissela Bok wrote: "Whatever matters to human beings, [trust is the atmosphere](#) in which it thrives." People, communities, corporations, markets, politics: everything. If we can figure out the optimal societal pressures to induce cooperation, we can reduce murder, terrorism, bank fraud, industrial pollution, and all the rest.

If we get pressures wrong, the murder rate skyrockets, terrorists run amok, employees routinely embezzle from their employers, and corporations lie and cheat at every turn. In extreme cases, an untrusting society breaks down. If we get them wrong in the other direction, no one speaks out about institutional injustice, no one deviates from established corporate procedure, and no one popularizes new inventions that disrupt the status quo—an oppressed society stagnates. The very fact that the most extreme failures rarely happen in the modern industrial world is proof that we've largely gotten societal pressures right. The failures that we've had show we have a lot further to go.

Also, as we'll see, evolution has left us with intuitions about trust better suited to life as a savannah-dwelling primate than as a modern human in a global high-tech society. That flawed intuition is vulnerable to exploitation by companies, con men, politicians, and crooks. The *only* defense is a rational understanding of what trust in society is, how it works, and why it succeeds or fails.



This book is divided into four parts. In Part I, I'll explore the background sciences of the book. Several fields of research—some closely related—will help us understand these topics: experimental psychology, evolutionary psychology, sociology, economics, behavioral economics, evolutionary biology, neuroscience, game theory, systems dynamics, anthropology, archaeology, history, political science, law, philosophy, theology, cognitive science, and computer security.

All these fields have something to teach us about trust and security.⁷ There's a lot here, and delving into any of these areas of research could easily fill several books. This book attempts to gather and synthesize decades, and sometimes centuries, of thinking, research, and experimentation from a broad swath of academic disciplines. It will, by necessity, be largely a cursory overview; often, the hardest part was figuring out what *not* to include. My goal is to show where the broad arcs of research are pointing, rather than explain the details—though they're fascinating—of any individual piece of research.⁸

In the last chapter of Part I, I will introduce societal dilemmas. I'll explain a thought experiment called the Prisoner's Dilemma, and its generalization to societal dilemmas. Societal dilemmas describe the situations that require intra-group trust, and therefore use societal pressures to ensure cooperation: they're the central paradigm of my model. Societal dilemmas illustrate how society keeps defectors from taking advantage, taking over, and completely ruining society for everyone. It illustrates how society ensures that its members forsake their own interests when they run counter to society's interest. Societal dilemmas have many names in the literature: collective action problem, Tragedy of the Commons, free-rider problem, arms race. We'll use them all.

Part II fully develops my model. Trust is essential for society to function, and societal pressures are how we achieve it. There are four basic categories of societal pressure that can induce cooperation in societal dilemmas:

- *Moral pressure.* A lot of societal pressure comes from inside our own heads. Most of us don't steal, and it's not because there are armed guards

and alarms protecting piles of stuff. We don't steal because we believe it's wrong, or we'll feel guilty if we do, or we want to follow the rules.

- *Reputational pressure.* A wholly different, and much stronger, type of pressure comes from how others respond to our actions. Reputational pressure can be very powerful; both individuals and organizations feel a lot of pressure to follow the group norms because they don't want a bad reputation.
- *Institutional pressure.* Institutions have rules and laws. These are norms that are codified, and whose enactment and enforcement is generally delegated. Institutional pressure induces people to behave according to the group norm by imposing sanctions on those who don't, and occasionally by rewarding those who do.
- *Security systems.* Security systems are another form of societal pressure. This includes any security mechanism designed to induce cooperation, prevent defection, induce trust, and compel compliance. It includes things that work to prevent defectors, like door locks and tall fences; things that interdict defectors, like alarm systems and guards; things that only work after the fact, like forensic and audit systems; and mitigation systems that help the victim recover faster and care less that the defection occurred.

Part III applies the model to the more complex dilemmas that arise in the real world. First I'll look at the full complexity of competing interests. It's not just group interest versus self-interest; people have a variety of competing interests. Also, while it's easy to look at societal dilemmas as isolated decisions, it's common for people to have conflicts of interest: multiple group interests and multiple societal dilemmas are generally operating at any one time. And the effectiveness of societal pressures often depends on why someone is considering defecting.

Then, I'll look at groups as actors in societal dilemmas: organizations in general, corporations, and then institutions. Groups have different competing interests, and societal pressures work differently when applied to them. This is an important complication, especially in the modern world of complex corporations and government agencies. Institutions are also different. In today's world, it's rare that we implement societal pressures directly. More often, we delegate someone to do it for us. For example, we delegate our elected officials to pass laws, and they delegate some government agency to implement those laws.

In Part IV, I'll talk about the different ways societal pressures fail. I'll look at how changes in technology affect societal pressures, particularly security. Then I'll look at the particular characteristics of today's society—the Information Society—and explain why that changes societal pressures. I'll sketch what the future of societal pressures is likely to be, and close with the social consequences of too much societal pressure.

This book represents my attempt to develop a full-fledged theory of coercion and how it enables compliance and trust within groups. My goal is to suggest some new questions and provide a new framework for analysis. I offer new perspectives, and a broader spectrum of what's possible. Perspectives frame thinking, and sometimes asking new questions is the catalyst to greater understanding. It's my hope that this book can give people an illuminating new framework with which to help understand how the world works.



Before we start, I need to define my terms. We talk about trust and security all the time, and the words we use tend to be overloaded with meaning. We're going to have to be more precise...and temporarily suspend our emotional responses to what otherwise might seem like loaded, value-laden, even disparaging, words.

The word *society*, as used in this book, isn't limited to traditional societies, but is any group of people with a loose common interest. It applies to societies of circumstance, like a neighborhood, a country, everyone on a particular bus, or an ethnicity or social class. It applies to societies of choice, like a group of friends, any membership organization, or a professional society. It applies to societies that are some of each: a religion, a criminal gang, or all employees of a corporation. It applies to societies of all sizes, from a family to the entire planet. All of humanity is a society, and everyone is a member of multiple societies. Some are based on birth, and some are freely chosen. Some we can join, and to some we must be invited. Some may be good, some may be bad—terrorist organizations, criminal gangs, a political party you don't agree with—and most are somewhere in between. For our purposes, a society is just a group of interacting *actors* organized around a common attribute.

I said actors, not people. Most societies are made up of people, but sometimes they're made up of groups of people. All the countries on the planet are a society. All corporations in a particular industry are a society. We're going to be talking about both societies of individuals and societies of groups.

Societies have a collection of *group interests*. These are the goals, or directions, of the society. They're decided by the society in some way: perhaps formally—either democratically or autocratically—perhaps informally by the group. International trade can be in the group interest. So can sharing food, obeying traffic laws, and keeping slaves (assuming those slaves are not considered to be part of the group). Corporations, families, communities, and terrorist groups all have their own group interests. Each of these group interests corresponds to one or more norms, which is what each member of that society is supposed to do. For example, it is in the group interest that everyone respect everyone else's property rights. Therefore, the group norm is not to steal (at least, not from other members of the group⁹).

Every person in a society potentially has one or more *competing interests* that conflict with the group interest, and *competing norms* that conflict with the group norm. Someone in that we-don't-steal society might really want to steal. He might be starving, and need to steal food to survive. He just might want other people's stuff. These are examples of *self-interest*. He might have some competing *relational interest*. He might be a member of a criminal gang, and need to steal to prove his loyalty to the group; here, the competing interest might be the group interest of another group. Or he might want to steal for some higher moral reason: a competing *moral interest*—the Robin Hood archetype, for example.

A *societal dilemma* is the choice every actor has to make between group interest and his or her competing interests. It's the choice we make when we decide whether or not to follow the group norm. Those who do *cooperate*, and those who do not *defect*. Those are both loaded terms, but I mean them to refer only to the action as a result of the dilemma.

Defectors—the liars and outliers of the book's title—are the people within a group who don't go along with the norms of that group. The term isn't defined according to any absolute morals, but instead in opposition to whatever the group interest and the group norm is. Defectors steal in a society that has declared that stealing is wrong, but they also help slaves escape in a society where tolerating slavery is the norm. Defectors change as society changes; defection is in the eye of the beholder. Or, more specifically, it is in the eyes of everyone else. Someone who was a defector under the former East German government was no longer in that group after the fall of the Berlin Wall. But those who followed the societal norms of East Germany, like the Stasi, were—all of a sudden—viewed as defectors within the new united Germany.




Figure 1: The Terms Used in the Book, and Their Relationships

Criminals are defectors, obviously, but that answer is too facile. Everyone defects at least some of the time. It's both dynamic and situational. People can cooperate about some things and defect about others. People can cooperate with one group they're in and defect from another. People can cooperate today and defect tomorrow, or cooperate when they're thinking clearly and defect when

they're reacting in a panic. People can cooperate when their needs are cared for, and defect when they're starving.

When four black North Carolina college students staged a sit-in at a whites-only lunch counter inside a Woolworth's five-and-dime store in Greensboro, in 1960, they were criminals. So are women who drive cars in Saudi Arabia. Or homosexuals in Iran. Or the 2011 protesters in Egypt, who sought to end their country's political regime. Conversely, child brides in Pakistan are not criminalized and neither are their parents, even though in some cases they marry off five-year-old girls. The Nicaraguan rebels who fought the Sandinistas were criminals, terrorists, insurgents, or freedom fighters, depending on which side you supported and how you viewed the conflict. Pot smokers and dealers in the U.S. are officially criminals, but in the Netherlands those offenses are ignored by the police. Those who share copyrighted movies and music are breaking the law, even if they have moral justifications for their actions.

Defecting doesn't necessarily mean breaking government-imposed laws. An orthodox Jew who eats a ham and cheese sandwich is violating the rules of his religion. A Mafioso who snitches on his colleagues is violating *omertà*, the code of silence. A relief worker who indulges in a long, hot shower after a tiring journey, and thereby depletes an entire village's hot water supply, unwittingly puts his own self-interest ahead of the interest of the people he intends to help.

What we're concerned with is the overall *scope of defection*. I mean this term to be general, comprising the number of defectors, the rate of their defection, the frequency of their defection, and the intensity (the amount of damage) of their defection. Just as we're interested in the general level of trust within the group, we're interested in the general scope of defection within the group.

Societal pressures are how society ensures that people follow the group norms, as opposed to some competing norms. The term is meant to encompass everything society does to protect itself: both from fellow members of society, and non-societal members who live within and amongst the society. More generally, it's how society enforces intra-group trust.

The terms *attacker* and *defender* are pretty obvious. The predator is the attacker, the prey is the defender. It's all intertwined, and sometimes these terms can get a bit muddied. Watch a martial arts match, and you'll see each person defending against his opponent's attacks while at the same time hoping his own attacks get around his opponent's defenses. In war, both sides attack and defend at the tactical level, even though one side might be attacking and the other defending at the political level. These terms are value-neutral. Attackers can be criminals trying to break into a home, superheroes raiding a criminal

mastermind's stronghold, or cancer cells metastasizing their way through a hapless human host. Defenders can be a family protecting its home from invasion, the criminal mastermind protecting his lair from the superheroes, or a posse of leukocytes engulfing opportunistic pathogens they encounter.

These definitions are important to remember as you read this book. It's easy for us to bring our own emotional baggage into discussions about security, but most of the time we're just trying to understand the underlying mechanisms at play, and those mechanisms are the same, regardless of the underlying moral context.

Sometimes we need the dispassionate lens of history to judge famous defectors like Oliver North, Oskar Schindler, and Vladimir Lenin.

PART I

The Science of Trust





2

A Natural History of Security

Our exploration of trust is going to start and end with security, because security is what you need when you don't have any trust and—as we'll see—security is ultimately how we induce trust in society. It's what brings risk down to tolerable levels, allowing trust to fill in the remaining gaps.

You can learn a lot about security from watching the natural world.

- Lions seeking to protect their turf will raise their voices in a “[territorial chorus](#),” their cooperation reducing the risk of encroachment by other predators for the local food supply.
- When [hornworms](#) start eating a particular species of sagebrush, the plant responds by emitting a molecule that warns any wild tobacco plants growing nearby that hornworms are around. In response, the tobacco plants deploy chemical defenses that repel the hornworms, to the benefit of both plants.
- Some types of [plasmids secrete](#) a toxin that kills the bacteria that carry them. Luckily for the bacteria, the plasmids also emit an antidote; and as long as a plasmid secretes both, the host bacterium survives. But if the plasmid dies, the antidote decays faster than the toxin, and the bacterium dies. This acts as an insurance policy for the plasmids, ensuring that bacteria don't evolve ways to kill them.

In the beginning of life on this planet, some 3.8 billion years ago, an organism's only job was to reproduce. That meant growing, and growing required energy. [Heat and light](#) were the obvious sources—photosynthesis appeared 3 billion years ago; chemosynthesis is at least a half a billion years older than that—but consuming the other living things floating around in the primordial ocean worked just as well. So life discovered predation.

We don't know what that [first animal predator](#) was, but it was likely a simple marine organism somewhere between 500 million and 550 million years ago. Initially, the only defense a species had against being eaten was to have so many individuals floating around the primordial seas that enough individuals were left to reproduce, so that the constant attrition didn't matter. But then life realized it might be able to avoid being eaten. So it evolved defenses. And predators evolved better ways to catch and eat.

Thus security was born, the planet's fourth oldest activity after eating, eliminating, and reproducing.

Okay, that's a pretty gross simplification, and it would get me booted out of any evolutionary biology class. When talking about evolution and natural selection, it's easy to say that organisms make explicit decisions about their genetic future. They don't. There's nothing purposeful or teleological about the evolutionary process, and I shouldn't anthropomorphize it. Species don't realize anything. They don't discover anything, either. They don't decide to evolve, or try genetic options. It's tempting to talk about evolution as if there's some outside intelligence directing it. We say "prehistoric lungfish first learned how to breathe air," or "monarch butterflies learned to store plant toxins in their bodies to make themselves taste bad to predators," but it doesn't work that way. Random mutation provides the material upon which natural selection acts. It is through this process that individuals of a species change subtly from their parents, effectively "trying out" new features. Those innovations that turn out to be beneficial—air breathing—give the individuals a competitive advantage and might potentially propagate through the species (there's still a lot of randomness in this process). Those that turn out to be detrimental—the overwhelming majority of them—kill or otherwise disadvantage the individual and die out.

By "beneficial," I mean something very specific: increasing an organism's ability to survive long enough to successfully pass its genes on to future generations. Or, to use Richard Dawkins's perspective from *The Selfish Gene*, genes that helped their host individuals—or other individuals with that gene—successfully reproduce tended to persist in higher numbers in populations.

If we were designing a life form, as we might do in a computer game, we would try to figure out what sort of security it needed and give it abilities accordingly. Real-world species don't have that luxury. Instead, they try new attributes randomly. So instead of an external designer optimizing a species' abilities based on its needs, evolution randomly walks through the solution space and stops at the first solution that works—even if just barely. Then it climbs upwards in the

fitness landscape until it reaches a local optimum. You get a lot of weird security that way.

You get teeth, claws, group dispersing behavior, feigning injury and playing dead, hunting in packs, defending in groups (flocking and schooling and living in herds), setting sentinels, digging burrows, flying, mimicry by both predators and prey, alarm calls, shells, intelligence, noxious odors, tool using (both offensive and defensive),¹ planning (again, both offensive and defensive), and a whole lot more.² And this is just in largish animals; we haven't even listed the security solutions insects have come up with. Or plants. Or microbes.

It has been convincingly argued that one of the reasons sexual reproduction evolved about 1.2 billion years ago was to **defend against** biological parasites. The argument is subtle. Basically, parasites reproduce so quickly that they overwhelm any individual host defense. The value of DNA recombination, which is what you get in sexual reproduction, is that it continuously rearranges a species' defenses so parasites can't get the upper hand. For this reason, a member of a species that reproduces sexually is much more likely to survive than a species that clones itself asexually—even though such a species will pass twice as many of its genes to its offspring as a sexually reproducing species would.

Life evolved two other methods of defending itself against parasites. One is to grow and divide quickly, something that both bacteria and just-fertilized mammalian embryos do. The other is to have an immune system. Evolutionarily, this is a relatively new development; it first appeared in **jawed fish** about 300 million years ago.³

A surprising number of evolutionary adaptations are related to security. Take vision, for example. Most animals are more adept at spotting movement than picking out details of stationary objects; it's called the **orienting response**.⁴ That's because things that move may be predators that attack, or prey that needs to be attacked. The human visual system is **particularly good** at spotting animals.⁵ The human ability, unique on the planet, to **throw things** long distances is another security adaptation. Related is what's called the **size-weight misperception**: the illusion that easier-to-throw rocks are perceived to be lighter than they are. It's related to our ability to choose good projectiles. **Similar stories** could be told about many human attributes.⁶

The predator/prey relationship isn't the only pressure that drives evolution. As soon as there was competition for resources, organisms had to develop security to defend their own resources and attack the resources of others. Whether it's plants competing with each other for access to the sun, predators fighting

over hunting territory, or animals competing for potential mates, organisms had to develop security against others of the same species. And again, evolution resulted in all sorts of *weird security*. And it works amazingly well.

Security on Earth went on more or less like this for 500 million years. It's a continual arms race. A rabbit that can run away at 30 miles per hour—in short bursts, of course—is at an evolutionary advantage when the weasels and stoats can only run 28 mph, but at an evolutionary disadvantage once predators can run 32 mph.




Figure 2: The Red Queen Effect in Action

It's different when the evolutionary advantage is against nature. A polar bear has thick fur because it's cold in the Arctic. And it's thick to a point, because the Arctic doesn't get colder in response to the polar bear's changes. But that same polar bear has fur that appears white so as to better sneak up on seals. But a better camouflaged polar bear means that only more wary seals survive and

reproduce, which means that the polar bears need to be even better at camouflage to eat, which means that the seals need to be more wary, and on and on and on up to some physical upper limit on camouflage and wariness.

This only-relative evolutionary arms race is known as the **Red Queen Effect**, after Lewis Carroll's race in *Through the Looking-Glass*: "It takes all the running you can do, to keep in the same place." Predators develop all sorts of new tricks to catch prey, and prey develop all sorts of new tricks to evade predators. The prey get more poisonous, so their predators get more poison-resistant, so the prey get even more poisonous. A species has to **continuously improve** just to survive, and any species that can't keep up—or bumps up against physiological or environmental constraints—becomes extinct.




Figure 3: The Red Queen Effect Feedback Loop

Along with becoming faster, more poisonous, and bitier, some organisms became smarter. At first, a little smarts went a long way. Intelligence allows individuals to adapt their behaviors, moment by moment, to suit their environment and circumstances. It allows them to remember the past and learn from experience. It lets them be individually adaptive. No one has a date, but vertebrates first appeared about 525 million years ago—and continued to improve on various branches of the tree of life: mammals (215 million years ago), birds (75 million years ago), primates (60 million years ago), the genus *Homo* (2.5 million years ago), and then humans (somewhere between 200,000 and 450,000 years ago, depending on whose evidence you believe). When it comes to security, as with so many things, humans changed everything.

Let's pause for a second. This isn't a book about animal intelligence, and I don't want to start an argument about which animals can be considered intelligent, or what about human intelligence is unique, or even how to define the

word “intelligence.” It’s definitely a fascinating subject, and we can learn a lot about our own intelligence by studying the intelligence of [other animals](#). Even my neat intelligence progression from the previous paragraph might be wrong: flatworms can be trained, and some cephalopods are surprisingly smart. But those topics aren’t really central to this book, so I’m going to elide them. For my purposes, it’s enough to say that there is a uniquely human intelligence.⁷

And humans take their intelligence seriously. The brain only represents 3% of total body mass, but uses 20% of the body’s total blood supply and 25% of its oxygen. And—unlike other primates, even—we’ll [supply our brains](#) with blood and oxygen at the expense of other body parts.

One of the things intelligence makes possible is [cultural evolution](#). Instead of needing to wait for genetic changes, humans are able to improve their survivability through the direct transmission of skills and ideas. These memes can be taught from generation to generation, with the more survivable ideas propagating and the bad ones dying out. Humans are not the only species that teaches its young, but humans have taken this to a new level.⁸ This caused a flowering of security ideas: deception and concealment; weapons, armor, and shields; coordinated attack and defense tactics; locks and their continuous improvement over the centuries; gunpowder, explosives, guns, cruise missiles, and everything else that goes “bang” or “boom”; paid security guards and soldiers and policemen; professional criminals; forensic databases of fingerprints, tire tracks, shoe prints, and DNA samples; and so on.

It’s not just intelligence that makes humans different. One of the things that’s unique about humans is the extent of our socialization. Yes, there are other social species: other primates, most mammals and some birds.⁹ But humans have taken sociality to a completely different level. And with that socialization came all sorts of new security considerations: concern for an ever-widening group of individuals, concern about potential deception and the need to detect it, concern about one’s own and others’ reputations, concern about rival groups of attackers and the corresponding need to develop groups of defenders, recognition of the need to take preemptive security measures against potential attacks, and after-the-fact responses to already-occurred attacks for the purpose of deterring others in the future.¹⁰

Some scientists believe that this increased socialization actually spurred the development of human intelligence.¹¹ Machiavellian Intelligence Theory—you might also see this called the Social Brain Hypothesis—holds that we evolved intelligence primarily in order to detect deception by other humans. Although the “Machiavellian” term came later, the idea first came from psychologist [Nicholas Humphrey](#). Humphrey observed that wild gorillas led a pretty simple

existence, with abundant and easily harvested food, few predators, and not much else to do but eat, sleep, and play. This was in contrast to gorillas in the laboratory, which demonstrated impressive powers of creative reasoning. So the obvious question is: what's the evolutionary advantage of being intelligent and clever if it's not required in order to survive in the wild? Humphrey proposed that the primary role of primate intelligence and creativity was to deal with the complexities of living with other primates. In other words, we evolved smarts not to outsmart the world, but to outsmart each other.

It's more than that. As we became more social, we needed to learn how to get along with each other: both cooperating with each other and ensuring everyone else cooperates, too. It involves understanding each other. Psychologist [Daniel Gilbert](#) describes it very well:

We are social mammals whose brains are highly specialized for thinking about others. Understanding what others are up to—what they know and want, what they are doing and planning—has been so crucial to the survival of our species that our brains have developed an obsession with all things human. We think about people and their intentions; talk about them; look for and remember them.

This makes evolutionary sense. Intelligence is a valuable survival trait when you have to deal with the threats from the natural world. But intelligence is an even more valuable survival trait when you have to deal with the threats from other intelligent individuals. An intelligent adversary is a different animal, so to speak, than an unintelligent adversary. An intelligent attacker is adaptive. An intelligent attacker can learn about its prey. An intelligent attacker can make long-term plans. An intelligent adversary can predict your defenses and incorporate them into his plans. If you're being attacked by an intelligent human, your most useful defense is to also be an intelligent human. Our ancestors grew smarter because those around them grew smarter, and the only way to keep up was to become even smarter.¹² It's a Red Queen Effect in action.

In primates, the frequency of deception is directly proportional to the size of a species' neocortex: the "thinking" part of the mammalian brain. That is, the bigger the brain, the greater the [capacity for deception](#). The human brain has a neocortex that's four times the size of its nearest evolutionary relative. Eighty percent of our brain is neocortex, compared to 50% in our nearest existing relative and 10% to 40% in [non-primate mammals](#).¹³

And as our neocortex grew, the complexity of our social interactions grew as well. Primatologist Robin Dunbar has studied primate group sizes. Dunbar

examined 38 different primate genera, and found that the volume of the **neocortex correlates** with the size of the troop. He established that the **mean human group size** is 150.¹⁴ This is the Dunbar number: the number of people with whom we can have explicit and personal encounters, whose history we can remember, and with whom we can experience some level of intimacy.¹⁵ Of course, it's an average. You personally might be able to keep track of more or fewer. This **number appears regularly** in human society: it's the estimated size of a Neolithic farming village; the size at which Hittite settlements split; and it's a basic unit in professional armies, from Roman times to the present day. It's the average size of people's Christmas card lists. It's a common department size in modern corporations.

So as our ancestors got smarter, their social groups got larger. Chimpanzees live in groups of approximately 60 individuals. *Australopithecus*—our ancestor from 4.5 million years ago—had an average group size of 70 individuals. When our first tool-using ancestors appeared 2 million years ago, the group size grew to 80. *Homo erectus* had a **mean group size** of 110, and Neanderthals 140. *Homo sapiens*: 150.

One hundred and fifty people is a lot to keep track of, especially if they're all clever, sneaky, duplicitous, and—as it turns out—murderous. There is a lot of evidence—both from the anthropological record and from ethnographic studies of contemporary primitive cultures—that humans are innately quite violent, and that intertribal warfare was endemic in primitive society. Several studies estimate that 15–25% of prehistoric males **died in warfare**.¹⁶

Economist **Paul Seabright** postulates that intelligence and murderousness are mutually reinforcing. The more murderous a species is, the greater the selective benefit of intelligence; smarter people are more likely to survive their human adversaries. And the smarter someone is, the more an adversary wants to kill him—and not just make him submit, as **other species** do.

Looking at the average weight of humans and extrapolating from other animals, humans **should primarily hunt** medium-sized rodents; indeed, early humans primarily hunted small game. And hunting small game is much more efficient for a bunch of reasons.¹⁷ Even so, **all primitive societies** hunt large game: antelopes, walrus, and so on. The theory is that although large-game hunting is less efficient, the skill set is the same as what's required for intertribal warfare. The groups that excelled at large-game hunting were more likely to survive the endemic warfare that existed in our evolutionary past. Group hunting also reinforced social bonds, which are a useful group survival trait.

A male killing another male of the same species—especially an unrelated male—eliminates a sexual rival. If you have fewer sexual rivals, you have more

of your own offspring. Natural selection favors murderousness. On the other hand, attempting to murder another individual of the same species is dangerous; you might get yourself killed in the process. This means fewer offspring, which implies a counterbalancing natural selection against murderousness.

It's another Red Queen Effect, this one involving murder. Evolutionary psychologist [David Buss](#) writes:

As the motivations to murder evolved in our minds, a set of counter-inclinations also developed. Killing is a risky business. It can be dangerous and inflict horrible costs on the victim. Because it's so bad to be dead, evolution has fashioned ruthless defenses to prevent being killed, including killing the killer. Potential victims are therefore quite dangerous themselves. In the evolutionary arms race, homicide victims have played a critical and unappreciated role—they pave the way for the evolution of anti-homicide defenses.

There is considerable debate about how violent we really are, with the majority opinion coming down on the “[quite violent](#)” side, especially among males from ages 16 to 24. On the other hand, [some argue](#) that human violence has declined over the millennia, primarily due to the changing circumstances that come with civilization. We do know it's been traditionally very hard to convince soldiers to [kill in war](#), and our experience with post-traumatic stress disorder shows that it has long-lasting ill effects. Our violence may be innate, but it depends a lot on context. We're comparable [with other primates](#).¹⁸

But if we are so naturally murderous, how did our prehistoric ancestors come to trust each other? We know they did, because if they hadn't, society would never have developed. People would never have gathered into groups that extended past immediate family, let alone into villages and towns and cities. Division of labor would have never evolved, because people couldn't trust others to do their parts. We would never have established trade with the strangers we occasionally encountered, let alone with companies based halfway across the planet. Friendships wouldn't exist. Societies based on either geography or interest would be impossible. Any sort of governmental structure: forget it. It doesn't matter how big your neocortex is or how abstractly you can reason: unless you can trust others, your species will forever remain stuck in the Stone Age.

The answer to that question will make use of the concepts presented in this chapter—the Red Queen Effect, the Dunbar number, our natural intelligence and murderousness—and it will make use of security. It turns out that trust in society isn't easy, and that we're still getting it wrong.



3

The Evolution of Cooperation

Two of the most successful species on the planet are humans and leafcutter ants of Brazil. Evolutionary biologist Edward O. Wilson has spent much of his career studying the ants, and argues that their success is due to [division of labor](#).¹ There are four different kinds of leafcutter workers: gardeners, defenders, foragers, and soldiers. Each type of ant is specialized to its task, and together the colony does much better than colonies of non-specialized ant species.

Humans specialize too, and—even better—we can adapt our specialization to the situation. A leafcutter ant is born to a particular role; we get to decide our specialization in both the long and short term, and change it if it's not working out for us.²

Division of labor is an exercise in trust. A gardener leafcutter ant has to trust that the forager leafcutter ants will bring leaf fragments back to the nest. I, specializing right now in book writing, have to trust that my publisher is going to print this book and bookstores are going to sell it. And that someone is going to grow food that I can buy with my royalty check. If I couldn't trust literally millions of nameless, faceless other people, I couldn't specialize.

Brazilian leafcutter ant colonies evolved trust and cooperation because they're all siblings. We had to evolve it the hard way.



We all employ both cooperating and defecting strategies. Most of the time our self-interest and group interest coincide, and we act in accordance with the group norm. Only sometimes do we act in some competing norm. It depends on circumstance, and it depends on who we are. Some of us are more cooperative, more honest, more altruistic, and fairer. And some of us are less so. There isn't

one dominant survival strategy that evolution has handed down to us; we have the flexibility to switch between different strategies.

One way to think of the relationship between society as a whole and its defectors is as a parasitic relationship. Take the human body as an example. Only 10% of the total number of cells in our human bodies are us—human cells with our particular genome. The other 90% are symbionts, genetically unrelated organisms.³ Our relationship with them ranges from mutualism (we both benefit) to commensalism (one benefits) to parasitism (one benefits and the other is harmed). The society of our bodies needs the cooperators to survive, and at the same time spends a lot of energy defending itself against the defectors.

Extending the analogy even further, our social systems are filled with parasites as well. Parasites steal stuff instead of buying it. They take more than their share in a communal situation. They overstay their welcome on their Aunt Faye's couch. They incur unsustainable debt, confident that bankruptcy laws—or some expensive lawyers—will enable them to bail out on their creditors when the going gets tough.

Parasites are all over the Internet. Crime is a huge business. Spammers are parasitic on e-mail. Grieffers in online games are parasitic on more conventional players. File sharers copy music instead of paying for it; they're parasitic on the music industry, getting the benefit of commercial music without giving back any money in return.

Excepting the smallest and simplest cases, every society has parasites living inside it. And there is an evolutionary advantage to being a parasite as long as there aren't too many of them and they aren't too good at it.

Being a parasite is a balancing act. Biological parasites do best if they don't immediately kill their hosts, but instead let them survive long enough for the parasites to spread to additional hosts. Ebola is too successful, so it fails as a species. The common cold does a much better job of spreading itself; it infects, and in the end kills, far more people by being much less "effective." Predators do best if they don't kill enough prey to wipe out the entire species. Spammers do better if they don't clog e-mail to the point where no one uses it anymore, and rogue banks are more profitable if they don't crash the entire economy. All parasites do better if they don't destroy whatever system they've latched themselves onto. Parasites thrive only if they don't thrive *too well*.

There's a clever model from game theory that illustrates this: the [Hawk-Dove game](#). It was invented by geneticists John Maynard Smith and George R. Price in 1971 to explain conflicts between animals of the same species. Like most game

theory models, it's pretty simplistic. But what it illuminates about the real world is profound.

The game works like this. Assume a population of individuals with differing survival strategies. Some cooperate and some defect. In the language of the game, the defectors are hawks. They're aggressive; they attack other individuals, and fight back if attacked. The cooperators are doves. They're pacific; they share with other doves, and retreat when attacked. You can think about this in terms of animals competing for food. When two doves meet, they cooperate and share food. When a hawk meets a dove, the hawk takes food from the dove. When two hawks meet, they fight and one of them randomly gets the food and the other has some probability of dying from injury.⁴

Set some initial parameters in the simulation: the value of sharing, the chance and severity of harm if two hawks fight each other, and so on. Program this model into a computer, set proportions for the initial population—50% hawks and 50% doves, for example—and let individuals interact with each other over multiple iterations.

What's interesting about this simulation is that neither strategy is guaranteed to dominate. Both hawks and doves can be successful, depending on the initial parameters. If the value of the food stolen is greater than the risk of death, the whole population becomes hawks. That is, if everyone is starving, people take what they can from each other without worrying about the consequences. Add a single dove, and it immediately starves. But as food gets less valuable (e.g., more plentiful) or fighting gets more dangerous, the population stabilizes into a mixture of hawks and doves. The more dangerous fighting is, the fewer hawks there will be. If food is reasonably plentiful and fighting reasonably dangerous, the population stabilizes into a mixture of mostly doves and fewer hawks. But unless you plug some really unrealistic numbers into the simulation—like starting out with a population entirely of doves—there will always be at least a few hawks in the mix.

This makes sense. Imagine a society made up entirely of cooperative doves. They share food whenever they meet each other, never stealing from one another. Now add a single hawk to the society. He does great. He steals food from all the doves, and since no one ever fights back, he has no risk of dying. It's the best survival strategy ever.

Now add a second hawk. The strategy is still pretty effective; if the population is large enough, the two hawks will never even meet. But as the number of hawks grows, the chance of two of them encountering each other—and one of

them dying in the resultant fight—increases. At some point, and the exact point depends on the parameters, there are enough other hawks around that being a hawk is as dangerous as being a dove has become. That's the stable percentage of hawks in the population.

Aside from making fighting more deadly or food less valuable, there are other ways to affect the percentages of hawks and doves. If doves can recognize hawks and refuse to engage, the population will have fewer hawks. If doves can survive hawk attacks without losing their food—by developing defenses, by learning to be sneaky—the population will have fewer hawks. If there is a way for doves to punish hawks, the population will have fewer hawks. If there is a way for doves to do even better if they work together, the population will have fewer hawks. If hawks can gang up on doves profitably, the population will have more hawks. In general, we get fewer hawks if we increase the benefits of being a dove and/or raise the costs of being a hawk, and we get more hawks if we do the reverse. All of this makes intuitive sense, and shouldn't come as a surprise.

And while a population consisting entirely of doves is stable, you can only get there if you start the game out that way. And if you assume that individuals in the game can think strategically and change their strategies as people can—doves can become hawks, and hawks can become doves—then an all-dove population is no longer stable. A physicist would describe an all-dove population as an unstable equilibrium. Given how easily a dove can become a hawk, it's very unstable. There will always be at least a minority of hawks.

The Hawk-Dove game is a model, and not intended to explain how cooperation evolved. However, several lessons can be learned by extrapolating the Hawk-Dove game into the real world. Any society will have a mix of people who cooperate and share, and people who defect and steal. But as the penalty, or cost, for attempting to steal, and failing, increases—it could be dying, it could be being jailed, it could be something else—there will be fewer defectors. Similarly, as the benefit of stealing increases—either in the value of what the thief gets, or in the probability he'll succeed in stealing—there will be more thieves.

In the real world, there are gradations of hawkishness. One person might murder someone to take his money; another might rob a person but let him live. A third might just shortchange him in some business transaction, or take an unfair share at the family dinner. Those are all hawkish behaviors, but they're not the same. Also, no one is 100% hawk or 100% dove; they're individual mixtures, depending on circumstance.⁵

If the benefit of being a hawk is greater than the risk of being a hawk, then hawks become the dominant strategy. Doves can't survive, and everyone

becomes a hawk. That’s anarchy: Hobbes’s “war of all against all.” In human terms, society falls apart. If we want to maintain a society based on cooperation, we need to ensure that the rate of defection stays small enough to allow society to remain cohesive.




Figure 4: Metaphorical Knobs to Control a Hawk-Dove Game

You can think of these parameters as knobs that control the rate of defection. We might not think of it in those terms, but it’s what we do all the time in the real world. Want fewer burglars? Increase the prison term for burglary, put more policemen on the street, or subsidize burglar alarms. Willing to live with more burglars? Understaff police departments, make it easier for burglars to fence stolen merchandise, or convince people to keep more cash at home.⁶ These are all societal pressures. So are increasing or decreasing social inequality, and teaching respect for other people’s property in school.

In our world, the costs and benefits of being a defector vary over time. As we develop new security technologies, and as the defectors develop new ways around them, society stabilizes with a different scope of defection. Similarly, as we develop new systems—Internet banking, for example—and defectors develop new ways to attack them, society stabilizes with a still different scope of defection. If the police force gets better at arresting speeders, there will be fewer of them. If someone invents a radar detector or if cars handle better at higher speeds, there will be more speeders.⁷

We’ll talk about this more in later chapters. The important point for right now is that no matter how hard we make life for the hawks among us—shunning them, removing them from society completely, making it less likely they will profit from their aggressive tactics—we will never be able to get the hawk percentage down to zero. Yes, we can make it very unprofitable to be a hawk, but if the percentage drops too low, being a hawk will become a more advantageous strategy. And because we humans are intelligent and adaptable, someone will figure that out and switch strategies.

Defectors are endemic to all complex systems. This is one of the dominant paradigms of life. We need to recognize that all of our complex human systems, whether they are millennia-old social systems or modern socio-technical systems, will always have parasites. There will always be a group of people who will try to take without giving back. The best we can hope for is to do what our bodies do, and what every natural ecosystem does: keep those parasites down to a tolerable level.

It's not even clear that natural selection favors the society with a minimum of hawk-like behavior. Hawks have value to society. In fact, if societies are in conflict with each other, it is evolutionarily advantageous to have some aggressive individuals. When war breaks out, the society with more hawks is likely to prevail. Again, think back to the primitive world in which we evolved. If you assume, as many anthropologists do, that tribal warfare was endemic among human societies, then having a substantial percentage of hawks around was valuable. Yes, they took advantage of the doves in peacetime, but they ensured the survival of those doves in wartime. Of course, we're now stuck with too many hawks because of the evolutionary pressures of 100,000 years ago.



I'm about to lump a lot of human traits together: cooperation, altruism, kindness, trustworthiness, and fairness. They're different, but all prosocial behaviors—behaviors intended to help others—and they're the glue that holds human society together. While psychologists put fine distinctions on them, considering them as facets of a whole is more useful for our purposes. They are all precursors of trust, and what allowed us to take the concept of specialization to a level unprecedented on our planet.

Figuring out how these traits evolved is an open question. Sure, they're great for our species as a whole, but that doesn't affect evolution. What matters for evolution is whether a particular characteristic helps the reproductive success of individuals with that characteristic. Kindness might be useful for society, but if it didn't result in kind people reproducing more successfully than unkind people, it would be bred out of the species pretty quickly.

There is an obvious evolutionary advantage in trusting kin: people with whom you share genetic material. If you have a gene, then your close relatives are likely to have that same gene. A gene that, on balance, makes it more likely for you to help your close relatives pass their genes on to future generations

will also be more likely to be passed on to future generations—assuming, of course, that the help it provides outweighs the cost to provide it. For example, if a lioness is genetically predisposed to suckle her sister's offspring, there's a good chance that her nieces and nephews share the genes responsible for that behavior, and will pass them on to their own offspring.

The natural world is filled with examples of animals trusting, helping, and behaving altruistically with each other. Not just ants: many insects defend their nests or hives with their lives. Some animals who live in groups and fear predation—prairie dogs, ground squirrels, some monkeys, assorted herd animals, and many birds—alert the group with an alarm call if they spot a predator. Other animals hunt in groups. Most of these examples turn out to be kin helping kin.⁸

Extending this tendency towards non-kin is much more difficult.⁹ Archaeologists have a four-stage model of the human process. Stage one happened 6 million years ago, when empathy and a motivation to help others developed in a common ancestor of humans and chimpanzees. Stage two began 1.8 million years ago; compassion can be seen in both short-term caring for sick individuals and special treatment for the dead. Stage three is much more recent; around 500,000 or 400,000 years ago, humans became dependent on group hunting, and started exhibiting long-term care for the injured and the infirm. Stage four occurred in modern humans starting 120,000 years ago, when [compassion extended](#) to strangers, animals, and sometimes even objects: religious objects, antiques, family heirlooms, etc. It probably didn't extend much past groups bigger than the Dunbar number of 150 until the [invention of agriculture](#), about 10,000 years ago—I guess that's a fifth stage.

Still, that doesn't tell us how or why it eventually did.

There are two basic types of non-kin cooperation. The first is mutualism.¹⁰ In some species, unrelated individuals cooperate because together they can perform tasks they couldn't do by themselves. A pack might hunt together because it can kill larger prey than the members could individually. [Unrelated elephants](#) help each other move objects they could not move alone.

Within a species, there's a tendency for individuals to cooperate by limiting their behavior. In many species, males fight each other for the prize of mating with a female. Primates fight to determine who is in charge of the tribe. In my house, the two cats fight to determine who gets to sit in the sunny chair. All these fights are serious, but tend to be non-injurious and are governed by ritual: roaring contests in red deer, claw-waving in male fiddler crabs, shell-rapping in hermit crabs. This is because these [ritualized battles](#) are often more about

getting information about the other individual than actually fighting, and a non-lethal battle is often a more survivable strategy. The Hawk-Dove game can model these types of conflicts: if the risk of being a hawk is great enough, it makes evolutionary sense to be a dove even if your opponent is a hawk, because it's more survivable to retreat than to fight.¹¹

So maybe we became smart enough to realize that cooperation usually beat defection as a survival skill, and modified our behavior accordingly. Those who could make that trade-off were more likely to pass their genes on to the next generation. This cooperation extended slowly outwards, from the immediate family group to more distant relatives to kith to familiar strangers—and over time, to unfamiliar strangers. And that cooperation slowly turned into trust.

Intelligence alone doesn't explain our trust of non-kin, though. Raw intelligence makes people calculating, but not necessarily honest or compassionate.¹² The missing ingredient is called *reciprocal altruism*. This is the second basic type of non-kin cooperation, and means that we tend to treat people as we have been treated.

Reciprocal altruism isn't limited to humans. [Vampire bats](#) must ingest blood every 60 hours or they'll die. If a bat can't find its own meal, a non-kin bat will often regurgitate some of its undigested blood and feed it to the hungry bat, knowing that another bat will regurgitate food for it at some later time. Then, the bats pay attention. They have [large frontal lobes](#) in their brains that they use to remember which other bats have shared blood with them in the past. A bat is more likely to share blood with a bat that has shared blood with it previously. Similarly, animals such as dogs, cats, horses, and some birds remember who was nice to them.

Think about our ancestors and their relationship with others living in their community. Cheating is valuable to the individual in the short term. But a person living in that community had an additional incentive not to cheat: if he did, he squandered his chance at future cooperation with his victim, and risked his reputation with the community. If the benefits of future cooperation are great enough, it makes evolutionary sense for non-kin to [help each other](#) if they can be reasonably sure they will be repaid at a later date.

A reasonable question, then, is whether altruism in the purest sense of the word really exists, or if it's all based on some anticipated reward or punishment. Perhaps Mother Teresa wasn't really altruistic; she expected her reward in Heaven. Perhaps our instinct to protect our children isn't really altruistic; it's because we expect them to care for us in our old age. We don't consider vampire

bats altruistic; they expect repayment at some future date. Even the mother who sacrifices her life for her child might just be ensuring that her genes survive.

If we simplify, the psychological theory of [transactional analysis](#) holds that people expect some sort of return—either emotional or material—from their apparent altruism and kindness. So we rescue a stranger from a burning building because we expect to survive and be praised, and we give money to charity because it makes us feel virtuous. You can argue that whenever we act in the group interest, it's because we know we're better off when we do.

There's even an alternate theory that explains altruistic behavior without any need for pure, selfless altruism. Biologist Amotz Zahavi's handicap principle explains costly "signals" within species. If you're an individual of above-average fitness, it makes evolutionary sense to spend some of that surplus on costly and hard-to-fake signals to advertise that fact to a [potential mate](#). This holds true for a peacock's tail and a stag's antlers, as well as for a human's apparently [altruistic acts](#). So the man who rescues a stranger from a burning building is advertising his kindness and physical prowess, and the woman who gives money to charity is advertising her wealth. We do know that agreeableness is a trait desired by others in a mate; [kind people](#) are more likely to reproduce.

This seems an irrelevant exercise, rather like debating whether or not there is such a thing as free will.¹³ [George Price](#), one of the inventors of the Hawk-Dove game, was unable to accept altruism's selfish basis, and spent much of his later life trying to demonstrate how wrong his mathematical model was. He gave his money away to strangers, let the homeless live in his house, and eventually committed suicide from depression. I think a more optimistic viewpoint is in order. People behave in ways that are altruistic, empathic, kind, trustworthy, fair, and cooperative. We do these things even though we don't have to. Yes, we have evolved into a species that finds these traits desirable. Yes, this is primarily reciprocal. Yes, we are also intelligent and calculating, but this is precisely the point. We have the ability to decide whether to be prosocial or not, and most of us, most of the time, decide positively. And we call these behaviors "altruism," "kindness," and "cooperation." We trust because others are trustworthy.

Humans seem to have evolved along these lines, overcoming the murderousness that accompanied our increasing intelligence. There is an enormous amount of laboratory research on altruism, fairness, cooperation, and trust. Experimenters have subjects play a variety of [bargaining games](#) where they divide a pot of money amongst themselves, with different outcomes depending on whether or not they act in the group interest or in self-interest. These have names like the

Ultimatum game¹⁴, the Dictator game¹⁵, the Trust game¹⁶, and the Public Goods game¹⁷, all with many different variants designed to tease out a particular aspect of human prosocial behavior.¹⁸ The general results seem to be that:

- People tend to be fair-minded.¹⁹ They routinely reduce their own rewards in order to be fair to other players.
- People tend to want to punish unfairness, even at their own personal expense.²⁰ We have a sense of justice and responsibility, and we react negatively to those who act contrary to that sense. In many instances people also reduce their own reward in order to punish someone whom they perceive to be as acting unfairly.
- People tend to follow social or cultural norms with respect to these prosocial behaviors.²¹ Definitions of fairness are cultural. People are more likely to be altruistic in a game that emphasizes altruism, and selfish in a game that emphasizes selfishness. Levels of trust and trustworthiness vary across cultures.
- People tend to be more trusting and altruistic with people they think they know and can identify with—even just a little bit—than with anonymous strangers.²²
- External factors matter a lot. In experiments, people were kinder after they [found a coin](#), traveled up an escalator (as opposed to traveling down), or watched a video of [flying through clouds](#) (as opposed to watching a video of driving on the ground).

Of course—and this is important to remember—these are typical results, and there is a wide variety of behavior among individual people.²³ This matches our experience in the world.

Neuroscience may also help explain altruism, most recently using [mirror neurons](#). These are neurons in our brain that fire both when we perform an action²⁴ and when we observe someone else performing the same action. First discovered in 1992, mirror neurons are theorized to be critical in imitation and learning, language acquisition, developing a theory of mind, empathy, and a variety of other prosocial behaviors.

Additionally, a large body of neuroscience research supports the notion that we are [altruistic innately](#), even if we receive no direct benefit, because at a deep level we want to be. Studies using functional magnetic resonance imaging (fMRI) show that the amygdala, the primitive part of the brain associated with

fear and anger, is involved in decisions about [fairness and justice](#). And it's probably not an unrelated side-effect that people who observe others acting either fairly or unfairly rate the fair people as significantly more agreeable, likeable, [and attractive](#) than the unfair people. We treat each other altruistically because it gives us pleasure to do so.

We not only innately trust, but we want to be trusted. A lot of this is intellectually calculated, but it goes deeper than that. Our need to be trusted is innate. There's even a biological feedback loop. Researchers have [found that oxytocin](#)—a hormone released during social bonding—naturally increases in a person who perceives that he is trusted by others. Similarly, artificially increasing [someone's oxytocin](#) level makes her more trusting.

The philosopher and economist [Adam Smith](#) expressed a similar sentiment 300 years ago:

How selfish so ever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it.

Of course, human trust isn't all-or-nothing. It's contextual, calibrated by our ability to calculate costs and benefits. A lot of our willingness to trust non-kin is calibrated by the society we live in. If we live in a polite society where trust is generally returned, we're at ease trusting first. If we live in a violent society where strangers are hostile and untrustworthy, we don't trust so easily and require further evidence that our trust will be reciprocated.

Our trust rules can be sloppy. We're more likely to trust people who are similar to us: look like us, dress like us, and speak the same language. In general, we're more likely to trust in familiar situations. We also generalize: if we have a good experience with people of a particular nationality or a particular profession, we are likely to trust others of the same type. And if we have a bad experience, we're likely to carry that mistrust to others of the same type.²⁵ These rules of thumb might not make logical sense in today's diverse world, but they seem to have been good ideas in our evolutionary past.

This is all good, but we have a chicken-and-egg problem. Until people start trusting non-kin, there is no evolutionary advantage to trusting non-kin. And until there's an evolutionary advantage to trusting non-kin, people won't be predisposed to trust non-kin. Just as a single hawk in a Hawk-Dove game can take advantage of everybody, a single dove in a Hawk-Dove game gets taken

advantage of by everybody. That is, the first trusting person who engages with a group of untrustworthy people isn't going to do very well.

It turns out that cooperative behavior can overcome these problems. Mathematical biologist Martin A. Nowak has explored the evolution of cooperation using mathematics, computer models, and experiments, and has found four different mechanisms by which altruistic behavior can spontaneously evolve in [non-kin groups](#):

- *Direct reciprocity*. Being altruistic towards you now is a good strategy because you'll be altruistic towards me later.
- *Indirect reciprocity*. Being altruistic towards you now is a good strategy because my reputation as an altruistic individual will increase, and someone else will be altruistic towards me later.²⁶
- *Network reciprocity*. Being altruistic towards you now is a good strategy because we are both in a group whose members are altruistic to each other, and being part of that group means that someone else will be altruistic towards me later.
- *Group selection*. Being altruistic towards you now is a good strategy because we're both part of a group whose members are altruistic to each other, and our group of altruists is more likely to survive than a group of non-altruists.²⁷

What methods work depend on how much it costs for one individual to help another, how beneficial the help is, and how likely it is that helpful individuals meet and recognize each other in the future. And, depending on details, there are several plausible [biological models](#) of how this sort of thing might have jump-started itself. Exactly how this evolved in humans is debated.²⁸ Philosopher [Patricia Churchland](#) suggests four coexistent characteristics of our pre-human ancestors that make all of Nowak's mechanisms likely: "loose hierarchy and related easygoing temperament, cooperative parenting extending to cooperating with the group, sexual selection, and lethal intergroup competition." The last one is especially interesting; [our murderousness](#) helped make us cooperative.

What's likely is that all six mechanisms—Nowak's four, kin selection, and Zahavi's handicap principle—were working at the same time. Also that there was a strong positive-feedback loop, as we became smarter and more social. Each individual mechanism contributes a bit towards the evolution of cooperation, which makes resultant individuals better able to pass their genes on to the next generation, which selects for a little more contribution from each mechanism,

which makes resultant individuals even better able to pass their genes on, and so on. [And these processes](#), especially group selection, work on both the genetic and cultural levels.

We became trustworthy, well...most of the time. We trusted others, well...most of the time. And, as we'll see, we used security to fill in the gaps where otherwise trust would fail. In a way, humans domesticated themselves.²⁹



4

A Social History of Trust

Trust is rare on this planet. Here's primatologist [Robert Sapolsky](#):

When baboons hunt together they'd love to get as much meat as possible, but they're not very good at it. The baboon is a much more successful hunter when he hunts by himself than when he hunts in a group because they screw up every time they're in a group. Say three of them are running as fast as possible after a gazelle, and they're gaining on it, and they're deadly. But something goes on in one of their minds—I'm anthropomorphizing here—and he says to himself, "What am I doing here? I have no idea whatsoever, but I'm running as fast as possible, and this guy is running as fast as possible right behind me, and we had one hell of a fight about three months ago. I don't quite know why we're running so fast right now, but I'd better just stop and slash him in the face before he gets me." The baboon suddenly stops and turns around, and they go rolling over each other like Keystone cops and the gazelle is long gone because the baboons just became disinhibited. They get crazed around each other at every juncture.

We're not like that. Not only do we cooperate with people we know, we cooperate with people we've never even met. We treat strangers fairly, altruistically sometimes. We put group interest ahead of our own selfishness. More importantly, we control other people's selfish behaviors.

We do this through a combination of our own prosocial impulses and the societal pressures that keep us all in line. This is what allowed for the hunter-gatherer societies of prehistory, the civilization of history, and today's globalization.

But while our cultures evolved, our brains did not. As different as our lives are from those of the primitive hunter-gatherers who lived in Africa 100,000 years ago, genetically we have barely changed at all.¹ There simply hasn't been enough time. As [Matt Ridley](#) writes in *The Red Queen*:

Inside my skull is a brain that was designed to exploit the conditions of an African savanna between 3 million and 100,000 years ago. When my ancestors moved into Europe (I am a white European by descent) about 100,000 years ago, they quickly evolved a set of physiological features to suit the sunless climate of northern latitudes: pale skin to prevent rickets, male beards, and a circulation relatively resistant to frostbite. But little else changed. Skull size, body proportions, and teeth are all much the same as they are in a San tribesman from southern Africa. And there is little reason to believe that the grey matter inside the skull changed much, either. For a start, 100,000 years is only three thousand generations, a mere eye blink in evolution, equivalent to a day and a half in the life of bacteria. Moreover, until very recently the life of a European was essentially the same as that of an African. Both hunted meat and gathered plants. Both lived in social groups. Both had children dependent on their parents until their late teens. Both passed wisdom down with complex languages. Such evolutionary novelties as agriculture, metal, and writing arrived less than three hundred generations ago, far too recently to have left much imprint on my mind.

It is this disconnect between the speed of cultural evolution and memes—intragenerationally fast—and the speed of genetic evolution—glacially slow, literally—that make trust and security hard. We've evolved for the trust problem endemic to living in small family groups in the East African highlands in 100,000 BC. It's 21st century New York City that gives us problems.²

Our brains are sufficiently neuroplastic that we can adapt to today's world, but vestiges of our evolutionary past remain. These cognitive biases affect how we respond to fear, how we perceive risks (there's a whole list of them in Chapter 15), and how we weigh short-term versus long-term costs and benefits. That last one is particularly relevant to decisions about cooperation and defection. Psychological studies show that we have what's called a [hyperbolic discounting](#) rate: we often prefer lower payoffs sooner to higher payoffs later. As we saw in the previous chapter, decisions to cooperate often involve putting our long-term interests ahead of our short-term interests. In some ways, this is unnatural for us.



As we saw in the previous chapter, any system of cooperators also includes some defectors. So as we as a species became more cooperative, we evolved strategies for dealing with defectors.

Making this happen isn't free. We have evolved a variety of different mechanisms to induce cooperation, the societal pressures I'll discuss in Chapters 6 through 10. [Francis Fukuyama](#) wrote: "Widespread distrust in society...imposes a kind of tax on all forms of economic activity, a tax that high-trust societies do not have to pay." It's a tax on the honest. It's a tax imposed on ourselves by ourselves, because, human nature being what it is, too many of us would otherwise become hawks and take advantage of the rest of us. And it's an expensive tax.³

[James Madison](#) famously wrote: "If men were angels, no government would be necessary." If men were angels, no security would be necessary. Door locks, razor wire, tall fences, and burglar alarms wouldn't be necessary. Angels never go where they're not supposed to go. Police forces wouldn't be necessary. Armies? Countries of angels would be able to resolve their differences peacefully, and military expenses would be unnecessary.

Currency, that paper stuff that's deliberately made hard to counterfeit, wouldn't be necessary, as people could just write down how much money they had.⁴ Angels never cheat, so nothing more would be required. Every security measure that isn't designed to be effective against accident, animals, forgetfulness, or legitimate differences between scrupulously honest angels could be dispensed with.

We wouldn't need police, judges, courtrooms, jails, and probation officers. Disputes would still need resolving, but we could get rid of everything associated with investigating, prosecuting, and punishing crime. Fraud detection would be unnecessary: the parts of our welfare and healthcare system that make sure people fairly benefit from those services and don't abuse them; and all of the anti-shoplifting systems in retail stores.

Entire industries would be unnecessary, like private security guards, security cameras, locksmithing, burglar alarms, automobile anti-theft, computer security, corporate security, airport security, and so on. And those are just the obvious ones; financial auditing, document authentication, and many other things would also be unnecessary.

Not being angels is expensive.

We don't pay a lot of these costs directly. The vast majority of them are hidden in the price of the things we buy. Groceries cost more because some people shoplift. Plane tickets cost more because some people try to blow planes up. Banks pay out lower interest rates because of fraud. Everything we do or buy costs more because some sort of security is required to deliver it.

Even greater are the non-monetary costs: less autonomy, reduced freedom, ceding of authority, lost privacy, and so on. These trade-offs are subjective, of

course, and some people value them more than others. But it's these costs that lead to social collapse if they get too high.

Security isn't just a tax on the honest, it's a very expensive tax on the honest. If all men were angels, just think of the savings!



It wasn't always like this. Security used to be cheap. Societal pressures used to be an incidental cost of society itself. Many of our societal pressures evolved far back in human prehistory, well before we had any societies larger than extended family groups. We touched on these mechanisms in the previous chapter: both the moral mechanisms in our brains that internally regulate our behavior, and the reputational mechanisms we all use to regulate each other's behavior.

Morals and reputation comprise our prehistoric toolbox of societal pressures. They are informal, and operate at both conscious and subconscious levels in our brains: I refer to the pair of them, unenhanced by technology, as *social pressures*. They evolved together, and as such are closely related and intertwined in our brains and societies. From a biological or behaviorist perspective, there's a reasonable argument that my distinction between moral and reputational systems is both arbitrary and illusory, and that differentiating the two doesn't make much sense. But from our perspective of inducing trust, they are very different.

Despite the prevalence of war, violence, and general deceptiveness throughout human history—and the enormous amount of damage wrought by defectors—these ancient moral and reputational systems have worked amazingly well. Most of us try not to treat others unfairly, both because it makes us feel bad and because we know they'll treat us badly in return. Most of us don't steal, both because we feel guilty when we do and because there are consequences if we get caught. Most of us are trustworthy towards strangers—within the realistic constraints of the society we live in—because we recognize it's in our long-term interest. And we trust strangers because we recognize it is in their interest to act trustworthily. We don't want a reputation as an untrustworthy, or an untrusting, person.

Here's an example from early human prehistory: two opposing tendencies that would cause society to fall apart if individuals couldn't trust each other. On one hand, we formed pair bonds for the purpose of child-rearing. On the other hand, we had a primarily gender-based division of labor that forced men and women to separate as they went about their different hunting and gathering tasks. This meant that primitive humans needed to trust that everyone honored

the integrity of these pair bonds, since individuals often couldn't be around to police them directly. The difficulty in resolving those opposing tendencies is known as Deacon's Paradox.⁵

No, anthropologists don't have unrealistic views on the sanctity of marriage. They know that illicit affairs go on all the time.⁶ But they also realize that such indiscretions occur with much less frequency than they would if mating weren't largely based on pair-bonding.⁷ Most people are honest most of the time, and most pair bonds are respected most of the time. Deacon singled out one particular human capability—the ability to form symbolic contracts—as the particular mechanism that polices sexual fidelity. This isn't just about two people deciding to cohabit, share food, and produce and raise offspring. It's about two people making a public declaration of commitment in marriage ceremonies, and enlisting other members of the community to simultaneously recognize and promote the stability of their pair bond. Because everyone has a stake in supporting sexual fidelity within the community, everyone keeps an eye on everyone else and punishes illicit matings.

This is an example of a social pressure. It's informal and ad hoc, but it protects society as a whole against the potentially destabilizing individual actions of its members. It protects society from defectors, not by making them disappear, but by keeping their successes down to a manageable rate. Without it, primitive humans wouldn't have trusted each other enough to establish gender-based division of labor and, consequently, could never have coalesced into communities of both kith and kin.

Other examples include being praised for good behavior, being gossiped about and snubbed socially for bad behavior, being shamed, shunned, killed, and—this is much the same as being killed—ostracized and cast out of the group.

I'm omitting a lot of detail, and there are all sorts of open research questions. How did these various social pressures evolve? When did they first appear, and how did their emergence separate us from the other primates—and other protohumans?⁸ How did trust affect intelligence, and how did intelligence affect trust? For our purposes, it's enough to say that they evolved to overcome our increased deceptiveness and murderousness.

In a primitive society, these social pressures are good enough. When you're living in a small community, and objects are few and hard to make, it's pretty easy to deal with the problem of theft. If Alice loses a bowl at the same time Bob shows up with an identical bowl, everyone in the community knows that Bob stole it from Alice and can then punish Bob. The problem is that these mechanisms don't scale. As communities grow larger, as they get more complex,

as social ties weaken and anonymity proliferates, this system of theft prevention—morals keeping most people honest, and informal detection, followed by punishment, leading to deterrence to keep the rest honest—starts to fail.

Remember the Dunbar number? Actually, Dunbar proposed several natural human group sizes that increase by a factor of approximately three: 5, 15, 50, 150, 500, and 1,500—although, really, the numbers aren't as precise as all that. The layers relate to both the intensity and intimacy of relationships, and the frequency of contact.

The smallest, three to five, is a *clique*: the number of people from whom you would seek help in times of severe emotional distress. The 12-to-20 person group is the *sympathy group*: people with whom you have a particularly close relationship. After that, 30 to 50 is the typical size of hunter-gatherer overnight camps, generally drawn from a single pool of 150 people. The 500-person group is the *megaband*, and the 1,500-person group is the *tribe*; both terms are common in ethnographic literature. Fifteen hundred is roughly the number of faces we can recognize, and the typical size of a hunter-gatherer society.⁹

Evolutionary psychologists are still debating Dunbar's findings, and whether there are as [many distinct levels](#) as Dunbar postulates. Regardless of how this all shakes out, for our purposes it's enough to notice that as we move from smaller group sizes to larger ones, our informal social pressures [begin to fail](#), necessitating the development of more formal ones. A family doesn't need formal rules for sharing food, but a larger group in a communal dining hall will. Small communities don't need birth registration procedures, marriages certified by an authority, laws of inheritance, or rules governing real-estate transfer; larger communities do. Small companies don't need employee name badges, because everyone already knows everyone else; larger companies need them and many other rules besides.

To put it another way, our trust needs are a function of scale. As the number of people we dealt with increased, we no longer knew them well enough to be able to trust their intentions, so our prehistoric trust toolbox started failing. As we developed agriculture and needed to trust more people over increased distance—physical distance, temporal distance, emotional distance—we needed additional societal pressures to elicit trustworthiness at this new scale. As the number of those interactions increased, and as the potential damage the group could do to the individual increased, we needed even more. If humans were incapable of developing these more formal societal pressures, societies either would have stopped growing or would have disintegrated entirely.

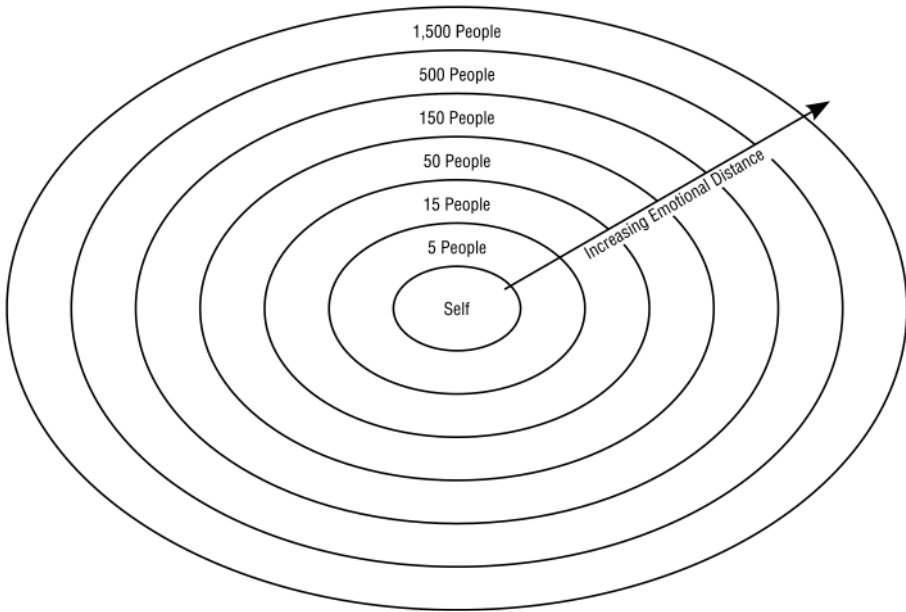


Figure 5: Dunbar Numbers

Agriculture required protecting resources, through violence if necessary. Luckily, two things happened. We invented institutions—government, basically—and we developed technology. Both of them allowed for human societies to grow larger without tearing themselves apart.

Institutions formalized reputational pressure. With government came laws, enforcement, and formal punishment. I'm not implying that the original purpose of government was facilitating trust, only that part of what these formal institutions did was codify the existing societal norms. This codification is a trust mechanism.

History has forgotten all of these early institutions. Some were undoubtedly civil. Some were religious.¹⁰ No one knows the details of how our ancestors made the transition from an extended family to a tribe of several extended families, because they happened thousands of years before anyone got around to inventing writing. Certainly there's overlap between formal reputational and early institutional pressures. It's enough to say that we made the transition, and that we augmented moral and reputational pressures with institutional pressures.¹¹

This was a critical development, one that gives us the ability to trust people's actions, even if we can't trust their intentions. We reinforced our informal

recognition of pair bonds with formal marriage through religious and civil institutions. We added laws about theft, and prescribed specific punishments. A lot of this, at least initially, is formalizing reputational mechanisms so that they scale to larger groups. But something important happens in the transition: institutional pressures require institutions to implement them. Society has to designate a subset of individuals to enforce the laws. Think of elders, guards, police forces, and judicial systems; priests also take on this role.

Institutions also enabled the formation of groups of groups, and subgroups within groups. So small tribes could become part of larger political groupings. Individual churches could become part of a larger religious organization. Companies could have divisions. Government organizations could have departments. Empires could form, and remain stable over many generations. Institutions scale in a way that morals and reputation do not, and this has allowed societies to grow at a rate never before seen on the planet.

The second force that allowed society to scale was technology—both technology in general and security technology specifically. Security systems are the final way we induce trust. Early security mechanisms included building earthen berms, wearing animal skins as camouflage, and digging pit traps. In one sense, security isn't anything new; we learned in Chapter 2 that it's been around almost as long as life itself. That primitive sort of security is what you might call natural defenses, focused on the individual. But when societies got together and realized they could, as a group, implement security systems, security became a form of societal pressure. In a sense, security technologies allow natural defenses to scale to protect against intra-group defection.

Technology also allowed informal social pressures to scale and become what I call societal pressures. Morals could be written down and passed from generation to generation. Reputation could similarly be recorded, and transferred from one person to another. This sort of thing depends a lot on technology: from the Bible and letters of introduction, to online debates on morality, to entries on the Angie's List database.

A good way to think about it is that both institutional pressure and security systems allow us to overcome the limitations of the Dunbar numbers by enabling people to trust systems instead of people. Instead of having to trust individual merchants, people can trust the laws that regulate merchants. Instead of having to evaluate the trustworthiness of individual borrowers, banks and other lending institutions can trust the credit rating system. Instead of trusting that people won't try to rob my house, I can trust the locks on my doors and—if I want to turn it on—my burglar alarm.

Technology changes the efficacy of societal pressures in another way as well. As soon as the different systems of societal pressure themselves need to be secured, it becomes possible for a defector to attack those security systems directly. Once there's a technologically enhanced system of reputational pressure, that system needs to be protected with security technology. So we need signet rings and wax seals to secure letters of introduction, and computer security measures that keep the Angie's List database from being hacked. Similarly, once forensic measures exist to help enforce laws, those forensic measures can be directly targeted. So burglars wear gloves to keep from leaving fingerprints, and file down VINs to prevent stolen cars from being tracked.

There's a bigger change that results from society's increased scale. As society moved from informal social pressures to more formal societal pressures—whether institutional pressures and security systems, or technologically enhanced moral and reputational pressures—the nature of trust changed. Recall our two definitions of trust from Chapter 1: trust of intentions and trust of actions. In smaller societies, we are usually concerned with trust in the first definition. We're intimately familiar with the people we're interacting with, and have a good idea about their intentions. The social pressures induce cooperation in specific instances, but are also concerned with their overall intentions. As society grows and social ties weaken, we lose this intimacy and become more concerned with trust in the second definition. We don't know who we're interacting with, and have no idea about their intentions, so we concern ourselves with their actions. Societal pressures become more about inducing specific actions: compliance.

Compliance isn't as good as actual trustworthiness, but it's good enough. Both elicit trust.

Also through history, technology allowed specialization, which encouraged larger group sizes. For example, a single farmer could grow enough to sustain more people, permitting even greater specialization. In this and other ways, general technological innovations enabled society to grow even larger and more complex. Dunbar's numbers remain constant, but postal services, telegraph, radio, telephone, television, and now the Internet have allowed us to interact with more people than ever before. Travel has grown increasingly fast and increasingly long distance over the millennia, and has allowed us to meet more people face-to-face. Countries have gotten larger, and there are multinational quasi-governmental organizations. Governments have grown more sophisticated. Organizations have grown larger and more geographically dispersed. Businesses have gotten larger; now there are multinational corporations employing

hundreds of thousands of people controlling assets across several continents. If your Facebook account has substantially more than 150 friends, you probably only have a superficial connection with many of them.¹²

Technology also increases societal complexity. Automation and mass distribution mean one person can affect more people and more of the planet. Long-distance communication, transport, and travel mean that people can affect each other even if they're far away. Governments have gotten larger, both in terms of geographical area and level of complexity. Computer and networking technology mean that things happen faster, and information that might once have been restricted to specialists can be made available to a worldwide audience. These further increases in scale have a major effect on societal pressures, as we'll see in Chapter 16.

None of this is to say that societal pressures result in a fair, equitable, or moral society. The [Code of Hammurabi](#) from 1700 BC, the first written code of laws in human history, contains some laws that would be considered barbaric today: if a son strikes his father, his hands shall be hewn off; if anyone commits a robbery and is caught, he shall be put to death; if anyone brings an accusation of any crime before the elders, and does not prove what he has charged, he shall, if a capital offense is charged, be put to death. And world history is filled with governments that have oppressed, persecuted, and even massacred their own people. There's nothing to guarantee that the majority actually approves of these laws; societal interest and societal norms might be dictated by an authoritarian ruler. The only thing societal pressures guarantee is that, in the short run at least, society doesn't fall apart.

Furthermore, societal pressures protect a society from change: bad, good, and indeterminate. Cooperators are, by definition, those who follow the group norm. They put the group interest ahead of any competing interests, and by doing so, make it harder for the group norm to change. If the group norm is unsustainable, this can fatally harm society in the long run.

Remember that society as a whole isn't the only group we're concerned about here. Societal pressures can be found in any group situation. They're how a group of friends protect themselves from greedy people at communal dinners. They enable criminal organizations to protect themselves from loose cannons and potential turncoats within their own ranks. And they're how a military protects itself from deserters and insubordinates, and how corporations protect themselves from embezzling employees.



5

Societal Dilemmas

Here are some questions related to trust:

- During a natural disaster, should I steal a big-screen TV? What about food to feed my family?
- As a kamikaze pilot, should I refuse to kill myself? What if I'm just a foot soldier being ordered to attack a heavily armed enemy bunker?
- As a company employee, should I work hard or slack off? What if the people around me are slacking off and getting away with it? What if my job is critical and, by slacking off, I harm everyone else's year-end bonuses?

There's a *risk trade-off* at the heart of every one of these questions. When deciding whether to cooperate and follow the group norm, or defect and follow some competing norm, an individual has to weigh the costs and benefits of each option. I'm going to use a construct I call a societal dilemma to capture the tension between group interest and a competing interest.

What makes something a societal dilemma, and not just an individual's free choice to do whatever he wants and risk the consequences, is that there are societal repercussions of the trade-off. Society as a whole cares about the dilemma, because if enough people defect, something extreme happens. It might be bad, like widespread famine, or it might be good, like civil rights. But since a societal dilemma is from the point of view of societal norms, by definition it's in society's collective best interest to ensure widespread cooperation.

Let's start with the smallest society possible: two people. Another model from game theory works here. It's called the Prisoner's Dilemma.¹

Alice and Bob are partners in crime, and they've both been arrested for burglary.² The police don't have enough evidence to convict either of them, so they bring them into separate interrogation rooms and offer each one a deal. "If you betray your accomplice and agree to testify against her," the policeman says, "I'll let you go free and your partner will get ten years in jail. If you both betray each other, we don't need your testimony, and you'll each get six years in jail. But if you cooperate with your partner and both refuse to say anything, I can only convict you on a minor charge—one year in jail."³

Neither Alice nor Bob is fully in charge of his or her own destiny, since the outcome for each depends on the other's decision. Neither has any way of knowing, or influencing, the other's decision; and they don't trust each other.

Imagine Alice evaluating her two options: "If Bob stays silent," she thinks, "then it would be in my best interest to testify against him. It's a choice between no jail time versus one year in jail, and that's an easy choice. Similarly, if Bob rats on me, it's also in my interest to testify. That's a choice between six years in jail versus ten years in jail. Because I have no control over Bob's decision, testifying gives me the better outcome, regardless of what he chooses to do. It's obviously in my best interest to betray Bob: to confess and agree to testify against him." That's what she decides to do.

Bob, in a holding cell down the hall, is evaluating the same options. He goes through the same reasoning—he doesn't care about Alice any more than she cares about him—and arrives at the same conclusion.

So both Alice and Bob confess. The police no longer need either one to testify against the other, and each spends six years in jail. But here's the rub: if they had both remained silent, each would have spent only one year in jail.

Societal Dilemma: Prisoners confessing.	
Society: A group of two prisoners.	
Group interest: Minimize total jail time for all involved.	Competing interest: Minimize individual jail time.
Group norm: To cooperate with the other prisoner and remain silent.	Corresponding defection: Testify against the other.

The Prisoner's Dilemma encapsulates the conflict between group interest and self-interest. As a pair, Alice and Bob are best off if they both remain silent and spend only one year in jail. But by each following his or her own self-interest, they both end up with worse outcomes individually.

The only way they can end up with the best outcome—one year in jail, as opposed to six, or ten—is by acting in their group interest. Of course, that only makes sense if each can trust the other to do the same. But Alice and Bob can't.

Borrowing a term from economics, the other prisoner's jail time is an *externality*. That is, it's an effect of a decision not borne by the decision maker. To Alice, Bob's jail time is an externality. And to Bob, Alice's jail time is an externality.

I like the prisoner story because it's a reminder that cooperation doesn't imply anything moral; it just means going along with the group norm. Similarly, defection doesn't necessarily imply anything immoral; it just means putting some competing interest ahead of the group interest.

Basic commerce is another type of Prisoner's Dilemma, although you might not have thought about it that way before. Cognitive scientist [Douglas Hofstadter](#) liked this story better than prisoners, confessions, and jail time.

Two people meet and exchange closed bags, with the understanding that one of them contains money, and the other contains a purchase. Either player can choose to honor the deal by putting into his or her bag what he or she agreed, or he or she can defect by handing over an empty bag.

It's easy to see one trust mechanism that keeps merchants from cheating: their reputations as merchants. It's also easy to see a measure that keeps customers from cheating: they're likely to be arrested or at least barred from the store. These are examples of societal pressures, and they'll return in the next chapters.

This example illustrates something else that's important: societal dilemmas are not always symmetrical. The merchant and a customer have different roles, and different options for cooperating and defecting. They also have different incentives to defect, and different competing norms.

Here's a societal dilemma involving two companies. They are selling identical products at identical prices, with identical customer service and everything else. Sunoco and Amoco gasoline, perhaps. They are the only companies selling those products, and there's a fixed number of customers for them to divvy up. The only way they can increase their market share is by advertising, and any increase in market share comes at the expense of the other company's. For simplicity's sake, assume that each can spend either a fixed amount on advertising or none at all; there isn't a variable amount of advertising spending that they can do. Also assume that if one advertises and the other does not, the company that advertises gains an increase in market share that more than makes up for the advertising investment. If both advertise, their investments cancel each other out and market share stays the same for each. Here's the question: advertise or not?

It's the same risk trade-off as before. From Alice's perspective, if she advertises and Bob does not, she increases her market share. But if she doesn't advertise and Bob does, she loses market share. She's better off advertising, regardless of what Bob does. Bob makes the same trade-off, so they both end up advertising and see no change in market share, when they would both have been better off saving their money.

Societal Dilemma: Advertising.	
Society: Two companies selling the same product.	
Group interest: Maximize profits. Group norm: To not engage in a costly and fruitless advertising arms race, and not advertise.	Competing interest: Maximize profits at the expense of the other company. Corresponding defection: Advertise.

This is your basic arms race, in which the various factions expend effort just to stay in the same place relative to each other. The USA and the USSR did this during the Cold War. Rival political parties do it, too.

If you assume the individuals can switch between strategies and you set the parameters right, the Hawk-Dove game is a Prisoner's Dilemma. When pairs of individuals interact, they each have the choice of cooperating (being a dove) or defecting (being a hawk). Both individuals know that cooperating is the best strategy for them as a pair, but that individually they're each better off being a hawk.

Not every interaction between two people involves a Prisoner's Dilemma. Imagine two drivers who are both stuck because a tree is blocking the road. The tree is too heavy for one person to move on his own, but it can be moved if they work together. Here, there's no conflict. It is in both their selfish interest and their group interest to move the tree together. But Prisoner's Dilemmas are common, and once you're primed to notice them, you'll start seeing them everywhere.⁴



The basic Prisoner's Dilemma formula involves two people who must decide between their own self-interest and the interest of their two-person group. This is interesting—and has been studied extensively⁵—but it's too simplistic for our purposes. We are more concerned with scenarios involving larger groups, with dozens, hundreds, thousands, even millions of people in a single dilemma.

Here's a classic societal dilemma: overfishing. As long as you don't catch too many fish in any area, the remaining fish can breed fast enough to keep up with demand. But if you start taking too many fish out of the water, the remaining fish can't breed fast enough and the whole population collapses.

If there were only one fisher, she could decide how much fish to catch based on both her immediate and long-term interests. She could catch all the fish she was able to in one year, and make a lot of money. Or she could catch fewer fish this year, making less money, but ensuring herself an income for years to come. It's a pretty easy decision to make—assuming she's not engaged in subsistence fishing—and you can imagine that in most instances, the fisher would not sacrifice her future livelihood for a short-term gain.

But as soon as there's more than one boat in the water, things become more complicated. Each fisher not only has to worry about overfishing the waters herself, but whether the other fishers are doing the same. There's a societal dilemma at the core of each one of their decisions.

Societal Dilemma: Overfishing.	
Society: A group of fishers all fishing out of the same waters.	
Group interest: The productivity of the fishing waters over the long term.	Competing interest: Short-term profit.
Group norm: To limit individual catches.	Corresponding defection: Take more than your share of fish.

Fisher Alice's trade-off includes the same elements as Prisoner Alice's trade-off. Alice can either act in her short-term self-interest and catch a lot of fish, or act in the group interest of all the local fishers and catch fewer fish. If everyone else acts in the group interest, then Alice is better off acting in her own selfish interest. She'll catch more fish, and fishing stocks will remain strong because she's the only one overfishing. But if Alice acts in the group interest while others act in their self-interest, she'll have sacrificed her own short-term gain for nothing: she'll catch fewer fish, and the fishing stocks will still collapse due to everyone else's overfishing.

Her analysis leads to the decision to overfish. That makes sense, but—of course—if everyone acts according to the same analysis, they'll end up collapsing the fishing stocks and ruining the industry for everyone. This is called a *Tragedy of the Commons*, and was first described by the ecologist Garrett Hardin in 1968.⁶

A Tragedy of the Commons occurs whenever a group shares a limited resource: not just fisheries, but grazing lands, water rights, time on a piece of shared

exercise equipment at a gym, an unguarded plate of cookies in the kitchen. In a forest, you can cut everything down for maximum short-term profit, or selectively harvest for sustainability. Someone who owns the forest can make the trade-off for himself, but when an unorganized group together owns the forest there's no one to limit the harvest, and a Tragedy of the Commons can result.

A Tragedy of the Commons is more complicated than a two-person Prisoner's Dilemma, because the other fishers aren't making this decision collectively. Instead, each individual fisher decides for himself what to do. In the two-person dilemma, Alice had to try to predict what Bob would do. In this larger dilemma, many more outcomes are possible.

Assume there are 100 fishers in total. Any number from 0 through 100 could act in their selfish interest and overfish. Harm to the group would increase as the scope of overfishing increases, regardless of what Alice does. Alice would probably not be harmed at all by 1 fisher overfishing, and she would be significantly harmed if all 99 chose to do the same. Fifty overfishers would cause some amount of harm; 20, a lesser amount. There are degrees of overfishing. Twenty fishers who each overfish by a small amount might do less damage to the fish stocks than 5 who take everything they can out of the water. What matters here is the scope of defection: the number of overfishers, but also the frequency of overfishing, and the magnitude of each overfishing incident.

At some scope of defection, stocks will be so depleted that everyone's catch in future years will be jeopardized. There's more at stake than whether Alice gets her fair share. In game theory, this is called a non-zero-sum game because wins and losses don't add up to zero: there are outcomes where everyone loses, and loses big.⁷ A fishery is non-zero-sum. Other societal dilemmas might seem like zero-sum games with a finite resource: if one person takes more, others get less. But even in these instances, there is a potential for catastrophe in widespread defection. If a community can't share a common water resource, everyone's crops will die because farmers can't plan on water use. If a few people constantly hog the exercise equipment, others won't come to the gym, which will lose membership and close. If someone consistently takes all the cookies, Mother will stop baking them. Remember: it's a bad parasite that kills its host.

The non-zero-sum property is an essential aspect of a societal dilemma. The group result barely depends on any single person's actions. Alice's cooperation or defection doesn't appreciably change the number of overfishers, nor is it likely to collapse the fishing stocks. It's the actions of the group that determine the overall result; at some point, the effects of the overfishers on the group will change from nothing to irreversible damage.

It's also possible that the group will not reach that point, even if all the members take as much fish as they want. There might not be enough fishers in the waters, or fishing technology might not be efficient enough. All the members of the group might be able to fish as much as they possibly can without affecting each other or future fishing stocks. But at some point, either the waters will get crowded enough or the fishers will get technologically advanced enough that the Tragedy of the Commons dilemma will occur.

The disconnect between Alice's individual actions and the effect of the group's actions as a whole makes societal dilemmas even harder to solve in larger groups. Under a rational economic analysis, it makes no sense for Alice to cooperate. The group will do whatever it does, regardless of her actions, and her individual cooperation or defection won't change that. All she's really deciding is whether to seize or forgo the short-term benefits of defecting.

Societal Dilemma: Tragedy of the Commons.	
Society: Some group of people, either a society of interest or a society of circumstance.	
Group interest: That the common resource not run out, and be available for all.	Competing interest: Get as much of that resource as possible in the short term.
Group norm: Cooperate and share that resource within its sustainability limits..	Corresponding defection: Take as much of that shared resource as you can.

In a Tragedy of the Commons, people acting in their self-interest harm the group interest. There's another type of societal dilemma, where people can receive the benefit of those who act in the group interest without having to act in the group interest themselves. It's called the *free-rider problem*

Whooping cough (otherwise known as pertussis) is a good example. It's both almost entirely preventable and almost entirely untreatable. Early in the 20th century, before the establishment of widespread vaccination programs, it was one of the most feared illnesses, and it remains a significant cause of death in developing countries. Compared to other vaccines, the pertussis vaccine isn't actually very effective at conferring immunity to any one individual. The standard infant schedule calls for four shots. After the first shot, about 30% become immune; after two, 50%; and even after all four shots have been administered, only about 90% of individuals have enough antibodies to fight off the disease.

What's more, vaccination is not without risk. The original pertussis vaccine carried a small risk of neurological damage. It has since been replaced with a

safer vaccine, but a minuscule risk of adverse reactions still persists, as it does with any vaccine. In the late 1960s and early 1970s, adverse vaccine reactions received a lot of attention in the media, most notably in Sweden, Japan, and the UK. Parents began to refuse vaccinations for their children, and doctors were often powerless to persuade them that the benefits outweighed the risks.

One of the primary benefits of vaccination is herd immunity. If almost everyone is vaccinated against a particular disease, there's no way for it to take hold in the community. Even if someone gets the disease, it's unlikely he will be able to infect others. Parents who refuse to have their children vaccinated do not only endanger their own children; they increase the risk of infection for everyone in the community. This increases, of course, as more parents opt out of vaccination programs. And while this is true for any vaccinated disease, the danger is particularly acute for whooping cough because the vaccine doesn't confer complete immunity and isn't recommended for the youngest infants or for those who are immune-compromised.

Between 1974 and 1979, the rate of pertussis vaccination among Swedish infants dropped precipitously, from 90% to 12%. Correspondingly, the incidence of whooping cough in Swedish children under four skyrocketed from 0.05% in 1975—effectively zero—to 3.4% by 1983. Sweden went from a country that had all but eradicated whooping cough to a country with a 1 in 30 infection rate.

When parents decide whether or not to immunize their child, they are faced with a societal dilemma. They can choose to cooperate and vaccinate their child, or they can choose to defect and refuse. As long as most children are vaccinated, a child is better off not being immunized: he avoids the chance of adverse effects, but reaps the benefit of herd immunity. But if there are too many defectors, everyone suffers the increased risk of epidemics. And it's a non-zero-sum game; there's a point where epidemics suddenly become much more likely.

Societal Dilemma: Vaccination.	
Society: Society as a whole.	
Group interest: No epidemics.	Competing interest: Avoid the small risk of adverse side effects (encephalopathy, allergic or autoimmune reactions, or—in extreme cases—contracting the disease from the vaccination).
Group norm: Vaccinate.	
Corresponding defection: Avoid vaccination.	

A free rider receives the benefit of everyone else's cooperation without having to cooperate himself. Think of a single person in the community who

doesn't pay his taxes; he gets all the benefits of the public institutions those taxes pay for—police and fire departments, road construction and maintenance, regulations to keep his food and workplace safe, a military—without having to actually pay for them.

But as more and more people stop paying their taxes, the government can provide fewer and fewer of those services—services that would be much more expensive or impossible for individuals to provide on their own—and the benefit of free riding is reduced. In the extreme, the whole system collapses.

Imagine a condominium without smoke detectors. The first tenant to install one is a sucker, because even though he pays for his detector, the building can burn down from a fire started elsewhere. The last tenant to install one is a fool, because he already receives the benefits of everyone else's detectors without having to pay anything.



It's easy to dismiss those original two-person examples as the responsibility of the two people alone. Alice and Bob can decide whether to rat on each other in jail, or whether to cheat each other when they buy and sell sealed bags. No one else needs to get involved. There's certainly no reason for society to get involved. Let the buyer and seller beware.

Society becomes involved because a broader societal dilemma emerges from Alice's and Bob's decisions. Let's look at the sealed bag exchange, focusing on customer Alice. She can either cooperate by paying for her purchase, or defect by defrauding merchant Bob. Yes, that decision most directly affects Bob, but—thinking more broadly about theft and society—it affects everyone.

Societal Dilemma: Defrauding merchants.	
Society: Those who buy and sell goods.	
Group interest: For commerce to operate smoothly.	Competing interest: Get stuff without having to pay for it.
Group norm: Don't defraud merchants.	Corresponding defection: Defraud merchants.

It's not that society cares about any particular thief; rather, society wants property rights to be respected. Note that it doesn't matter what sort of property rights deserve respect. There could be communal property, there could be personal property, and there could be everything in-between. What's important for

society is for everyone to respect what society decides are the property rules that make collective life work, and then for everybody to be able to trust that those rules will be followed.⁸

Similarly, if we focus on merchant Bob, we can see that he is in a corresponding societal dilemma with the society composed of all the other merchants: he can either treat his customers fairly or he can defraud them. Society doesn't want dishonest merchants; not only because we don't want to be defrauded, but also because we know that our entire system of commerce hinges on trust.

The alternative just wouldn't work. Merchants would stop doing transactions with all customers, not just with Alice. And customers would stop doing transactions with all merchants. Or they could both implement expensive and time-consuming bag-checking procedures that require them to each hire someone to help them perform transactions. And so on. Without trust, commerce collapses.

Even prisoners can have a broader community with a stake in whether or not prisoners confess. A criminal organization won't be concerned with Alice or Bob personally, but with members' loyalty to the organization. The organization as a whole benefits if it is viewed by individual members as an association in which they can trust others to keep their secrets, even at great personal cost.

Societal Dilemma: Criminals testifying against each other.	
Society: The criminal organization.	
Group interest: To minimize the amount of jail time for the society.	Competing interest: To minimize personal jail time.
Group norm: Don't testify against each other.	Corresponding defection: Testify against each other in exchange for reduced jail time.

The interesting thing about these dilemmas is that, looking at them in isolation, there's no logical solution. Thinking back to the prisoners, there is no analysis by which cooperation makes sense. Because they can't trust each other, they both end up confessing. This is the fundamental problem with cooperation: trust is unnatural, and it's not in the individual's short-term self-interest. This problem is why cooperation is so rare in the natural world, why it took so long to develop in humans, and why we have developed societal pressures as a way to enforce cooperation and hold society together.

PART II

A Model of Trust





6

Societal Pressures

In game theory, Prisoner's Dilemmas have no solution. Because the two prisoners, or the merchant and customer, can't trust each other, they both end up defecting. The larger societal dilemmas—the arms race, the Tragedy of the Commons, and the free-rider problem—are similarly unsolvable. Defecting is the only course that makes logical sense, even though the end result will be disastrous for the entire group.

But that's not how people generally operate. We cooperate all the time. We engage in honest commerce, although Enron and AIG and Countrywide are some pretty spectacular exceptions. Most of us [don't overfish](#), even though the few of us who do have depleted the ocean's stocks. We mostly vaccinate our children, despite the minor risk of an adverse reaction. Sometimes, even, we don't rat on each other in prison.¹

Prisoner's Dilemmas involve a risk trade-off between group interest and self-interest, but it's generally only a dilemma if you look at it very narrowly. For most people, most of the time, there's no actual dilemma. We don't stand at the checkout line at a store thinking: "Either the merchant is selling me a big screen TV, or this box is secretly filled with rocks. If it's rocks, I'm better off giving him counterfeit money. And if it has a TV, I'm still better off giving him counterfeit money." Generally, we just pay for the TV, put it in our car, and drive home. And if we're professional check forgers, we don't think through the dilemma, either. We pay for the TV with a bad check, put it in our car—I suppose it's a getaway car this time—and drive back to our lair.

The problem isn't with people; the problem is with the dilemma.² Societal dilemmas are choices between group interest and some competing individual interest. It assumes the individuals are only trying to minimize their jail time, or maximize their fishing catch or short-term profits. But in the real world, people

are more complicated than that. Our competing interests are more nuanced and varied, and they're subjective and situational. We try to maximize things other than our selfish self-interest. And our societal dilemmas are part of our ongoing relationships with other people.

Society solves societal dilemmas by making it in people's best interest to act in the group interest. We do this so naturally and so easily that we don't even notice the dilemma. Because of laws and police, it's not obviously better to steal a big screen TV than go without. Because of how everyone will react, it's not obviously smarter to betray a friend. Sure, no jail time is better than risking six years in jail, and catching more fish is better than catching fewer fish, but even those assessments fail to capture the richness of human emotion. Is no jail time but a reputation as a stool pigeon better than six years in jail? Is catching more fish but contributing to the degradation of the oceans better than catching fewer fish, even if everyone else is catching more than you? It depends. It depends on who you are. It depends on what you are. It depends on where you are.

Another famous dilemma illustrates this. The Stag Hunt was first formulated by [Jean-Jacques Rousseau](#) in 1754. In his scenario, a small group of hunters—it could be two and it could be more; it doesn't matter—are hunting a stag together. As would be obvious to readers of his day, everyone needs to work together in order to pull this off.

If it was a matter of hunting deer, everyone well realized that he must remain at his post; but if a rabbit happened to pass within reach of one of them, we cannot doubt that he would have gone off in pursuit of it without scruple and, having caught his own prey, he would have cared very little about having caused his companions to lose theirs.

What makes this different than the Prisoner's Dilemma is that the benefit of cooperation is more than the benefit of defection: a stag is much more food, even divided a few ways, than a rabbit. It would seem there's no actual dilemma; for all players, cooperate—cooperate is better than any other option. In the real world, however, [defections happen](#) in this sort of cooperative game all the time. It seems to make no sense.

Rousseau, too, ignored the variety and subjectivity of the hunters' competing interests. It's not obvious—for all people all the time—that a share of a stag is better than a whole rabbit. Sure, it's more meat, but that's not the only consideration. First of all, the stag isn't a done deal. The choice is between a guaranteed

rabbit—they're small and easy to catch—and the possibility, maybe even the probability, of a share of a stag. Is our intrepid hunter Alice an optimist or a pessimist? Does she want to stalk stag for hours, or does she want to snare her rabbit, go home, and do something she really enjoys with the rest of her day? Maybe she's tired. Maybe she's bored. Maybe she doesn't even like the taste of stag, and has a great rabbit stew recipe she's been dying to try. (Me, I like the one in Julia Child's *The Way to Cook*.) Maybe she is happy to forgo the rabbit for a stag, but doesn't trust that her fellow hunters will do the same. The point is that it's not for Rousseau to conclude which of these considerations matter to the hunters; the hunters get to decide for themselves. And they're all going to decide differently.

Another dilemma is called the [Snowdrift Dilemma](#), sometimes called Chicken.³ Two drivers are trapped by a snowdrift; each can either cooperate by shoveling or defect by remaining in his own car. If both remain in their cars, both remain stuck. If at least one of them shovels, both are freed; and two shovelers will get the job done much faster and more reliably than one. But unlike a Prisoner's Dilemma, it's in each driver's best interest to cooperate, even if the other defects.⁴

It turns out there are several different dilemmas⁵—generally called [social dilemmas](#) or coordination games—whose differences depend on the relative value of the various outcomes. Those nuances make a huge difference to game theorists, but are less important to everyday people. We make trade-offs based on what we want to do.⁶

When you look at the details of players' competing interests, motivations, and priorities, you often realize they might not be playing the same game. What might be a Prisoner's Dilemma for Alice could be a Snowdrift for Bob. What might be a Snowdrift for Alice might be a Stag Hunt for Bob. For Alice, cooperating might be the obviously smart thing to do. She might feel bad about herself if she defected. She might be afraid of what her friends would think if she defected. There might be a law against defecting, and she might not want to risk the jail time. She'll have her own particular trade-off: her own subjective values about cooperating and defecting. Bob might cooperate or defect for completely different reasons. And even if Bob and Alice are playing the same game today, they might each play a different game tomorrow. The complexities of these societal dilemmas are [much more complicated](#) than simple game theory models.

Think back to the baboon story at the start of Chapter 4. Notice the societal dilemma:

Societal Dilemma: Gazelle hunting.	
Society: Society of baboons.	
Group interest: Tasty gazelle meat for everyone.	Competing interest: Gaining an advantage over a fellow baboon.
Group norm: Hunt cooperatively.	Corresponding defection: Attack a fellow baboon during the hunt.

One of the great achievements of our species is our ability to solve societal dilemmas. In a way, we solve them by cheating. That is, we don't solve them within the parameters of the game. Instead, we modify the game to eliminate the dilemma. Recall the two drivers stuck behind a fallen tree that neither one can move by himself. They're not in a Prisoner's Dilemma. They're not even in a Snowdrift Dilemma. In their situation, their selfish interest coincides with the group interest—they're going to move the tree and get on with their lives. The trick to solving societal dilemmas is make them look like that. That's what societal pressures do: they're how society puts its thumb on the scales.



Solving societal dilemmas often means considering the people involved and their situations more broadly. The sealed-bag exchange is no longer a Prisoner's Dilemma if we assume the people involved have a sufficiently strong conscience.

Alice might be thinking: "If I assume Bob will cooperate, I have two choices. If I cooperate, I'll get my purchase and feel good about cooperating with Bob. If I defect, I'll get my purchase for free but I'll feel guilty about cheating Bob. That guilty feeling is worse than giving up the money, so it makes sense for me to cooperate. On the other hand, if I assume Bob will cheat me, my two choices look like this: If I cooperate, Bob will take my money and I'll feel stupid and angry for cooperating with a cheat. If I defect, I won't get my purchase and will feel guilty for trying to cheat Bob. That stupid feeling for being cheated is a little worse than the guilty feeling for trying to cheat Bob—who turned out to be a cheat himself. But Bob is making this same analysis, and he doesn't want to feel guilty about cheating me, either. So he's not going to defect."

And indeed, Bob makes the same analysis and also cooperates, although—most likely—they both don't consciously decide anything and both just behave honestly and trust each other to do the same. Maybe I have the emotions wrong—they could be motivated by a moral compass, by a sense of fairness, or by altruism towards the other person. In any case, dilemma solved.

Those guilty feelings come from inside our heads. Feelings of guilt are a societal pressure, one that works to varying degrees in each of us.

Moral pressure isn't the only thing we use to solve societal dilemmas. All of the considerations that make cooperation more attractive and defection less attractive are societal pressures. These include the rewards society directs towards cooperators and the penalties it directs towards defectors, the legal punishments society metes out to defectors, and the security measures that make defecting difficult to pull off and even more difficult to get away with.⁷

Societal Dilemma: Stealing.	
Society: Society as a whole.	
Group interest: Respect property rights. Group norm: Don't steal.	Competing interest: Get stuff without having to pay for it. Corresponding defection: Steal.
To encourage people to act in the group interest, society implements these societal pressures:	
Moral: People feel good about being honest and bad about stealing. People have been taught religious admonitions like "Thou shalt not steal."	
Reputational: Society shuns people who have a reputation for being thieves.	
Institutional: Stealing is illegal, and society punishes thieves.	
Security: Door locks, burglar alarms, and so on.	

Of course, there's a lot more going on, and I'll discuss that in later chapters. The real world isn't this simplistic; any analysis of human interaction must take circumstances into account. If Alice is a tourist in a foreign country, Bob might cheat her anyway. If the dollar value of cheating is high enough, either Alice or Bob might decide that cheating is worth more than the negative feelings that result from cheating. In Chapter 3, I said that trust is contextual; all of that analysis applies here.

For most of us, it is more worthwhile to cooperate than to defect. It can be a better strategy for us, given what we know about the people who share in our dilemma.⁸ And, for different and equally valid reasons, some of us find defection to be more valuable than cooperation. Not universally, not all of the time, but at that moment for that person and that particular trade-off. There are no actual dilemmas; there are just individual subjective risk trade-offs.

Here are six different ways societal pressures can reduce the scope of defection—which I'll illustrate using the example of Alice potentially cheating a merchant.

- *Pressures that increase the actual or perceived difficulty of defecting.* Actual commerce usually doesn't happen inside sealed bags. Bob takes various additional security precautions to minimize the risk that Alice might cheat. Bob requires her to pay with hard-to-forge currency, or runs her credit card through a third-party authentication system. Window bars and burglar alarms make it harder for Alice to steal from Bob.
- *Pressures that raise the consequences of defecting.* These would be largely implemented after the fact; think prison terms, fines, cutting off a thief's hand,⁹ and social ostracism. Even if they never catch anyone, the police can make it difficult and expensive to commit a crime; every heist movie demonstrates this entertainingly.
- *Pressures that reduce the actual or expected benefits of defecting.* Exploding ink cartridges can make stolen garments less useful to thieves, and daily ATM withdrawal limits restrict how much a thief can steal.
- *Pressures that limit the damage caused by the defections that happen.* Bob won't keep a lot of cash in his store. He might even store some of his expensive inventory elsewhere. He'll also have an insurance policy that will help him resume normal business quickly after a theft.
- *Pressures that increase the benefits of cooperating.* Reputation serves this function; Alice derives value from being known in society as honest and honorable in her business dealings, more so if she is part of the same society as the merchant she patronizes. Certainly Alice's credit rating is a part of her reputation. We also have a powerful [need to conform](#) to the group.
- *Pressures that lower the costs of cooperating.* Society makes it easy to cooperate. Stores make check-out stands easy to find. Unforgeable paper money and credit cards make it easy to conduct commerce, as opposed to a barter system, or needing to lug around a sackful of gold and silver. Or think of the iTunes store, which makes it easy to buy music legitimately online.

There's a lot of overlap here, and many of these techniques are tightly coupled. When you reduce the benefits of defecting, you almost certainly reduce the frequency of defecting.



Figure 6: Societal Pressure Knobs

Think back to the Hawk-Dove game, and the knobs society can use to set the initial parameters. The categories in that figure are all individual knobs, and societal pressures provide a mechanism for the group to control those knobs. In theory, if the knobs are calibrated perfectly, society will get the exact scope of defection it's willing to tolerate.



There are many ways to sort societal pressures. The system I'm using sorts them by origin: moral pressures, reputational pressures, institutional pressures, and security systems.¹⁰ These are categories you've certainly felt yourself. We feel moral pressure to do the right thing or—at least—to not do the wrong thing. Reputational pressure is more commonly known as peer pressure, but I mean any incentives to cooperate that stem from other people. Institutional pressure is broader and more general: the group using rules to induce cooperation. Security systems comprise a weird hybrid: it's both a separate category, and it enhances the other three categories.

The most important difference among these four categories is the scale at which they operate.

- *Moral pressure works best in small groups.* Yes, our morals can affect our interactions with strangers on the other side of the planet, but in general, they work best with people we know well.
- *Reputational pressure works well in small- and medium-sized groups.* If we're not at least somewhat familiar with other people, we're not going to be

able to know their reputations. And the better we know them, the more accurately we will know their reputations.

- *Institutional pressure works best in larger-sized groups.* It often makes no sense in small groups; you're unlikely to call the police if your kid sister steals your bicycle, for example. It can scale to very large groups—even globally—but with difficulty.
- *Security systems can act as societal pressures at a variety of scales.* They can be up close and personal, like a suit of armor. They can be global, like the systems to detect international money laundering. They can be anything in between.

I'm being deliberately vague about group sizes here, but there definitely is a scale consideration with societal pressures. And because the increasing scale of our society is one of the primary reasons our societal pressure systems are failing, it's important to keep these considerations in mind.

Another difference between the categories of societal pressure is that they operate at distinct times during a security event. Moral pressure can operate either before, during, or after an individual defects. Reputational, as well as most institutional, pressure operates after the defection, although some institutional pressure operates during. Security can operate before, during, or after.

Any measures that operate during or after the event affect the trade-off through a feedback loop. Someone who knows of the negative outcome—perhaps ostracism due to a bad reputation, or a jail sentence—either through direct knowledge or through seeing it happen to someone else, might refrain from defecting in order to avoid it. This is deterrence.

All of this, and more, is illustrated in the complicated block diagram below. Along the bottom axis is the timeline: before, during, and after defection. Along the left are the different categories of societal pressure: moral and reputational (considered together), institutional, and security systems. The traits/tendencies box represents the physical and emotional aspects of people that make them more or less likely to defect. Natural defenses are aspects of targets that make them more or less difficult to attack. Neither of these are societal pressure systems, but I include them for the sake of completeness.

An example might be useful here. Alice is deciding whether to burglarize a house. The group interest is for her not to burglarize the house, but she has some competing interest—it doesn't matter what it is—that makes her want to burglarize the house. Different pressures affect her risk trade-off in different ways.

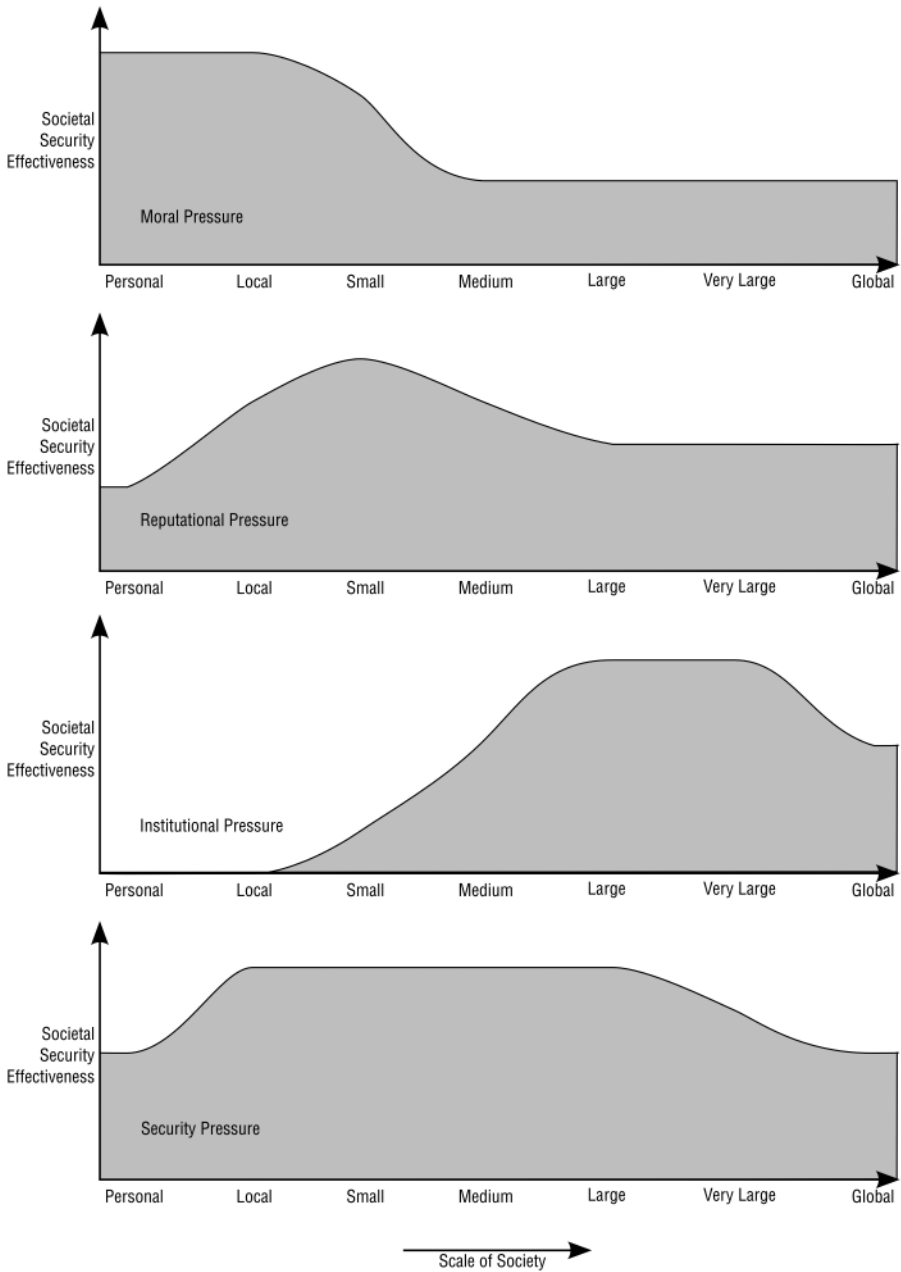


Figure 7: The Scale of Different Societal Pressures




Figure 8: How Societal Pressures Influence the Risk Trade-Off

- Traits/tendencies. If Alice is afraid of heights, she won't try to break in through a second-story window. If she has a broken leg, she probably won't try to break in at all. These considerations operate before defection, at the point of the risk trade-off, when she's deciding whether or not to burglarize the house. (Note that this is not societal pressure.)

- Natural defenses operate during the burglary. Maybe the owner of the house is home and might tackle Alice. (Again, note that this is not societal pressure.)
- Most moral pressures operate at the point of the risk trade-off, or decision: Alice's sense of fairness, sense of right and wrong, and desire to obey the law. Some operate during the actual burglary: feelings of empathy, for example. Some operate after she's committed her crime: guilt, shame, and so on.
- Reputational pressures, assuming she's caught, operate after she's done burglarizing the house. They stem from the reactions and responses of others.
- Institutional pressures also operate after she's done burglarizing the house. Think of laws and mechanisms to punish the guilty in this case.
- Security systems can operate before, during, or after. Preemptive interventions, including incarcerating Alice before she commits the crime or forcing her to take some mood-altering medication that makes her not want to burglarize houses, operate before. Defenses operate during: door locks and window bars make it harder for her to burglarize the house. Detection systems can operate during or after: a burglar alarm calls a response that may or may not come in time. Interventions, like camouflage and authentication systems, operate during as well. Forensic systems operate afterwards, and may identify Alice as the burglar. There's one more type of security system: recovery systems that operate after a burglary can provide a perverse incentive to those aware that the consequences of their misbehavior can be mitigated at no cost to themselves. If Alice knows the house owner can easily recover from a burglary—maybe he has a lot of money, or good insurance—she's more likely to burglarize him.

Systems that work during or after the burglary usually have a deterrence effect. Alice is less likely to burglarize a house if she knows the police are diligent and jail sentences are severe. Or if she knows there's a homeowner who is skilled at karate, or has a burglar alarm.

These categories are not meant to be rigid. There are going to be societal pressures that don't fit neatly into any one category. That's okay; these categories are more meant to be a way of broadly understanding the topic as a whole than a formal taxonomy.

In the next four chapters, I'll outline each type of societal pressure in turn. I'll talk about how they work, and how they fail.

Societal pressure failures occur when the scope of defection is either too high or too low: either there are too many burglaries, or we're spending too much

money on security to prevent burglary. This is not the same as individual burglaries; if someone's house was burglarized, it's not necessarily a societal security failure. Remember, we can never get the number of hawks down to zero; and sooner or later, further reducing their number becomes prohibitively expensive.

In some ways, societal pressures are like a group's immune system. Like antibodies, T cells, and B cells, they defend society as a whole against internal threats without being particularly concerned about harm to individual members of the group. The protection is not perfect, but having several different mechanisms that target different threats in different ways—much as an immune system does—makes it stronger.



7

Moral Pressures

Looking back at all the elections I've had the opportunity to vote in, there has never been one whose outcome I affected in any way. My voting has never even changed the vote percentages in any perceptible way. If I decided to never vote again, democracy wouldn't notice. It would certainly be in my best interest **not to vote**. Voting is a pain. I have to drive to the polling place, stand in line, then drive home.¹ I'm a busy guy.

Voting is a societal dilemma. For any single individual, there are no benefits to voting. Yes, your vote counts—it just doesn't matter. The rare examples of small elections decided by one vote don't change the central point: voting isn't worth the trouble. But if no one voted, democracy wouldn't work.

Still, **people vote**. It makes sense if 1) the voters see a difference between the two candidates, and 2) they care at least a little bit about the welfare of their fellow citizens. Studies with **actual voters** bear this out.²

Societal Dilemma: Voting.	
Society: Society as a whole.	
Group interest: A robust democracy. Group norm: Vote.	Competing interest: Do what you want to do on election day. Corresponding defection: Don't bother voting.
To encourage people to act in the group interest, society implements these societal pressures: Moral: People tend to feel good when they vote and bad when they don't vote, because they care about their welfare and that of their fellow citizens.	

Caring about the welfare of your fellow citizens is an example of a moral pressure.³ To further increase voter turnout, society can directly appeal to morals. We impress upon citizens the importance of the issues at stake in the election; we even frame some of the issues in explicitly moral terms. We instill in them a sense of voting as a civic duty. We evoke their sense of group membership, and remind them that their peers are voting. We even scare them, warning that if they don't vote, the remaining voters—who probably don't agree with them politically—will decide the election.

Murder is another societal dilemma. There might be times when it is in our individual self-interest to kill someone else, but it's definitely in the group interest that murder not run rampant. To help prevent this from happening, society has evolved explicit moral prohibitions against murder, such as the Sixth (or Fifth, in the Roman Catholic and Lutheran traditional number) Commandment, "Thou shalt not kill."

Morality is a complex concept, and the subject of thousands of years' worth of philosophical and theological debate. Although the word "moral" often refers to an individual's values—with "moral" meaning "good," and "immoral" meaning "bad"—I am using the term "morals" here very generally, to mean any innate or cultural guidelines that inform people's decision-making processes as they evaluate potential trade-offs.⁴ These encompass conscious and unconscious processes, explicit rules and gut feelings, deliberate thoughts, and automatic reactions. These also encompass internal reward mechanisms, for both cooperation and defection. Looking back at Figure 8 in Chapter 6, there is going to be considerable overlap between morals and what I called "traits/tendencies"—I hope to ignore that as well. As we saw in Chapter 3, all sorts of physiological processes make us prone to prosocial behaviors like cooperation and altruism. I'm lumping all of these under the rubric of morals.

And while morals can play a large part in someone's risk trade-off, in this chapter I am just focusing on how morals act as a societal pressure system to reduce the scope of defection. In Chapter 11, I'll discuss how morals affect the decision to cooperate or defect more broadly.

Beliefs that voting is the right thing to do, and that murdering someone is wrong, are examples of moral pressure: a mechanism designed to engage people's sense of right and wrong. Natural selection has [modified our brains](#) so that trust, altruism, and cooperation feel good, but—as we well know—that doesn't mean we're always trustworthy, altruistic, and cooperative. In both of the above examples, voting and murder, morals aren't very effective. With voting, a large defection rate isn't too much of a problem. U.S. presidential elections are decided by

about half the pool of eligible voters.⁵ Elections for lesser offices are decided by an even smaller percentage. But while this is certainly a social issue, the harm non-voters cause is minimal.

With murder, the number of defectors is both smaller and more harmful. In 2010, [the murder rate](#) in the U.S. was 5.0 per 100,000 people. [Elsewhere in the world](#), it ranges from 0.39 per 100,000 in the relatively murder-free nation of Singapore, to an astonishing 58 per 100,000 in El Salvador.



Morals affect societal dilemmas in a variety of ways. They can affect us at the time of the risk trade-off, by making us feel good or bad about a particular cooperate/defect decision. They can affect us after we've made the decision and during the actual defection: empathy for our victims, for example. And they can affect us after the defection, through feelings of guilt, shame, pride, satisfaction, and so on. Anticipating moral feelings that may arise during and after defection provides either an incentive to cooperate or a deterrent from defection.

At the risk of opening a large philosophical can of worms, I'll venture to say that morals are unique in being the only societal pressure that makes people "want to" behave in the group interest. The other three mechanisms make them "have to."

There are two basic types of morals that affect risk trade-offs, one general and the other specific. First, the general. The human evolutionary tendencies toward trust and cooperation discussed in Chapter 3 are reflected in our moral, ethical, and religious codes. These codes vary wildly, but all emphasize prosocial behaviors like altruism, fairness, cooperation, and trust. The most general of these is the [Golden Rule](#): "Do unto others what you would have them do unto you." Different groups have their own unique spin on the Golden Rule, but it's basically an admonition to cooperate with others. It really is the closest thing to a universal moral principle our species has.⁶

Moral reminders don't have to be anything formal. They can be as informal as the proverbs—which are anything but simplistic—that we use to teach our children; and cultures everywhere have proverbs about altruism, diligence, fidelity, and being a cooperating member of the community.⁷ The models of the world we learn sometimes have moral components, too.

This can go several ways. You might learn not to defecate upstream of your village because you've been taught about cholera, or because you've been taught that doing so will make the river god angry. You might be convinced not to

throw down your weapon and leave your fellow pikemen to face the charge alone either by a love for your comrades-at-arms like that for your brothers, or by knowledge that pikemen who run from charging horsemen get lanced in the back—not to mention what happens to deserters.

Traditionally, religion was the context where society codified and taught its moral rules. The Judeo-Christian tradition has the Ten Commandments, and Buddhism the Four Noble Truths. Muslims have the Five Pillars of Islam. All of these faiths call for the indoctrination of children in their teachings. Religion even exerts a subtle influence on the non-religious. In one experiment, theists and atheists alike gave more money away—to an anonymous stranger, not to charity—when they were first asked to unscramble a jumbled sentence containing words [associated with religion](#) than when the sentence contained only religion-neutral words. Another found less cheating when they were asked to recall the [Ten Commandments](#). A third experiment measured cheating behavior as a function of belief in a deity. They found no difference in cheating behavior between believers and non-believers, but found that people who conceived of a loving, caring, and forgiving God were much more likely to cheat than those who conceived of a [harsh, punitive, vengeful](#), and punishing God.⁸

Often, morals are not so much prescriptions of specific behaviors as they are meta rules. That is, they are more about intention than action, and rarely dictate actual behaviors. Is the Golden Rule something jurors should follow? Should it dictate how soldiers ought to treat enemy armies? Many ethicists have long argued that the Golden Rule is [pretty much useless](#) for figuring out what to do in any given situation.

Our moral decisions are contextual. Even something as basic as “Thou shalt not kill” isn’t actually all that basic. What does it mean to kill someone? A more modern translation from the original Hebrew is “Thou shalt not murder,” but that just begs the question. What is the definition of murder? When is killing not murder? Can we kill in self-defense, either during everyday life or in war-time? Can we kill as punishment? What about abortion? Is euthanasia moral? Is assisted suicide? Can we [kill animals](#)? The devil is in the details. [This is the stuff](#) that philosophers and moral theologians grapple with, and for the most part, we can leave it to them. For our purposes, it’s enough to note that general moral reminders are a coarse societal pressure.

Contextualities are everywhere. Prosocial behaviors like altruism and fairness may be universal, but they’re expressed differently at different times in each culture. This is important. While morals are internal, that doesn’t mean we develop them naturally, like the ability to walk or grasp. Morals are taught.

They're memes and they do evolve, subject to the rules of natural selection, but they're not genetically determined.

Or maybe some are. There's a theory that we have a [moral instinct](#) that's analogous to our language instinct. Across cultures and throughout history, all moral codes have rules in common; the Golden Rule is an example. [Others relevant](#) to this book include a sense of fairness, a sense of justice, admiration of generosity, prohibition against murder and general violence, and punishment for wrongs against the community. Psychologist and animal behaviorist Marc Hauser even goes so far as to propose that humans have [specific brain functions](#) for morals, similar to our language centers.⁹ And psychologist Jonathan Haidt proposes [five fundamental systems](#) that underlie human morality.

- *Harm/care systems.* As discussed in Chapter 3, we are naturally predisposed to care for others. From mirror neurons and empathy to oxytocin, our brains have evolved to exhibit altruism.
- *Fairness/reciprocity systems.* Also discussed in Chapter 3, we have natural notions of fairness and reciprocity.
- *In-group/loyalty systems.* Humans have a [strong tendency](#) to divide people into two categories, those in our group ("us") and those not in our group ("them"). This has serious security ramifications, which we'll talk about in the section on group norms later in the chapter, and in the next chapter about group membership.
- *Authority/respect systems.* Humans have a tendency to [defer to authority](#) and will follow orders simply because they're told to by an authority.
- *Purity/sanctity systems.* This is probably the aspect of morality that has the least to do with security, although patriarchal societies have used it to police all sorts of female behaviors. Mary Douglas's *Purity and Danger* talks about how notions of [purity and sanctity](#) operate as stand-ins for concepts of unhealthy and dangerous, and this certainly influences morals.

You can think of these systems as moral receptors, and compare them to taste and touch receptors. Haidt claims an evolutionary basis for his categories, although the evidence is scant. While there may be an innate morality associated with them, they're also strongly influenced by culture. In any case, they're a useful partitioning of our moral system, and they all affect risk trade-offs.

These five fundamental systems are also influenced by external events. [Spontaneous cooperation](#) is a common response among those affected by a natural disaster or other crisis. For example, there is a marked [increase in solidarity](#)

during and immediately after a conflict with another group, and the U.S. exhibited that solidarity after the terrorist attacks of 9/11. This included a general **increase in prosocial** behaviors such as spending time with children, engaging in religious activities, and donating money. **Crime in New York** dropped. There was also an increase in **in-group/out-group** divisions, as evidenced by an **increase in hate crimes** against Muslims and other minorities throughout the country.

Some of our moral pressure is very strong. **Kin aversion**, the particular disgust we have for the idea of mating with people we grew up with, works without any prompting or ancillary security.¹⁰ So does our tendency to feel **protective impulses** towards children, which can extend to **small animals** and **even dolls**. This makes sense. We avoided incest and looked after our offspring for millions of years before we became human. These strong moral inclinations can be deliberately tapped. Think of evocations of kin relationships to foster cooperation: blood brothers, brothers in Christ, and so on. Or how cartoon animals are so often drawn with the big-head-big-eyes look of babies in an attempt to make them more universally likeable.

A lot of our morals are cultural. For example, while fairness is a universal human trait, **notions of fairness** differ among groups, based on variables like community size and religious participation. Psychologist Joseph Henrich used a cooperation game to study notions of fairness, altruism, and trust among the Machiguenga tribesmen deep in the Peruvian Amazon. While Westerners tended to share a lucky find with someone else, the tribesmen were more likely to keep it to themselves. In both instances, the actions were perceived as fair by others of the same culture.¹¹

Think back to the various societal dilemmas we've discussed so far. Many of them have a moral component that encourages people to cooperate. We're taught—or conditioned, depending on what social science theory you believe—that stealing and fraud are wrong, although different cultures have different definitions of property. We're taught that taking more than our fair share is wrong: whatever “fair” means in our culture. We're taught that sitting idly by while others do all the work is wrong, although no one accuses the incapacitated, the infirm, the elderly, or infants of being immoral. Even criminals have moral codes that prohibit ratting on each other.

Of course, the effectiveness of these rules depends largely on individual circumstances, and some of them—such as “honor among thieves,” or the politeness rule against taking the last item on a communal plate of food—are notoriously weak. But there would be even less honor among thieves if the phrase didn't exist to remind them of their moral obligation to the group.

Moral pressure can do better, though. In addition to general admonitions to cooperative behavior, other measures specifically remind people of their moral obligations to the group, such as the obligation to vote. For example, think about the signs in restaurant bathrooms that read, “Employees must wash hands before returning to work.”¹² Another example are the signs that remind people not to litter. My favorite is “[Don’t mess with Texas](#),” one of the best advertising slogans ever.

Of course, signs warning “No shoplifting,” or “Shoplifting is a crime,” primarily remind shoplifters that they run the risk of getting caught and going to jail; more about that in a couple of chapters. These reminders nonetheless have an unmistakable moral component. And, more precisely, public service announcements that deliberately invoke people’s feelings of guilt and shame have been shown to be [effective in changing](#) behavior.

One afternoon at a historic monument in Rome, I saw a sign advising visitors: “This is your history. Please don’t graffiti it.” That sign was an artifact of a societal dilemma: it’s fun to carve one’s name into the rock wall, but if everyone does that, historic monuments will soon look ugly. The difference between selfish interest and group interest is small in this case; for some people, a simple reminder is enough to tip the scales in favor of the group.



Group norms are themselves a form of moral societal pressure. Voting turnout rates range from as high as 92% in Austria to as low as 48% in the United States. Yes, some countries make voting mandatory and use other categories of social pressure to get people to vote, and we’ll talk about those in the next chapter; but these are rates in countries where voting is entirely optional.

The “Don’t mess with Texas” slogan is so good because it doesn’t just remind people not to litter. It reinforces the group identity of Texans both as people who don’t leave messes and who are not to be messed with.

We not only absorb our moral codes and definitions of right and wrong from the group; the group also transmits cues about cooperation and defection and what it means to act in a trustworthy manner. People are more likely to suppress their self-interest in favor of the group interest if they feel that others are doing so as well, and they’re less likely to do so if they feel that others are [taking advantage](#) of them. The psychological mechanism for this is unclear, but certainly it is related to our innate sense of fairness. We generally don’t mind sacrificing for the group, as long as we’re all sacrificing fairly. But if we feel like

we're being taken advantage of by others who are defecting, we're more likely to defect as well.

If you see your neighbor watering his lawn during a drought restriction and getting away with it, your sense of fairness is offended. To restore fairness, you have two options: you can turn him in, or you can take the same benefit for yourself. You have to live with your neighbor, so defecting is easier. (Recall the phrase "if you can't beat 'em, join 'em.")

Psychologist [Andrew Colman](#) called this the Bad Apple Effect.¹³ Any large group is likely to contain a few bad apples who will defect at the expense of the group interest and inspire others to do likewise. If someone is speeding, or littering, or watering his lawn in spite of water-use restrictions, others around him are more likely to do the same.

This can occasionally create a positive-feedback loop, driven by individual differences in [how people evaluate](#) their risk trade-off. The first defectors provide a small additional incentive for everyone else to defect. Because there are always some people who are predisposed to cooperate, but just barely, that incentive may push them over to the defecting side. This, in turn, can result in an even greater incentive for everyone else to defect—a cascade that can sometimes lead to mass defection and even a mob mentality.

Both experiment and observation bear this out. Littering is a societal dilemma: it is in everyone's self-interest to drop his or her own trash on the ground—carrying it to a can is bothersome—but if everyone did that, the streets would be a mess. People are [more likely to litter](#) if there is a small amount of litter already on the ground, and two or three times more likely to litter if there is a lot. Just seeing a single person litter, or [seeing someone](#) pick up litter, modifies behavior. In a recent book, James B. Stewart points to the current epidemic of lying by public figures, and blames it for the [general breakdown](#) of ethics in America: when lying is believed to be normal, more people lie. In psychological experiments, a single [unpunished free rider](#) in a group can cause the entire group to spiral towards less and less cooperation. These patterns reflect the human tendency to adhere not only to social norms, but to moral norms. [In Islam](#), announcing that you've sinned is itself a sin.¹⁴

This effect can motivate cooperation, too. For years, society has tried to encourage people to conserve energy. It's another societal dilemma: we're better off collectively if we conserve our natural resources, but each of us individually is better off if we use as much as we want. Even more selfishly, if I use as much as I want and everyone else conserves, I get all the benefits of conservation without

actually having to do anything. Awareness campaigns have worked somewhat to mitigate this problem, but not enough.

Every month, included in my electric bill, is a chart comparing my electricity usage to my neighbors' average usage. It tells me if I'm using more, or less, electricity than average. On the face of it, why should I care? Electricity isn't free. The more I use, the more I pay. And aside from the savings from lower bills—which exist even without the chart—I get no personal benefit for conserving, and incur no penalty for not conserving.

But the chart works. People use less energy when they can compare their energy usage with that of their neighbors.¹⁵ That's because there actually is a benefit and a penalty, albeit entirely inside the heads of those receiving the bill: their competitive nature, their desire to conform to group norms, and so on.

Similarly, people are more likely to pay their taxes if they think others are paying their taxes as well. People are more likely to vote, less likely to overfish, more likely to get immunized, and less likely to defraud their customers if they think these practices are the group norm. This isn't peer pressure; in these cases, the risk trade-off is made in secret. Of course, this sort of thing works even better when the group knows whether or not you've cooperated or defected, but that's the subject of the next chapter.

Morals can be influenced by a powerful ruler, or a ruling class, or a priestly class. Especially if you can manipulate people's in-group/out-group designations, some awful things can be done in the name of morality: slavery and genocide are two examples.¹⁶ Interestingly, genocide is often precipitated by [propaganda campaigns](#) that paint the victims as vermin or otherwise less than human: underscoring of the moral predispositions people have towards other people.

Psychologist Simon Baron-Cohen maintains that in psychopaths, cruelty and evil stem from a [failure or absence](#) of empathy. Extending this notion into our model, a person is more inclined to cooperate if he feels empathy with the other people in the group, and is more inclined to defect if he doesn't feel that empathy. Both general moral rules and specific moral reminders serve to enhance empathy to the group, by reminding people of both their moral principles and the group interest.¹⁷



For more than ten years, economist Paul Feldman brought bagels into his workplace and sold them on the honor system. He posted prices that people were expected to pay, securing the system by nothing more than the bagel-eaters'

morals.¹⁸ Although it was easy to take a bagel without paying, Feldman succeeded in collecting about 90% of the posted price, resulting in much more profit than if he had to pay someone to sell the bagels and guard the money. He eventually turned this into a full-time business, selling **food on the honor** system to 140 companies in the Washington, DC, area.

Societal pressure based on morals largely succeeds because of who we are as human beings. When we meet someone for the first time, we tend to cooperate. We act trustworthy because we know it's right, and we similarly extend some amount of trust. We tip in restaurants. We pay for our bagels. We **follow social norms** simply because they *are* social norms. This is all contextual, of course, and we're not stupid about it. But it is our nature.

Philosopher **Emmanuel Levinas** said that morality is grounded in face-to-face interactions. In general, moral pressure works best at close range. It works best with family, friends, and other intimate groups: people whose intentions we can trust. It works well when the groups are close in both space and time. It works well when it's immediate: in crises and other times of stress. It works well with groups whose members are like each other, whether ethnically, in sharing an interest, or some other trait. Even having a common enemy works in this regard.

Think about the chart that shows my energy use compared to my neighbors'. It doesn't compare me to the rest of the world, or even to my country. It compares me to my neighbors, the people most like me.

Morals sometimes work at long range. After the 2011 earthquake and tsunami in Japan, **people turned in** thousands of wallets and safes found in the rubble, filled with \$78 million in cash. People regularly protest working conditions at factories, or give to relief efforts, or fight social injustices, in other countries halfway across the planet. People have moral beliefs that encompass all of humanity, or all animals, or all living creatures. We are a species that is capable of profound morality.

We are also a species capable of profound immorality. And while moral pressure works, it also regularly fails. When it does, it fails for several specific reasons:

People vary in their individual behavior. Sure, most people will cooperate most of the time, but some people will defect some of the time, and almost everyone will defect once in a while.

Morals often conflict. We'll talk about this more in Chapter 11. Sometimes defectors are people whose morals lead to different imperatives than those

reached by the cooperators. These people will be largely unaffected by societal moral pressure. Society will have an easier time convincing a potential thief that stealing is wrong than it will have convincing an abolitionist that slavery is good.

Morals often overreach. It's relatively easy to use morals to enforce basic prosocial behaviors, because those are aligned with what's already in our brains. Enforcing arbitrary moral codes is much harder. If the group norm goes against any of Haidt's five fundamental moral systems, more people will have conflicting morals, and more will defect.

Throughout history, totalitarian regimes have attempted to impose moral codes on their citizens, suppressing some heretofore acceptable behaviors and inventing new obligations. Perhaps the most well-known modern example of an authoritarian attempt to reorient popular moral sensibilities was the Soviet Union's unsuccessful prohibition on the [practice of religion](#), which threatened to undermine materialist Communist ideology. This kind of thing isn't rare. When I visited Myanmar in 1991, I saw large billboards everywhere, courtesy of the government's "People's Desire Campaign," exhorting the populace to believe and act in all sorts of pro-government ways. These campaigns often just drive the behaviors underground.

Morals can be manipulated. [Confidence tricksters](#), in particular, manipulate the very same traits that make us cooperate: kindness, altruism, fairness.

Morals scale badly. They fail as societies become larger and the moral ties that bind their members weaken.

Remember Baron-Cohen's theory of empathy. As the group gets larger and more anonymous, there's less empathy. Joseph Stalin said, "the death of one man is a tragedy, the death of millions is a statistic"; similarly, [we have trouble](#) thinking about large groups in the same moral way we think of the people closest to us.

All of these reasons make morals the weakest of the societal pressures. Morals are the societal pressure that works "when no one is watching." They determine whether we keep a wallet that no one saw us find and pick up, whether we litter on a deserted street, whether we conserve energy or crank up the air conditioning, and whether we help ourselves to a bagel on a tray in an empty break room.

When nothing other than moral pressure influences a societal dilemma, the number of defectors will be at its largest. However, opportunities for individuals to make moral choices when they are unobserved represent only a small portion of societal dilemmas. We humans are a social species, and more often than not someone is watching. And that makes an enormous difference.



8

Reputational Pressures

From the perspective of trust, societal dilemmas involve a Red Queen Effect. On one hand, defectors should evolve to be better able to fool cooperators. And on the other, cooperators should evolve to better recognize defectors. It's a race between the [ability to deceive](#) and the ability to detect deception.

There's a lot of research on detecting deception, and humans seem not to be very good at it. There are exceptions, and people can learn to be better at it—but in general, we can't tell liars from truth-tellers. Like the Lake Wobegon children who are all above average, most of us think we're much better at detecting deception than we actually are. We're better, but still not great, at predicting cooperators and defectors.¹

This is surprising. The Red Queen Effect means both sides improve in order to stay in place, yet in this case, defectors have the upper hand. A possible reason is that we have developed another method for figuring out who to trust and who not to. We're a social species, and in our evolutionary past we interacted with the same people over and over again. We don't have to be that good at predicting bad behavior, because we're really good at detecting it after the fact—and using reputation to punish it.²

In fact, our brains have specially evolved to deal with cheating after the fact. Perhaps the most vivid demonstration of this can be seen with the [Wason Selection Task](#). The test compares people's ability to solve a generic logical reasoning problem with their ability to solve the same problem presented in a framework of detecting cheaters: for example, "if Alice went to Boston, she took the train" versus "if Alice is served alcohol, she is over 21." People are generally much [better at solving](#) the latter. Additionally, [fMRI scans of the brains](#) performing this task show that we have specific brain circuitry for cheater detection.

Think back to how contrived and artificial the sealed bag exchange from Chapter 5 seemed. That's because there's more going on than short-term decision making. Commerce isn't a one-time event. It happens again and again, day after day, often between the same people. We know the individuals and companies with whom we interact, maybe personally, maybe casually, maybe by their brand. Everyone has a reputation, and it's important. While morals are part of the reason we cooperate with each other, the preponderance of the evidence—both observational and experimental—supports the hypothesis that **we cooperate primarily** because we crave reward (engagement) and fear punishment (exclusion) from other members of our group.³

Bob depends on his reputation as an honest merchant. If he cheats Alice, she won't do business with him again. Even worse, she'll tell her friends.⁴ Bob couldn't survive as a merchant if he had a reputation as a cheater. If we assume that the cost to Bob's reputation if he defects is greater than the value of the item being purchased, he has no dilemma. He is better off cooperating, regardless of what Alice does. Reputation is such a major factor for Bob that he almost certainly allows Alice to reverse the transaction after the fact, a process commonly known as returning the purchase.⁵ This is the fundamental threat of damage to your reputation. A business works to make its customers happy, because it knows its reputation will be damaged if it doesn't deliver. Customers, knowing this is true, are more willing to trust the business.⁶

Societal Dilemma: Cheating customers.	
Society: Group of merchants/society as a whole.	
Group interest: Merchants are trusted. Group norm: Don't cheat customers.	Competing interest: Maximize short-term profits. Corresponding defection: Cheat customers when possible.
To encourage people to act in the group interest, society implements these societal pressures: <ul style="list-style-type: none"> Moral: Guilt, shame, sense of fairness, kindness, etc. Reputational: Merchants want to be seen as trustworthy. Customers share their experiences with merchants, making merchants less likely to cheat customers so as to retain their good reputation. 	

Customer reputation used to be a bigger deal than it is today. When commerce was transacted entirely in local markets with local merchants, in situations where buyers and sellers knew each other and knew they would need to do business with each other many times in the future, reputation mattered just as much to the customer as it did to the merchant. In today's world of global commerce, where potential customers may be located a half a world away, customer reputation matters much less than merchant reputation. Means for ascertaining the integrity of potential customers are coming back, though. Online reputation systems, like eBay's feedback mechanism, gave both merchants and customers [reputation information](#) about each other. (In 2008, [eBay changed this](#), and no longer allows merchants to give feedback on customers, citing abuse of the process by merchants.)

We take our reputations very seriously, and spend a lot of time and effort maintaining them, sometimes defending them to the point of death.⁷ We go to these extremes because we recognize that if we want others to trust us and cooperate with us, we need a good reputation. So we keep our reputation clean, cover up blemishes, or fake our reputation completely.

Tellingly, psychological and brain research both show that we [remember negative](#) information about people more vividly, with more detail, and for a longer time than positive information. It seems that knowing who will defect is more important than knowing whom to trust.

We're also good at keeping up with the reputation of others. There's a [theory that gossip](#) originated as a mechanism for learning about the reputations of others and helping us know whom to trust. Of course, gossip requires language. Humans are unique on the planet for our ability to gossip,⁸ and humans everywhere on the planet are enthusiastic about it. It tells us who is likely to be cooperative and who is not, so we know [whom to interact](#) with. It helps [establish group interests](#) and group norms. It works as a societal pressure system, too; both observational studies and experiments show that [gossip helps keep](#) people in line. Social networking sites are the most modern manifestations of these ancient needs.

Reputation is a common mechanism for raising the costs of defecting and increasing the benefits of cooperating. Buskers generally don't disrupt each other's acts because they don't want a bad reputation among their peers. [Diamond merchants](#) generally pay their debts promptly, don't pocket other people's

diamonds, and don't substitute worthless stones for valuable ones, because they don't want to jeopardize their reputation within the community. And prisoners sometimes won't testify against each other because they don't want to be known as stool pigeons by the other criminals in town.

Here's how finely tuned we are to others watching our actions. The coffee room at the Division of Psychology at the University of Newcastle in Australia works on the honor system, just like Feldman's bagel business. Researchers found that if they put a sign above the pay box with a picture of a [pair of eyes](#)—not an entire face, just a pair of eyes—people put almost three times as much money in the box as they did when the sign had an image of flowers. Similarly, children who were told to take only one piece of Halloween candy but were left alone with a full bowl defected less when the bowl was placed in front of a mirror. And they defected even less when they were asked their [names and addresses](#) before being given the same opportunity. Along the same lines, religion often provides a universal observer. God is omniscient and the arbiter of one's final reputation, and a calculating believer behaves accordingly.⁹

Not only do we guard our reputation against blemishes, we also take pains to advertise our good reputation. This can be as grandiose as a company touting its customer satisfaction ratings or product quality awards, or as mundane as those small "I Voted" stickers that many polling places in the United States give to voters to wear for the rest of Election Day. The effect is both reputational and moral; voters can publicly demonstrate that they behaved in the group interest and voted, and simultaneously remind others of their civic responsibility. There's even a German expression, "*Tu Gutes und rede darüber*": "Do good and talk about it."¹⁰

We need several more pieces to make a reputational pressure system work. We don't always have perfect information about what other people are doing. Maybe they cooperated when we thought they defected, or vice versa. Or they might have done the wrong thing accidentally or because they weren't thinking clearly. Our reputational systems have to work despite the occasional mistake. This requires two things: contrition and forgiveness. If you defect by accident, apologize, make amends, and then return to cooperating. And if someone does that to you, forgive and return to cooperating.

I'm glossing over a lot of subtleties here. Forgiveness is a complicated emotion, and there's a fine line between being forgiving and being a sucker, and between being contrite and being a doormat. [There's a great word](#) from the Tshiluba language, spoken in southeastern Democratic Republic of Congo, that regularly appears on impossible-to-translate word lists. *Ilunga* means someone who forgives any abuse the first time it occurs, tolerates it the second time, and

neither forgives nor tolerates it the third time. The English saying is snappier: “Fool me once, shame on you; fool me twice, shame on me.”



Throughout most of history, commerce was a local phenomenon. Reputation made it work, and reputation was local. The emergence of long-distance commerce in the Western world was aided in great part by the involvement of European Quakers, who earned a reputation for dealing honorably with their business partners. Prior to the mid-17th century, European traders ran a significant risk that trading partners from other countries would act in their own self-interest and renege on promises they had made; overseas contracts were often unenforceable, so the potential for profit often outweighed the likelihood of punishment. However, the Quakers’ religious commitment to integrity and simple living, and their belief in the essential worth of every individual, informed all of their business dealings. Being upright with God was more important to them than making a fast buck. The moral benefit they experienced from acting in accord with their consciences, and the ensuing reputational benefits within both their religious and business communities, outweighed any short-term financial gains that might have come from shady dealing. A Quaker found to have dealt with others dishonestly ran the risk not only of losing business opportunities, but of being expelled from his religious community. As a result, [Quakers would cooperate](#) even if it went against their self-interest, and—as they consolidated their positions in industry—there was a gradual increase of trust in them among overseas traders.

The Quakers were an exception. The problem with reputation is that it doesn’t naturally scale well. Recall Dunbar’s numbers. We can recognize 1,500 faces, but the number of people we know enough about to know their reputation is much lower—maybe 500 or even 150. Once our societies get larger than that, we need other mechanisms by which to infer reputation than direct knowledge of the other person. And, as you’d expect, we have developed several of these.

One of the ways to scale reputation is to generalize based on group membership. So we might believe that people with a particular skin color, or who speak a particular language, or who worship a particular God, are untrustworthy. We might believe that a Quaker is trustworthy.

During the mid-17th century, being a Quaker meant something to the general community. In different periods of history, so did being a Freemason, or a member of the Medici family. In the 12th century, you could take a Templar letter

of credit issued in England all the way to Jerusalem. In the 11th century, the [Maghribi traders](#) of the medieval Mediterranean had a reputation similar to the Quakers. A thousand years earlier, Roman letters of introduction were similarly trusted throughout the empire.

Political scientist [Robert Putnam](#) has argued that mistrust increases in a community as ethnic diversity increases. Evidence of this effect comes from sources as diverse as studies of carpooling, Peruvian micro-credit cooperatives, and Civil War deserters. Even worse, this inherent mistrust of those in other ethnic groups isn't offset by an increase in trust of those in one's own ethnic group; trust across the board weakens in more ethnically diverse communities.¹¹

So it should come as no surprise that we have an enormous number of membership markers that we use to determine who is like us: language, dress, ethnicity, gang tags, haircuts, tattoos, jewelry, T-shirt slogans, food choices, gestures, secret handshakes, turns of phrase in speech, formal membership credentials, and so on. We generalize based on profession, city of residence, political affiliation, religious affiliation, sexual orientation, interests, and pretty much any other category you can think of. The theory is that all of these are [vestigial remnants](#) of prehistoric kin recognition mechanisms. But while these might have worked better in our evolutionary past than they do today, our brains are still stuck on them.

Take appearance, for example. Numerous experiments indicate that we are more likely to trust [people who look like](#) us. The phenomenon goes well beyond race; experimenters have digitally manipulated images of faces to more or less resemble those of their subjects and found that a variety of prosocial behaviors are correlated with facial similarity.

Dialect is a particularly interesting marker of group membership. With the nationalization and globalization of mass media, both accents and dialects are fading, but for most of human history, they were localized.¹² They're hard to fake, unless you're a rare gifted mimic, and they're generally set by adolescence. There is a lot of evidence, worldwide, that people are predisposed to cooperate with someone who speaks the same dialect they do. For instance, in one experiment, subjects were more likely to trust people with the [same accent](#) they had.¹³ And we naturally [change our patterns](#) of speech and body language to mimic those around us, unconsciously trying to fit into the group.¹⁴ Of course, the flip side of this is that we're less likely to trust people who [don't sound like](#) we do. Again, dialect preference seems to be a vestigial kin-recognition system.

It's worth noting that membership markers are harder to acquire and fake for groups that involve long-term—even inter-generational—cooperation and trust,

than they are for groups involving more near-term cooperation and trust. It's much easier to learn the knowledge and skills to be a member of the community of football fans, or stamp collectors, or a particular church, than it is to acquire a new facial feature like an epicanthic fold or a dialect.

Gauging reputation by group membership is a lousy way to prejudice someone—another name for the practice is “stereotyping.” But it's not an unreasonable cognitive shortcut, given our inability to interact meaningfully with more than 150 people, or even to put names to more than 1,500 faces. Historically, as the number of people we interacted with grew, we had to develop these shortcuts. Identifying someone as a member of a particular community, whether an ethnic community or a community of choice such as a professional association, gives us some indication about whether she is likely to cooperate with us or defect. If she's a member of the same community as us, we know she's likely to share the same set of ethical rules we do.

Recall the Golden Rule. It's not enough to want to cooperate. You also need to know **how to cooperate** according to your society's particular definition, so others can know you're reliably cooperative. One popular business-success book tried to “improve” on the Golden Rule, creating what it called the **Platinum Rule**: do unto others as they would want you to do unto them. That sounds even more altruistic, but it's not what has been encoded in our brains. Figuring out what someone else wants is easy to get wrong. It's much easier to assume that another person wants what you want. Of course, that works best if you only deal with people who are like you, and are likely to want the same things you want.

Social norms tell us how to cooperate. This is one of the reasons societies have tended to be homogeneous in their morals: it's advantageous. When people with different morals interact, they may have different default assumptions about what it means to cooperate. Remember the Machiguenga tribesmen in Chapter 7? They use a different definition of “fair” than Westerners do. Cooperation works better if we all agree on what it means to cooperate.¹⁵

Of course, just because our brains are hard-wired for this sort of in-group/out-group division doesn't mean it's the right thing to do.¹⁶ There are all sorts of reasons why stereotyping is a bad system for judging individual people, and for those reasons we should strive to get beyond our more base instincts.

A substitute that can help **reputation scale** is commitment. By committing ourselves to an action in a way that we cannot undo, we can make up for a lack of reputation. Consider the coordination problem between a prostitute and a prospective client. The two have met in a bar and have agreed to meet upstairs in his hotel room later in exchange for \$100. She wants him to pay her

in advance, because she doesn't trust that he will pay her in his hotel room as promised. Similarly, he is concerned that, once having received the money, she won't follow through and meet him later. If the two could trust each other, this would be easy to solve. But they don't.

One solution is to tear the \$100 bill in half, one piece for each of them. In the U.S. at least, half a \$100 bill has no value, so neither party has the money. Now both parties have effectively committed to the rendezvous: so she can receive the other half of the \$100 bill and he can receive the service. If either one of them defects and misses the meeting, neither gets the money.¹⁷ eBay escrow services serve the same function; they facilitate trust by forcing the buyer and seller into a commitment they can't get out of easily.

A similar mechanism is to deliberately cut off your escape routes, so you have no choice but to follow through on your commitment. This could mean literally burning your bridges behind you. In 1519, when Hernán Cortés invaded what today is Veracruz, Mexico, he scuttled the ships he arrived on, signaling to both the Aztecs waiting for him and his own men that there would be no reneging on his commitment.

A second way to demonstrate commitment is to move in steps. When I hired a contractor to perform renovation work on my home, the contract stipulated several partial payments at different milestones during the project. This step-by-step approach—me paying the contractor partially, him doing some of the work, me paying some more, him doing some more work, etc.—helped both of us trust each other during the entire project because the severity of defection was lessened.

This was also the general idea behind the Cold War doctrine of mutually assured destruction. Both the U.S. and the USSR worked to convince the other that they were committed to massive retaliation in the event of a first strike. The result was that neither side was willing to use nuclear weapons; the two countries might not have trusted each other in general, but they both trusted that the other side was crazy enough to follow through on its commitment.

A third way to signal commitment is ritual. This could be a handshake to seal a commercial deal, a ceremony to seal a marriage, or an Eagle Scout induction ceremony. Rituals work because 1) reputation is at stake, and 2) society provides sanctions against anyone who reneges. Of course, these only work if everyone understands what the ritual is and what it means.

Zahavi's handicap signals from Chapter 3 are another way to scale reputation: costly and hard-to-fake demonstrations of our reputation. These include publicly attending religious services to demonstrate our morality,¹⁸ ostentatiously

spending money to demonstrate our social class, and engaging in particular activities to demonstrate our political or cultural proclivities.¹⁹ Nobilities have complex displays of etiquette. Banks spend some of their money on imposing buildings to show off their financial health. **Criminals have signals**, too, to advertise their “good” reputation as a career criminal: prison time (that fellow criminals vouch for), tattoos, and deliberate physical self-harm.

Branding is yet another way to make reputation scale, similar to group membership. In many cases, we interact with organizations as groups rather than as individuals. That is, the corporate reputation of McDonald’s is more important to our decision about whether or not to trust it than the individual reputation of any of the stores or the individual employees.

Branding isn’t necessarily about quality; it’s about sameness. Chain restaurants don’t necessarily promise the best food, they promise consistency in all of their restaurants. So when you sit down at a McDonald’s or a Cheesecake Factory, you know what you’re going to get and how much you’re going to pay for it. Their reputation reduces uncertainty.

Advertising can be about persuading consumers to associate a certain brand with a certain reputation. **Shared brand names** serve as means of aggregating individual reputations into an overarching group reputation, which—if it’s maintained in good standing—benefits all members of the coalition. Companies call attention to their age, their size, and the quality of their products and services, all in an effort to enhance their reputation. Witness the ubiquity of advertising boasting about firms’ positions on environmental and workplace issues or contributions to worthy causes. Or the effort the principals of the **Saudi Binladin Group** construction company have spent trying to differentiate themselves from their terrorist relative.

In ascertaining quality, consumers will often rely on the cognitive shortcut provided by a brand name, and will even **pay a premium** for products with brand names they associate with a reputation for quality. One study of Bordeaux wines found that customers will pay a premium for bottles from a **more reputable producer**, even if the wine is no better. Notions of branding have leaked into individual reputation as well. Career counselors now advise professionals to “**cultivate their brand.**”

A final way to make reputation scale is to systemize it, so that instead of having to trust a person or company, we can trust the system. A professional police force and judiciary means that you don’t have to trust individual policemen, you can trust the criminal justice system. A credit bureau means that lenders don’t have to decide whether or not to trust individual borrowers, they can trust the

credit rating system. A credit card relieves merchants from having to figure out whether a particular customer is able to pay later; the system does that work for them. Dunbar's number tells us there is a limit to the number of individuals we can know well enough to decide whether or not to trust; a single trust decision about a system can serve as a proxy for millions of individual trust decisions.

We have a lot of experience with this kind of thing online: ratings of sellers on eBay, reviews of restaurants on sites like Yelp, reviews of contractors on sites like Angie's List, reviews of doctors, accountants, travel agencies...pretty much everything you can think of. Social networking sites systemize reputation, showing us whom we might want to trust because we have friends in common.

This is an enormous development in societal pressure, one that has allowed society to scale globally. It used to be that companies could ignore the complaints of a smallish portion of their customers, because their advertising outweighed the word-of-mouth reputational harm. But on the Internet, this isn't necessarily true. [A small complaint](#) that goes viral can have an enormous effect on a company's reputation.

On the other hand, while these reputational systems have been an enormous success, they have brought with them a new type of trust failure. Because potential defectors can now attack the reputational systems, they have to be secured. We'll talk about this in Chapter 10.



Reputation [isn't an effective](#) societal pressure system unless it has consequences, and we both reward cooperators and punish defectors.

We reward cooperators all the time incidentally through our actions. We choose to do business with merchants who have proven to be trustworthy. We spend time with people who have demonstrated that they're trustworthy. We try to hire employees who have good reputations, and we promote and give bonuses to employees who cooperate. From a security perspective, friendships are mutual reward systems for cooperating.

The common thread in all of these rewards is participation. Humans are a social species, and we reward by allowing others to participate in the group: whatever it is doing, whatever benefits it is accruing, whatever status and credibility it has achieved. Our brains are hard-wired to need to participate; we crave the approval of the group.

We also punish defectors. And if participation is the canonical reward, exclusion is the corresponding punishment. In our evolutionary past, the most severe

punishment was banishment from the group. As interdependent as humans were, this punishment was tantamount to death.

We still banish people today. We tell them we're no longer their friends and that they shouldn't come around anymore. We cut all ties with certain relatives, kick trolls out of online communities, and unfriend people on Facebook. On a different scale, someone with a destroyed credit rating is pretty much banished from the lending community.

Other punishments are less severe: physical violence, property damage, and so on. Sometimes we call this sort of thing "revenge." Here's how [Maine lobstermen](#) deal with one of their group violating traditional territories:

Ordinarily, repeated violation of territorial boundaries will lead to destruction of the offender's gear. It is usual for one man operating completely on his own to first warn an interloper. In some places this is done by tying two half hitches around the spindle of the offending buoys; in other places by damaging the traps slightly. At this point, most intruders will move their traps. If they are not moved, they will be "cut off." This means cutting off the buoy and warp line from the trap, which then sinks to the bottom where the owner has no chance of finding it.... A man who violates a boundary is ordinarily never verbally confronted with the fact of his intrusion. And the man who destroys his gear will traditionally never admit to it.²⁰

Most punishments are even less extreme. We may still hang around with some friends, but not rely on them as much or not tell them our intimate secrets. We may still invite *those* relatives to the family's holiday party, but not talk to them much.

[Shame is a common](#) reputational punishment, and—as a result—an important social emotion. Much of this is the informal kind of shaming we've all experienced amongst our friends and colleagues. More formal examples include police blotter reports, IRS quarterly listings of Americans who renounce their citizenship, public disclosure of [excessive CEO pay](#), televised arrests, deadbeat dads in the media, and TV shows like *America's Most Wanted*. Of course, these all have a technological component, and some might be more properly put into the category of institutional pressure.

Informal punishments are so common we can miss them if we're not paying attention. An employee who has frequent conflicts with his colleagues may find himself shuffled into a dead-end position, or assigned the graveyard shift. A husband spied flirting with the housekeeper may be told by his friends and family that such behavior is unacceptable. An entertainer who espouses unpopular

political positions may encounter a dip in his popularity at the box office as moviegoers boycott his films.

Informal punishments are common in our society. They're certainly prevalent through childhood, from early play amongst small children through social ostracism in school. Some groups define themselves through the exclusion of others.

Remember the Bad Apple Effect from the previous chapter? As you might expect, the effects of those **bad apples diminishes** if punishment is threatened.

There is an old idea that punishment can be transferred from one person to another. In many traditions, God punishes a person's relatives in addition to punishing the transgressor.²¹ In some societies, **if Alice kills Bob**, one of Bob's relatives is allowed to kill one of Alice's relatives in retribution. The **Nazis instituted this** as government policy; it was called "*sippenhaft*." This practice is a form of societal pressure. If Alice is considering defecting from some group norm—killing another person, committing adultery, whatever—she not only has to worry about human or divine retribution against her personally, but also retribution against members of her family. And threatening Alice's family should she defect both raises her perceived cost of defection and enlists her family members in persuading her not to defect in the first place. The Israeli government's current practice of bulldozing the homes of suicide bombers' families is an example. Of course, sometimes this goes very badly. Think of "honor killings" of rape victims, or blood feuds in **various cultures** throughout history, like the **Hatfields versus the McCoys**.



There's a variant of the Hawk-Dove game that demonstrates how **reputation can solve** a societal dilemma. It's designed so doves are more likely to interact with doves. When this happens, hawks can be isolated and their numbers reduced.

Compared with the basic Hawk-Dove game, cooperation turns out to be an even better strategy. Stable populations have even fewer hawks because doves, by preferring to interact with other doves, can effectively isolate them. Left to fight amongst themselves, hawks tend to kill each other off. Returning to human society, we are at our most cooperative when we seek out other cooperative people and avoid those who would take advantage of us. We learned this in Chapter 3 when we looked at the evolution of cooperation: cooperators do better when they can recognize each other. Reputation not only encourages cooperation, but also marginalizes defectors to the point where there ends up being fewer of them to deal with.

This is important. We've been talking about societal dilemmas as if they're always decisions to either cooperate or defect. In the real world, we often have a choice of people with whom to interact. We don't walk into stores randomly, wondering if the merchant will cheat us. We only walk into stores where we believe the merchant will not cheat us. Instead of defecting and cheating the merchant as punishment, we prefer to shop elsewhere.



In Chapter 3, we learned that two things are required for cooperation: reciprocal altruism and a calculating intelligence. Morals and reputation, the two things I've been calling our primitive toolbox of social pressures, provide that reciprocal altruism. Even so, reputational societal pressure can fail in many ways.

Defectors take steps to hide facts that can harm their reputation, or manipulate facts to help their reputation. Recall that in the mid-1970s, [John Wayne Gacy](#) managed to rape and kill 33 young men. All the while, his Chicago neighbors and colleagues on civic and charitable committees never suspected “Pogo the Clown” Gacy was involved in any work more diabolical than entertaining children for good causes. In the UK, [Dr. Harold Shipman](#) had a similar story. Described as “a pillar of the community” by his neighbors, he killed at least 250 people, mostly elderly widows, before he was caught. Most examples are less extreme. A politician might go to church and publicly pray, to encourage people to think he's honest: a whited sepulchre. An American trekking through Europe might sew a Canadian flag on his backpack.

Confidence tricksters spend a lot of time manipulating reputation signals. They employ all sorts of props, façades, and other actors—shills—to convince their victims that they have a good reputation by appearing authentic, building confidence, and encouraging trust. Corporations and political candidates both do similar things; they use paid supporters to deliberately spread artificial reputational information about them. This is becoming even more prevalent and effective on the Internet. Hired hands write fake blog posts, blog comments, tweets, Facebook comments, and so on. Scammers on eBay create fake feedback, giving themselves a better reputation. There are even companies that will give you fake Facebook friends, making you seem more popular with attractive people than you actually are.

Defectors try to minimize the effects of their bad reputation. People get new friends, move to another city, or—in extreme cases—change their names, get plastic surgery, or steal someone else's identity. Philip Morris renamed itself

Altria, because who would want to buy their Kraft Mac and Cheese from a cigarette company? ValuJet, its brand ruined after Flight 592 crashed in the Everglades in 1996, now operates as AirTran Airways. Blackwater, the defense contractor notorious for numerous Iraq war abuses, became Xe Services and then Academi. The School of the Americas, implicated in training many human rights-abusing military staff in Latin America, rebranded itself as the Western Hemisphere Institute for Security Cooperation.

Corporations work to minimize the effects of negative reputation on their brands through advertising and public relations. Multinational food and consumer product companies like Unilever and Procter & Gamble deliberately downplay their corporate brand and focus attention on their sub-brands. There are a lot of marketing reasons to engage in sub-branding, but the security thinking is that if there is a problem with one of their brands, the negative publicity won't spread across the company's entire product line.

Some people simply don't care about reputation. Like our individual morals, our individual concern about reputation varies—from person to person as well as from situation to situation. Some of us care a lot; others, not so much. Of course, this is contextual. We all have different reputations in different groups with respect to different personal attributes.

Some people end up with the wrong reputation. Even if someone does nothing wrong, there's no guarantee that his reputation is accurate. Untrue stories can circulate by mistake. Someone else might lie to give him a bad reputation. We all know people who have reputations they don't deserve, both good and bad.

Defectors band together in subgroups that have different reputational rules. Gang members thrive in groups. Sure, they have a terrible reputation in the broader community, but they care primarily about their reputation within their gangs. This dynamic is also true for defectors who have a different moral system from the dominant culture: a lone pot smoker in a pot-free community is going to have a lot harder time than one who finds other pot smokers in the vicinity. His friends will help him defect. In effect, he will choose to cooperate with the smaller society of defectors, rather than with the pot-abstaining majority. The same is true for those worshipping in secret out of fear: early Christians in the Roman Empire, pagans afterwards, Jews in post-expulsion Spain, devout Russian Orthodox in the former Soviet Union. We'll talk about this more in Chapters 11 and 12.

The value of defecting might be worth the reputational damage. Maybe it's a single large transaction, and the merchant is willing to sacrifice her reputation for

the money. Or maybe it's a situation where the merchant can outrun his reputation.²² We've all heard stories of home remodeling contractors that score a big contract, and then either don't do the work or do a quick, shoddy job, and disappear with the money. "Fly by night," it's called. They've made the risk trade-off and decided their reputation wasn't that valuable. If they're career scammers, a big payoff may even enhance their reputation among their fellow scammers. A restaurant owner in a tourist area could serve lousy food, confident that the reputational damage matters less when there's no repeat clientele. A corporate CEO might decide his company's ability to repair reputational damage allows it to get away with misdeeds he wouldn't have authorized if he didn't have such an effective public relations department.

*The most important reason reputational pressure starts to fail is that groups get too large.*²³ Assisted by technology, reputational pressure can scale globally. Think of the reputations of public figures and celebrities, companies and brands, or individuals on the Internet. Think of eBay's reputation system, review sites like Yelp, or how we can make friends on shared-interest websites. Think of the FBI's criminal databases, the information about you kept by credit bureaus, or Google's database of your interests. Think of passports, driver's licenses, or employee badges. These are all reputational systems, and all serve to apply reputational pressure in different risk trade-offs.

But these systems can have all sorts of inaccuracies. What we know about celebrities, corporations, and people in faraway places doesn't always match reality. It's not only the natural errors that creep into any large-scale process, it's that these systems can be manipulated and the technologies used to support them can be attacked. In order for reputation to scale, we need to trust these reputational systems, but sometimes that trust is not well-founded. We'll talk about this more in Chapter 10.

Reputational pressure works best within a group of people who know each other: a group of friends or coworkers in an office, compared to a bunch of strangers on a bus or a city full of people. Neighbors are good at [settling disputes](#); people who don't live so close to each other are less good at it.

However, once the group size grows larger and the social ties between people weaken, reputation alone doesn't cut it.

Commenting on Hardin's original Tragedy of the Commons paper, psychologist Julian Edney wrote that "the upper limit for a simple, self-contained, sustaining, well-functioning commons may be as low as 150 people."²⁴

Eleven years later, [Dunbar wrote](#):

The Hutterites, a group of contemporary North American religious fundamentalists who live and farm communally, regard 150 as the maximum size for their communities. What is interesting is the reason they give for splitting communities at this size. They find that when there are more than about 150 individuals, they cannot control the behaviour of members by peer pressure alone.

Commenting on the Hutterites, [Hardin suggested](#), “[Perhaps](#) we should say a community below 150 really is managed—managed by conscience.”

I read somewhere once that police officers represent a failure of the underlying social system.²⁵ The social system should be self-policing, and formal rules and rule enforcement should not be required. But it’s not self-policing, and not just because we’re wary of vigilantism. It’s simply a natural effect of increasing the scale of the underlying social group.



Institutional Pressures

Store owners generally get to set their own hours. If their customers tend to shop early, the store opens early. If their customers tend to sleep in, the store doesn't open until late morning. Nights, weekends, holidays: a smart store owner is going to match his store's hours to his customers' needs.

This isn't true of stores in a shopping mall. Shopping malls have preset hours, and if you have a store in the mall, you have to adhere to those hours. It doesn't matter who your customers are. Called a "continuous operations clause," it's written into most mall leases.

This solves a societal dilemma: stores are individually better off if they can set their hours to suit their business, but the stores are collectively better off if everyone shares the same hours so customers know that everything will be open when they go. To ensure that stores follow the group interest, mall operators enforce continuous operations clauses through steep fines.

Societal Dilemma: Mall hours.	
Society: Group of merchants.	
Group interest: Mall stores all have uniform hours.	Competing interest: Maximize short-term profits.
Group norm: Stay open during agreed upon hours.	Corresponding defection: Open and close when it makes financial sense.
To encourage people to act in the group interest, society implements these societal pressures:	
Institutional: The group fines stores that close during common hours.	

As we saw in the previous chapter, solving a Prisoner's Dilemma involves changing the costs and benefits of acting in the person's selfish interest versus acting in the group interest. Shopping malls solve their Prisoner's Dilemma by using fines. A fine raises the cost of a store owner acting in his self-interest. Raise that cost high enough, and owners will open and close their stores in unison.

The common mall hours, and the fines for violating them, are an example of an institutional societal pressure. It's a rule established by the institution that owns the mall—it might even be a cooperative institution consisting of all the stores—that the society of store owners all agree to.

Political philosophers have long argued that informal societal pressures aren't enough for a successful human society. Thomas Hobbes, writing in the mid-17th century, believed individuals couldn't be trusted, and the opportunities to defect were simply too tempting. In this "state of nature"—that's anarchy, although he never used the word—our lives would be "solitary, poor, nasty, brutish, and short." [Martin Luther](#) said the same thing, a century earlier.

[Immanuel Kant](#) put it this way at the end of the 18th century:

The problem of organizing a state, however hard it may seem, can be solved even for a race of devils, if only they are intelligent. The problem is: "Given a multitude of rational beings requiring universal laws for their preservation, but each of whom is secretly inclined to exempt himself from them, to establish a constitution in such a way that, although their private intentions conflict, they check each other, with the result that their public conduct is the same as if they had no such intentions."

The result is Social Contract Theory, which posits that people willingly grant government power that compels them to subordinate their immediate self-interest to the long-term group interest in order to protect themselves and their fellow citizens from harm. Thomas Hobbes, [John Locke](#), [Jean-Jacques Rousseau](#), and the 20th-century philosopher [John Rawls](#) all proposed different flavors of this idea. Their conclusions about the ideal way to achieve social order vary, but all maintain that it is both necessary and moral to forcibly limit individual freedoms, reasoning that without a government enforcing laws, defectors would take over, to the detriment of all.

At its basest form, it's an argument we've seen in the previous chapter: fear of punishment is what keeps the tempted honest. In [Plato's Republic](#), Glaucon argues that if you remove that fear, the righteous will behave no differently than the wicked: "Mankind censure injustice, fearing that they may be the victims of it and not because they shrink from committing it." During the Italian

Renaissance, [Niccolò Machiavelli](#) built an entire political philosophy around this principle.

Men never act well except through necessity: but where choice abounds and where license may be used, everything is quickly filled with confusion and disorder. It is said therefore that Hunger and Poverty make men industrious, and Laws make them good.

Of course, that's not precisely true. The righteous aren't really just calculating scoundrels behaving well only because they fear that someone else—or perhaps God—will punish them if they step out of line. Reciprocal altruism works, and most people are honest most of the time. It's the defectors that Machiavelli was talking about, and for them he got it mostly right.

Laws, regulations, and rules in general are all institutional societal pressures. They're similar to reputational pressure, but formalized. We all agree to comply with all sorts of institutional pressures as a precondition of being part of a group, the most common of which are the laws by which we agree to be bound as a condition of being part of whatever political units we're part of. (It's certainly debatable whether individuals “agree to be bound by” all of the rules that end up being applicable to them, but that's generally how political philosophers look at it.)

It's not always clear exactly when informal social mores become rules. The social pathologists make a distinction between codified and explicit norms established by the government and non-formal norms agreed upon by the group, but that leaves a large grey area for less-official groups. Still, codifying our reputational pressure into laws was a big step for the development of society, and it allowed larger and more complex social groupings—like cities.

[Garrett Hardin](#), who created the phrase “the Tragedy of the Commons,” later wished he'd called it “the tragedy of the unmanaged commons.” The point of his paper was not that defectors will inevitably ruin things for the group, but that unless things are managed properly, they will. He was stressing the need for institutional pressure.

Institutional pressure requires an institution for implementation and enforcement; I mean the term very broadly. Institutions include governments of all sizes, but also religious institutions, corporations, criminal organizations, and so on. These institutions implement rules, laws, edicts—there are several terms—and sanctions for disobeying them and possibly incentives to obey them.

Burglary has costs that exceed the value of the goods stolen. Burglary costs in the time and effort to replace what's been stolen, the psychological effect

of having one’s home violated, the cost to the community of investigating the crime and prosecuting the accused, and even the cost of defending the suspect if he happens to be indigent. Sometimes the costs to the burglarized far exceed the value to the burglar: think of someone who steals copper wire out of a data center to sell as scrap metal, or destroys a building to get at valuables inside. But these costs are not borne by the burglar. They are externalities to him.

A well-written law combined with proper enforcement raises the costs to the burglar to the point where he is forced to bear the full costs of his actions. It could even raise the costs to the point where breaking, entering, and stealing is a worse trade-off than buying the same things legitimately.

Voting is another example. In the U.S., voter turnout is so low in part because there’s no legal requirement to vote. In countries where **voting is required** by law—Australia, Belgium, Bolivia, etc.—turnout is much higher. This is also true in countries that don’t have explicit voting laws, but have laws that raise the cost of not voting in other ways. For example, in Greece, it’s harder for non-voters to get a passport or driver’s license. If you don’t vote in Singapore, you’re removed from the electoral rolls and must provide a reason when you reapply. In Peru, your stamped voting card is necessary to obtain some government services. And in Mexico and Italy, there are informal consequences of not voting, harking back to the previous chapter. These “innocuous sanctions,” as they’re called in Italy, make it—for example—harder to get day care for your child.

Deacon’s Paradox is another example. The societal dilemma looks like this:

Societal Dilemma: Respecting pair bonds.	
Society: Society as a whole.	
Group interest: Everyone trusts each other enough to go about daily tasks away from their long-term partners. Group norm: Respect each other’s pair bonds.	Competing interest: Maximize personal pleasure, maximize gene propagation. Corresponding defection: Have sex with whomever you want.
To encourage people to act in the group interest, society implements these societal pressures: Moral: Teaching that adultery is wrong. The occasional commandment. Reputational: Public shaming of people who break their marriage vows. Institutional: Legal marriage contracts. Adultery laws.	

Initially, [marriage rites were informal](#) and reputational; both religious and civil institutions formalized them as we developed rules about property and inheritance. Of course, this isn't perfect. Philandering is as old as human society; rules are generally only selectively enforced, and friends of a philanderer will always be tempted to look the other way. But formalized marriage rules have been in effect throughout history, and they're largely effective.

Gridlock is another example. If you've ever driven in a crowded city center, you know the problem. Drivers stay as close as possible to the car in front of them, so no one will be able to cut in front of them and they will get where they're going as quickly as possible. The inevitable result of this strategy is that cars get stuck in the middle of an intersection when the light turns red, and cars going the other way can't pass. This is both inconvenient and a danger to public safety as emergency vehicles become unable to pass through the congestion. In extreme cases, gridlock can tie up traffic for hours. Everyone would do better if no one entered the intersection until the car was able to completely clear it on the other side, but unless everyone shows restraint, those who do are penalized. The solution: in many cities, it's now illegal to enter an intersection if you are unable to pass completely through without blocking cross-traffic.

Some societal dilemmas are particularly resistant to institutional pressure. Kidnapping and piracy are two examples. The dilemma is obvious. Kidnapping and piracy are bad for society, so paying ransom is bad because it makes these crimes profitable and emboldens those who commit them. Nonetheless, each and every one of us wants an exception to be made if we, our loved ones, or our cargo are held for ransom. So people follow their self-interest, their self-preservation interest, or their relational interest and pay up. This practice has made kidnapping profitable in many countries, most notably Mexico, Colombia, and Iraq, and has contributed to the escalation of piracy in Southeast Asia and off the coast of Somalia. All of these countries could pass laws making it illegal to pay kidnapping ransoms, but those would be hard to enforce. Both parties to these transactions want to hide them from the authorities. It's not enough to declare kidnapping illegal; enforcement matters, and most high-kidnapping and high-piracy countries have ineffective police forces at best, or corrupt police serving as [accomplices at worst](#). Piracy has an additional externality; the costs are not borne by the country that hosts the pirates.¹ In countries like the United States, harsh enforcement has made kidnapping for ransom a very rare crime, and piracy nearly nonexistent. In other countries, like Somalia, paying ransoms is common, even though the government [occasionally jails](#) those who do so.

Compare this to bribery. Like kidnapping, bribery of public officials is a societal dilemma. Society is much better off without bribery, but when individuals are faced with a recalcitrant government official, they can be easily motivated to ignore that and pay up. Where bribery is illegal for both the giver and the receiver, both parties have an incentive to hide the bribe from the police, which makes enforcement of anti-bribery laws difficult. (The fact that it's sometimes the police who have to be bribed makes it even worse.) India's chief economic advisor recently argued that, for some classes of bribes, [offering a bribe](#) should be decriminalized. The rationale here is that if the bribe giver is not treated as a criminal, he will be more willing to help prosecute public employees who demand bribes. Of course, this only works for one-time bribes, where an official is demanding payment for a service that the recipient should normally receive. It doesn't work for bribes in which an official is being asked to do something he shouldn't normally do, or for a series of bribes over time. In all cases, the bribe payer would not want to make his actions public, regardless of the law. But in the more normal case of a government official trying to line his pockets through a one-off transaction, decriminalizing the bribe giver's actions would make it more likely for him to go public.²

Similarly, while it's bad policy to negotiate with terrorists, it's [easy to make exceptions](#). At the height of the IRA's bombing campaign in the UK, Prime Minister Thatcher was publicly affirming that her government would never negotiate with terrorists while at the same time conducting secret back-channel negotiations with senior IRA figures. This was in addition to the negotiations the non-militarist wing of the IRA was conducting with the British government.



Just like reputational pressure, institutional pressure requires consequences to work. The difference is that while reputational consequences are informal, institutional consequences are formal, codified, and tangible. These can be punishments, more properly called sanctions, or rewards, better called incentives.

Think back to Bob and Alice in their respective prison cells, making their own risk trade-offs. Not implicating others enhances the reputation of a criminal; additionally, criminal organizations hunt down and punish those who don't keep silent.

Societal Dilemma: Criminals testifying against each other.	
Society: The criminal organization.	
Group interest: Minimize the amount of punishment for the society.	Competing interest: Minimize personal punishment.
Group norm: Don't testify against each other.	Corresponding defection: Testify against each other in exchange for reduced punishment.
<p>To encourage people to act in the group interest, the society implements a variety of trust mechanisms.</p> <p style="padding-left: 40px;">Moral: People feel good when they support other members of their group, and bad when they let them down.</p> <p style="padding-left: 40px;">Reputational: Those who testify against their fellow criminals are shunned.</p> <p style="padding-left: 40px;">Institutional: The criminal organization severely punishes stool pigeons.</p>	

You could argue whether the criminal code of silence—and the practice of killing police informants—belongs in this chapter or the previous one. I suppose it depends on how formal the rules are. Certainly it goes far beyond shunning.

Sanctions serve several purposes. Modern penologists hold that prisons are primarily intended to reeducate and reform, minimizing recidivism. Financial sanctions serve as a penalty, raising the financial cost of defecting. And, unfortunately, both have an aspect of revenge about them—another formalization of reputational pressure.³ But the part of it that matters most for societal pressure is the deterrent effect. A rule or law will encourage some people to cooperate simply based on their innate moral tendency to obey authority and follow the rules, but primarily—like reputational systems—laws rely on punishment as a deterrent to defection. Unlike reputational systems, though, imposing sanctions is more formalized.⁴ This doesn't necessarily mean something that has been written down and agreed to like a legal code, although it generally is. Sanctions reduce the number of defections. And recalling the Bad Apple Effect from Chapter 7, they prevent further defections.

The general idea of such rewards is to formalize coercion. Even prohibitive laws have some aspect of this; they're prescriptive as well as punitive. They operate both before the decision about whether to cooperate or defect occurs, by providing guidelines for acceptable behavior and prior notification of any penalties, and afterwards, through enforcement.

Laws are only as good as society's ability to enforce them. It's not enough to pass a law requiring people to pay their taxes, or banning child labor, or limiting the amount of insect parts in your breakfast cereal; if you don't also sanction defectors, the laws will not act as much of a deterrent. Fines have to be assessed and collected. Jail time has to be served. And all of this has to be implemented with an eye towards solving the societal dilemma.

Alexander Hamilton said as much in *The Federalist 15*:

It is essential to the idea of a law, that it be attended with a sanction; or, in other words, a penalty or punishment for disobedience. If there be no penalty annexed to disobedience, the resolutions or commands which pretend to be laws will, in fact, amount to nothing more than advice or recommendation.

Sanctions fall into three basic categories: confiscation of resources or possessions, shaming, or physical penalties. Fines and forced servitude fall into the first category, and the last category further breaks down into incarceration, physical harm, and execution. Shaming and physical harm were more common historically; the stocks are a good example of both, as people restrained by them could be abused by the community. Sex offender registries are a common modern shaming sanction, but others—such as requiring an offender to stand in a public place wearing a sign that broadcasts the nature of his offense—are slowly making a comeback, in spite of persuasive arguments that they are *immoral*, *ineffective*, and degrade the public as much as those subjected to them. House arrest, monitored by an electronic bracelet, has a shaming aspect too. So does community service, if it's obvious and in public.

Most modern sanctions consist of either incarceration or financial penalties. Incarceration removes the defector from society for a period of time, and prevents him from committing further defections. Done right, jail is a place to reform criminals. Done wrong, jail is a place where criminals learn how to be better criminals.

Financial penalties can be tricky to implement, and are therefore worthy of a longer discussion. Speeding is a risk trade-off. There are risks to speeding—accidents—but there are also rewards, such as getting to your destination sooner or the adrenaline rush that comes with driving faster. There are all sorts of pathologies in the trade-off—the rewards are immediate and constant, but the risks are nebulous and only happen occasionally—and one might think there's no reason society can't just let people make the trade-offs by themselves.

The problem is that when Alice speeds, she also increases the risk to everyone around her.⁵ So there is a societal dilemma at work, and if you want Alice not to

speed you're going to have to make it illegal and penalize her for doing it. Studies show that [fines reduce speeding](#) overall, even though they don't deter habitual speeders. Drunk driving laws and their enforcement are a similar example.

Societal Dilemma: Speeding.	
Society: Society as a whole.	
Group interest: Minimize automobile deaths.	Competing interest: Minimize travel time.
Group norm: Obey speed limits.	Corresponding defection: Speed.
<p>To encourage people to act in the group interest, the society implements a variety of trust mechanisms.</p> <p style="padding-left: 40px;">Moral: It's moral to drive in a way that doesn't endanger others. Also, it's moral to follow the rules.</p> <p style="padding-left: 40px;">Reputational: There is some social pressure, in some circles, not to be known as a speeder or a reckless driver.</p> <p style="padding-left: 40px;">Institutional: Speed limits.</p>	

It's vital for the financial penalties to be high enough to make the behavior unprofitable. For example, if customs has a 10% chance of catching a smuggler, then the penalty for smuggling needs to be at least ten times the value of the goods—otherwise it would make financial sense to smuggle. One report demonstrated that [uninsured drivers](#) in the UK are capable of doing the math, and will remain uninsured if the expected penalty for doing so is less than the cost of insurance. This is even more important when dealing with corporations; there are many examples of fines being so small as to be viewed as an incidental cost of doing business. We'll talk about this more in Chapter 13.

Fixed financial penalties are regressive. Like everything else about the speeding trade-off, the cost of a speeding ticket is relative. If you're poor, a \$100 speeding ticket is a lot of money. If you're rich, it's a minor cost of the trip.⁶ Finland, Denmark, and Switzerland address the problem by [basing traffic fines](#) on the offender's income. Wealthy people in those countries have regularly been issued speeding tickets costing over \$100,000. You might disagree with the system as a matter of policy, but it certainly is a more broadly effective societal pressure. Jail time for speeders accomplishes much the same thing.

There are two basic ways the law can prescribe financial penalties. It can pass a direct law, or it can institutionalize liabilities. If the affected individuals can sue the defectors and win sufficient punitive damages, that will also increase the

cost of defecting. In both cases, laws remove the externality by making sure the defector pays the cost of defection. Watch how it works:

- *Overfishing.* Pass and enforce a law fining (or even jailing) those who overfish, and the dilemma goes away. Assuming the cost of the fine multiplied by the probability of getting caught—that’s “cost” defined broadly, in terms of money, jail time, social stigma, whatever—is greater than the value of the additional fish, it changes Alice’s risk trade-off.
- *Polluting the river.* Allow people living downstream from the polluter to sue. Assuming the court system works properly, the cost of the lawsuits to the polluter will be greater than the cost not to pollute the river.
- *Unsafe food handling.* Consumer protection laws raise the cost of ignoring food safety—presumably to save money—by imposing financial penalties on those who engage in it.

All this assumes a system where both the plaintiff and the defendant can afford the same quality of legal representation. The trade-off changes when the river polluters are corporations with deep pockets, and the people affected don’t have the means to pay for lawyers. Or when the bad behavior occurs when foreign companies import food into another country, and the probability of getting caught is low. Again, we’ll return to these considerations in Chapter 13.



Taxes can be another type of institutional pressure. It’s weird, because it doesn’t actually prohibit anything. But if the goal is to reduce the scope of defection, charging people for their marginal defection is one way to do it. Like fines, taxes increase the cost of defecting. But unlike fines, they operate during and not after the defection.⁷ For example, a sanction for littering requires the authorities to detect the crime and then assess the penalty. This happens after the littering occurs, and there’s always the chance of not getting caught. A tax on excess trash occurs at the time of trash pickup, although the person may pay the tax later.

The societal dilemma surrounding antibiotics mirrors the one surrounding vaccination: overuse vs. underuse. It’s in everyone’s immediate self-interest to use antibiotics to treat conditions that respond to them, but if they’re overused, bacteria develop resistance, making them ineffective for everyone. A big part of the problem is the wholesale use of antibiotics in agriculture: administering antibiotics to livestock in order to produce faster growth, regardless of whether

they are needed to treat disease. The problem of antibiotic-resistant supergerms is an externality. But doctors also contribute significantly to the problem; they frequently prescribe antibiotics in cases where they're not really necessary. But here again, use of antibiotics makes sense from the perspective of the doctor, who reasons that they won't hurt and might help the immediate patient, whose patients regularly ask for them, and to whom the larger social costs are an externality as well.

One solution, proposed by the Infectious Diseases Society of America, is to [tax the use of antibiotics](#). This is a societal pressure, increasing the cost of using antibiotics as a way to remove the externality. We can debate the effectiveness of the measure: it'll definitely help in agricultural uses, but how much it will reduce superfluous doctor prescriptions will depend on who pays the tax and how. Not to mention how easy it would be to smuggle in untaxed antibiotics.



The converse of penalties are incentives: rewarding someone for cooperating. There is a whole class of institutional pressure systems designed to reward cooperative behavior. Examples include:

- Tax deductions or tax credits for certain behaviors.
- Faster tax refunds for people who file their returns electronically.⁸
- Time off a prison sentence for good behavior.
- Employee bonuses.
- Bounties and rewards for turning in wanted fugitives.
- Deferred or non-prosecution of SEC violations as an incentive to provide evidence against other, larger, violators.
- Certifications, both coercive ones (FDA approval for new drugs) and optional ones (LEED certifications for buildings).
- Whistle-blower statutes, where the whistle-blower gets a percentage of the fraud found.

The problem with rewarding cooperators via an institutional mechanism is that it's expensive. If we assume that the majority will cooperate regardless of the reward, then a lot of people will get a reward for doing what they were going to do already. Either the reward will have to be very small and not much of an additional incentive to cooperate, or the total cost of rewarding everyone will be

very expensive. In general, it's more efficient to spend that money going after the minority of defectors.

Financial incentives and penalties interact weirdly with other categories of societal pressures. It's easy to regard societal pressures as cumulative—and to assume that moral plus institutional pressure will be more effective than morals alone—but our moral systems are more complicated than that.

In one experiment, participants were presented with a societal dilemma: they were in charge of a manufacturing plant that emitted toxic gas from its smokestacks. They could either spend more money to clean up a lot of the toxin, or spend less money to clean up a little bit of the toxin. The dilemma came from the fact that pending government legislation—a bad thing in the experiment's scenario—depended on how much cleaning up the manufacturing plants did collectively. It's a free-rider problem: a subject could either cooperate and clean up his share, or defect and hope enough others cleaned up enough to forestall legislation.

What makes this experiment particularly interesting is that half of the subjects were also told that the industry would be inspecting smokestacks to verify compliance and fining defectors. It wasn't a big risk; both the chance of inspection and the cost of noncompliance were low. Still, inspections are a societal pressure, and you'd expect they would have some positive effect on compliance rates. Unexpectedly, they had a negative effect: subjects were more likely to cooperate if there were no noncompliance fines than if there were. The addition of money made it a financial rather than a moral decision. Paradoxically, financial penalties intended to discourage harmful behavior can have the reverse effect.

For this reason, signs featuring anti-littering slogans like "Don't Mess with Texas" are more effective than signs that only warn, "Penalty for Littering: \$100"; and "smoking in hotel rooms is prohibited" signs are more effective than signs that read "\$250 cleaning penalty if you smoke." In one experiment with day care providers, researchers found that when they instituted a fine for parents picking their children up late, late pickups increased. The fine became a fee, which parents could decide to pay and assuage any moral resistance to defection.

More generally, the very existence of rules or laws can counter moral and reputational pressure. Some towns are experimenting with eliminating all traffic laws and signs. The idea is that drivers who must follow the rules pay less attention to the world around them than drivers with no rules to follow.

Financial rewards have the same effect that financial penalties do; they engage the brain's greed system and disengage the moral system. A fascinating incident in Switzerland illustrates this. Trying to figure out where to put a nuclear

waste dump, researchers polled residents of several small towns about how they would feel about it being located near them. This was 1993, and a lot of fear surrounded the issue; nonetheless, slightly more than half of the residents agreed to take the risk, for the good of the country.

In order to motivate the other half, the researchers offered money in exchange for siting the nuclear dump near them: about \$2,000 per person per year. Instead of enticing more residents to accept the dump, it reduced their number by half. The researchers doubled and then tripled the amount offered, but it didn't make a difference. When they simply asked nicely, the researchers stimulated the altruistic part of the residents' brains—and, in many cases, they decided it was the right thing to do. Again, the addition of money can increase the defection rate.⁹

Financial advisors exhibit this unconscious bias in favor of their clients. In one experiment, analysts [gave different weights](#) to the same information, depending on what the client wanted to hear. An obvious societal pressure system to address this problem would be to [require advisors](#) to disclose any conflicts of interest; but this can have the reverse effect of increasing the number of defectors. By disclosing their conflicts, financial advisors may feel they have been granted a moral license to pursue their own self-interest, and may feel partially absolved of their professional obligation to be objective.



Elinor Ostrom received a Nobel Prize in 2009 for studying how societies deal with Tragedies of the Commons: grazing rights in the Swiss Alps, fishing rights off the coast of Turkey, irrigation communities in the Philippines. She's studied commons around the world, and has a list of rules for successfully managing them.¹⁰ Generalizing them to our broad spectrum of societal dilemmas, they serve as a primer for effective institutional pressure:

1. Everyone must understand the group interest and know what the group norm is.
2. The group norm must be something that the group actually wants.
3. The group must be able to modify the norm.
4. Any institution delegated with enforcing the group norm must be accountable to the group, so it's effectively self-regulated. We'll discuss these institutions in Chapter 14.
5. The penalties for defecting must be commensurate with the seriousness of the defection.

6. The system for assessing penalties must be consistent, fair, efficient, and relatively cheap.
7. The group must be able to develop its own institutional pressure and not have it imposed from outside the group.
8. If there are larger groups and larger group interests, then the groups need to be scaled properly and nested in multiple layers—each operating along these same lines.

Ostrom's rules may very well be the closest model we have to our species' first successful set of institutional pressures. They're not imposed from above; they grow organically from the group. Societies of resource users are [able to self-regulate](#) if they follow these rules, and that self-regulation is stable over the long term. It's generally when [outsiders come in](#) and institutionalize a resource-management system that things start to fail.

I mentioned institutional pressure as a formalization of reputational pressure. This works in several ways. Laws formalize reputation itself. In Chapter 8, we talked about group membership as a substitute for individual reputation. As societies grow, laws formalize some group memberships.

For example, doctors need a license to practice. So do architects, engineers, private investigators, plumbers, and real estate agents. Restaurants need licenses and regular inspections by health officials to operate. The basic idea is that these official certifications provide a basis for people to trust these doctors, private investigators, and restaurants without knowing anything about their reputations. Certification informs potential clients that a businessperson has at least the minimum formal education and skill required to safely and competently perform the service in question, and that the businessperson is accountable to someone other than the customer: a licensing body, a trade organization, and so on. Handicap license plates are another formalized reputational system. Not all certifications are controlled by the government; some come from private institutions, such as Underwriter's Laboratories' certifications, the Good Housekeeping Seal of Approval, Consumer Reports rankings, and a variety of computer standards.

Other formal memberships that serve as reputation substitutes include academic degrees, bar associations for lawyers, the Better Business Bureau, food products' labels of origin—*appellation d'origine contrôlée* in France, and U.S. counterparts like “Wisconsin cheese” and “Made in Vermont”—USDA Organic certification, consumer credit ratings and reports, bonding, accreditation of educational institutions.

Negative reputation can also be institutionalized: public sex-offender registries, the DHS terrorist “no fly” list, blacklists for union organizers or suspected

Communists, and designations on driver’s licenses of a felony conviction. [The scarlet letter](#) is an older example, and the yellow star the Nazis required Jews to wear is a particularly despicable one.

Laws also formalize commitment. Legal contracts are probably the best example. Marriage licenses, curfew laws, and laws that enforce parents’ commitment to raise their children are others.

Societal Dilemma: Following contracts.	
Society: Society as a whole.	
Group interest: Effectively formalize agreements.	Competing interest: Maximize some self-interest.
Group norm: Follow contracts.	Corresponding defection: Break contracts.
<p>To encourage people to act in the group interest, the society implements a variety of trust mechanisms.</p> <p>Moral: We feel good about keeping our word.</p> <p>Reputational: No one does business with individuals and companies with a reputation for breaking contracts.</p> <p>Institutional: There are all sorts of laws regarding the legality of contracts, and sanctions for breaking them.</p>	

Finally, laws formalize societal norms that reputation traditionally enforced: anti-incest laws and age-of-consent laws, minimum drinking ages, bans on false advertising, blue laws, public indecency/intoxication laws, city lawn and weed ordinances, noise ordinances, libel and slander laws, zoning regulations, laws against impersonating police officers, and—in a perverse way—laws prohibiting people from criticizing the government. Employment applications that ask if you have ever been convicted of a felony are a formalization of reputation.

All of these institutional pressures allow reputation to scale, by giving people a system to trust so they don’t have to necessarily trust individuals. If I trust the system of government-issued identification cards and driver’s licenses, I don’t have to wonder whether to trust each and every a person when he tells me he’s old enough to drink in my bar.



There are many ways institutional pressure fails:

There is too little or too much of it. We've seen how institutional pressure is required to augment moral and reputational pressures in large and complex societies. Too little institutional pressure and the scope of defection is too great. For example, there's more tax evasion if the crime goes unpunished.

But more institutional pressure isn't always better. Gary Becker won a Nobel Prize in economics in part for his work in criminology. He asked the obvious question, what's the optimal level of crime? The naïve answer is zero, but that is unattainable and requires so much institutional pressure that society falls apart. Too much institutional pressure, and you get a country that looks like North Korea or the former East Germany: police states with a good part of the population working for the police. The other extreme—no police—doesn't work, either. You get lawless countries like Afghanistan and Somalia. Somewhere in the middle is the optimal scope of defection and the optimal level of enforcement.

In a lot of ways, this is similar to how evolution solves security problems. Antelopes don't need perfect protection against lions, and such protection would be too expensive in evolutionary terms. Instead, they accept the cost of losing the occasional antelope from the herd and increase their reproductive efficiency to compensate.

Similarly, we can never ensure perfect security against terrorism. All this talk of terrorism as an [existential threat](#) to society is nonsense. As long as terrorism is rare enough, and most people survive, society will survive. Unfortunately, it's not politically viable to come out and say that. We're collectively in a pathological state where people expect perfect protection against a host of problems—not just terrorism—and are unwilling to accept that that is not a reasonable goal.

Laws don't always have their intended effect. They can be a blunt tool, especially when it comes to violent crime and disaffected populations. There isn't a clean cause-and-effect relationship between incentives and behavior; more often than not, incentives are emotional, and are far more compelling than a rational consideration of even the most severe sanction.¹¹ There's a lot of research in this area, with many counterintuitive—and sometimes contradictory—results. We know that, in general, spending more money on police reduces crime somewhat. On the other hand, there are studies that demonstrate that the [death penalty reduces](#) murders as well as studies that [demonstrate it doesn't](#). While it's easy for politicians to be “tough on crime,” it's not always obvious that that's the best solution. An increase in the severity of punishment often doesn't translate into a drop in crime; an increase in the [probability of punishment](#) often does.¹² Often the [societal causes of crime](#) are what's important, and changes in the law do very little to help.

Laws have a clearer effect on more calculating crimes. Increasing penalties against tax fraud reduces tax fraud, at least for a little while. Increasing penalties on corporate crimes reduces those crimes. In those instances, potential defectors have plenty of time to make a rational risk trade-off.¹³

It's not always possible to enforce a law. International law, for example, only matters to the extent that the countries are willing to observe it or are able to enforce it on each other. Viktor Bout was an international arms dealer for about twenty years before his arrest in 2008. He was able to ship weapons to every conflict region imaginable, even those under UN embargo. He benefited from the lack of international law addressing transnational criminal activity, deliberately slack customs enforcement in countries seeking to attract business, and nations that found it convenient to let him do their dirty work.

Laws are open to interpretation, and that interpretation process can be expensive. Earlier I talked about solving the societal dilemma of pollution with a legal security measure: allowing people downstream from the polluter to sue. This is good in theory, but can be problematic in practice. The polluter can hire a team of lawyers skilled in the art of legal delay. If the cost of the lawyers is less than the cost of cleaning up the pollution, or if the polluter can outspend his legal opponents, he can neutralize their ability to raise the cost of defecting. This kind of expensive legal defense can also work against government regulations, tying the case up in the courts until the government gives up. In the state anti-trust suits against Microsoft, almost all states settled before trial.

Laws can have loopholes. This can happen by accident, when laws are linguistically ambiguous, contain simple errors, or fail to anticipate some new technological development. It can also happen deliberately, when laws are miswritten to enable the skillful few to evade them.

Examples of accidental loopholes are the “Double Irish” and “Dutch Sandwich” loopholes that allow multinational corporations to avoid U.S.—and other—taxes.¹⁴ It's how [Google pays](#) only 2.8% of profits in tax. One estimate claims the U.S. loses \$60 billion per year in taxes this way. Another loophole allows [large paper mills](#) to claim \$6 billion in tax credits per year for mixing diesel fuel in with a wood byproduct they already burn; the law with the loophole was intended to reduce the consumption of fossil fuels.¹⁵ A variety of loopholes make [video games](#) one of the most highly subsidized industries in the U.S. And, so as not to entirely pick on the U.S., the International Whaling Commission's loophole for research that [Japan exploits](#) to hunt whales commercially is another example.

Although it's hard to prove, there are many examples of laws believed to be deliberately written with loopholes to benefit someone. The UN Convention on

the [Law of the Sea](#) provisions on international fisheries are deliberately ambiguous, making much of it impossible to enforce. Also at the UN, [Security Council Resolution 1441](#)—used to justify invading Iraq—seems to have been designed to be ambiguous enough to both support and oppose the use of force.

More generally, loopholes are ways institutional pressure is subverted by defectors to do things it wasn't originally intended to do. Think of patent law, originally intended to protect inventors but now used by corporations to attack other corporations, or by patent trolls to extort money out of corporations. Or the legal profession, originally intended to serve justice but now used as an offensive weapon. Or stocks, originally intended to provide capital for companies but now used for all sorts of unintended purposes: weird derivatives, indexes, short-term trading, and so on. These are all defections. Either the law should be effective, or it shouldn't exist. A law with a loophole is the worst of both.

Laws can be applied inconsistently. If laws aren't objective, common, and universally applied, they are seen as unfair; and unfairness can exacerbate the Bad Apple Effect.

[Judge Gordon Hewart](#) put it best:

There is no doubt, as has been said in a long line of cases, that it is not merely of some importance, but of fundamental importance, that justice should both be done and be manifestly seen to be done.

Laws try to outlaw legitimate and moral behavior. Sometimes it's perfectly legitimate for someone to follow her individual self-interest, regardless of the group interest. There's no inherent dividing line, and different people—and societies—will draw it differently.

[Invasive species](#) are at best a serious problem, and at worst an [ecological disaster](#). They also pose a societal dilemma in which even a single defector can cause the group severe harm. All it took was one farmer releasing [silver carp](#) into the natural waterways of North America for it to invade everywhere, one flight accidentally carrying a [pregnant brown tree snake](#) to decimate the ecosystem of Guam, and one boat with [zebra mussel larvae](#) in its ballast water or milfoil clinging to its hull to overwhelm a previously pristine lake. As such, there need to be some pretty strong societal pressures in place to deal with this problem.

Some invasive species are easy to define as pests, but others are not. [Monk parakeets](#) are an invasive species in the U.S., thought to have been first released by pet owners either accidentally or as an easy way to get rid of them. The main harm they cause is [crop damage](#), although they [also cause fires](#) and blackouts by building massive, elevated nests in electrical equipment, and they outcompete

indigenous birds. On the other hand, they make cute pets and a lot of people like them. This results in a legal mess: the Wild Bird Conservation Act of 1992 prohibits importing them into the U.S., but [state laws vary wildly](#), with some states banning them, while others have no laws whatsoever.

One of the most useful things a court system does is strike a balance between polarities of interest. How should society balance my individual right to play loud music with my neighbors' right to peace and quiet? Or my right to run a tannery versus my neighbors' right to an unsmelly environment? How should society balance my individual desire to keep a parakeet as a pet with the community's need to minimize the dangers posed by feral birds? Laws that try to outlaw legitimate and moral behavior are less likely to succeed.

Laws don't affect every type of defector equally. In addition to those who can afford to fight and those who can't, there are three broad types of defectors when it comes to laws. The first are the individuals who know the law, believe the law is good (or at least that they don't want these things happening to *them*), and choose to break it anyway: burglars, muggers, kidnappers, murderers, speeders, and people in desperate straits. The second are individuals who know the law, believe the law is wrong, and choose to break it: pot smokers, some parakeet and ferret owners, and members of the Underground Railway who helped escaped slaves from the American South flee to safety in Canada. There is also a third category: those who don't know they're breaking the law, or don't realize how their actions affect the group. People might speed because they legitimately didn't see the speed limit sign, or they might not realize that certain sexual practices are against the law. These three groups will react differently to different laws, sanctions, and incentives.

Sometimes and for some people, laws aren't enough. Sometimes the incentives to defect are worth the risk. That's where security technologies come in.

10

Security Systems

Security systems are all around us, filling in the gaps where moral, reputational, and institutional pressures aren't effective enough. They include the door locks and burglar alarms in our homes, the anti-counterfeiting technologies in major world currencies, and the system of registering handguns and taking ballistic prints. They can be high-tech, like automatic face recognition systems, or low-tech, like defensive berms and castle walls. They don't even have to be physical systems; they can be procedural systems like neighborhood watches, customs interviews, and police pat-downs.

Theft of hotel towels isn't high in the hierarchy of world problems, but it can be expensive for hotels. Moral prohibitions against stealing prevent most people from stealing towels. Many hotels put their name or logo on their towels. That works as a reputational pressure system; most people don't want their friends to see obviously stolen hotel towels in their bathrooms. Sometimes, though, this has the opposite effect: [making towels souvenirs](#) of the hotel and more desirable to steal. It's against the law to steal hotel towels, of course, but with the exception of large-scale thefts, the crime will never be prosecuted.¹ The result is that the scope of defection is higher than hotels want. And large, fluffy towels from better hotels are expensive to replace.

The only thing left for hotels to do is take security into their own hands. One system that has become increasingly common is to set prices for towels and other items, and automatically charge the guest for them if they disappear from the rooms. This works with things like bathrobes, but it's too easy for the hotel to lose track of how many towels a guest has in his room, especially if piles of them are available at the pool or can easily be taken from a housekeeper's cart in the hallway.

A newer system, still not widespread, is to embed washable computer chips into the towels and track their movement around the hotel electronically. One anonymous Hawaii hotel claims they've **reduced towel theft** from 4,000 a month to 750, saving \$16,000 monthly in replacement costs. Assuming the RFID tags are inexpensive and don't wear out too quickly, that's a pretty good security system.

Let's go back to our two prisoners. They are morally inclined not to betray each other. Their reputation in the underworld depends on them not betraying their fellow criminal. And the criminal organization they're part of has unwritten but very real sanctions against betraying other criminals to the police. That's probably enough for most criminals, but not all. And—depending on the country—the police can be very persuasive.

What some organizations do—terrorists and spies come to mind—is add a security system. They organize themselves in cells so that each member of the criminal organization only knows a few other members: the members of his cell and maybe one or two others. There are a lot of ways to do this, and the organizational structure of the World War II French Resistance wasn't the same as Al Qaeda. If he's arrested or otherwise captured and interrogated, there's only so much damage he can do if he defects. This doesn't help the two captured prisoners, of course, but it does protect the rest of the criminal organization.

Societal Dilemma: Criminals testifying against each other.	
Society: The criminal organization.	
Group interest: Minimize the amount of jail time for the society.	Competing interest: Minimize personal jail time.
Group norm: Don't testify against each other.	Corresponding defection: Testify against each other in exchange for reduced jail time.
<p>To encourage people to act in the group interest, the society implements a variety of trust mechanisms.</p> <p>Moral: People feel bad when they let members of their group down.</p> <p>Reputational: Those who testify against their fellow criminals are shunned.</p> <p>Institutional: The criminal organization punishes stool pigeons.</p> <p>Security: The criminal organization limits the amount of damage a defecting criminal can inflict.</p>	

Of course, there are some good reasons not to run an organization like this. Imagine how much less effective a corporate worker would be if he only knew the five people in his department, and only communicated with his supervisor

using dead drops and the occasional voice-disguised conversation from constantly changing pay phone locations. But sometimes security wins out over clean lines of communication and managerial open-door policies.

In Chapter 6's Figure 8, I broke out several different types of security systems:

- *Defenses*. This is what you normally think of as security: weapons, armor, door locks, bulletproof vests, guard dogs, anti-virus software, speed bumps, bicycle locks, prison walls, panic rooms, chastity belts, and traffic cones. The common aspect of all these things is they try to physically stop potential defectors from doing whatever they're trying to do.
- *Interventions*. These are other security measures that happen during the defection that either make defection harder or cooperation easier. To make defection harder, think of obfuscation and misdirection measures, security cameras in casinos, guard patrols, and authentication systems. To make cooperation easier, think of automatic face-recognition systems, uniforms, those automatic road-sign radar guns that tell you what speed you're going, and road signs that inform you of the rules.
- *Detection/response systems*. These include burglar alarms, sensors in smokestacks to detect pollutants, RFID tags attached to store merchandise—or hotel towels—and detectors at the doorways, intrusion-detection systems in computer networks, and a UV light to detect if your hotel's bed sheets are clean.
- *Audit/forensic systems*. These are primarily enhancements to institutional societal pressure. They include fingerprint- and DNA-matching technology and the expert systems that analyze credit card spending, looking for [patterns of fraud](#).
- *Recovery systems*. These are security measures that make it easier for the victim to recover from an attack. Examples are a credit monitoring service or an insurance plan. What's interesting about these measures is that they don't directly influence the risk trade-off. If anything, they make someone more likely to defect, because he can more easily rationalize that the victim won't be hurt by his actions.
- *Preemptive interventions*. These operate before the attack, and directly affect the risk trade-off. Think of things like forced castration (chemical or otherwise), mandatory drug therapy to alter the personality of a career criminal, or a frontal lobotomy. Yes, these are often punishments after an attack, but they can prevent a future attack, too. Incarceration is also a preemptive intervention as well as a punishment; there are entire categories of crimes that someone in jail simply can't commit. So is execution, for the

same reason. Also in this category are [predictive policing programs](#) that increase police presence at times and places where crimes are likely to occur.

I'd be the first to admit this classification isn't perfect, and there are probably examples that don't fit neatly into my different boxes. That's okay; I'm less interested in precisely categorizing all possible security countermeasures, and more interested in looking at the breadth of security systems we use every day for societal pressures—many without even realizing it.

Security systems comprise a unique category of societal pressure. They're the last layer of defense—and the most scalable—against defection. You can view them as a way to technologically enhance natural defenses. Even if humans were complete loners and had never formed society, never worried about societal dilemmas, and never invented societal pressures, security systems could still protect individuals.

As a technological analog to natural defenses, they're the only societal pressure that actually puts physical constraints on behavior. Everything else we've discussed so far affects the risk trade-off, either directly, such as moral pressure, or through feedback, such as reputational pressure. Security can work this way as well, but it can also stop someone who decides to defect. A burglar might not have any moral qualms about breaking into a jewelry store, and he might not be worried about his reputation or getting caught—but he won't be able to steal anything unless he can pick the door lock and open the safe. Security might constrain him technically (the ability to pick the lock), financially (the cost to buy an oxyacetylene torch capable of cutting open the safe), or temporally (the time required to cut open the safe). Sometimes the constraints are relative, and sometimes they're absolute. This is what makes security systems so powerful and scalable. Security systems can work even if a defector doesn't realize that he's defecting. For example, a locked gate will stop someone who doesn't realize he's entering private property.

Also as an analog to natural defenses, security systems aren't always used as societal pressures. That distinction depends on who implements the security system and why. Think back to the sealed-bag exchange: the merchant could implement a variety of security systems to prevent his customers from shoplifting, cheating, or otherwise defrauding him. He could install security cameras and put anti-theft tags on his merchandise. He could buy a device that detects counterfeit money. He could use a service that verifies checks. All of this is the merchant's decision and the merchant's doing, and none of it is related to intra-group trust.

If a storeowner installs a camera behind his cash register, it's not societal pressure; if a city installs cameras on every street corner, it is. And if the police

use all the individually installed cameras in the area to track a suspect—as was done with [Timothy McVeigh's van](#)—then it's societal pressure. If society decides to subsidize in-store cameras, that's also societal pressure.

If I carry a gun for self-defense, it's not societal pressure; if we as a society collectively arm our policemen, it is. You could argue there is no societal dilemma involved in the hotel's towel-security decision. This is certainly true, and illustrates that the boundary between individual security and security as societal pressure can be fuzzy. The same security measure—a camera, for example—might be individual in one instance and societal in another. There are also going to be security measures that are some of both. I'm less concerned with the hard-to-classify edge cases than I am with the general categories.

Even if a security system is implemented entirely by individuals, that doesn't mean it can't also serve as societal pressure. A security camera is more likely to displace crime than reduce it; a potential thief can just go to another store instead. But if enough stores install hidden cameras, potential burglars might decide that the overall risk is too great. Lojack, widely deployed, will [reduce car theft](#) (and will increase car theft in neighboring regions that don't have the same system). Various computer security systems can have a similar result. If a security system becomes prevalent enough, potential defectors might go elsewhere because the value of defection is reduced.

Of course, society often limits what sort of security systems someone can implement. It may be illegal for a store to install security cameras in dressing rooms, even if it would reduce shoplifting. And I'm not allowed to bury land mines in my front yard, even if I think it would deter burglars.

Our security systems are also limited by our own abilities. Carrying a gun for self-defense makes less sense if you don't know how to use one. And I don't have the time to test every piece of food I eat for poison, even if I wanted to. A more realistic example: a store might have a policy to test if large bills are counterfeit, but not bother with smaller bills. (Of course, defectors take advantage of this: it's why \$20 bills are counterfeited more often than \$100 bills.)

Security systems are both their own category of societal pressure and augment the other three categories, allowing them to scale better. Quite a lot of the societal pressures we've talked about in the previous three chapters have a security component. Examples include:

- *Security-augmented moral pressure.* Something as simple as a sign stating “Employees must wash hands after using the restroom” can be viewed as a security system. Measures that make compliance easier are another

way to enhance morals, such as the electronic completion and filing of tax returns, photography to put a human face on victims and potential victims, and recycling bins in prominent locations. Other, more modern, technologies directly affect moral societal pressures: psychiatric therapies, personality-altering drugs, and brain-altering surgeries.

- *Security-augmented reputational pressure.* The eBay feedback mechanism is a reputational system that requires security to ensure the system can't be hacked and manipulated by unscrupulous merchants. Other examples are letters of introduction, tribal clothing, employee background checks, sex offender databases, diplomas posted on walls, and U.S. State Department travel advisories. Informal online reviews of doctors allow us to trust people we don't know anything about, with our health. Online reputational systems allow us to trust unknown products on Amazon, unknown commenters on Slashdot, and unknown "friends" on Facebook. Credit-rating systems codify reputation. In online games, security systems are less of an enhancement to, and more of a replacement of, moral and reputational pressures for ensuring game fairness.
- *Security-augmented institutional pressure.* A community might install cameras to help enforce speed limits. Or a government might use correlation software to analyze millions of tax returns, looking for evidence of cheating. Other examples include alarm systems that summon the police, surveillance systems that allow the police to track suspects, and forensic technologies that help prove guilt. Also time-lock safes, anti-shoplifting tags, cash register tapes, hard-to-forge currency, time cards and time clocks, credit card PIN pads, formal licensing of doctors, and the entire legal profession.

Let's put this all together. Think about an employee traveling for company business on an expense account. He can either live frugally, or enjoy the most expensive hotels, restaurants, and so on. It's a societal dilemma:

Here are some more societal dilemmas, and corresponding security systems that act as societal pressures.

- *Gridlock.* Security measures include traffic cops to keep cars moving, specially striped intersections to demarcate off-limits areas, and cameras to assist enforcement at gridlock-prone intersections.
- *Vaccines.* There is ongoing research on how to rebuild public [confidence in vaccines](#) and reduce defection. Tactics could include ad campaigns and other types of marketing. Also, inhalable vaccines make it easier to cooperate.

- *Cheating at games.* It's more fun for the group if everyone plays fairly, but it's sometimes more fun for the individual to cheat and win. To help combat cheating, the **new version of Monopoly** comes with an electronic gadget that keeps track of everyone's money and makes sure they go to the right square—no cheating.

Societal Dilemma: Corporate expenses.	
Society: The corporation.	
<p>Group interest: Minimize corporate expenses.</p> <p>Group norm: Spend the corporation's money frugally.</p>	<p>Competing interest: More enjoyable corporate travel.</p> <p>Corresponding defection: Spend a lot on hotel, meals, and so on.</p>
<p>To encourage people to act in the group interest, the society implements a variety of trust mechanisms.</p> <p style="padding-left: 40px;">Moral: A company-wide belief that frivolous expenses are tantamount to stealing.</p> <p style="padding-left: 40px;">Reputational: Praising people who save the company money. Publicly chastising people who spend lavishly.</p> <p style="padding-left: 40px;">Institutional: Corporate travel policies, including per diem systems and daily spending limits.</p> <p style="padding-left: 40px;">Security: E-mail reminders that people should be parsimonious with the company's money (enhances moral pressure).</p> <p style="padding-left: 40px;">Requiring employees to submit for approval estimates of how much they'll spend beforehand, and making it difficult to get additional expenses reimbursed (enhances both moral and reputational pressure).</p> <p style="padding-left: 40px;">Putting everyone's travel expenses on a website that everyone in the company can see (enhances reputational pressure).</p> <p style="padding-left: 40px;">Requiring booking of airfare and hotels through a dedicated travel agent, who enforces the corporate policies (enhances institutional pressure).</p> <p style="padding-left: 40px;">Auditing of travel expenses, with overspenders being forced to reimburse the company (enhances institutional pressure).</p>	

A lot of those might not feel like security systems, but they are. The breadth of security systems is vast. This chart—from criminal justice professor Ronald V. Clarke—illustrates just how diverse security can be.

Crime Prevention Techniques					
Increase the Effort	<p>1. Target harden</p> <ul style="list-style-type: none"> Steering column locks and ignition immobilizers Anti-robbery screens Tamper-proof packaging 	<p>2. Control access to facilities</p> <ul style="list-style-type: none"> Entry phones Electronic card access Baggage screening 	<p>3. Screen exits</p> <ul style="list-style-type: none"> Ticket needed for exit Export documents Electronic merchandise tags 	<p>4. Deflect offenders</p> <ul style="list-style-type: none"> Street closures Separate bathrooms for women Disperse pubs 	<p>5. Control tools/weapons</p> <ul style="list-style-type: none"> "Smart" guns Restrict spray paint sales to juveniles Toughened beer glasses
Increase the Risks	<p>6. Extend guardianship</p> <ul style="list-style-type: none"> Go out in a group at night Leave signs of occupancy Carry cell phone 	<p>7. Assist natural surveillance</p> <ul style="list-style-type: none"> Improved street lighting Defensible space design Support whistle-blowers 	<p>8. Reduce anonymity</p> <ul style="list-style-type: none"> Taxi driver IDs "How's my driving?" decals School uniforms 	<p>9. Use place managers</p> <ul style="list-style-type: none"> CCTV for double-deck buses Two clerks for convenience stores Reward vigilance 	<p>10. Strengthen formal surveillance</p> <ul style="list-style-type: none"> Red light cameras Burglar alarms Security guards
Reduce the Rewards	<p>11. Conceal targets</p> <ul style="list-style-type: none"> Off-street parking Gender-neutral phone directories Unmarked armored trucks 	<p>12. Remove targets</p> <ul style="list-style-type: none"> Removable car radio Women's shelters Pre-paid cards for pay phones 	<p>13. Identify property</p> <ul style="list-style-type: none"> Property marking Vehicle licensing and parts marking Cattle branding 	<p>14. Disrupt markets</p> <ul style="list-style-type: none"> Monitor pawn shops Controls on classified ads License street vendors 	<p>15. Deny benefits</p> <ul style="list-style-type: none"> Ink merchandise tags Graffiti cleaning Disabling stolen cell phones
Reduce Provocations	<p>16. Reduce frustrations and stress</p> <ul style="list-style-type: none"> Efficient lines Polite service Expanded seating Soothing music/muted lights 	<p>17. Avoid disputes</p> <ul style="list-style-type: none"> Separate seating for rival soccer fans Reduce crowding in bars Fixed cab fares 	<p>18. Reduce temptation and arousal</p> <ul style="list-style-type: none"> Controls on violent pornography Enforce good behavior on soccer field Prohibit racial slurs 	<p>19. Neutralize peer pressure</p> <ul style="list-style-type: none"> "Idiots drink and drive" "It's OK to say No" Disperse trouble makers at school 	<p>20. Discourage imitation</p> <ul style="list-style-type: none"> Rapid repair of vandalism V-chips in TVs Censor details of modus operandi
Remove Excuses	<p>21. Set rules</p> <ul style="list-style-type: none"> Rental agreements Harassment codes Hotel registration 	<p>22. Post instructions</p> <ul style="list-style-type: none"> "No Parking" "Private Property" "Extinguish camp fires" 	<p>23. Alert conscience</p> <ul style="list-style-type: none"> Roadside speed display boards Signatures for customs declarations "Shoplifting is stealing" 	<p>24. Assist compliance</p> <ul style="list-style-type: none"> Easy library checkout Public lavatories Litter receptacles 	<p>25. Control drugs and alcohol</p> <ul style="list-style-type: none"> Breathalyzers in bars Server intervention programs Alcohol-free events

In fact, one way to look at societal pressures is that *everything* I've written about in these past four chapters is a security system. Morals act as a preemptive intervention system. Reputation is a detection and response system; so are laws and sanctions. Taxes and incentives are interventions. And so on. While that may be true—and as a security guy that's really how I think of it all—it's more useful to think of security as its own thing.

I'm not going to talk more about specific security systems, both because such discussions can quickly get very technical, and because there are shelves full of books already written on the subject.



The use of performance-enhancing drugs in professional sports is a societal dilemma, and a good example of how security systems fail as a societal pressure.²

Societal Dilemma: Doping in professional sports.	
Society: All the athletes in the sport.	
Group interest: A safe and fair sport. Group norm: Don't take performance-enhancing drugs.	Competing interest: Winning and making a lot of money. Corresponding defection: Take performance-enhancing drugs.
To encourage people to act in the group interest, the society implements a variety of trust mechanisms.	
Moral: Guilt at not winning fair and square. Reminders that athletes are role models, and appeals to "think of the children."	
Reputational: Keep fans and endorsements by maintaining the reputation of a fair player.	
Institutional: Bans on performance-enhancing drugs.	
Security: Drug testing for specific performance-enhancing drugs.	

That's the idea, at least.³ It turns out that **enforcing anti-doping** rules is very difficult. The problem is while the intent of the rules is to ban performance-enhancing drugs in general, the temptation to ignore the group interest and take these drugs is enormous. Here's a quote from professional cyclist [Alex Zülle](#):

I've been in this business for a long time. I know what goes on. And not just me, everyone knows. The riders, the team leaders, the organizers, the

officials, the journalists. As a rider you feel tied into this system. It's like being on the highway. The law says there's a speed limit of 65, but everyone is driving 70 or faster. Why should I be the one who obeys the speed limit? So I had two alternatives: either fit in and go along with the others or go back to being a house painter. And who in my situation would have done that?

Before the sport started paying attention, distance cyclists used stimulants such as caffeine, cocaine, nitroglycerine, amphetamines, and painkillers to improve their endurance. It's a classic arms race—everyone had to partake in order to keep up—and many athletes suffered catastrophic health effects from long-term use. Morals and reputation aren't going to work in situations like this, and the only effective measures are institutional rules enforced by security systems: tests for specific drugs. France passed the first anti-doping laws in 1965; testers found that **almost a third** of the participants in the Tour de France the next year tested positive for amphetamines. Over the decades, each new potentially performance-enhancing substance was countered with a ban and then a test.⁴ Blacklists now encompass hundreds of substances.

Yet inconsistencies among various regulatory bodies' blacklists have led to the occasional sanction against athletes who **never intended** to break the rules.⁵ At the 2000 Olympics, Romanian gymnast Andreea Răducan was **stripped of her gold medal** because she tested positive for pseudoephedrine; she had taken two pills of an over-the-counter cold medicine prescribed by her team doctor.

Security systems fail for several broad reasons.

They don't work as well as advertised. Technologies are often developed and sold by companies that tout their value, even if there's no real evidence to support it. So municipalities **install security cameras** in a mistaken belief that they prevent crime, the TSA buys **full-body scanners** in a mistaken belief that they prevent terrorism, and the **military spends billions** on weapons systems like the Sgt. York air defense gun that don't work. In previous centuries, physiognomy (facial features) and phrenology (skull measurements) were both believed to be useful in identifying criminal personalities.

Attackers develop ways around the technologies. Attackers are always trying to figure out ways around security systems, and some of them succeed. Every anti-counterfeiting measure is eventually successfully overcome by counterfeiters.⁶ (Not just paper money; improvements in metallurgy result in **better slugs**.) No matter how many tax loopholes are closed, there is enough complexity in the tax code—and enough legislators willing to slip in provisions to benefit special interests—that unscrupulous companies can always find more. There are many ways to break the security of **door locks and safes**.

Major technologies change in ways that affect the security technologies. We'll talk about this extensively in Chapter 16. The Internet has given us an endless series of lessons in previously stable systems that failed when they moved online. For example, the security measures against impersonation fraud—identity theft—that worked in the world of face-to-face meetings and telephone conversations failed completely with online commerce. Computers make paper documents easier to forge, and [fax machines](#) make forgeries easier still. [Electronic voting machines](#) are considerably less secure than their predecessors. Modern [electronics in cars](#) bring with them new security risks. [Networked medical devices](#) can be hacked. There are hundreds of examples like this.

Sometimes the technological changes have absolutely nothing to do with the societal dilemma being secured. Between the ubiquity of keyboards and the tendency for teachers to focus on standardized tests, [cursive is not being taught](#) as much in schools. The result is that signatures are more likely to be either printed text or illegible scrawls, both easier to forge.

Security systems that augment other societal pressures, opening new avenues for attack. An example will illustrate.

In a small town, everyone knows each other, and lenders can make decisions about whom to loan money to, based on reputation (like in the movie *It's a Wonderful Life*). The system isn't perfect; there are ways for defectors to hide their reputations and cheat lenders. The real problem, though, is that the system doesn't scale. In order to enable lending on a larger scale, we enhanced reputation with technological security systems. For instance, credit reports and scores are a security-augmented reputational pressure. This system works well, and lending has exploded in our society in part because of it. But the new reputational pressure system can be attacked technologically. A defector could hack the credit bureau's database and enhance her reputation by boosting her credit score. Or she could steal someone else's reputation. All sorts of attacks that just weren't possible with a wholly personal reputational system become possible against a system that works on reputation plus a security system.

Even worse, many people don't realize that adding technological security to a reputational system makes such a difference, and continue to assume that it's a wholly reputational system. This adds to the risks. Some examples:

- Licensing is an institutional—formalized reputational—pressure system. When it is augmented with physical or electronic credentials, forging them becomes a way to attack it.

- Bank payment systems once had a combination of reputational and institutional pressure systems. Today it's primarily technological, and attackable through that technology.
- We traditionally used physical cues to assess the reputation of a business: the cleanliness of a restaurant, the impressiveness of a bank's building, and so on. Today we get a lot of those same cues from websites, where they are much easier to fake.⁷ More generally, our learned abilities to [read trust signals](#) are continually being overtaken by technology.
- Universal ID systems can make impersonation fraud more profitable, because a single stolen ID can be used in many more places. Sometimes, a [harder-to-forge ID](#) is even riskier, because it is that much more profitable to forge.

There's a more general change afoot. We're moving a lot of our interactions with other people from evolved social systems into deliberately created socio-technical systems. Instead of having a conversation face-to-face or voice-to-voice, we have it via text or e-mail. Instead of showing our friends our vacation pictures over drinks, we publish them on Flickr. Instead of sharing the intimacies of our life in person, we do it on Facebook. Instead of hanging out with our friends in bars or even street corners, we meet in massive multi-player games with a social component like World of Warcraft and Eve Online. This is an important change. In many of these systems, the technology fades to the background—that's the point, after all—and our brains primarily focus on the social aspects. As a result, we focus on the moral and reputational pressures endemic to the human interactions and ignore the technological part. So we forget that text conversations can be stored forever, retrieved later, and shared with other people. We forget there are people reading our Facebook comments who are not generally privy to the intimacies our life. We forget that Eve Online isn't the same as a face-to-face get-together. The technology changes how our social interactions work, but it's easy to forget that.

In this way, our traditional intuition of trust and security fails. There's a fundamental difference between handing a friend your photo album and allowing him to look through it and giving her access to your Flickr account. In the latter case, you're implicitly giving her permission to make copies of your photos that she can keep forever or give to other people.

Our intuitions about trust are contextual. We meet someone, possibly introduced by a mutual friend, and grow to trust her incrementally and over time. This

sort of process happens very differently in online communities, and our intuitions aren't necessarily in synch with the new reality. Instead, we are often forced to set explicit rules about trust—whom we allow to see posts, what circles different “friends” are in, whether the whole world can see our photos or only selected people, and so on. Because this is unnatural for people, it's easy to get wrong.

Science is about to give us a completely new way security-augmented reputational pressure can fail. In the next ten years, there's going to be an explosion of results in genetic determinism. We are starting to learn which genes are correlated with which traits, and this will almost certainly be misreported and misunderstood. People may use these genetic markers as a form of reputation. Who knows how this will fall out—whether we'll live in a world like that of the movie *Gattaca*, where a person's genes determine his or her life, or a world where this sort of research is banned, or somewhere in-between. But it's going to be interesting to watch.

I don't mean to imply that it is somehow wrong to use technological security systems to scale societal pressures, or wrong to use security to protect those technological systems. These systems provide us with enormous value, and our society couldn't have grown to its present size or complexity without them. But we have to realize that, like any category of societal pressure, security systems are not perfect, and will allow for some scope of defection. We just need to watch our dependence on the various categories of societal pressure, and ensure that by scaling one particular system and implementing security to protect it, we don't accidentally make the scope of defection worse.

Expenditures on security systems can outweigh the benefits. Security systems can get very expensive, and there's a point of diminishing returns where you spend increasingly more money and effort on security and get less and less additional security in return.⁸ Given a choice between a \$20 lock and a \$50 lock, the more expensive lock will probably be more secure, and in many cases worth the additional cost. A \$100 lock will be even more secure, and might be worth it in some situations. But a \$500 lock isn't going to be ten times more secure than a \$50 lock. There's going to be a point where the more expensive lock will only be slightly more secure, and not worth the additional cost. There'll also be a point where the burglar will ignore the \$500 lock and break the \$50 window. But even if you increase the security of your windows and everything else in your house, there's a point where you start to get diminishing returns for your security dollar.

The same analysis works more generally. In the ten years since 9/11, the U.S. has [spent about \\$1 trillion](#) fighting terrorism. This [doesn't count the wars](#) in Iraq and Afghanistan, which total well over \$3 trillion. For all of that money, we have not increased our security from terrorism proportionally. If we double

our security budget, we won't reduce our terrorism risk by another 50%. And if we increase the budget by ten times, we won't get anywhere near ten times the security. Similarly, if we halve our counterterrorism budget, we won't double our risk. There's a point—and it'll be different for every one of us—where spending more money isn't worth the risk reduction. A [cost-benefit analysis](#) demonstrates that it's smart to allow a limited amount of criminal activity, just as we observed that you can never get to an all-dove population.

There can be too much security. Even if technologies were close to perfect, all they could do would be to raise the cost of defection in general. Note that this cost isn't just money, it's freedom, liberty, individualism, time, convenience, and so on. Too much security system pressure lands you in a police state.




Figure 9: Security's Diminishing Returns

It's impossible to have enough security that every person in every circumstance will cooperate rather than defect. Everybody will make an individual risk trade-off. And since these trade-offs are subjective, and there is so much variation in both individuals and individual situations, the defection rate will never get down to zero. We might possibly, in some science-fiction future, raise the cost of defecting in every particular circumstance to be so high that the benefit of cooperating exceeds that of defecting for any rational actor, but we can never raise it high enough to dissuade all irrational actors. Crimes of passion, for example, are ones where the cost of the crime far outweighs the benefits, so they occur only when [passion overrides rationality](#).

PART III

The Real World



11

Competing Interests

In a societal dilemma, an individual makes a risk trade-off between the group interest and some competing interest. Until now, we've ignored those competing interests: it's been mostly selfish interests, with the occasional competing moral interest. It's time to fully explore those competing interests.

In general, there are a variety of competing interests that can cause someone to defect and not act according to the group norm:

- *Selfish self-interest.* This is the person who cheats, defrauds, steals, and otherwise puts his own selfish interest ahead of the group interest. In extreme cases, he might be a sociopath.
- *Self-preservation interest.* Someone who is motivated by self-preservation—fear, for example—is much more likely to behave according to her own interest than to adhere to the group norm. For instance, someone might defect because she's being blackmailed. Or she might have a drug addiction, or heavy debts. Jean Valjean from *Les Miserables*, stealing food to feed himself and his family, is a very sympathetic defector.
- *Ego-preservation interest.* There are a lot of things people do because they want to preserve a vision of who they are as a person. Someone might defect because he believes—rightly or wrongly—that others are already defecting at his expense and he can't stand being seen as a sucker. Broker [Rhonda Breard embezzled](#) \$11.4 million from her clients, driven both by greed and the need to appear rich.
- *Other psychological motivations.* This is a catch-all category for personal interests that don't fit anywhere else. It includes fears, anxieties, poor impulse control, genuine laziness, and temporary—or permanent—insanity. [Envy can motivate](#) deception.¹ So can greed or sloth. People do [things out](#)

of anger that they wouldn't otherwise do. Some pretty [heinous behavior](#) can result from a chronic deprivation of basic human needs. And there's a lot we're still learning about how people make risk trade-offs, especially in [extreme situations](#).

- *Relational interest.* Remaining true to another person is a powerful motivation. Someone might defect from a group in order to protect a friend, relative, lover, or partner.
- *Group interest of another group.* It's not uncommon for someone to be in two different groups, and for the groups' interests—and norms—to be in conflict. The person has to decide which group to cooperate with and which to defect from. We'll talk about this extensively later in this chapter.
- *Competing moral interest.* A person's individual morals don't always conform to those of the group, and a person might be morally opposed to the group norm; someone might defect because he believes it is the right thing to do. There are two basic categories here: those who consider a particular moral rule valid in general but believe they have some kind of special reason to override it, and those who believe the rule to be invalid per se. Robin Hood is an example of a defector with a competing moral interest. An extreme example of people with a competing moral interest are suicide bombers, who are convinced that their actions are serving some greater good—one paper calls them "[lethal altruists](#)."
- *Ignorance.* A person might not even realize he's defecting. He might take something, not knowing it is owned by someone. (This is somewhat of a special case, because the person isn't making a risk trade-off.)

An individual might have several simultaneous competing interests, some of them pressuring him towards the group norm and some away from it. In 1943, [Abraham Maslow](#) ordered human needs in a hierarchy, from the most fundamental to least fundamental: physiological needs, safety, love and belonging, self-esteem, self-actualization, and self-transcendence. Some of those needs advocate cooperation, and others advocate defection.

Figuring out whether to cooperate or defect—and then what norm to follow—means taking all of this into account. I'm not trying to say that people use some conscious calculus to decide when to cooperate and when to defect. This sort of idea is the basis for the [Rational Choice Theory](#) of economics, which holds that people make trade-offs that are somehow optimal. Their decisions are "rational" not in the sense that they are based solely on reason or profit maximization, but in the much more narrow sense that they minimize costs

and maximize benefits, taking risks into account. For example, a burglar would trade off the prospective benefits of robbing a home against the prospective risks and costs of getting caught. A homeowner would likewise trade off the benefits of a burglar alarm system against the costs—both in money and in inconvenience—of installing one.

This mechanistic view of decision making is crumbling in the face of new psychological research into the psychology of decision making. It's being replaced by models of what's called **Bounded Rationality**, which provide a much more realistic picture of how people make these sorts of decisions. For example, we know that much of the **trade-off process** happens in the unconscious part of the brain; people decide in their gut and then construct a conscious rationalization for that decision. These gut decisions often have strong biases shaped by evolution, but we know that a lot of assessment goes into that gut decision and that there are all sorts of contextual effects.




Figure 10: Competing Interests in a Societal Dilemma

This all gets very complicated very quickly. In 1958, psychologist Lawrence Kohlberg outlined six stages of moral development. Depending on which stage a person is reasoning from, he will make a **different type of trade-off**. The stage of moral reasoning won't determine whether a person will cooperate or defect, but instead will determine what moral arguments he is likely to use to decide.²

More generally, there are several counterbalancing pressures on a person as she makes her trade-off. We can organize pressures from the person outwards: self, kith and kin, less intimate friends, various larger and potentially overlapping

groups, society as a whole (however we define that), all of humanity (a special case of society as a whole), and some higher moral system (religion, political or life philosophy, or whatever). Sometimes the pressures come entirely from a person's own head, as with the various self-interests. The rest of the time, they come from other people or groups.

Kohlberg's Stages of Morality³	
Level 1: Preconventional Morality Right and wrong determined by rewards/punishment.	Stage 1: Punishment-avoidance and obedience Makes moral decisions strictly on the basis of self-interest. Disobeys rules, if possible without getting caught.
	Stage 2: Exchange of favors Recognizes that others have needs, but makes satisfaction of own needs a higher priority.
Level 2: Conventional Morality Other's views matter. Avoidance of blame; seeking of approval.	Stage 3: Good boy/good girl Makes decisions on the basis of what will please others. Concerned about maintaining interpersonal relations.
	Stage 4: Law and order Looks to society as a whole for guidelines about behavior. Thinks of rules as inflexible, unchangeable.
Level 3: Postconventional Morality Abstract notions of justice. Rights of others can override obedience to laws/rules.	Stage 5: Social contract Recognizes that rules are social agreements that can be changed when necessary.
	Stage 6: Universal ethical principles Adheres to a small number of abstract principles that transcend specific, concrete rules. Answers to an inner moral compass.

This is important, because the stronger the competing pressure is, the easier it becomes to defect from the group interest. Self-preservation interests can be strong, as can relationship interests. Moral interests can be strong in some people and not in others. Psychological motivations like fears and phobias can be very strong. The group interests of other groups can also be strong, especially if those groups are smaller and more intimate.⁴ Scale and emotional distance

matter a lot. The diagram gives some feel for this, but—of course—it's very simplistic. Individuals might have different emotional distances to different levels, or a different ordering.

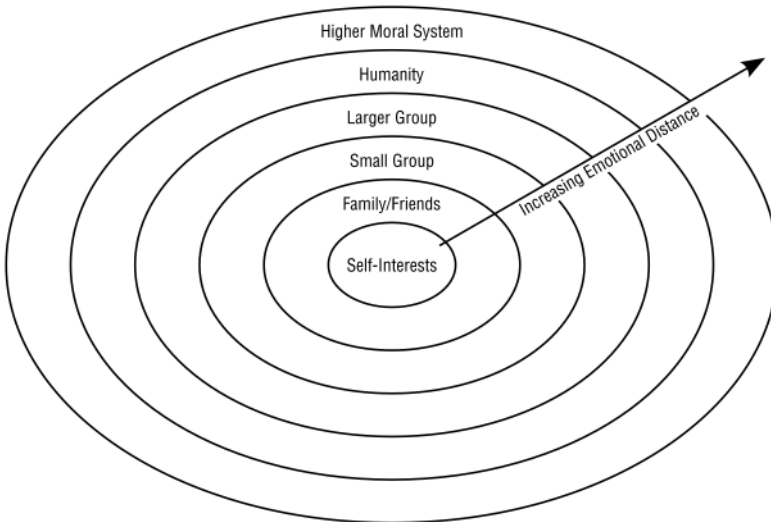


Figure 11: Scale of Competing Interests

Emily Dickinson wrote that people choose their own society, then “shut the door” on everyone else.

Competing interests, and therefore competing pressures, can get stronger once defectors start to organize in their own groups. It's one thing for Alice to refuse to cooperate with the police because she believes they're acting immorally. But it's far easier for her to defect once she joins a group of activists who feel the same way. The group provides her with moral arguments she can use to justify their actions, a smaller group she can personally identify with as fellow defectors, advice on how to properly and most effectively defect, and emotional support once she decides to defect. And scale matters here, too. Social pressures work better in small groups, so it's more likely that the morals of a small group trump those of a larger one than the other way round. In a sense, defectors are organizing in a smaller subgroup where cooperating with them means defecting from the larger group.

Depending on their competing interests, people may be more or less invested in the group norm. The selfish interests tend to come from people who are not invested in the group norm, and competing moral interests can come from people who are strongly invested in the group norm while also strongly invested

in other norms. Because of these additional investments, they have to explicitly wrestle with the corresponding competing interests, and the trade-offs can become very complicated.

Someone with criminal tendencies might have a simple risk trade-off to make: should I steal or not? But someone who is both moral and invested in the group norm—Jean Valjean or Robin Hood—has a much harder choice. He has to weigh his self-preservation needs, the morality of his actions, the needs of others he's helping, the morality of those he's stealing from, and so on. Of course, there's a lot of individual variation. Some people consider their **morality to be central** to their self-identity, while others consider it to be more peripheral. René Girard uses the term “**spiritual geniuses**” to describe the most moral of people. We also describe many of them as martyrs; being differently moral can be dangerous.⁵ Society, of course, wants the group interest to prevail.

Ralph Waldo Emerson wrote:

Society is a joint-stock company, in which the members agree, for the better securing of his bread to each shareholder, to surrender the liberty and culture of the eater. The virtue in most request is conformity. Self-reliance is its aversion. It loves not realities and creators, but names and customs.

Henry David Thoreau talks about how he went along with the group norm, despite what his morals told him:

The greater part of what my neighbors call good I believe in my soul to be bad, and if I repent of anything, it is very likely to be my good behavior. What demon possessed me that I behaved so well?

When historian Laurel Thatcher Ulrich wrote “Well-behaved women seldom make history,” she was referring to defecting.

Socrates's morals pointed him in the other direction, choosing to cooperate and drink poison rather than defect and escape, even though he knew his sentence was unjust.

We accept that people absorb and live according to the morals of their culture—even to the point of absolution from culpability for actions we now consider immoral—because we examine culpability in light of the commonly available moral standards within the culture *at that time*.⁶

This all might seem unrelated to this book; however, it's anything but. **Misunderstanding the defector** is a common way societal pressures fail, something we'll talk about more in Chapter 15. Think of the risk trade-off as a balance. When Alice is deciding whether to cooperate or defect, she's weighing the costs

and benefits of each side. Societal pressures are how the group puts its thumb on the scales and tries to make cooperating the more attractive option. If you think Alice is defecting because she's selfish (she's in it for the money) or concerned about her ego (she wants to look cool in front of her friends) when she actually has a competing moral interest, you're going to get societal pressures wrong. The details are different for every dilemma, but they're almost always important.

There's another important reason to understand the competing interests: you might get a different type of defection, depending on the competing interest. To illustrate this, let's use a more subtle societal dilemma: whether Alice should cooperate with the police.⁷ This is important, because whether and to what extent members of society report crime and assist the police greatly influences how well laws against those crimes work. In the absence of 100% automated burglar alarms connected to the police station, a monitored security camera in every niche and nook, or police patrols tailing every citizen 24/7, the likelihood that a burglar is going to get caught depends mostly on the willingness of bystanders to take action: either by calling the police, or by tackling the burglar and then calling the police. The more people who report illegal activity—both crimes in progress and crimes after the fact—the better institutional pressure works.

That's the group interest. Competing interests for not reporting include:

- *Selfish self-interest.* Alice might simply not care enough about society to cooperate. She might be too busy with other things in her life, and not have the time to get involved. She might have concluded in a risk–reward calculation that her time and the hassle of reporting a theft outweighs her benefit from reporting it.
- *Self-preservation interest.* Alice might be scared to cooperate with the police for any of several reasons. 1) She might be a criminal herself, and would rather not have anything to do with the police. And even though the police often give protection to lesser criminals who help prosecute their more powerful bosses, that protection is irregularly applied, and [there's no guarantee](#) a particular criminal witness will be adequately protected. 2) The police might be a danger to her. It's not universally true that the police are benevolent and helpful. There are people who won't willingly interact with the police out of legitimate fear. 3) She might fear retaliation from the criminals or the criminals' compatriots. Criminal organizations stoke this fear of retribution to allow themselves to commit crimes in a community with relative impunity. There was even a "[Stop](#)

Snitching” campaign, including a DVD produced by the Baltimore underworld, designed to intimidate people into not reporting crimes.

- *Ego-preserving interest.* She might be invested in a self-image that emphasizes keeping one’s head down and not borrowing trouble. She might, as a victim, be embarrassed and not want to admit it.
- *Other psychological motivation.* She might have an irrational fear of authority figures, severe anxiety, or pathological shyness.
- *Relational interest.* She might know the criminal in question, and would rather protect that person than assist the police.
- *Group interest of another group.* She might be part of, or sympathetic to, the group committing the crime, and decide to cooperate with the group rather than society as a whole.⁸ For instance, she might notice her employer committing a crime and decide not to report it. Or she might be a cop watching another cop abusing a prisoner, and she feels loyalty to her fellow officers trumps her moral obligation to report crime.
- *Competing moral interest.* She might not believe in the law. Many people in our society would never even think of calling the police when they see an illegal alien (“that law is immoral”), or discover that someone downloads copyrighted music off the Internet (“that law is ridiculous”). She might think the police behave immorally, or that the victim of the crime deserved it.

Even when someone is the victim of a crime, he might choose not to report it. Examples include crimes like rape (which can be demeaning to the victim to prosecute), some kinds of fraud (which carry a social stigma with them), small-scale crimes where it is unlikely that the police can help, and instances where the victim has reason to fear the police. Would a prostitute call the police after being raped? When my wife’s pocket was picked on the Budapest subway a decade ago, we didn’t bother reporting it to the police because we didn’t think they could do anything. Internet crimes can fall into this category, too. Quite a bit of credit card fraud isn’t reported to the police because the amount is too small for the police to worry about. In fact, a fraudster can make a good living stealing small amounts of money from large numbers of people because it’s not worth anyone’s effort to pursue him.

As a side note, people have lots of reasons for not reporting crime. Sometimes crimes are simply too hard to report. International crimes, made easier by globalization and the Internet, fall into this category. Internet scam victims fleeced by criminals in Nigeria probably have no idea whom to call—and the

unfortunate realization that no one can help. **Con artists** try to ensure that their victims don't call the police, because they thought they themselves did something illegal or because they're too embarrassed at being suckered.

Societal Dilemma: Cooperating with the police.	
Society: Society as a whole.	
<p>Group interest: Effective law enforcement.</p> <p>Group norm: Cooperate with the police.</p>	<p>Competing interest: Laziness.</p> <p>Competing norm: Ignore the police.</p>
	<p>Competing interest: Self-preservation; that is, a legitimate fear of the police or criminals.</p> <p>Competing norm: Avoid the police.</p>
	<p>Competing interest: Ego-preservation as someone who doesn't get involved in others' affairs.</p> <p>Competing norm: Don't get involved.</p>
	<p>Competing interest: Friend or relative of the person the police are investigating.</p> <p>Competing norm: Mislead the police, either actively by lying or passively by remaining silent.</p>
	<p>Competing interest: Member of a group that opposes the police.</p> <p>Competing norm: Several possible, depending on the group norm of the group.</p>
	<p>Competing interest: Believes that the police are not morally justified in their actions.</p> <p>Competing norm: Avoid, obstruct, or mislead the police.</p>

Not being aware of the crime is a problem with a lot of Internet fraud. **Fake anti-virus software** scams trick users into believing they have a virus, and charge them \$25, or \$50, or more for software to "remove" it. It's a multi-million-dollar industry, and most of the victims never realize they were scammed. There are **Internet money laundering** schemes that work the same way.

Competing interests are normal, and our society recognizes that people have them. Sometimes we even have mechanisms for dealing with these conflicts of interest. Judges are supposed to recuse themselves from cases in which they have a potential competing interest. Many governments exempt conscientious objectors from compulsory military service. Newspaper columnists, academic researchers, and others are supposed to declare any competing interests so their readers can understand their biases. Certain laws have religious exemptions.

Mostly, we're all better off because of these mechanisms: recusal makes it less likely that judges will issue decisions that reflect a personal bias; conscientious objector status makes it less likely that soldiers will have to rely on unwilling comrades to defend them in battle. But public health is *not* better off because there are religious exemptions to vaccination laws.

We even recognize the validity of certain competing interests in the law through the [doctrine of necessity](#). Something as straightforward as prohibitions against murder have exceptions for things like self-defense, a self-preservation competing interest. But note that the onus is on the person to demonstrate the validity of that competing interest. If Alice shoots and kills Bob, the presumption—and by this I mean social presumption, not legal presumption—will be that she committed murder, unless she can demonstrate otherwise.

Another point: morals are complicated, and societal dilemmas can disappear because people don't recognize a particular moral claim and corresponding competing interest. Overfishing is not a societal dilemma if you're unconcerned about the long-term sustainability of the seas. You might not even notice as fishers deplete the oceans, because there will probably still be fish in the grocery store as long as you're alive.⁹ Slavery isn't a problem if you don't believe the slave class has the same rights as the rest of the community. Even genocide isn't a societal dilemma if you have sufficiently dehumanized those you are slaughtering.

Of course, there's a lot more here that other books cover, and I recommend reading the literature on competing morals for some insights into how people should make trade-offs among a variety of [competing interests](#). For our purposes, it's enough to recognize that people have many competing interests, the details of which affect the efficacy of societal pressures as well as the means of defection. And for societal pressures to work, we need mechanisms that address the motivation for defection as well as the means.



We're all members of many formal and informal groups. These are the societies in our societal dilemmas. For most people, humanity is the largest one.

Society as a whole—whether we define it as our town, our country, or all of humanity—is a group. The company we work for is a group. Our political party is a group. Our city of residence is a group. These groups might have subgroups: the particular department in our company, the particular local political organization of the national party, our neighborhood in our city. Extended families can also be considered groups, and they have lots of different subgroups. Large corporations have many levels of subgroups; so do militaries and some religious groups.

These groups and subgroups often come into conflict with each other. We regularly have to make risk trade-offs in societal dilemmas where the interest of one group is in opposition with the interest of another group, and where cooperating with one group means defecting from the other, and vice versa. The rest of this chapter and the next three chapters discuss group interests and competing group allegiances. This is essential to understand how cooperation and defection work in the real world.

Recall our prisoner, in a societal dilemma with the rest of his criminal organization. That's a fine story, but real life is more complicated.

For the early years of his life, Sean O'Callaghan was a domestic terrorist. He joined the Provisional IRA when he was fifteen, and over the next five years, participated in nearly seventy attacks against the British, including two murders. In 1976, he had a change of heart. For the next ten years, he was a police informant intent on sabotaging the IRA. He thwarted several bombing attempts, including one against the Prince and Princess of Wales, and disrupted the delivery of tons of weapons and explosives. He also [publicly confessed](#) to his own crimes, and testified against many other IRA members.

[Defecting from the IRA](#) was a very dangerous thing to do. He did so—and risked retribution—because of a competing moral interest, and also because of another group interest: that of his community. “I realised that there was only one way in which I could help damage the ruthless killing capacity of the IRA: by handing myself up to the RUC [Royal Ulster Constabulary] and giving evidence against as many people as possible.” He was just as much cooperating with the larger group as he was defecting from the smaller group.

O'Callaghan faced a pair of societal dilemmas, both of which we've gone over in detail: criminals either cooperating with or defecting from their criminal organization and citizens assisting the police. These two societal dilemmas were in conflict. O'Callaghan had to make a choice: cooperate with the IRA and defect from society as a whole, or cooperate with society as a whole and defect from the IRA. The table below just lists the two competing societies and ignores the various other competing interests from previous tables.

Societal Dilemma: Cooperating with the police against the IRA.	
Society: Society as a whole.	Competing society: the IRA.
<p>Group interest: A peaceful divided Ireland. Group norm: Cooperate with the police.</p>	<p>Competing group interest: A free united Ireland, and the well-being of the IRA. Corresponding group norm: Don't cooperate with the police.</p>
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: Society teaches people to value peace over freedom and to help convict IRA terrorists.</p> <p>Reputational: Society praises people who help the police catch criminals. We give them awards, write articles about them, host them on television shows, and so on.</p> <p>Institutional: Laws against deliberately withholding evidence from the police, or actively misleading the police.</p> <p>Security: Hotlines that allow people to report crime anonymously. Witness protection programs.</p>	<p>To encourage people to act in the competing group interest, the society implements a variety of societal pressures.</p> <p>Moral: IRA teaches people to value freedom over peace and not to let fellow IRA members down.</p> <p>Reputational: Those who testify against their fellow criminals are shunned, or worse.</p> <p>Institutional: The criminal organization punishes police informants.</p> <p>Security: The criminal organization limits the amount of damage a defecting criminal can inflict.</p>

Competing societal dilemmas represent the normal state of affairs. Rarely is the real world so tidy as to isolate a single societal dilemma from everything else. Group interests are often in conflict, and cooperating in one necessitates defecting in another.

Nepotism is a problem in many organizations: companies, governments, and so on.¹⁰ For example, President [Ulysses S. Grant](#) found jobs for many of his relatives. He appointed his brother-in-law as minister to Denmark, his cousin as minister to Guatemala. Another brother-in-law was made counsel to Leipzig, and a third became the White House usher.

It's a pair of societal dilemmas. Grant was a citizen of the United States, and bound by its laws and customs. He was also a member of his own family, and endeavored to further its interests.

Societal Dilemma: Nepotism	
Society: The organization.	Competing society: The family. (Other competing interests not listed.)
<p>Group interest: Hiring the best people for the job.</p> <p>Group norm: Not showing any favoritism.</p>	<p>Competing group interest: Making sure your relatives do as well as they can.</p> <p>Corresponding group norm: Showing favoritism towards relatives.</p>
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: It feels good to put the best interests of the organization ahead of personal interests.</p> <p style="padding-left: 40px;">Reputational: The rest of the organization will react badly to charges of nepotism.</p> <p style="padding-left: 40px;">Institutional: Anti-nepotism laws.</p> <p style="padding-left: 40px;">Security: A free press that exposes nepotism.</p>	<p>To encourage people to act in the competing group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: It feels good to assist relatives. It's moral to take care of your family.</p> <p style="padding-left: 40px;">Reputational: The rest of the family will react well to showing favoritism and badly to not showing favoritism.</p> <p style="padding-left: 40px;">Institutional: You might be required to support your unemployed relatives.</p> <p style="padding-left: 40px;">Security: None.</p>

Other examples of competing societal dilemmas include:

- A politician who is a member of a political party and also a resident of the district that elected him.
- An employee of a department who is also an employee of the whole corporation.
- A corporate employee who is also a member of a union of workers of that corporation.
- Someone who is an employee of Company A working as a contractor to Company B. Or, similarly, someone who is an employee of Defense Contractor A working for the military of Country B. Think of employees of private security firms working for the U.S. military in Iraq.
- A schoolteacher who is both a member of a teacher's union and a parent who sends her children to the school she teaches at.
- Married lawyers representing opposing parties,¹¹ or judges related to lawyers appearing before them.

- Legislators who have worked for companies they write laws to regulate, or who expect jobs with those companies after their legislative careers.
- Multiple people serving on the boards of directors of the same corporations, known as an [interlocking directorate](#).

Some of these competing groups are nested, meaning that the members of the smaller group are all also members of the larger group. Others are overlapping, meaning that only some members of each group are also in the other group.

Most instances of competing interests are not societal dilemmas, and many are not relevant to security. That voter who lives in one district and votes in another will have to balance her competing interests when voting, but there are no security implications. That changes when the members of one group are expected to conform to some norm, and the members of the other group are expected to conform to some conflicting norm.

Another way to view moral interests is as a group interest. We are all part of the group of “everybody,” whether we define it as society as a whole, the human race, all life on this planet, or whatever. Many people associate morals with the responsibilities arising from that group membership. Relational ethicists do this at all scales, but the moral concepts of “universal justice” and “human rights” are implicitly connected to membership in the human race. People differ on what those terms mean, of course.

When multiple groups have competing interests, it can feel like a battle to see which group has a stronger tie with the individual. Often it’s the smaller, more intimate group. This is by no means definite—people don’t blindly commit crimes just because their friends are doing so—but it is a tendency. In general, it takes more societal pressure to get someone to defect from a smaller group interest in favor of a larger group interest than it does to get someone to do the reverse.

Think of the societal dilemma that someone who is a member of society as a whole and also a member of a criminal gang might have about cooperating with the police. There are benefits to being a member of society, and there are also benefits to being a member of the criminal gang—especially for people who live in dangerous communities in which the police offer little protection, or even pose an additional threat. Defecting from the broader society is a whole lot easier than defecting from a tight-knit and violent gang. And as long as the criminal gang maintains social ties that are stronger than those of general society, members are more likely to cooperate with the gang than with society in

general. This sort of dynamic plays out in organizations as diverse as corporations and terrorist organizations. Such organizations want their members to have a stronger loyalty to them than to society as a whole.

Benefits of group membership are important. In Chapter 12, we'll talk about instances where employees cooperate with their employers in breaking the law—defecting from society as a whole. It can be hard to go against your boss, even when you know that the company's actions are wrong. The benefits of being part of the group—things like the reputational approval of your colleagues and a continued paycheck—are powerful motivators to cooperate with the organization.

Over the millennia, we have developed a variety of measures that enhance group loyalty: initiation rites, rituals, shared markers of group membership, and so on. We talked about these in Chapter 8, as societal pressures to ensure commitment. When there are two groups in opposition, these measures become even more important.

It's not just formal organizations. When [Larry Froistad confessed](#) to killing his daughter on an online support group, Lisa DeCarlo alerted the authorities. She was vilified by the other members of the group.

[Police solidarity](#) provides another example. Called the “Blue Wall” or the “Blue Code of Silence,” it can be very difficult to get police officers to incriminate each other. For example, during the Toronto G20 Summit meeting in 2010, several witnesses reported a [policeman beating](#) a protester. The officer was eventually identified. During the investigation, however, a pretty good photograph of the then-anonymous policeman was shown to his fellow officers, and no one admitted being able to identify him, including eight policemen who were in his immediate vicinity during the incident, and one who was his roommate during the summit.

A huge variety of informal groups can have competing interests in societal dilemmas: social groups, ethnic groups, class groups, and racial groups; and there are many societal dilemmas that involve protecting the rights of minorities.

Families are filled with competing interests. Do you save money for your child's college tuition, or do you help your aging parents with their medical expenses? Do you side with your spouse or your parents? After your parents divorce and one of them remarries, who do you spend holidays with? If you are close friends with both members of a couple and you find out one of them is having an affair, do you tell or remain silent? When your father's driving skills begin to deteriorate, do you take away his car keys? Sometimes security is at stake in family situations: if your parents abused you as a child, do you let your

children sleep over at their house? In such cases, the societal pressures are primarily informal: morals and reputation. Only for extreme cases of abuse, or very dysfunctional families, is the legal system brought into play.

Since individuals are members of multiple competing groups, there are often redundant or conflicting societal pressures operating simultaneously. The societal pressures for one group don't always transfer to the other, and the pressures invoked at any time depend on the group that is being secured.

For example, the rules about what you can do in your family are different from the rules about what you can do out in the real world. We don't call the police if one of our children steals from the other, just as we don't call a burglar's parents if we catch him breaking into our house. Sometimes incidents of employee misconduct are dealt with by the corporation, but sometimes the police are called in. Groups of nations have their own organizations to call upon to deal with rogue nations. It all depends on scale.

When the scale changes, it can be confusing to know which rules to follow. We see this in schools, where some teachers and principals have begun calling the police for infractions that they previously would have called parents about. The scale has increased; the rules for dealing with these infractions are more often coming from the school district or larger community than from inside the classroom. This means the more informal moral and reputational systems stop working, and teachers feel the need to shift towards an institutional model. In general, as the size of the group grows, more formal societal pressures are required. And switching scales is messy, because the new systems are unfamiliar and often require new ways of thinking, and initially aren't good enough to work well.



In the next three chapters, we're going to talk about societal dilemmas surrounding organizations: both dilemmas facing groups and dilemmas within the groups. First, we'll talk about organizations in general, and then about specific types of organizations: corporations and (primarily government) institutions. We're going to see a lot of competing interests and societal dilemmas, and some pretty complicated societal pressures.



12

Organizations

So far, we've been talking primarily about how societal pressures affect individuals. Organizations—groups as small as several people and as large as hundreds of millions of people—also behave as individual actors. These organizations can be part of larger groups, just as individuals can, and those groups have group interests and corresponding group norms that affect those organizations. And just as Alice has to decide whether or not to steal, break a contract, or cooperate with the police, so do organizations.

It can be hard to think about organizations as a collective object. We often use individual metaphors when we talk about groups, and that results in us trying to use our understanding of individuals when we try to understand groups. We say things like “al Qaeda hates America,” “Google is trying to control the Internet,” and “China wants a strong dollar” as if those groups could have psychological states. It's metonymy, and while there's value to these generalizations, they also have their hazards. We're going to try to navigate those hazards.

Organizations are of course made up of individuals, who bring with them the sorts of societal dilemmas we've already discussed: both the dilemmas between the organizational interest and the individual's own competing interests, and the societal dilemmas that come from the individual being a member of the organization and a member of society as a whole. But we often treat organizations as if they actually were individuals, assuming that societal pressures work on them in the same way they do on individuals. This doesn't work, and results in some pretty bad trust failures, and high scopes of defection.



Organizations' competing interests include:

- *Selfish interest.* Organizations can have selfish interests, just like individuals. Depending on the organization, it might be profit, power, authority, influence, notoriety, or some combination of those things.
- *Self-preservation interest.* Organizations have strong self-preservation interests. They want to survive, just like individuals.
- *Ego-preserving interest.* Organizations have an analogue of self-image, and do things to preserve that image. For example, some organizations have a mission statement and go to great lengths to make sure their actions are consistent with their words. (Google's "don't be evil" motto is a good example.) Some organizations have particular reputations they want to preserve, for being honorable, ruthless, quick, and so on. Other organizations take pride in their geographic origins or in how long they've been in business. Still others have charitable foundations.
- *Other psychological motivations.* Organizations don't have psychologies, but they do have cultures. Examples are the not-invented-here syndrome, where companies become reluctant to adopt solutions from outside the organization; a "CYA," or "cover your ass" mentality, which predisposes an organization towards some solutions and away from others; dysfunctional communications, which lead to defection at one level that other levels don't know about; a caste system that can breed resentment in one group and lead to sabotaging behavior; or a skunk works dynamic, where a group inside the organization operates autonomously and in secret for a while.

Organizations also have competing group interests with other groups: rival organizations; groups of organizations, such as industry associations or geographical chambers of commerce; or society as a whole. Multinational organizations have potentially competing interests with a variety of countries.

An example of organizational interest is the March of Dimes. It started out as an organization to raise money to fight polio. After the polio vaccine was developed and polio almost eradicated, the March of Dimes didn't have a big party and wind up its accounts. Instead, it reconstituted itself as an organization to prevent birth defects in general, which should keep it going roughly forever.

Even though organizations have interests, the societal pressures we've already talked about work differently on organizations than they do on people.

- *Moral pressure.* Organizations are not people; they don't have brains, and they don't have morals. They can have group interests that are analogous to morals, though. Charities can have lofty mission statements, and a corporate mission statement like "don't be evil" is effectively a moral.
- *Reputational pressure.* For groups, reputation works differently than for individuals. Organizations care about their reputation just as individuals do: possibly more, due to size. They also have more control over it. Organizations can spend money to repair their reputations by undertaking advertising and public relations campaigns, making over their images, and so on—options that are simply unavailable to most individuals. On the other hand, because organizations are larger, their reputations are more valuable, and can be significantly harmed by the actions of a few of its members.
- *Institutional pressure.* Laws can be effective, but organizations cannot have sanctions imposed on them the way people can. They can't be put in jail or executed. In the U.S., there are occasional instances of physical-like punishments to corporations—the breakup of Standard Oil in 1911 comes to mind—and sometimes political parties are outlawed, such as Iraq's Ba'ath party in 2003. In extreme cases, individuals within organizations are jailed in punishment for organizational activity, but those are exceptions. Sometimes organizations are prohibited from certain actions by law as a punishment. For the most part, however, financial penalties are the only sanctions organizations face, which leads to all the issues of financial interests taking precedence over other moral interests we talked about in Chapter 9.
- *Security systems.* Security works differently against organizations than it does against people, primarily because they're *not* people. For the most part, security works against individuals inside the organizations rather than on the organization as a whole.



Organizations can be actors in all of the societal dilemmas discussed above. They have to decide whether to cooperate with the police and defraud people they do business with. They are affected by societal pressure. In addition, though, people within organizations have their own societal dilemmas with the organizations.

An example: **Carol Braun** was described as a “dedicated, dependable, competent and conscientious” 27-year employee of Goodwill Industries of North Central Wisconsin. She must have had some pretty good skills at reputation management, because over seven years, she used her position as comptroller to embezzle more than half a million dollars. Her actions were discovered when auditors found a \$77,000 discrepancy and conducted a comprehensive fraud investigation. Braun pleaded no contest to a single charge of embezzlement in 2003, and was sentenced to five years in prison and another five years of extended supervision. Braun’s actions resulted not only in significant financial loss to the Wisconsin Goodwill, but also in financial loss to her colleagues, whose pay had to be cut to make up the budget shortfall, and reputational damage to the agency.

That’s a particularly egregious example, but organizations teem with societal dilemmas. We often don’t notice them because we’re intuitively adept at dealing with groups of people. We understand hierarchies and authority, and the difference among superiors, colleagues, and subordinates. We’re facile at office politics because we’ve evolved to deal with social groups. But the societal dilemmas are still there, and sometimes it only takes a little nudge to bring them to the surface.

Every employee of an organization is faced with a societal dilemma: should he do what he wants, or should he do what the organization wants him to do? Stripped of context, it looks like this:

Societal Dilemma: Working within an organization.	
Society: The organization.	
Group interest: Maximize organizational interest.	Competing interest: Maximize personal interest.
Group norm: Do what the organization tells you to.	Corresponding defection: Do what you want.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: Work ethic, pride in a job well done, etc.</p> <p>Reputational: In some organizations, people who are perceived to work harder are treated better by their peers. In most organizations, they’re treated better by their superiors.</p> <p>Institutional: Organizations have all sorts of rules about employee behavior. Employees are supervised. Firing, promotions, and raises are all tied to performance—at least in theory.</p> <p>Security: Time cards, auditing, employee monitoring, formal performance reviews, and so on.</p>	

In economics, this is known as the *principal–agent problem*: the principal (in this case the organization) hires an agent (the employee) to pursue the principal's interests, but because the competing interests of the principal and the agent are different, it can be difficult to get the agent to cooperate.

Defection isn't all-or-nothing, either. Defections can be as diverse as coming in late, not working very hard, venting, whining, passive-aggressive behavior with coworkers, stealing paper clips from the office supply closet, or large-scale embezzlement. Remember the employee traveling for business from Chapter 10. He can cooperate with the organization and limit his expenses, or he can put his own self-interest first and spend wildly—or anything in-between.

We've all had experience with these sorts of defectors. Whether it's company employees, government employees, or members of any type of organization, there are always people who simply don't do the job they're supposed to.

There's another kind of defecting employee: someone who doesn't think of his employer's best interest while doing his job. Think of the officious employee who cares more about the minutiae of his procedures than the job he's actually supposed to do, or the employee who spends more time on office politics than actually working. The comic strip Dilbert is all about the dynamics of defecting employees and their defecting managers.

The fact that organizations almost never stop functioning because of defecting employees is a testament to how well societal pressures work in these situations. Organizations pay their employees, but there's a lot more than just salary keeping people doing their jobs. People feel good about what they do. They like being part of a team, and work to maintain their good reputation at work. They respond to authority, and generally do what their superiors want them to do. There's also a self-selection process going on; companies tend to hire and retain people who set aside their personal interests in favor of their employer's interests, and individuals tend to apply to work at companies that share their own balance between corporate and personal interests. And if those incentives aren't enough, corporations regularly fire employees who don't do what they're paid to do—or employees quit when they don't like their working conditions. There are also other financial incentives to cooperation in the workplace: commissions, profit sharing, stock options, efficiency wages, and rewards based on performance.

The poorer the job is—the less well-paying, the less personally satisfying, the more unpleasant, etc.—the more restrictive the security measures tend to be. Minimum-wage employees are often subject to rigorous supervision, and punitive penalties if they defect. Higher-level employees are often given more latitude and autonomy to do their job, which comes with a greater ability to defect.

This means that the ability to defect, and the stakes of defection, generally increase the higher up someone is within an organization. The overall trade-off is probably good for the organization, even though the occasional high-ranking defecting employee can do more damage before being discovered and realigned or fired than some misbehaving staff on the bottom rung. A senior executive can modify the organizational interest to be more in line with his own. And since he is in charge of implementing societal pressures to ensure that employees act in the organizational group interest, he can design solutions that make employees more likely to cooperate while still leaving him room to defect. He can build in loopholes. Additionally, because he can implement societal pressures to limit defections among the other employees, he can minimize the Bad Apple Effect that would magnify the adverse effects of his defection to the organization. In extreme cases, a CEO can run the company into bankruptcy for his personal profit, a ploy called “corporate looting” or “control fraud.” His power makes it possible for him to impose his personal agenda on top of the organizational agenda, so the organization becomes—at least in part—his personal agent.

This kind of thing doesn't have to be as extreme as fraud. Think of a CEO whose salary depends on the company's stock price on a particular date. That CEO can either cooperate with the group interest by doing what's best for the company, or defect in favor of his self-interest and do whatever is necessary to drive the stock price as high as possible on that date—even if it hurts the company in the long run.¹

[Sambo's restaurants](#) had an odd incentive scheme called “fraction of the action” that let managers buy a 10% interest in individual restaurants: not only the ones they worked at, but others as well. This enabled rapid early expansion for the chain, since it both helped finance new openings and gave managers a huge incentive to make restaurants prosper. But as the chain grew, people all over the hierarchy had individual financial interests that conflicted with their loyalty to the chain as a whole. People responsible for getting food to a whole region were able to favor specific restaurants, for instance.

On the other hand, executives have a lot of societal pressures focused on them that's supposed to limit this sort of behavior. In the U.S., Sarbanes-Oxley was passed precisely for this purpose. And the inherent restraints of their roles prevent most of them from being brazen about it. But there are exceptions, and some of those are what we read about in the newspapers.



We've mentioned organizations as individual actors in societal dilemmas, and we've talked about individuals in societal dilemmas inside organizations. Now let's put them together.

Think back to the overfishing societal dilemma from Chapter 5, but instead of Fisherman Bob deciding whether to cooperate and limit his catch or defect and overfish, it's the Robert Fish Corporation. Fisherman Bob and the Robert Fish Corporation face the same societal dilemma, but a corporation isn't actually a person; it's an organization of people in a hierarchy. Let's go through it step by step.

The Robert Fish Corporation has to decide whether or not to overfish. The scale is certainly different than the simpler example—the Robert Fish Corporation might have dozens of large fishing boats all over the world—but the idea is the same. The corporation will collect whatever information it needs via its employees, and some person or group within that corporation will decide whether to cooperate or defect. That trade-off will be made based on the corporation's competing interests and whatever societal pressures are in place.

For the moment, let's assume that the corporation decides to defect. For whatever reason, it is official policy of the Robert Fish Corporation to follow its short-term self-interest at the expense of the group interest of society as a whole. In this case, that means catching as much fish as possible, whenever possible, regardless of whether that depletes the stock.

Alice is the Vice President of Operations of the Robert Fish Corporation. Her job is to implement that corporate policy. Alice is in charge of overfishing. As an employee, Alice has a societal dilemma to address. She can either cooperate and implement corporate policy to overfish, or defect and undermine her employer's goals. But in addition to her role as a Robert Fish Corporation employee, Alice is a member of society as a whole. And as a member of society, she has a second societal dilemma: she can either cooperate and ensure that her company fishes responsibly, or defect and allow it to overfish.

Those two societal dilemmas conflict, just like O'Callaghan's two societal dilemmas. Cooperating in one means defecting in the other. So when Alice decides whether or not her company is going to overfish, she is caught between two societal dilemmas. A whole gamut of corporate rules will pressure her to implement corporate policy, and laws against overfishing will pressure in the opposite direction.

Societal Dilemma: Overfishing	
Society: The Robert Fish Corporation.	Competing society: Society as a whole. (Other competing interests not listed.)
Group interest: Follow the corporate norms.	Competing group interest: Ensure long-term viability of fishing stocks.
Group norm: Overfish.	Corresponding group norm: Don't overfish.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: It feels good to put the best interests of the organization ahead of personal interests.</p> <p style="padding-left: 40px;">Reputational: The rest of the organization will react badly to someone who doesn't act in the organizational interest.</p> <p style="padding-left: 40px;">Institutional: Specific corporate overfishing policy regulating behavior. Raises and promotions tied to amount of fish caught.</p> <p style="padding-left: 40px;">Security: Super-efficient fishing technology that is optimized to maximize the catch.</p>	<p>To encourage people to act in the competing group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: Good stewardship of earth's resources, being a good global citizen are valorized.</p> <p style="padding-left: 40px;">Reputational: Environmental groups report on company behavior and organize letter writing campaigns or boycotts of defectors.</p> <p style="padding-left: 40px;">Institutional: Laws prohibiting overfishing.</p> <p style="padding-left: 40px;">Security: Possibly government monitoring of fishing. Pesky protest boats.</p>

There is also the normal gamut of competing interests that Alice might have. Alice might be morally predisposed to respect the authority of her bosses and go along with her group. She might believe that overfishing is morally wrong. She probably has some specialized knowledge of the life cycle of fish and the effects of overfishing. Concerns about her reputation as a good employee or a team player will make her more likely to cooperate with her employer. Her self-regard and her reputation as a moral individual might make her more likely to cooperate with society. Her self-preservation interest—she might be fired if she disobeys the corporate policy—comes into play as well. And remember that emotional distance is important: if Alice has stronger ties to her employer than to society, she's more likely to cooperate with her employer and defect from society. Organizations try to keep their employees loyal for this reason.²

Clearly Alice has a tough choice to make. Here are some examples of how that choice has played out in the real world. There is a lot of research in decision

making within groups, especially corporations. We've already seen in Chapter 9 how financial considerations dampen moral considerations. There is considerable evidence, both observational and experimental, that the group dynamics of a hierarchical organizational structure, especially a corporate one, dampen [moral considerations](#) as well. There are [many reasons for this](#), and it seems to increase as organizations grow in size.

From 1978 to 1982, the [Beech-Nut Corporation](#) sold millions of bottles labeled as apple juice, intended for babies, that contained no actual apple products. If you read the story of how this happened, and how it kept on happening for so long, you can watch as the senior executives wrestled with their two societal dilemmas. They could cooperate with society and not sell phony apple juice, but that would mean defecting from their corporation. Or they could cooperate with their corporation, first by not questioning how this "juice" supplier could be 25% cheaper than anyone else, and then by continuing to sell the product even after they knew it was phony; but that would mean defecting from society. In the end, the economic and social ties they had with their company won out over any ties with greater society, and it wasn't until an independent laboratory discovered their deception that they stopped the practice. In 1987, they were tried in federal court, and eventually agreed to pay a \$2 million fine—at the time, the largest ever paid to the U.S. Food and Drug Administration. This is also one of the rare occasions that individuals within a corporation were jailed.

Since the mid-1980s, a growing docket of complaints, criminal prosecutions, and civil suits in the United States, Europe, and elsewhere has revealed that, since at least 1950, Roman Catholic bishops knowingly transferred thousands of priests accused of child molestation into unsuspecting parishes and dioceses, rather than diminish the ranks and reputation of the priesthood and expose the church to scandal. By 2011, allegations had been made against nearly 5,000 U.S. priests, and over 15,000 U.S. residents had testified to [being victimized](#). (Estimates of the [actual number of victims](#) range as high as 280,000.) In a 2002 tally, approximately two-thirds of sitting U.S. bishops were alleged to have either [retained accused priests](#) in their then-current positions or moved them to new assignments. This was in keeping with the [Vatican's exhortations](#) to investigate cases of sexual abuse in secret, so they would remain bound only by canon law.

What happened inside the church can be explained as a pair of societal dilemmas. The larger one was within society as a whole: we are definitely all better

off if people don't sexually molest minors, and we have implemented a variety of societal pressures—moral, reputational, and legal—to keep that particular defec-tion down to a minimum. We even have a variety of security mechanisms to detect child porn on the Internet and determine who is taking and trading those pictures.

Meanwhile, a smaller societal dilemma unfolded within the Roman Catholic Church. Of course pedophilia and ephebophilia aren't the societal norm within the church; pedophile priests are just as much defectors from the church as they are from society as whole. But the church hierarchy (the bishops and the Vati-can) decided that its ability to function as a trustworthy religious institution depended on reputation. This is known as the "doctrine of scandal," and means that its reputation was more important than justice—or preventing further transgressions. So the church systematically worked to keep secret the problem and the identities of the perpetrators.³ The church has some pretty strong soci-et al pressures at its disposal—primarily moral and reputational—which is why this scandal took decades to become public. In some cases, it even forced the victims to sign non-disclosure agreements (an institutional pressure).

Societal Dilemma: Protecting Pedophiles	
Society: The Roman Catholic Church.	Competing society: Society as a whole. (Other competing interests not listed.)
Group interest: A scandal-free church. Group norm: Protect pedophile priests from exposure and prosecution.	Competing group interest: Protecting minors. Corresponding group norm: Arrest, convict, and punish pedophiles.
To encourage people to act in the group interest, the society implements a variety of societal pressures. Moral: Exposing the church is seen as a sin against it. Reputational: Praising people who kept quiet and punishing those who exposed the church. Institutional: Imposing sanctions against those who exposed the church. Non-disclosure agreements. Security: None.	To encourage people to act in the competing group interest, the society implements a variety of societal pressures. Moral: Child molestation is bad. Protecting minors, and punishing sex offenders, is paramount in our society. Reputational: People are rewarded, either emotionally or physically, for exposing pedophiles. Pedophiles are ostracized. Institutional: Laws against pedophilia. Rewards for turning in pedophiles. Security: Chemical castration, actual castration.

In the end, this backfired massively. Unfortunately, cover-ups are not uncommon, as organizations try to protect their own reputation—and their profits from cheaper products. It happens within corporations. It happens within governments. It can happen within any type of organization.

On the other hand, there's a new trend that cuts in the opposite direction. One theory of corporate damage control advocates full disclosure, acknowledgement, and public displays of contrition, in hopes of a quick reputational resurrection. Lots of politicians have been taking this tactic with their tearful public confessions, resignations, treatment center visits, and then quick return to public life, problem supposedly solved.



Whistle-blowers are an extreme example of someone defecting from an organization to cooperate with society as a whole. When WorldCom's Vice President of Internal Audit, Cynthia Cooper, first expressed her concerns about bookkeeping anomalies she had discovered, she was met with hostility from her supervisor and apathy from the company's auditors. Despite this, she unilaterally conducted a full-scale financial audit of the company. What she discovered was that [top WorldCom executives](#) had routinely misidentified operating costs as capital expenditures, ultimately preventing \$11 billion from being subtracted from the company's bottom line, and thereby misrepresenting the company's value to its board and investors. Cooper's discovery led to an SEC investigation, bankruptcy and reorganization of the company, and criminal convictions of WorldCom's top executives and accountants. It also brought into renewed focus the need for public companies to implement internal societal pressures to protect themselves and the public from defectors in their ranks.

Along similar—but not nearly as extreme—lines as Sean O'Callaghan, Cooper put herself at considerable personal risk by becoming a police informant.

This is a complicated risk trade-off, one that includes both the group interests of WorldCom and society as a whole, as well as Cooper's various self-interests.

Societal Dilemma: Whistle-blowing		
Society: The organization.	Competing society: Society as a whole.	Other competing interests.
<p>Group interest: The best interest of the organization.</p> <p>Group norm: Organizational loyalty; do what the organization expects you to do, regardless of competing interests.</p>	<p>Competing group interest: Lawfulness.</p> <p>Competing group norm: Cooperate with the police and expose organizational wrongdoing.</p>	<p>Competing interest: Keep your job.</p> <p>Competing norm: Do what the organization wants.</p> <p>Competing interest: Do what's morally right.</p> <p>Competing norm: Expose and help prosecute crime.</p> <p>Competing interest: Don't get involved.</p> <p>Competing norm: Quit the job and don't say anything.</p>
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: Acting in the best interest of the organization is "the right thing to do."</p> <p>Reputational: People who act in the best interest of the organization are seen as good and loyal employees.</p> <p>Institutional: People who act in the best interest of the organization are rewarded, both financially and with advancement.</p> <p>Security: Employee monitoring, indoctrination procedures.</p>	<p>To encourage people to act in the competing group interest, the society implements a variety of societal pressures.</p> <p>Moral: Protecting the greater community is "the right thing to do."</p> <p>Reputational: People who protect the greater community are rewarded with the admiration of the media and the public.</p> <p>Institutional: Laws protecting whistle-blowers from retaliation.</p> <p>Security: Cameras, photocopies, and other recording devices make evidence gathering easier.</p>	

I don't know the exact dimensions of the trade-off—likely the full range of competing interests includes everything related to cooperating with the police—but you get the general idea. And it's not just employees; **corporate board members** face a similar pair of societal dilemmas. Cooper had a variety of competing interests, and the full force of WorldCom's societal pressures fighting her.

Organizations can muster considerable societal pressures to prevent and punish whistle-blowing defections. Some extreme examples:

- National Security Agency analyst [Thomas Drake](#), alarmed by the agency's initiation of warrantless domestic electronic surveillance after the September 11 attacks, first expressed his concerns to his superiors, then supplied classified information on the program to the House Permanent Select Committee on Intelligence and oversight committees investigating 9/11-related intelligence failures, and finally shared what he believed was unclassified information about the NSA with a reporter. A year and a half later, federal agents raided Drake's home, confiscated his computers, books, and papers, and accused him of participating in a conspiracy to violate the Espionage Act—a law originally aimed at those who aid the enemy and harm national security, rather than those whose disclosures serve the public interest.⁴
- The [1962 Vatican Instruction](#) "*Crimen Sollicitationis*" prescribed excommunication of those who violated the oath of secrecy imposed on parties to investigations of sexual misconduct by priests, including pedophilia investigations.
- As a research physician, Nancy Fern Olivieri was part of a group conducting a clinical trial of a drug for the pharmaceutical company Apotex. When she came to believe that the [drug was ineffective](#) and possibly toxic, Apotex threatened all sorts of legal action against her if she took her concerns public.
- Detective Jeff Baird [exposed misconduct](#) in the New York City Police Department's Internal Affairs division. As a result, he was shunned and harassed by his fellow officers, gratuitously transferred to different units, and received anonymous death threats.

It's no wonder so few people become whistle-blowers, the consequences can be so devastating. Imagine you're in the middle of a Madoff-like pyramid scheme. Do you expose the scheme and risk prosecution or retaliation, feign naïveté and try to get out, or actively participate for greater rewards and greater risk?

An even more extreme example is military desertion in wartime. Militaries need strict hierarchies to function effectively. It's important that soldiers obey the orders of their superiors, and be able to give orders to their subordinates. But since these orders might be otherwise pretty abhorrent to individuals, the military implements a lot of societal pressure to make it all work. This is why military training uses substantial social pressures around strict obedience and group cohesion. In addition, militaries have strict rules about obeying orders, with serious sanctions for breaching them. Throughout much of history, desertion was

punishable by death with less than due process, because it was just too important to the group preservation interest to allow for individual self-preservation.

This can change when the military is ordered to take action against the very people it believes it is protecting. In 2011, two high-ranking Libyan **military pilots defected** rather than carry out orders to bomb protesters in the Libyan city of Benghazi. The pilots realized that they were in a pair of societal dilemmas, and chose to cooperate with their fellow countrymen against the government rather than cooperate with their fellow soldiers against the protesters.

Societal Dilemma: Military Desertion		
Society: The military.	Competing society: Society as a whole.	Other competing interests.
<p>Group interest: The best interest of the military.</p> <p>Group norm: Do whatever your superiors tell you to do.</p>	<p>Competing group interest: The best interest of the people in society.</p> <p>Competing group norm: Don't attack your fellow citizens.</p>	<p>Competing interest: Self-preservation.</p> <p>Competing norm: Don't put yourself in harm's way.</p>
		<p>Competing interest: Ego preservation.</p> <p>Competing norm: Don't let your fellow soldiers down.</p>
		<p>Competing interest: Do what's morally right.</p> <p>Competing norm: Don't kill people.</p>
<p>To encourage people to act in the group interest, the military implements a variety of societal pressures.</p> <p>Moral: Basic training instills a military morality.</p> <p>Reputational: Military units have strong group cohesion.</p> <p>Institutional: Disobeying orders is strictly punished.</p> <p>Security: A variety of security measures constrain soldiers.</p>	<p>To encourage people to act in the competing group interest, society implements a variety of societal pressures.</p> <p>Moral: Moral teaching not to harm others.</p> <p>Reputational: Society ostracizes those who turn against their own people.</p> <p>Institutional: Laws against war crimes.</p> <p>Security: None.</p>	

In 2005, Captain Ian Fishback exposed the [U.S.'s use of torture](#) in Iraq because of his religious convictions. Similarly, Bradley Manning had to deal with two competing societal dilemmas in 2010 when he allegedly became a whistle-blower and sent 250,000 secret State Department cables to the anti-secrecy group WikiLeaks, which made them public.⁵ Like the Libyan pilots, he chose to defect from the government and cooperate with what he perceived as the country as a whole. His subsequent treatment by the U.S. government—which incarcerated him, stripped him of due process, and tortured him—is in part a societal pressure by the government to prevent copycat defections. In previous eras, the king might have put his head on a pike for all to see.

Such anti-defection measures don't work perfectly, of course. Almost all corporate, government, and other institutional misdeeds become public eventually. All militaries have some level of insubordination and desertion. Historically, desertion was huge, mostly because there was no good way to enforce cooperation most of the time. These days, in most countries, it's generally kept at a low enough level that it doesn't harm the military organization as a whole.



It's not always the case that someone who defects from an organization hurts the organization. An individual member of the organization can defect against the desires of the organization but for the benefit of the organization.

This is easiest to explain with an example. Let's return to the Robert Fish Corporation. This time, the corporation decides it will not overfish. Alice, a fisher working for the corporation, has a societal dilemma as an employee: she can cooperate and implement the corporate policy, or she can defect and do what she wants. She also has the same dilemma as a member of society.

Like most employees, Alice generally cooperates and does what the corporation wants. The problem is that the corporation wants a lot of things, but only measures and pays attention to some of them. In our example, Alice's level of cooperation is measured by how much her actions affect the profitability of the corporation. She's rewarded for keeping revenues high and costs low, and penalized for doing the reverse.

Alice might overfish, even though the official corporate policy is not to. She defects in the societal dilemma with society as a whole, and also in the societal dilemma with the Robert Fish Corporation. But unless her management is specifically measuring her on overfishing, they're not going to realize that her increased revenues are coming from something that is against corporate policy.

And unless management penalizes her for doing so, she will be motivated to continue the practice.

This sort of dynamic is not uncommon in a corporate environment.

- In 2010, BP's [Deepwater Horizon](#) drilling rig exploded, killing 11 workers and injuring 17 others, then collapsed, and spilled 205 million gallons of oil and 225,000 tons of methane from the Macondo well into the Gulf of Mexico. Reading the reports on the Deepwater Horizon oil spill, it's obvious the company cut all sorts of safety corners and took undue risks. The [employees of BP](#) didn't do what was required by law. More importantly, BP's employees didn't do what was required by BP.⁶
- The "Big Dig," the massive highway project in downtown Boston from the 1990s, had a long list of defects resulting from shoddy business practices. Again, cutting costs and time was more important than doing the job right on a project already way over time and budget. And while in many cases the companies who did the substandard work were successfully sued, the individuals inside those companies who made the decisions were largely untouched.
- Before the 2008 financial crisis, there was an expression around Wall Street: "I'll be gone. You'll be gone." It was what self-interested investment bankers would say about worthless mortgage-backed securities, weird derivatives, or anything else that was more smoke and mirrors than real value. Yes, those who sold these financial instruments were going against the long-term interests of their employer by dealing in them. But by the time anyone would find out, they expected to be rich, retired, and beyond any reprisals from their bosses. It's only a small step removed from pump-and-dump stock scams.⁷

This isn't always a defection from the organization. Sometimes it's a defection in detail but not in spirit. Sometimes senior managers make sure they don't know the details of what's happening. Or they're perfectly aware corners are being cut and regulations violated, but make sure the facts never appear in a memo or e-mail. This gives them plausible deniability in the face of prosecution. In extreme cases, companies hire public relations people to lie to the public without realizing that they're lying. Of course, if someone gets caught doing this, the individual will be accused of not following company policy.

On the other hand, sometimes this is innocent and nothing more than the organization's failed societal pressure systems resulting in a too-high scope of

defection. In either case, in corporations where this sort of thing is prone to happen, the individuals who do it are the ones who will most likely be rewarded. So if there are five fishers, and one of them breaks the rules and secretly overfishes, she will bring in the most revenue to the company and get promoted to manager. The fishers that cooperated and didn't overfish will be passed over. Investment managers who sold the toxic securities were the ones who got the big bonuses.

Sometimes this is incremental. If your colleagues are all overfishing 2%, then overfishing 3% isn't a big deal. But then it becomes 5%, 7%, 10%, and so on. As long as the incentive structure rewards doing slightly better than your colleagues, the incentive to defect remains. You get what you reward.



Larger organizations are naturally nested: departments within corporations, agencies within the government, units within a larger military structure, states within a country, and so on. This nested structure regularly leads to societal dilemmas. They're much like the employee societal dilemmas—should he work hard for the group interest of the company, or slack off for his own self-interest—but a subgroup inside the organization is the actor, rather than an individual. Should an airport screener act in the best interest of the TSA or in the best interest of the federal government? Should an employee act in the best interest of his department, his office, or his company as a whole?

I once worked for a company that had rigid rules about controlling costs. Those rules were implemented by department, not company-wide. The idea, of course, was that cost minimization at the smaller level would translate to cost minimization across the entire company. But sometimes it didn't work that way. I remember several instances where I had a choice between an action that would cost my department more, and an action that would cost my department less but would—because of costs to other departments—cost the company more. For example, I could fly a multi-city itinerary on several more expensive tickets, each allocated to the department that was responsible for that particular city. Or I could fly on a single cheaper ticket. Of course, my boss told me to choose the option that cost our department less, because that's how he was rewarded.

There are other competing interests within organizations: profits, perks (use of the corporate jet, for example), the corporate brand, an alternate idea of what the corporate brand should be, and so on. There are lots of these sorts of conflicts of interest in the investment banking world, such as the [conflict between the group that takes companies public and the group that recommends stocks to investors](#).⁸ [A full discussion](#) of that would take an entire book.



13

Corporations

Everything we discussed in the previous chapter applies to corporations, and some of the examples we used in the previous chapter were corporations. But because they are actors in so many societal dilemmas—they're legal persons in some countries—they warrant separate discussion. But before examining how societal dilemmas affect corporations, we need first to understand the basic supply-and-demand mechanics of a market economy as a pair of societal dilemmas.

Suppose a local market has a group of sandwich merchants, each of whom needs to set a sale price for its sandwiches. A sandwich costs \$4 to make, and the minimum price a merchant can sell them at and stay in business is \$5. At a price of \$6 per sandwich, consumers will buy 100 of them—sales equally divided amongst the merchants. At a sale price of \$5 per sandwich, consumers will buy 150—again, equally divided. If one merchant's prices are lower than the others', the undercutter will get all the business.

The merchants face a societal dilemma, an Arms Race akin to the advertise-or-not example in Chapter 5. It's in their collective group interest for prices to remain high; they collectively make a greater profit if they all charge \$6 for a sandwich. But by keeping their prices high, each of them runs the risk of their competitors acting in their self-interest and undercutting them. And since they can't trust the others not to do that, they all preemptively lower their prices and all end up selling sandwiches at \$5 each. In economics this is known as the "race to the bottom."

Societal Dilemma: Setting prices.	
Society: All the merchants.	
Group interest: Make the most money as a group.	Competing interest: Make the most money individually, and in the short term.
Group norm: Keep prices high.	Corresponding defection: Undercut the competition.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: The group encourages loyalty.</p> <p>Reputational: The group reacts negatively to those who break the cartel.</p> <p>Institutional: Various price-fixing schemes.</p> <p>Security: Internet price-comparison sites.</p>	

This societal dilemma is in continuous force. Day after day, month after month, the merchants are under constant temptation to defect and lower their prices, not just down to \$5, but even lower, if possible. The end result is that all of them end up selling sandwiches as cheaply as they possibly can, to the benefit of all the customers.

It's obvious how to solve this: the merchants need to trust each other. Like the mall stores at the beginning of Chapter 9, they can collectively agree to sell sandwiches at a minimum price of \$6 because they know it benefits them as a group. This practice was common throughout history. The medieval guild system was a way for sellers to coerce each other into keeping prices high; it was illegal to engage in trade except through the guild, and the system was enforced by the king. Cartels are a more modern form of this; oligopolies are another. Another way is to convince the government to pass a law outlawing cheaper sandwiches. Whatever name you use, the **result is price-fixing**.

Merchants like doing this, because keeping prices high is profitable. As **Adam Smith** said, "People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices."

Price-fixing has had varying degrees of success throughout history.¹ Sometimes it lasts for a long time. De Beers has successfully controlled the diamond market and kept prices artificially high since the 1880s. And sometimes it collapses quickly—the **global citric acid cartel** lasted only four years and the DRAM

computer-memory cartel just three. Sometimes buyers, such as Gateway and Dell in the DRAM price-fixing case, have a hand in breaking cartels, but it's usually government. Similarly, it's usually government that helps support them. Smuggling and other commerce often take place outside the cartel, but the cartel still works as long as they're kept to a minimum.

That's not good enough for a modern market economy. It is a basic tenet of capitalism that competition—sellers competing for buyers—rather than cartels are what should set prices. Capitalist society wants universal defection amongst sellers, because we recognize that a constant downward pressure on prices benefits the economy as a whole.

What we realize is that there's another societal dilemma functioning simultaneously and competing with the first.

Societal Dilemma: Setting prices.	
Society: Society as a whole.	
Group interest: Competition. Group norm: Do not collude in setting prices.	Competing interest: Make the most money as a group. Competing norm: Keep prices high.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: The belief that price-fixing is wrong and that competition is good.</p> <p style="padding-left: 40px;">Reputational: Being known as the merchant with the lowest price gives you an advantage, and being known as a price-fixer makes you look sleazy.</p> <p style="padding-left: 40px;">Institutional: Anti-trust laws.</p> <p style="padding-left: 40px;">Security: Various price-comparison websites.</p>	

Each merchant is in a societal dilemma with all of the other sandwich sellers; they're also in a larger societal dilemma with all the rest of society, including all the other sandwich sellers. Cooperating in one means defecting in the other, and in a modern market economy, the latter dilemma takes precedence.²

This works to the buyer's advantage, although more in theory than in practice. The previous societal dilemma pushes prices down only when there are more salable goods than there are buyers, and sellers are competing for buyers.

In some cases, the buyers can get stuck in a societal dilemma as well, pushing prices up. This is the other half of a market economy: buyers competing with each other. Imagine that a sandwich seller has twenty sandwiches left, and there

are forty people who want to buy one—including customer Bob. The normal price for the sandwich is \$5, but the seller has raised his price to \$6.

Here's the new societal dilemma. Bob is actually willing to pay \$6 for the sandwich, but he'd rather get it for \$5. So would everyone else. If everyone cooperated and refused to pay \$6 for a sandwich, the seller would eventually be forced to lower his prices. But there's always the incentive to defect—and be sure of getting a sandwich—rather than cooperate so that everyone who gets a sandwich pays only \$5.

Societal Dilemma: Competing on to-buy prices.	
Society: All the customers.	
Group interest: Keep prices low.	Competing interest: Getting the item you want.
Group norm: Don't bid up the price of items.	Corresponding defection: Differing to pay more for an item.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: It's unfair to bid up merchandise.</p> <p>Reputational: There are negative reputational consequences for bidding up merchandise and for overpaying.</p> <p>Institutional: None.</p> <p>Security: None.</p>	

Of course, this kind of thing never happens at sandwich shops. But it regularly happens in real estate markets, when buyers bid amounts higher than the asking price in order to out-compete other buyers for properties. It also happens with popular concerts and sporting events, where scalpers create a secondary market with higher prices as more buyers compete for a limited number of seats.

Auctions are fueled by this societal dilemma. As long as there are more bidders who want an item than there are items, they'll compete with each other to push prices as high as possible. And auctions implement societal pressures to prevent buyer collusion. For example, eBay makes it difficult for buyers to contact each other and collude.

A similar mechanism occurs with clothing in department stores. All department stores eventually mark down their seasonal inventory to get rid of it. Selling it cheap, or even at a loss, is better than keeping it on the shelves or in a storeroom somewhere. If Alice finds something she wants to buy early in the

season, she is faced with a societal dilemma. If she cooperates with everyone else and refuses to buy the clothing at full price, eventually the entire inventory will be discounted—drastically. But she risks others defecting and buying the garments at full price, and there not being any left of what she wants at the end of the season for the store to discount. Some discount retailers such as Outnet.com explicitly make use of this societal dilemma in their sales techniques. A garment starts out at full price, and is discounted more each week, until it reaches a final—very large—discount. Shoppers are truly faced with a societal dilemma: buy now at the higher price, or wait for a lower price and potentially lose the garment to someone else.³ Many antique shops and consignment stores use this strategy, too. As long as multiple buyers want the same item, it works.⁴

On the other hand, traditional buying clubs allow buyers to cooperate and push prices down. In addition to minimizing distribution and presentation costs, Costco and Sam's Club negotiate lower prices on behalf of their members.

Both of these pairs of societal dilemmas assume that, within each subgroup, buyers, sellers, and sandwiches are interchangeable. But of course that's not the case. Humans are a species of innovators, and we're always looking for ways to sell more profitable sandwiches and buy cheaper ones. The seller has two basic options:

- Merchant Alice can sell a cheaper sandwich. If Merchant Alice can substitute cheaper ingredients or use a cheaper sandwich-making process, she can either sell her sandwiches more cheaply than the competition or sell them at the same price with a greater profit margin—both options making her more money. It might not work. If the customers notice that Alice's sandwiches are of poorer quality than Bob's, they'll value them less. But if the customers don't notice that the sandwiches are any worse, then Alice deserves the increased business. She's figured out a way to make sandwiches cheaper in a way that makes no difference to the customer.⁵
- Merchant Alice can sell a better sandwich. Maybe she finds more expensive but tastier ingredients, or uses a more complicated sandwich-making process. Or she could make the sandwich-buying experience better by serving it with a smile and remembering her regular customers' names. She can either sell that better sandwich at the same price, bringing her more customers and more profit, or she can sell the better sandwiches at a more expensive price—whatever price the customers think those new sandwiches and the premium experience are worth. Of course, this requires that the customers value this better sandwich more. If they do, then Alice also deserves the increased business.

Both of these things happen all the time. Innovation is one of the important things a market economy fuels. On the buyer's side, the ways for customers to innovate are more limited.

Yes, this is all basic supply-and-demand economics; but it's economics from the perspective of societal pressures. You can look at a market economy as two different pairs of competing societal dilemmas: one preventing sellers from colluding, and the other preventing buyers from colluding. On a local scale, moral and reputational pressure largely enforces all of this. As long as buyers know the prices sellers are selling at and the sellers know what buyers are willing to pay—and this is generally true in local public markets—competition works as a price-setting mechanism. And if there are enough sellers, it's hard for them to collude and fix prices; someone is bound to defect and undercut the group. Sellers can try to differentiate their products from each other—either by selling less-desirable variants at a cheaper price or more-desirable variants at a higher price—and buyers will compete against each other to set new prices. The best way to succeed in this marketplace is to offer the best products at the lowest prices: that is, to have the best reputation for quality and price. There need to be enough buyers and sellers to make the market fluid, and enough transparency that the buyers know what they're buying; but if those things are true, then it all works.

It's only when you scale things up that these systems start failing. Societal pressures don't work the same when the sellers are large corporations as they do when they're sole proprietors in a public market. They don't work the same when the products are complicated—like cell phone plans—as they do when the products are simple. They don't work the same when commerce becomes global. They don't work the same when technology allows those corporations to defect at a scope larger than their own net worth.



During the early years of Prohibition, there was an epidemic of paralysis in the American South and Midwest, caused by “Jamaica Ginger,” a popular patent medicine. It was mostly alcohol,⁶ but about 500,000 bottles were laced with what turned out to be a nerve poison. It's hard to imagine a reputational pressure system being effective enough to prevent this kind of thing from happening. Sure, the company that sold this product was vilified, but not before tens of thousands of people were affected. (The “United Victims of [Ginger Paralysis Association](#)” had 35,000 members.) And, in fact, [this incident led](#) to the passage

of the 1938 Food, Drug, and Cosmetic Act and the establishment of regulations requiring pre-market approval for drugs.

Corporations are organizations. They come in all sizes. The company that made all that Jamaican Ginger consisted of two guys and an office; many corporations employ more than 100,000 people; and Wal-Mart employs over 2,000,000. They have some of the same characteristics as individuals—they try to maximize their trade-offs, they have a self-preservation instinct, etc.—but they are not individuals. In some very important ways, they differ from individuals.

These differences may affect corporations' defection characteristics:

- *They have a single strong self-interest: the profit motive.* The case can be made that it's the only relevant interest a corporation has. A corporation is legally required to [follow its charter](#), which for a for-profit corporation means maximizing shareholder value. Individuals have many more competing motivations.
- *They try to hire people who will maximize their selfish interest.* The people who run corporations, as well as the people promoted within them, tend to be willing to put the corporation's selfish interest (and sometimes their own selfish interest) ahead of any larger group interest. Individuals can't hire arms and feet selected to meet their needs.
- *They can be very large in several dimensions.* They can have a lot of assets, products, sales, stores, and employees. This increases their potential scope of defection: they can defect with greater frequency, and each defection can have greater intensity.
- *They can spread themselves over a large geographical area, so much so that they become unmoored from any physical location.* This reduces the effectiveness of institutional pressure that's tied to physical location: laws. It also reduces moral and reputational pressure against senior executives in those corporations, as they can remain socially isolated from those they harm.
- *They can be complex, especially if they're large.* This creates more internal subgroups at varying scales and intimacies, and the competing interests within them can change what they do. This gives them more options for evading accountability. It can also make it more difficult for people acting locally to determine what the competing interests actually are. Sometimes a single corporation can encompass different business units that compete directly with each other.

- *They can be powerful.* The combination of money and size can make corporations very powerful, both politically and socially. They can influence national and local legislation.⁷
- *Millions of people depend on corporations for their livelihood.* When a major corporation has problems—or even if it makes strategic decisions about automating, outsourcing, shutting down or starting up new product lines, and so on—many people and their families are affected. Whole communities can be affected. This means there are unintended consequences to many societal pressure systems.
- *They can be difficult to punish.* Corporate employees or owners are not the same as the corporation. Also, punishing a corporation can have ripple effects through society, hurting those who were in no way responsible for the corporation's misdoings.
- *They can live forever.* They are not tied to their founders, or to any particular people. They can live far longer than human lifespans.
- *They have more to lose than individuals do.* A damaged reputation can have much larger effects on corporations than on individuals, especially the big ones. This makes them more conservative.

Because of these differences, societal pressures work differently. **Moral pressure** is dampened in corporations. We've already seen in Chapter 9 that adding financial incentives tends to trump moral considerations. At the extreme, by telescoping the complexities of human morality into a wholly financial risk trade-off, corporations can largely relieve themselves of moral considerations. We also saw in Chapter 12 that morals are dampened in hierarchical group settings. The research is pretty clear on this point.

The upshot, to paint with a broad brush, is that corporations' risk trade-offs are much more focused on making a financial profit than individuals' are.⁸ People are emotionally complicated, and will regularly forgo money in exchange for more subjective benefits. Corporations, because of their group nature, are simpler; they are far more likely to choose the more profitable trade-off. To take a familiar example, it's far easier for a chef/owner of a restaurant to forgo some profit to create the sort of restaurant that gives him the most creative satisfaction, while a corporate-owned restaurant chain will be more concerned about consistency and the bottom line.

Another example is a garment or shoe designer buying goods made in **overseas sweatshops** staffed with child labor. An individual might refuse to do that on moral grounds, recognizing that she is going to have to pay more for those

goods made elsewhere and deliberately forgo the extra profits. A corporation is more likely to buy the goods, as long as it's legal to do so. And, as we've seen in Chapter 12, the person who is in charge of making this decision will do better personally if he ignores his own moral considerations and cooperates with his employer. Even worse, if the corporation doesn't maximize profits, it risks a shareholder lawsuit.

Additionally, market competition encourages sellers to ignore moral pressure as much as they can. Imagine if you were in a corporate boardroom, discussing the Double Irish tax loophole and how it could save your company millions. After it has been explained how the maneuver is perfectly legal, and how other companies are doing it, how far do you think a “but it's immoral” argument is going to go? Even if you don't want to do it, if you don't and your competitors do, you'll be uncompetitive in the marketplace—reminiscent of the sports doping example from Chapter 10. Morals have nothing to do with it; this is business. Likewise, on a smaller scale, hospitals tend to [replace management](#) teams who don't exploit Medicare billing loopholes, or engage in illegal upcoding, with teams that do.

Even when a corporation engages in seemingly altruistic behavior—investing in the community, engaging in charitable activities, pledging to follow fair labor guidelines, and so on—it is primarily doing so because of the value of increasing its reputation. It's only a bit over the top to call corporations “[immortal sociopaths](#),” as attorney and writer Joel Baken did. For corporations, the closest thing they have to morals is law. The analogy is pretty precise. Morals tell people what's right and what's wrong; the law tells corporations what's right and what's wrong. If corporations behave morally, it's generally because they believe it is [good for their reputation](#), and to a lesser extent because it's good for employee morale. This is less likely to be true with smaller corporations run by individuals or small groups of individuals; there, the corporation is more likely an extension of the person.

Or as [Baron Thurlow](#), a Lord Chancellor of England, put it sometime before 1792: “Corporations have neither bodies to be punished, nor souls to be condemned, they therefore do as they like.” In more modern language, John Coffee wrote that corporations have “[no soul to damn](#); no body to kick.”⁹

Reputational pressure can also fail against corporations. There's a belief that the market's natural regulation systems are sufficient to provide societal pressure, and that institutional pressure—laws and regulations—are both unnecessary and have harmful side effects. From the perspective of this book, this is just another name for reputational pressure.

Let's take an example: toxins in bottled water. Assume there's no institutional pressures, only reputational. Consumers decide for themselves what sort of toxin levels they are willing to tolerate, and then either buy or don't buy the product. (The assumption here is that removing the toxins costs money, and will result in a more expensive bottle.) Companies that sell toxin-free water enjoy a good reputation. Companies that allow too much toxin in their bottled water face a diminished reputation, and as a result, will reduce those toxins in an attempt to repair their reputation. If this works, it effectively "regulates" the bottled water companies.

We already know how reputational pressures fail when arrayed against an individual, and those failures are even more likely in the case of corporate reputation.

- The corporation will try to manage its reputation. Just as a person tries to accentuate his good qualities and minimize his bad ones, corporations do the same. The difference is that corporations will employ people whose entire job is to do this. Corporate reputation management equals public relations, and corporations spend **a lot of money** on advertising—\$130 billion annually in the U.S. alone. The science of advertising has completely changed over the past couple of decades. Today, it's more like **psychological manipulation** with a healthy **dose of neuroscience**.¹⁰ As such, there can be a large difference between a corporation's behavior and what the public thinks is the corporation's behavior. It can be hard to remember the relative toxicity levels of different bottled water brands when the corporations are all engaged in advertising designed to make you believe you'll be more successful with the opposite sex if you would only drink their product.
- For reputation to work as a societal pressure system, there needs to be transparency. But consumers might not know enough about the relative toxicity levels to have it affect the reputation of the various companies. (They might not know what chemicals are in the water, they might not know at what concentrations those chemicals are toxic, and they might not know the toxic effects of those chemicals.) Corporations can be very private, especially about things that make them look bad. Sure, testing companies like Consumers Union can give consumers information about the various bottled water companies, but there seems to be very little demand for that sort of thing. Salience matters a lot, here. When you want a bottle of water, you're thinking about your thirst—not about independent

third-party evaluations of water quality. To give a real example, corporations have successfully fought the labeling of genetically modified foods, so consumers aren't able to decide for themselves whether to eat them.

- Corporations might co-opt the testing and rating process. Those “independent third-party evaluations” aren't always so independent, and without transparency, consumers won't know.
- The damage resulting from the bad behavior might be so severe that no reputational consequences would be enough. Imagine that the bottled water is toxic enough that people start dying. Sure, the company will be out of business. But that seems like an inadequate penalty for killing people. And while this is an extreme story, there are lots of real-world examples of corporate decisions resulting in long-term disease and even death. In 2007 and 2008, at least ten Chinese companies produced contaminated batches of the blood-thinning drug heparin, substituting a cheap synthetic ingredient for a costlier natural one. At least **150 people died** as a direct result of the contaminated drug; we may never know how many secondary deaths or related illnesses there were.
- There can be a long time lag between the bad behavior and the reputational consequences. If the toxin in the bottled water is slow-acting, people might not know about its effects for years or even decades. So a corporation could continue selling toxin-laced water for a long time before it suffered any reputational damage. Remember “I'll be gone, you'll be gone”? That's an economically rational self-interest strategy in that instance.
- Consumers might not be able to penalize the company that's making the bottled water. In an open-air market, customers know who their suppliers are. In the complex world of international outsourcing and subcontracting, it can be much harder. In 2011, **Cargill recalled** 36 million pounds of ground turkey because of salmonella risk. None of that turkey was sold under the Cargill name, making it difficult for customers to penalize Cargill. In 2005, the data broker **ChoicePoint allowed** a criminal group to steal the identifying information of 140,000 consumers. If consumers wanted to penalize the company by not doing business with them anymore, they couldn't—consumers aren't ChoicePoint's customers.
- The profit resulting from the bad behavior might be large enough that it'd be worth the reputational loss. If customers have no choice but to buy the bottled water—maybe there's no competition and the groundwater is even more toxic—then the corporation doesn't have to worry about what customers will think. Less-extreme versions of this scenario happen

all the time in the real world; many industries benefit from the difficulty customers have in switching to a competing product.¹¹

All this is made worse by the various substitutes people use in place of direct reputation when it comes to brands. **There's recognition**: people buy what is familiar to them. **There's social proof**: people buy what others buy. There's even something called **attribute substitution**: people buy the red bottle because they like the color red and don't have any other way of choosing. These are some of the reasons consumers can be manipulated so easily.

Reputation relies on transparency to work, but for many modern products, the seller knows a lot more than the buyer. There's a general economic theory about this, called **a lemons market**. Both experiment and observation demonstrate that in a lemons market, bad products drive out good products. That is, if one company is selling cheap toxic water—or cheap unhealthy sandwiches—and the buyer doesn't know the difference between the good products and the cheap ones, he'll buy the cheap ones, and competitors will be pressured to make their products equally cheap and equally bad.

What we know about reputational pressures is that they work best in small groups where there are strong social ties among the individuals. A sandwich seller in a local public market probably doesn't need a whole lot of institutional pressure. He's part of a community, and if his sandwiches start making people sick fast enough that they notice the connection, no one will buy them anymore. But just as this sort of security system doesn't scale for individuals as the community gets larger, social ties weaken, and the value of the items being bought and sold increases, it doesn't scale for corporations, either. **Globalization is making** the effects of reputational pressure weaker. As a result, the effects of defection are greater. Three examples:

- In 2011, the pharmaceutical giant **Glaxo Smith-Kline** was fined \$750 million for marketing drugs manufactured in a Puerto Rican plant whose managers ignored numerous FDA letters warning that products were likely contaminated.
- Hundreds of people in Haiti, Panama, and Nigeria **died of kidney failure** in the 1990s and 2000s after consuming medicinal syrups manufactured with toxic diethylene glycol—an industrial chemical used to make plastics. Economically minded manufacturers had secretly substituted the toxic chemical for the more expensive, but nontoxic, glycerin.
- Starting in the mid-1990s, the Ford Motor Company knew that its Explorer model was **prone to rollover**, but didn't do anything to fix the

problem until 2002. Until they did, there were 185 deaths and 700 injuries resulting from the problem.

Just as moral and reputational pressures can fail against corporations, so can institutional pressures. We've discussed some of the ways they fail against individuals in Chapter 9: interpretation, loopholes, lack of enforcement. These failures can be more severe in corporations, because corporations can afford more and better lawyers to figure out how to evade laws. And law enforcement is much more consumer-friendly when it comes to dealing with individual defectors. If someone steals your wallet, you know how to call the police. If a corporation breaks the law, whom do you call?

Fines can be an effective institutional penalty, but can fail if they're too small. The DeCoster family egg farms, responsible for the huge [salmonella outbreak](#) in 2010, had been repeatedly fined for health violations for over ten years. In 2011, the large pharmaceutical company Merck Serono agreed to pay a \$44.5 million fine for illegally marketing [the drug Rebif](#). That sounds like a lot, until you realize that the annual sales of the drug were \$2.5 billion and the misconduct occurred over an eight-year period. It's no wonder the firm was a repeat offender; the fines were just a cost of doing business. Another example: the penalties for [using child labor](#) are so small in some countries—\$59 to \$147 in Egypt, \$470 in India, \$70 in Kenya, \$47 to \$470 in Nicaragua, \$25 to \$253 in the Philippines—that it makes financial sense for Western companies to defect. In Chapter 11, I mentioned the fake anti-virus industry. One company largely ignored the Federal Trade Commission prosecution because it was making more money than the fine was likely to be.¹²

We discussed other societal pressure failures inside corporations in the previous chapter: employees of a corporation defecting from that corporation, employee loyalty that encourages cooperation with the corporation and defection from society as a whole, and employees defecting from a corporation to benefit that corporation. Additionally, two of the differences between corporations and people listed above—that millions of people depend on them for their livelihood and that punishing them can have ripple effects through society—mean that sometimes it's in society's best interest to not punish defecting corporations: a fact a smart corporation can use to its advantage.

There is one more societal pressure failure that is unique to large and powerful corporations: the co-option of institutional pressure to further their own self-interest.

Imagine a societal dilemma, one that affects a rich and powerful interest: probably a corporation or an industry, but maybe a person or group of people. It could

be the [oil industry wanting](#) government subsidies (in 2011, the U.S. effectively provided \$4.4 billion in tax breaks to this industry alone, not even counting the military costs to protect their supply chains); or the Walt Disney Corporation wanting the government to extend the period of copyright so Mickey Mouse doesn't fall into the public domain. The group interest is to resolve the dilemma fairly. The self-interest for the corporation is to resolve the dilemma in its favor.

Societal Dilemma: Getting public money for projects.	
Society: Society as a whole.	
Group interest: Distribute government money fairly and maintain a level playing field. Group norm: Play by the rules.	Competing interest: Get as much money as you can for your pet projects. Corresponding defection: Manipulate the rules.
To encourage people to act in the group interest, the society implements a variety of societal pressures.	
Moral: It can feel wrong to take too much from the government.	
Reputational: It can look bad to take too much from the government.	
Institutional: Laws determine what benefits different interests get, and prohibit any one interest from taking too much.	
Security: The Congressional Record provides evidence of some of this, assuming anyone actually reads it. There are now websites that try to track political donations.	

If a company can convince the government to resolve the dilemma in its favor, then its self-interest becomes the group interest. In this way, companies can defect in spirit by deliberately changing the laws so they are not defecting in practice—thereby circumventing or subverting societal pressures. So, for example, companies that make car seats, airbags, full-body scanners, compact fluorescent bulbs, car insurance, surveillance cameras, vaccines, radon detectors, and Internet filters for schools have had laws passed mandating—or at least encouraging—their use. And the healthcare industry got a [law passed limiting](#) its liability for care improperly delayed or denied.

In a sense, what corporations are doing here is reversing the principal-agent relationship. They're deliberately manipulating institutional pressures so they can directly benefit from them. In economics, changing laws to suit your desires without adding any value is known as rent-seeking.

One way to manipulate laws is through licensing requirements. Over the past several years, there have been debates in several states about [licensing interior](#)

designers. It's either a necessary measure to keep charlatans out of the business, or an onerous, pro-cartel, anti-competitive system. Another way is through public opinion. The political decision not to regulate the derivatives markets is a good example: not only did it involve lobbyists and campaign contributions to get laws changed, but also public relations to convince journalists and the public that keeping the markets unregulated was a good idea.

Here's another example. Hydraulic fracturing, or fracking, is a means of extracting oil and gas from subterranean reservoirs by forcing pressurized fluid into underground rock formations. The process was originally commercialized in 1949 and in its first few decades of use was primarily used to boost production of old wells. Recent advances in **horizontal drilling technology**, combined with hydraulic fracturing, have enabled the tapping of heretofore inaccessible reserves, and the recent rise in oil prices has made it economically viable. However, the procedure also poses environmental risks, most notably the risk that chemicals used in the process—including methanol, benzene, and diesel fuel—might contaminate ground water, degrade air quality, and migrate to the earth's surface; and that the resultant toxic wastewater might be impossible to decontaminate.¹³ This societal dilemma sounds a lot like the monk parakeet example from Chapter 9, and you'd expect society to figure out whether this procedure is worth it. But the companies that use the procedure—Halliburton is a big player here—lobbied successfully for a provision in the 2005 Bush administration energy bill exempting fracking from regulation by the U.S. Environmental Protection Agency under the Safe Drinking Water Act.¹⁴ That's the effect of reversing the principal-agent relationship: the government becomes the agent of the corporation.

One common way to do this is regulatory capture, which we'll talk about in the next chapter. Another way is to simply be unregulatable for political or economic reasons. Homebuilders have been sued repeatedly over the past decade for **shoddy building practices**, many of them illegal. "Too big to regulate" is how one source put it, making it impossible for homeowners to know they're getting a substandard house until it's too late. The banking industry is similarly trying very hard to be unregulatable, claiming that any regulations would damage the economy more than it would help it.



When it comes to organizations, size is proportional to power. Legislative bodies used to rule fewer people and smaller geographic areas. In the

United States, many laws that were passed by states in the 1800s became federal matters in the 1900s. There's nothing sinister about it; it's just that it now makes more sense to deal with these laws on that scale. Today, international legislative bodies have increasingly more power—simply because more things make sense to deal with on a multinational level.

This is especially true in corporations. Broadly speaking, there's a natural size of an organization based on the technology of its time. The average organization size [used to be smaller](#), became larger, and now is even larger. Historically, there have only been a few very large organizations: the Roman Empire, the Catholic Church, and so on. These worked because they were organizations of organizations. That's how countries work; the U.S. has federal, state, and municipal governments. That's also how feudalism, militaries, franchise stores, and large multinational corporations work.

It still works this way, but we're better at it now. Organizational size is restricted by the limits of moving information around. Different people within, and different parts of, an organization need to communicate with each other; and the larger an organization, the harder that is to do. Most organizations are hierarchical, making communications easier. And militaries have generally been examples of the largest-sized organization a particular technological level can produce. But there's a limit where the costs of communications outweigh the value of being part of one organization. Economist [Ronald Coase first](#) pointed this out in 1937. Called "Coase's limit" or "Coase's ceiling," it's the point of diminishing returns for a company: where adding another person to an organization doesn't actually add any value to the organization. You can think of an employee inside of an organization having two parts to his job: coordinating with people inside the organization and doing actual work that makes the company money. Some people are wholly focused inside the organization: the HR department, for example. Others do the actual work, but still have internal coordination roles. There's a point where adding an additional person to the organization increases the internal coordination for everyone else to a point that's greater than the additional actual work he does. So, the company actually loses money overall by hiring him.¹⁵ The ease of collecting, moving, compiling, analyzing, and disseminating information affects Coase's ceiling, and one of the effects of information technology is that it raises Coase's ceiling because the resultant efficiency increases.¹⁶

Larger size has several effects on societal dilemmas:

- Large corporations can do more damage by defecting. A single company, Enron, did \$11 billion worth of financial damage to the U.S. economy.

That much damage might previously have required ten smaller companies to defect. This means that as large corporations grow, fewer defectors can do even more damage. So society needs more security, to further reduce the amount of defection, in order to keep the potential damage constant.

- Individuals within a large corporation can defect from the corporation to a greater degree, for greater personal gain and to the greater detriment of the corporation. [Nick Leeson's](#) unauthorized trading while he worked for Barings Bank destroyed the entire company in 1995. Kenneth Lay, Jeffrey Skilling, and other senior Enron executives destroyed that company. [Kweku Adoboli](#) lost \$2.3 billion for the investment bank UBS in 2011.
- Large corporations have more power to deliberately manipulate societal pressures. This includes getting laws passed specifically to benefit them, and engaging in jurisdictional arbitrage by deliberately moving certain operations to certain countries in order to take advantage of local laws. [Different countries](#) have different, often conflicting, laws about price-fixing, and international companies have an easier time forming cartels. This sort of thing can be more local, too. Until recently, [Amazon.com used its](#) large national footprint and lack of physical stores to avoid having to charge sales tax in most states.
- Punishing a large corporation might result in so much cost or damage to society that it makes sense to let them get away with their wrongdoing. The ultimate expression of this is when a company is “too big to fail”: when the government is so afraid of the secondary effects of a company going under that they will bail the company out in order to prevent it.¹⁷
- Individuals within large corporations can be emotionally further away from the individuals they're affecting when they make decisions about whether to cooperate or defect. Remember that moral pressure decreases in effectiveness with emotional distance. The larger the corporation, the larger the tendency towards emotional distance.
- Larger corporations have more to lose by defecting. Their reputation is more valuable, and damage to it will have greater effects on the corporation. This serves to restrict what they're willing to do.

Large corporations can also play one societal dilemma off another. Remember our sandwich seller in the market. He's stuck in a societal dilemma with all the other sandwich sellers, and has to set his prices accordingly. In order to prevent the market's sandwich sellers from cooperating, society as a whole—as part of a larger societal dilemma—passes laws to prevent collusion and price-fixing. But

a larger sandwich seller has more options. He can expand his product offering across several dimensions:

- *Economies of scale.* He can buy his ingredients in bulk and streamline his production processes.
- *Depth.* More sandwich options.
- *Size.* Larger or smaller sandwiches.
- *Time.* Breakfast sandwiches or sandwiches for midnight snacks.
- *Scope.* Sandwich-like things, such as hot dogs, bagels, wraps, and muffins.
- *Accessories.* Chips and sodas, groceries.
- *Service.* Sandwich subscriptions, delivery, free wi-fi to go along with the sandwiches.

All this makes it much more difficult to enforce the basic societal dilemmas of a market economy. On the face of it, as a seller diversifies, he is now stuck in multiple different societal dilemmas: one with the other sandwich sellers in the market, and another with—for example—chip sellers. But by tying the two products together, perhaps selling a sandwich and chips together, or offering a once-a-week chip subscription with the purchase of a sandwich subscription, he is able to play the two societal dilemmas off each other, taking advantage of both.

We see this with various product schemes. Whether it's Citibank selling credit cards and consumer loans and anti-theft protection plans to go with those credit cards; or Apple selling computer hardware and software; or Verizon bundling telephone, cable, and Internet; product bundles and subscription services hide prices and make it harder for customers to make buying decisions. There's also a moral hazard here. The less Citibank spends on antifraud measures, the more protection plans it can sell; the higher its credit card interest rates, the more attractive its consumer loans are.

Large corporations can also use one revenue stream to subsidize another. So a big-box retail store can temporarily lower its prices so far that it's losing money, in order to drive out competition. Or an airline can do the same with airfares in certain markets to kill an upstart competitor.

Things get even more complicated when sellers have multiple revenue streams from different sources. Apple sells iPhones and iPads to customers, sells the ability to sell customer apps to app vendors, and sells the right to sell phone contracts to phone companies. Magazines sell both subscriptions and their subscription lists. This sort of thing is taken to the extreme by companies like Facebook, which don't even charge their users for their apps at all, and make

all their money selling information about those users to third parties.¹⁸ It turns out that offering a product or service for **free is very different** than offering it cheaply, and that “free” perturbs markets in ways no one fully understands. The optimal way to do business in an open-air market—offer the best products at the lowest prices—fails when there are other revenue streams available.

An additional complication arises with products and services that have high barriers to entry; it’s hard for competitors to emerge. In an open-air market, if the sandwich vendors all sell their sandwiches at too-high prices, someone else can always come in and start selling cheaper sandwiches. This is much harder to do with cell phone networks, or computer operating systems, or airline tickets, because of the huge upfront costs. And industries can play the meta-game to prevent competition, as when the automobile industry bought and then dismantled cities’ trolley networks, big agriculture lobbied government to impose draconian regulations on small farms, and so on.

There’s one more problem with the technological corporations that doesn’t really exist on the small scale of an open-air market: the risks of defection can be greater than the total value of the corporations themselves. An example will serve to explain.

Chemical plants are a terrorism risk. Toxins such as phosgene, chlorine, and ammonia could be dispersed in a terrorist attack against a chemical plant. And depending on whose numbers you believe, **hundreds of plants** threaten hundreds of thousands of people and some threaten millions. This isn’t meant to scare you; there’s a lot of debate on how realistic this sort of terrorist attack is right now.

In any case, the question remains of how best to secure chemical plants against this threat. Normally, we leave the security of something up to its owner. The basic idea is that the owner of each chemical plant best understands the risks, and is the one who loses out if security fails. Any outsider—in this case, a regulatory agency—is just going to get it wrong.

And chemical plants do have security. They have fences and guards. They have computer and network security. They have fail-safe mechanisms built into their operations.¹⁹ There are regulations they have to follow. **The problem is** that might not be enough. Any rational chemical-plant owner will only secure the plant up to its value to him. That is, if the plant is worth \$100 million, it makes no sense to spend \$200 million on securing it. If the odds of it being attacked are less than 1%, it doesn’t even make sense to spend \$1 million on securing it. The math is more complicated than this, because you have to factor in such things as the reputational cost of having your name splashed all over the media after an incident, but that’s the basic idea.

But to society, the cost of an actual attack could be much, much greater. If a terrorist blows up a particularly toxic plant in the middle of a densely populated area, deaths could be in the tens of thousands and damage could be in the hundreds of millions. Indirect economic damage could be in the billions. The owner of the chlorine plant would pay none of these costs; to him, they are externalities borne by society as a whole.

Sure, the owner could be sued. But he's not at risk for more than the value of his company, and the outcome of a lawsuit is by no means preordained. Expensive lawyers can work wonders, courts can be fickle, and the government could step in and bail him out (as it did with airlines after 9/11). And a smart company can often protect itself by spinning off the risky asset in a subsidiary company, or selling it off completely. Mining companies do this all the time.

The result of all this is that, by leaving the security to the owner, we don't get enough of it.

In general, the person responsible for a risk trade-off will make the trade-off that is most beneficial to *him*. So when society designates an agent to make a risk trade-off on its behest, society has to solve the principal-agent problem and ensure that the agent makes the same trade-off that society would. We'll see how this can fail with government institutions in the next chapter; in this case, it's failing with corporations.

Think back to the sandwich sellers in the local market. Merchant Alice is one of those sandwich sellers, and a dishonest, unscrupulous one at that. She has no moral—or reputational—issues with potentially poisoning her buyers. In fact, the only thing that's standing in the way of her doing so is the law. And she's going to do the math.

She has the opportunity of making her sandwiches using some substandard but cheaper process. Maybe she's buying ingredients that aren't as clean. Whatever she's doing, it's something that saves her money but is undetectable by her customers.

If her increased profit for selling potentially poisonous sandwiches is \$10,000, and the chance of her getting caught and fined is 10%, then any fine over \$100,000 will keep her cooperating (assuming she's rational and that losing \$100,000 matters to her).

Now consider a large sandwich corporation, ALICE Foods. Because ALICE Foods sells so many more sandwiches, its increased profit from defecting is \$1,000,000. With the same 10% probability of penalty, the fine has to be over \$10,000,000 to keep it from defecting. But there's another issue. ALICE Foods only has \$5,000,000 in assets. For it, the maximum possible fine is everything

the corporation has. Any penalty greater than \$5,000,000 can be treated as \$5,000,000. So ALICE Foods will rationally defect for any increased profit greater than \$500,000, regardless of what the fine is set at (again, assuming the same 10% chance of being fined and no semblance of conscience).

Think of it this way. Suppose ALICE Foods makes \$10,000,000 a year, but has a 5% chance of killing lots of people (or of encountering some other event that would bankrupt the company). Over the long run, this is a guaranteed loss-making business. But in the short term, management can expect ten years of profit. There is considerable incentive for the CEO to take the risk.

Of course, that incentive is counteracted by any laws that ascribe personal liability for those decisions. And the difficulty of doing the math means that many companies won't make these sorts of conscious decisions. But there always will be some defectors that will.

This problem occurs more frequently as the value of defecting increases with respect to the total value to the company. It's much easier for a large corporation to make many millions of dollars through breaking the law. But as long as the maximum possible penalty to the corporation is bankruptcy, there will be illegal activities that are perfectly rational to undertake as long as the probability of penalty is small enough.²⁰

Any company that is too big to fail—that the government will bail out rather than let fail—is the beneficiary of a free insurance policy underwritten by taxpayers. So while a normal-sized company would evaluate both the costs and benefits of defecting, a too-big-to-fail company knows that someone else will pick up the costs. This is a moral hazard that radically changes the risk trade-off, and limits the effectiveness of institutional pressure.

Of course, I'm not saying that all corporations will make these calculations and do whatever illegal activity is under consideration. There are still both moral and reputational pressures in place that keep both individuals and corporations from defecting. But the increasing power and scale of corporations is making this kind of failure more likely. If you assume that penalties are reasonably correlated with damages—and that a company can't buy insurance against this sort of malfeasance—then as companies can do more damaging things, the penalties against doing them become less effective as security measures. If a company can adversely affect the health of tens of millions of people, or cause large-scale environmental damage, the harm can easily dwarf the total value of the company. In a nutshell, the bigger the corporation, the greater the likelihood it could unleash a massive catastrophe on society.

In talking about group interests and group norms, I've mostly ignored the question of who determines the interests, sets the norms, and decides what scope of defection is acceptable and how much societal pressure is sufficient. It's easy to say "society decides," and from a broad enough viewpoint, it does. Society decides on its pair-bonding norms, and what sorts of societal requirements it needs to enforce them. Society decides how property works, and what sorts of societal pressures are required to enforce property rights. Society decides what "fair" means, and what the social norms are regarding taking more or doing less than your fair share. These aren't deliberate decisions; they're evolved social decisions. So just as our immune system "decides" which pathogens to defend the body against, societies decide what the group norms are and what constitutes defecting behavior. And just as our immune system implements defenses against those pathogens, society implements societal pressures against what it deems to be defection.

But many societal pressures are prescribed by those in power,¹ and while the informal group-consensus process I just described might explain most moral and reputational pressure, it certainly doesn't explain institutional pressure. Throughout most of our history, we have been ruled by autocrats—leaders of family groups, of tribes, or of people living in geographical boundaries ranging in size from very small to the Mongol Empire. These individuals had a lot of power—often absolute power—to decide what the group did. They might not have been able to dictate social norms, but they could make and enforce laws. And very often, those laws were immoral, unfair, and harmful to some, or even most, people in the group.

Throughout most of our history, people had no say in the laws that ruled them. Those who ruled did so by force, and imposed laws by force. If the monarch in

power decided that the country went to war, that's what the people did. The group interest was defined by what the king wanted, and those who ignored it and followed some competing interest were punished. It didn't matter if the majority agreed with the king; his word defined the group norm. "*L'État, c'est moi*" and all.²

I'm eliding a lot of nuance here. Few rulers, from tribal leaders to emperors, had—or have—absolute power. They had councils of elders, powerful nobles, military generals, or other interests they had to appease in order to stay in power. They were limited by their roles and constrained by the societies they lived in. Sometimes a charismatic and powerful ruler could radically change society, but more often he was ruled by society just as much as he ruled it. Sometimes group norms are decided by privileged classes in society, or famous and influential people, or subgroups that happen to be in the right place at the right time.

In parts of our history, laws and policy were decided not by one person but by a cohort: the ancient Roman Senate, the *Maggior Consiglio* in medieval Venice, the British Parliament since the Magna Carta. Modern constitutional democracies take this even further, giving everybody—more or less—the right to decide who rules them, and under what rules those rulers rule.

This dynamic isn't limited to government; it also plays out in other groups. Someone in charge decides what the group's norms are, constrained by the "rules" of his office. A CEO can be removed from office by the board of directors. A Mafia head can be deposed by a rival; criminal gangs and terrorist groups have their own organizational structures.

The deciders generally don't decide the details of the norms and societal pressures. For example, while the king might decide that the country will go to war and all able-bodied men are to be drafted into the army, he won't decide what sorts of security measures will be put in place to limit defectors. Society delegates the implementation of societal pressures to some subgroup of society. Generally these are institutions, which I'll broadly define as an organization delegated with implementing societal pressure. We've already discussed delegation and the principal-agent problem. We're now going to look at how that plays out with institutions.



In 2010, [full-body scanners](#) were rushed into airports following the underwear bomber's failed attempt to blow himself up along with an airplane. There are a lot of reasons why the devices shouldn't be used, most notably because they can't directly detect the particular explosive [the underwear bomber](#) used, and probably wouldn't have detected his underwear bomb. There have been several

court cases brought by people objecting to their use. One of them, filed by the Electronic Privacy Information Center, alleged the TSA didn't even follow its own rules when it fielded the devices. (Full disclosure: I was a plaintiff in that case.) I want to highlight an argument a Department of Homeland Security lawyer made in federal court. He contended that the agency has the legal authority to [strip-search every](#) air traveler, and that a mandatory strip-search rule could be instituted without any public comment or rulemaking. That is, he claimed that DHS was in charge of airline security in the U.S., and it could do anything—*anything*—it wanted to in that name.

After the [September 11 attacks](#), people became much more scared of airplane terrorism. The data didn't back up their increased fears—airplane terrorism was actually a much larger risk in the 1980s—but 9/11 was a huge emotional event and it really knocked people's feeling of security out of whack. So society, in the form of the government, tried to improve airport security. George W. Bush signed the Aviation and Transportation Security Act on November 19, 2001, creating the Transportation Security Administration.

Societal Dilemma: Airplane terrorism.	
Society: Society as a whole.	
<p>Group interest: Safe air travel.</p> <p>Group norm: Not to blow up airplanes.</p>	<p>Competing interest: Blowing up airplanes is believed to be an effective way to make a political point or advance a political agenda.³</p> <p>Corresponding defection: Blow up airplanes.</p>
<p>To encourage people to act in the group interest, society implements these societal pressures:</p> <p style="margin-left: 40px;">Moral: Our moral systems hold that murdering people and destroying property is wrong.</p> <p style="margin-left: 40px;">Reputational: Society punishes people who kill innocents, and even people who espouse doing that. In some cases, people are publicly vilified not because they themselves advocate violence, but because they aren't sufficiently critical of those who do.</p> <p style="margin-left: 40px;">Institutional: Nation states implement laws to fight airplane terrorism, including invasive passenger screening. We have severe punitive measures to deter terrorists, at least the non-suicide kind.</p> <p style="margin-left: 40px;">Security: Magnetometers, x-ray machines, swabs fed into machines that detect potential explosives, full-body scanners, shoe scanners, no-fly lists, behavioral profiling, and on and on.</p>	

The societal dilemma of airplane terrorism is a particularly dangerous one, because even a small number of defectors can cause thousands of deaths and billions of dollars in economic damage. People are legitimately concerned about this, and want strong societal pressures.⁴ Moral and reputational pressures aren't nearly enough, both because the [scale is too large](#) and the competing group interest is so strong. Institutional pressure is required, and the institution in the U.S. that has been delegated with this responsibility is the Transportation Security Administration.

There are actually several levels of delegation going on. The people delegate security to their leaders—Congress and the president—who delegate to the Department of Homeland Security, which delegates to the TSA, which delegates to individual TSA agents staffing security checkpoints.

Figure 12 illustrates how institutional pressure is delegated. Ultimately, institutions are put in charge of enforcement. These aren't always governments; they can be any subgroup of society given the power to enforce institutional pressure at any level, such as:

- The police, who implement societal pressures against a broad array of competing norms. (Okay, I admit it. That's an odd way to describe arresting people who commit crimes against people and property.)
- The judicial system, which 1) punishes criminals and provides deterrence against future defections, and 2) adjudicates civil disputes, providing societal pressures based on both formal and informal societal norms.
- Government regulatory agencies, such as the U.S.'s TSA, the Occupational Safety and Health Administration, the Federal Communications Commission, and the Food and Drug Administration.
- Industry organizations, which implement industry self-regulation. (This is often agreed to in order to forestall government regulation.)
- Corporate security offices, which implement the physical and data-security policies of a corporation.
- Corporate auditors, who 1) verify the same, and 2) verify the corporation's books, providing societal pressures against corporate financial malfeasance.
- An independent security company, hired by an organization to guard its buildings.

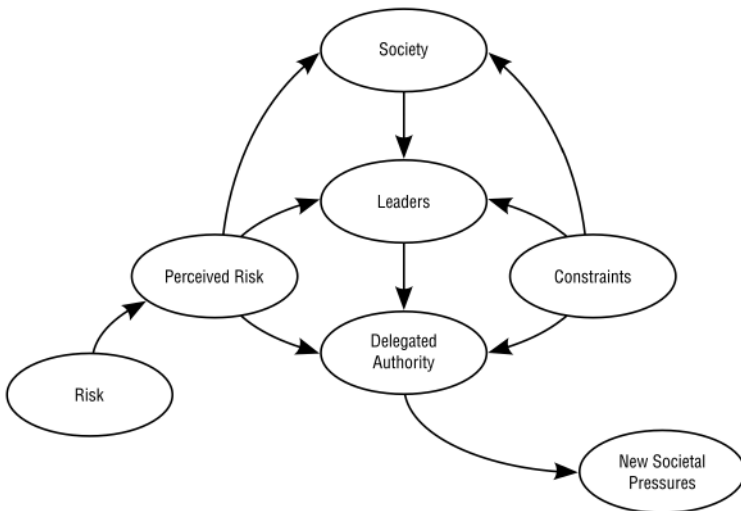


Figure 12: How Societal Pressures Are Delegated

The goal of delegation is for the institution to act as the group’s agent. That is, to implement societal pressures on behalf of, and in the name of, the group. But because of the principal–agent problem, that institution doesn’t have the same competing interests as the group as a whole—or even as any institution or subgroup above them. As a result, it won’t necessarily implement societal pressures to the same degree or in the same way as the group would like. And that’s an endless source of problems.



When it comes to terrorism and airplane security, those problems are legion. The TSA is a government institution with a mandate and funding from the U.S. government. It answers to the government. And the government has a mandate from, is funded by, and answers to, the people. Given all of that, you’d expect the people to have a lot of input into what the TSA does. Which is why it can seem so weird when it does things with absolutely no input from anyone. But it’s a natural effect of the principal–agent problem.

The TSA’s interests aren’t the same as those of any of the groups it’s an agent for: DHS, the government, or society as a whole.

For one, the TSA has a self-preservation interest. If it is seen as unnecessary—that is, if society as a whole believes there’s a sufficiently diminished terrorist threat—it might be disbanded. Or perhaps its function would be taken over by some

international security organization. In either case, like a person, the TSA is concerned about its own survival. (By the way, people working within the TSA are also concerned about their jobs, power, and reputation within the agency, and so on.)

For another, the TSA is concerned with its own reputation in the eyes of society. Yes, it wants to do a good job, but it also needs to be *seen* as doing a good job. If there's a terrorist attack, the TSA doesn't want to be blamed for not stopping the terrorists. So if a terrorist bombs a shopping mall instead of an airplane, it's a win for the TSA, even though the death toll might be the same.⁵ Even without an actual terrorist attack, if it is seen as doing a bad job—even if it's actually doing a good job—it will be penalized with less public support, less funding, and less power.⁶

Finally, the TSA is concerned about its relative power within the government. The more funding it has, and the closer it is to the president, the better job it can do and the more likely it is to survive.

Societal Dilemma: Implementing airplane security.	
Society: Society as a whole.	
	Competing interest: Selfish interest—garner as much power and prestige as it can. Corresponding defection: Get as much money for its budget as possible.
Group interest: Airplane security whose benefits exceeds the costs. Group norm: Implement airplane security at a reasonable level.	Competing interest: Self-preservation—ensure that it won't be disbanded by the government. Corresponding defection: Become an indispensable part of airplane security.
	Competing interest: ego preservation - ensure that if there is a terrorist attack, it won't be blamed. Corresponding defection: implement a greater level of airplane security than the risk trade-off warrants.
To encourage people to act in the group interest, society implements these societal pressures:	
Moral: We teach people to do the right thing. Reputational: Institutions that put their own survival ahead of their nominal missions aren't thought of very well. Institutional: Legislators and courts rein institutions in. Security: Auditors, inspectors, cameras, and monitoring.	

The TSA's competing interests are common in government agencies. You can see it with the police and other law-enforcement bodies. These institutions have been delegated responsibility for implementing institutional pressure on behalf of society as a whole, but because their interests are different, they end up implementing security at a greater or lesser level than society would have.

Exaggerating the threat, and oversecuring—or at least overspending—as a result of that exaggeration, is by far the most common outcome. The TSA, for instance, would never suggest returning airport security to pre-9/11 levels and giving the rest of its budget back so it could be spent on broader anti-terrorism measures that might make more sense, such as intelligence, investigation, and emergency response. It's a solution that goes against the interests of the TSA as an institution.

This dynamic is hardly limited to government institutions. For example, corporate security officers exhibit the same behavior. In Chapter 10, I described the problem of corporate travel expenses, and explained that many large corporations implement societal pressures to ensure employee compliance. This generally involves approval—either beforehand for things like airfare and hotels, or after-the-fact verification of receipts and auditing—of travel expenses. To do this, the corporation delegates approval authority to some department or group of people, which determines what sort of pressures to implement. That group's motivation becomes some combination of keeping corporate travel expenses down and justifying its own existence as a department within the corporation, so it overspends.

Recall the professional athletes engaging in an arms race with drug testers. It might be in the athletes' group interest for the sport of cycling to be drug-free, but the actual implementation of that ideal is in the hands of the sport's regulatory bodies. The World Anti-Doping Agency takes the attitude of “ban everything, the hell with the consequences.” It might better serve the athletes if the agency took more time and spent more money developing more accurate tests, was more transparent about its testing methodology, and had a straightforward redress procedure for athletes falsely accused—but it's not motivated to make that risk trade-off. And as long as it's in charge, it's going to do things its way.

Enforcing institutions have a number of other competing interests resulting from delegation. A common one has to do with how the enforcing institutions are measured and judged. We delegate to the police the enforcement of law, but individual policemen get reviewed and promoted based on their arrest and conviction rate. This can result in a variety of policing problems, including a police department's willingness to pursue an innocent person if it believes it can get a

conviction, and pushing for an easy conviction on a lesser charge rather than a harder conviction on a more accurate charge.

There's one competing interest that's unique to enforcing institutions, and that's the interest of the group the institution is supposed to watch over. If a government agency exists only because of the industry, then it is in its self-preservation interest to keep that industry flourishing. And unless there's some other career path, pretty much everyone with the expertise necessary to become a regulator will be either a former or future employee of the industry, with the obvious implicit and explicit conflicts. As a result, there is a tendency for institutions delegated with regulating a particular industry to start advocating the commercial and special interests of that industry. This is known as [regulatory capture](#), and there are many examples both in the U.S. and in other countries. U.S. examples include:

- The Minerals Management Service, whose former managers saw nothing wrong with [steering contracts](#) to ex-colleagues embarking on start-up private ventures, and having sexual relationships with and accepting gifts from oil and gas industry employees. In fact, the MMS was broken up in 2010 because this cozy relationship was blamed in part for the Deepwater Horizon oil spill.
- The [Federal Aviation Administration](#), whose managers' willingness to overlook or delay action on crucial safety problems contributed to the 1996 crash of a ValuJet Airlines DC-9 in the Everglades, and the 2011 sudden in-flight failure of a section of fuselage on a Southwest Airlines 737.
- The [Securities and Exchange](#) Commission, whose lawyers routinely move to government employment from the banking industry, and back after their term of service is over. One of the effects of this revolving door was a poorly regulated banking industry that caused the financial crisis of 2008.

One way to think about all this is as a battle between diffuse interests and concentrated interests. If you assume that specific regulations are a trade-off between costs and benefits, a regulatory institution will attempt to strike a balance. On one side is the industry, which is both powerful and very motivated to influence the regulators. On the other side is everyone else, each of whom has many different concerns as they go about their day and none of whom are particularly motivated to try to influence the regulators. In this way, even if the interests of society as a whole are greater than the interests of the industry, they're not as well-represented because they're so diffuse. And to the extent that the institution

is society's agent for implementing societal pressures, this becomes a colossal failure of societal interest. Moreover, each level of delegation introduces new competing interests, like a horribly conflicted game of telephone.



Institutions have power, and with that power comes the ability to defect. Throughout history, governments have acted in the self-interest of their rulers and not in the best interest of society. They can establish social norms and enforce those norms through laws and punishment. They can do this with or without the support of the people.

But there's a new type of potentially defecting institution, one that's made possible by the information age: corporations acting in the role of institutions. This can happen whenever public infrastructure moves into private hands. With the rise of the Internet as a communications system, and social networking sites in particular, corporations have become the designers, controllers, and arbiters of our social infrastructure. As such, they are assuming the role of institutions even though they really aren't. We talked in Chapter 10 about how combining reputational pressure with security systems gives defectors new avenues for bypassing societal pressures, like posting fake reviews on Yelp. Another effect is that the corporation that designs and owns the security mechanisms can facilitate defection at a much higher level.

Like an autocratic government, the company can set societal norms, determine what it means to cooperate, and enforce cooperation through the options on its site. It can take away legal and socially acceptable rights simply by not allowing them: think of how publishers have eroded fair use rights for music by not enabling copying options on digital players. And when the users of the site are not customers of the corporation, the competing interests are even stronger.

Take Facebook as an example. Facebook gets to decide what privacy options users have. It can allow users to keep certain things private if they want, and it can deny users the ability to keep other things private. It can grant users the ability to fine-tune their privacy settings, or it can only give users all-or-nothing options. It can make certain options easy to find and easy to use, and can make other options hard to find and even harder to use. And it will do or not do all of these things based on its business model of selling user information to other companies for marketing purposes. Facebook is the institution implicitly delegated by its users to implement societal pressures, but because it is a for-profit corporation and not a true agent for its users, it defects from society and acts in

its own self-interest, effectively reversing the principal–agent relationship. Of course, users can refuse to participate in Facebook. But as Facebook and other social networking sites become embedded in our culture and our socialization, opting out becomes less of a realistic option. As long as the users either have no choice or don't care, it can act against its users' interests with impunity.

It's not easy to implement societal pressures against institutions that put their competing interests ahead of the group interest. Like any other organization, institutions don't respond to moral pressure in the same way individuals do. They can become impervious to reputational pressure. Since people are often forced to interact with institutions, it often doesn't matter what people think of them. Yes, in a democracy, people can vote for legislators who will better delegate societal pressures to these institutions, but this is a slow and indirect process. You could decide to not use a credit card or a cell phone and therefore not do business with the companies that provide them, but often that's not a realistic alternative.

Sometimes the authorities are just plain unwilling to punish defecting institutions. No one in the U.S. government is interested in taking the National Security Agency to task for [illegally spying](#) on American citizens (spy agencies make bad enemies). Or in punishing anyone for [authorizing the torture](#) of—often innocent—terrorist suspects. Similarly, there's little questioning legislatively about President Obama's self-claimed [right to assassinate](#) Americans abroad without due process.

The most effective societal pressures against institutions are themselves institutional. An example is the lawsuit I talked about at the start of this chapter. EPIC sued the TSA over full-body scanners, claiming the agency didn't even follow its own rules when it fielded the devices. And while the [court rejected EPIC's](#) Fourth Amendment arguments and allowed the TSA to keep screening, it ordered the TSA to conduct notice-and-comment rulemaking. Not a complete victory by any means, but a partial one.

And there are many examples of government institutions being reined in by the court system. In the U.S., this includes judicial review, desegregating schools, legalizing abortion, striking down laws prohibiting interracial and now same-sex couples from marrying, establishing judicial oversight for wiretapping, and punishing trust fund mismanagement at the Bureau of Indian Affairs.

What's important here is accountability. It is important that these mechanisms are seen publicly, and that people are held accountable. If we're going to keep government from overstepping its bounds, it will be through separation of powers: checks and balances. But it's not just government that needs to be watched; it's corporations, non-government institutions, and individuals. It's everyone's responsibility to keep everyone else in check.

PART IV

Conclusions



15

How Societal Pressures Fail

Let's start our discussion of societal pressure failures with an example: taxes. Paying taxes is a classic free-rider problem; if almost everyone cooperates by paying taxes, defectors get all the benefits of whatever those taxes are paying for without having to suffer the financial penalties of actually paying.¹ There are laws and enforcement, but at least in the U.S., with the exception of payroll taxes, income tax is almost entirely enforced by voluntary compliance. It's not just a financial risk trade-off; there are two pieces of moral pressure at work here: [people paying taxes](#) because it's the right thing to do, and people paying taxes because it's the law and following the law is the right thing to do.

Still, there's [a lot of fraud](#) in the U.S. According to the IRS, in 2001—the most recent year I could find comprehensive numbers for—the difference between total taxes owed and total taxes paid was \$345 billion; about 19% of the total taxes due. A [third-party estimate](#) from 2008 tax returns also showed a 19% tax gap. Note that this gap is in the percentage of money owed, not the percentage of cheaters. By one estimate, 25% of individuals admit to [cheating on their taxes](#). On the other hand, a single corporation avoiding billions in taxes costs taxpayers vastly more money than many thousands of waiters lying about their tip income.

There are many [reasons people cheat](#) on their taxes, and they all point to failures of societal pressure. First, there is [very little enforcement](#). In 2007, for example, the [IRS examined less](#) than 1% of the 179 million tax returns filed, initiated criminal prosecutions in only 4,211 cases, and obtained indictments in only 2,322 cases. [Corporate audits are down](#), too, both in number and thoroughness. And while there's debate about whether [increasing the penalties](#) against tax evaders increases compliance, we do know that increasing the number of audits [increases compliance](#) and—of course—collects more of the taxes owed.

Aside from low-level cheating that can be easily detected by computer matching, cheating on your taxes is easy and you're not likely to get caught.

Second, it's profitable. These days, if you're making a 5% return on your investments, you're doing really well. With the top federal tax rate at 35%, the money you can save by cheating is a pretty strong motivation. These are not people who can't afford to pay taxes; the typical tax cheat is a male under 50 in a high tax bracket and with a complex return. (Poorer users, with all their income covered by payroll taxes, have less opportunity to cheat.) The current situation creates an incentive to cheat.

Third, people think that lots of other people do it. Remember the Bad Apple Effect? There's a 1998 survey showing people believe that 38% of their fellow taxpayers are **failing to declare** all their income and listing false deductions. And the high-profile tax cheats that make the news reinforce this belief.

And fourth, recent political rhetoric has **demonized taxes**. Cries that taxation equals theft, that the tax system is unfair, and that the government just wastes any money you give it gives people a different morality, which they use to justify underpayment. This weakens the original moral pressure to pay up.

All of **these reasons interact** with each other. One study looked at tax evasion over about 50 years, and found that it increases with income tax rates, the unemployment rate, and public dissatisfaction with government. Another **blamed income inequality**.

Despite all of this, the U.S. government collects 81% of all taxes owed. That's actually pretty impressive compared to some countries.

There's another aspect to this. In addition to illegal tax evasion, there's what's called tax avoidance: technically legal measures to reduce taxes that run contrary to the tax code's policy goals. We discussed tax loopholes at length in Chapter 9. There are a lot of creative companies figuring out ways to follow the letter of the tax law while completely ignoring the spirit. This is how companies can make billions in profits yet pay little in taxes. And make no mistake, industries, professions, and groups of wealthy people deliberately manipulate the legislative system by lobbying Congress to get special tax exemptions to benefit themselves. One example is the **carried-interest tax** loophole: the taxation of private equity fund and hedge fund manager compensation at the 15% long-term capital gains tax rate rather than as regular income. Another is the **investment tax credit**, intended to help building contractors, that people used to subsidize expensive SUVs. There's also tax flight—companies moving profits out of the country to reduce taxes.

Estimates of **lost federal revenue** due to legal tax avoidance and tax flight are about \$1 trillion. Adding tax evasion, the total amount of lost revenue is \$1.5 trillion, or 41% of total taxes that should be collected. Collecting these taxes would more than eliminate the federal deficit.

Okay, so maybe that's not so good.

There are a lot of societal pressure failures in all of this. Morals differ: people tend to **perceive tax evasion** negatively, tax flight—companies moving profits out of the country to reduce taxes—neutrally, and tax avoidance positively: it's legal and clever. Even so, a reasonable case can be made that tax avoidance is **just as immoral** as tax evasion. The reputational effects of being a public tax cheat are few, and can be positive towards people who are clever enough to find legal loopholes. Institutional pressure depends on enforcement, which is spotty. Security systems are ineffective against the more complex fraud.



Remember the goal of societal pressures. We want a high level of trust in society. Society is too complex for the intimate form of trust—we have to interact with too many people to know all of their intentions—so we're settling for cooperation and compliance. In order for people to cooperate, they need to believe that almost everyone else will cooperate too. We solve this chicken-and-egg problem with societal pressures. By inducing people to comply with social norms, we naturally raise the level of trust and induce more people to cooperate. This is the positive feedback loop we're trying to get.

Societal pressures operate on society as a whole. They don't enforce cooperation in all people in all circumstances. Instead, they induce an overall level of cooperation. Returning to the immune system analogy, no defense works in all circumstances. As long as the system of societal pressures protects society as a whole, individual harm isn't a concern. It's not a failure of societal pressure if someone trusts too much and gets harmed because of it, or trusts too little and functions poorly in society as a result. What does matter is that the overall scope of defection is low enough that the overall level of trust is high enough for society to survive and hopefully thrive.

This sounds callous, but it's true. In the U.S., we tolerate 16,000–18,000 murders a year, and a tax gap of \$1.5 trillion. By any of the mechanisms discussed in Chapter 14, society gets to decide what level of defection we're willing to tolerate, and those numbers have fluctuated over the years. These are

only failures of societal pressure if society thinks these numbers are either too high or too low.

In Chapter 6, I talked about societal pressures as a series of knobs. Depending on the particular societal dilemma, society determines the scope of defection it can tolerate and then—if it's all working properly—dials the societal pressure knobs to achieve that balance. Recall the Hawk-Dove game from Chapter 3; a variety of different initial parameters result in stable societies. If we want less murder, we increase societal pressures. If that ends up being too expensive and we can tolerate a higher murder rate, we decrease societal pressures.

That metaphor is basically correct, but it's simplistic. We don't have that level of accuracy when we implement societal pressures. In the real world, the knobs are poorly marked and badly calibrated, there's a delay after you turn one of them before you notice any effects, and there's so much else going on that it's hard to figure out what the effect actually is. Think of a bathtub with leaky unmarked faucets, where you can't directly see the water coming out of the spout...outside, in the rain. You sit in the tub, oscillating back and forth between the water being too hot and too cold, and eventually you give up and take an uncomfortable bath. That's a more accurate metaphor for the degree of control we have with societal pressures.

Figure 13 tries to capture all of this.² On the left is the main feedback loop, between new societal pressures and the scope of defection. New societal pressures cause a change in the scope of defections, which causes a change in both risk and perceived risk. Then, the new perceived risk causes calls for changes in societal pressures.

Notice the delay between implementing new societal pressures and seeing corresponding changes in the scope of defection. The delay comes from several sources. One, moral and reputational pressures are inherently slow. Anything that affects risk trade-offs through a deterrence effect will require time before you see any effects from it. Depending on the form of government, new institutional pressures can also be slow. So can security systems: time to procure, time to implement, time before they're used effectively.

For example, the [first people arrested](#) for writing computer viruses in the pre-Internet era went unpunished because there weren't any applicable laws to charge them with. Internet e-mail was not designed to provide sender authentication; the result was the emergence of spam, a problem we're still trying to solve today. And in the U.S., the FBI regularly complains that the [laws regulating surveillance](#) aren't keeping up with the rapidly changing pace of communications technology.

Two, it can take time for a societal pressure change to propagate through society. All of this makes it harder to fine-tune the system, because you don't know when you're seeing the full effects of the societal pressures currently in place. And three, it takes time to measure any changes in the scope of defection. Sometimes you need months or even years of statistical data before you know if things are getting better or worse.

The feedback is also inexact. To use a communications theory term, it's noisy. Often you can't know the exact effects of your societal pressures because there are so many other things affecting the scope of defection at the same time; in Figure 13, those are the "other considerations." For instance, in the late 20th century, the drop in the U.S. crime rate has been linked to the [legalization of abortion](#) 20 years previously. Additionally, society's perceptions of risks are hard to quantify, and contain a cultural component. I'll talk more about this later in the chapter.

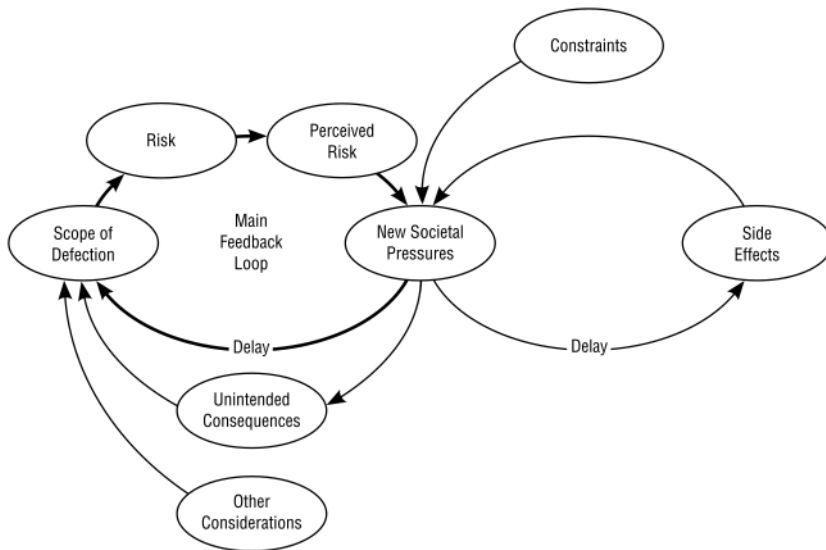


Figure 13: Societal Pressure's Feedback Loops

A related feedback loop, shown as the lower loop on the left in Figure 13, is also important. These are the [unintended consequences](#) of societal pressures that often directly affect the scope of defection. A large-scale example would be the effects on crime of Prohibition, or of incarcerating 16–25% of young black men in the U.S. A smaller-scale example is that hiring guards to prevent shoplifting may end up [increasing shoplifting](#), because regular employees now believe

that it's someone else's job to police the store and not theirs. Electronic sensor tags have a similar effect.

Security systems are complex, and will invariably have side effects on society. This is shown as the loop on the right side of Figure 13. For example, the U.S. incarceration rate has much broader social effects than simply locking up criminals. Prohibition did, too. A simple side effect is that some societal pressures, mostly security systems, cost money. More subtle side effects are [fewer bicycle riders](#) as a result of helmet laws, a [chilling effect on](#) computer-security research due to laws designed to prevent the digital copying of music and movies, and [increased violence](#) as a result of drug enforcement.

Decisions about whether to implement a new societal pressure require careful consideration of the [trade-off between](#) its costs and benefits—which are extremely difficult to predict.

Security systems are often what economists call an experiential good: something you don't understand the value of until you've already bought, installed, and experienced it.³ This holds true for other forms of societal pressure as well. If you're knowledgeable and experienced and perform a good analysis, you can make some good guesses, but it can be impossible to know the actual effects—or unintended consequences—of a particular societal pressure until you've already implemented it. This means that implementing societal pressures is always an iterative process. We try something, see how well it works, then fine-tune.

Any society—a family, a business, a government—is constantly balancing its need for security with the side effects, unintended consequences, and other considerations. Can we afford this particular societal pressure system? Are our fundamental freedoms and liberties more important than more security?⁴ More onerous ATM security will result in fewer ATM transactions, costing a bank more than the ATM fraud. A retail store that installs security cameras in its dressing rooms will probably have fewer customers as a result, with a greater loss of revenue than was saved by the decrease of shoplifting. Online retailers face similar choices, since complicated security measures reduce purchases. In Chapter 9, we talked specifically about how hard it is to get the security effects of laws right. It's hard for all categories of societal pressure.⁵

What all of this means is that it's easy to get societal pressures wrong. We implement more or less societal pressure than the risk warrants. We implement suboptimal, ineffective, or the wrong kind of security systems. Then, when we try to fix them, we get it wrong again. Many of the excesses in the War on Terror

can be chalked up to overcompensation for the security failures that led to the terrorist attacks of 9/11.



In Chapters 7 through 10 we talked about how specific types of societal pressure fail. Here, I am going to talk more generally about societal pressure failures. These failures can be broken into several broad categories. These categories aren't hard and fast, and there's going to be some overlap. The goal here is just to give a feeling for how societal pressures can go wrong.

Misunderstanding the actor. Potential defectors have many competing interests, ranging from selfish to moral; if you misunderstand them, you're likely to get security wrong. Defectors also have [different characteristics](#), such as motivation, skill, money, risk aversion, and so on.

It makes no sense to spend \$2 to forge an ID card worth \$1, right? That's true if the defector is in it for the money. But if he's a security researcher analyzing weaknesses in the production process, a competing company trying to damage the business, or a hacker just trying to understand how the stuff works, it might be. Similarly, if you think terrorists are all foreigners, you'll miss the homegrown ones.

We've also touched on the problem of organized defectors. Organization is common in crime—well-funded criminal organizations are far more effective than lone criminals—and in terrorism.⁶ It's also common among reform-minded defectors: abolitionists, animal rights activists, and so on. When defectors organize, societal pressures that worked in the past might not work as well. We talked about both of these problems in Chapter 11. A common misunderstanding is to assume that defectors are unorganized when they are—this happens often with crime—or to assume that defectors are organized when they are not, as happened with al Qaeda.

Misunderstanding the security incentives. Sometimes societal pressure can fail because it creates an incentive for the wrong competing norm. An example will help make this clear.

Convincing people to reduce their trash is a societal dilemma. Moral pressure only goes so far, and reputational pressure against having a lot of trash is generally pretty weak. By far the easiest institutional pressure is to charge people by the amount of trash they generate: by the bag, by the bin, by the pound. The idea is to tax marginal defection and encourage people to reduce their trash.

Societal Dilemma: Limiting personal trash.	
Society: Society as a whole.	
Group interest: Limit the use of landfills.	Competing interest: Laziness or apathy. Corresponding defection: Throw away as much trash as you want.
Group norm: Limit trash.	Competing interest: Minimize cost. Corresponding defection: Overstuff the trash can.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p style="padding-left: 40px;">Moral: Awareness campaigns that emphasize the immorality of polluting.</p> <p style="padding-left: 40px;">Reputational: Social pressure against people who put out a lot of trash.</p> <p style="padding-left: 40px;">Institutional: Charge residents extra, based on how much trash they produce.</p> <p style="padding-left: 40px;">Security: Garbage monitoring.⁷</p>	

However, a resident who wants to avoid the extra charges has several other options. He can stuff his trash more tightly into his bin. He can burn his trash to reduce the volume. He can dump his trash on the side of the road, or in the bin of a neighbor down the block. These options were always available to him, but before the extra trash collection fee, there was no reason to bother. As soon as you add societal pressures, some people will look for ways to get around them without having to cooperate in the original dilemma.

This isn't just theoretical. A study of nine municipalities showed exactly this sort of behavior—increases in trash burning and dumping—when **unit pricing** was implemented. **Stuffing more trash** in the bins, known as the “Seattle stomp” after the municipality where it was first noticed, is very common.

The failure here is the assumption that there is only one competing norm. In this case, there are a variety of ways to defect. And if the societal pressures only raise the cost of one competing norm, it could make the others more attractive. In this example, the trash fee didn't increase the cost of generating more trash; it merely increased the cost of generating more trash *and* putting that trash in trash cans. Directly targeting trash creation would be a better institutional pressure, but I can't think of any way a municipality could possibly make that work. On a larger scale, a disposal tax could be assessed when someone purchases a product. This would motivate product manufacturers to reduce packaging, or otherwise make their products more disposal-friendly, depending on the particulars

of the tax. Of course, administering that would be difficult, and society would have to balance that cost with the benefit.⁸

Misunderstanding the risk. We don't make risk trade-offs based on actual risk; as shown in Figure 13, we make them **based on perceived** risk. If we believe the scope of defection is higher or lower than it really is, we're not going to implement optimal societal pressures. And there are lots of ways we get risk wrong.

Natural Biases in Risk Perception	
We exaggerate risks that are...	We downplay risks that are...
Spectacular	Pedestrian
Rare	Common
Personified	Anonymous
Beyond our control	More under our control
Externally imposed	Taken willingly
Talked about	Not discussed
Intentional or man-made	Natural
Immediate	Long-term or diffuse
Sudden	Evolving slowly over time
Affecting us personally	Affecting others
New and unfamiliar	Familiar
Uncertain	Well understood
Directed against children	Directed against adults
Morally offensive	Morally desirable
Entirely without redeeming features	Associated with some ancillary benefit

This is all well-studied by psychologists. Current U.S. counterterrorism policy demonstrates these biases. Political scientist [John Mueller wrote](#):

Until 2001, far fewer Americans were killed in any grouping of years by all forms of international terrorism than were killed by lightning, and almost none of those terrorist deaths occurred within the United States itself. Even with the September 11 attacks included in the count, the number of Americans killed by international terrorism since the late 1960s (which is when the State Department began counting) is about the same as the number of Americans killed over the same period by lightning, accident-causing deer, or severe allergic reaction to peanuts.

But that's not the way people think. Terrorism is rare, spectacular, beyond our control, externally imposed, sudden, new and unfamiliar, uncertain, potentially directed against our children, offensive, and entirely without redeeming features. For these and other reasons, we [exaggerate the risk](#) and end up spending much too much on security to mitigate it.

Another example is computer crime. It's pedestrian, common, slowly evolving, affecting others, increasingly familiar, and (at least by techies) well-understood. So it makes sense that we understate the risks and underfund security.

There are cultural biases to risk as well. According to one study conducted in 23 countries, people have a higher risk tolerance in cultures that avoid uncertainty or are individualistic, and a lower risk tolerance in cultures that are egalitarian and harmonious. Also—and this is particularly interesting—the wealthier a country is, the lower its citizens' [tolerance for risk](#). Along similar lines, the greater the [income inequality](#) a society has, the less trusting its citizens are.

Creating a dilemma that encourages deception. Think back to the two prisoners for a minute. Throughout this entire book, we've assumed that Alice and Bob are both actually guilty. What if they're not? Now, what is Alice's best strategy?

Disturbingly, it may still be in her best interest to confess and testify against Bob. Follow me here: if Bob lies and testifies against Alice, she is looking at either six or ten years in jail. Lying and testifying against Bob is the better choice for Alice: six years is better than ten. And if Bob remains silent, she's looking at either freedom or one year in jail. Again, lying is the better choice for Alice: freedom is better than one year in jail. By this analysis, both Alice and Bob fare best if they confess to crimes they did not commit in an attempt to get leniency for themselves while falsely accusing the other. To make matters worse, assume that Bob is innocent and Alice is guilty. It's still in Alice's interest to falsely testify against Bob.

Of course, the risk trade-off is more complicated than that. Alice and Bob have to assess the prosecutor's case, and weigh the trade-off between their false confession and the hope that justice will prevail in the end. But as soon as the police offer Alice and Bob this deal, they increase the likelihood that one or both of them will confess to a crime they didn't commit. This is the reason that plea bargaining is illegal in many countries: it sets up perverse incentives. This can only be exacerbated by the surprising tendency of people to

make **false confessions**.⁹ Generalizing, we find that all sorts of unsavory people try to align themselves with the police in exchange for leniency for their own actions. This kind of thing can happen whenever people cooperate with a norm they don't believe in.

Accidentally making the costs of cooperation too high. Recall Chapter 11, where we talked about people assisting the police. One of Alice's potential competing interests is that cooperating with the police is too difficult, time-consuming, or dangerous. So even if Alice wants to cooperate, the cost is too high and she's forced to defect. This is the reason laws requiring the police to enforce immigration laws are a bad idea. The last thing you want is for someone to be afraid to assist the police out of fear that he will be deported. Another example is rape; if the cost of reporting a rape and helping prosecute the rapist is too emotionally high, women will not come forward. In general, there is a cost associated with cooperating. If we want to limit defections, we need to limit the costs—and/or increase the benefits—of cooperation.

Accidentally increasing the incentive to defect. The point of societal pressure is to induce cooperation. Sometimes the results are backwards, and societal pressure induces defection. Again, an example will explain this. Currently in the United States, standardized student testing has incredible influence over the future fates of students, teachers, and schools. Under a law called the No Child Left Behind Act, students have to pass certain tests; if they don't pass, their schools are penalized. In the District of Columbia, the school system offered teachers \$8,000 bonuses for improving test scores, and threatened them with termination for failing. Scores did increase significantly during the period, and the schools were held up as examples of how incentives affect teachers' behavior.

It turns out that a lot of those score increases were faked. In addition to teaching students, **teachers cheated on** their students' tests by changing wrong answers to correct ones.

There's a societal dilemma at work here. Teachers were always able to manipulate their students' test scores, but before the No Child Left Behind law, the competing interests were weak. People become teachers to teach, not to cheat... until their jobs depended on it. When the competing interests became stronger, the school districts should have increased societal pressures, probably security systems, to restore balance.¹⁰

Societal Dilemma: Cheating on students' tests.	
Society: Society as a whole.	
Group interest: Accurate testing of students.	Old competing interest: Selfish interest of having a star classroom. Old corresponding defection: Fake students' tests so they have a higher score.
Group norm: Allow students to take their own tests.	New competing interest: Financial reward, job retention. New corresponding defection: Fake students' tests so they have a higher score.
<p>To encourage people to act in the group interest, society implements these societal pressures:</p> <p>Moral: Teacher integrity.</p> <p>Reputational: Loss of reputation if caught cheating.</p> <p>Institutional: Changing answers on students' tests is fraud, and there are laws against it.</p> <p>Security: Secure handling of tests makes it harder for teachers to change answers.</p> <p>Statistical analysis of test data can show evidence of cheating.</p>	

There's a rule at work here. When you start measuring something and then judge people based on that measurement, you encourage people to game the measurement instead of doing whatever it is you wanted in the first place. If a company penalizes customer-support people for having long phone calls, they have an incentive to hang up on callers. If you reward software programmers for fixing bugs, they have an incentive to create buggy software to have bugs to fix instead of getting it right the first time.¹¹ [If you pay CEOs](#) based on stock price, they have an incentive to inflate the stock price at the expense of the company's long-term interest.

The incentive to defect can also be increased when the reason a thing is attacked changes. Driver's licenses are a great example. Originally, they were nothing more than proof that a person is legally allowed to drive a car. As such, there wasn't much of an incentive to forge them, and security around the licenses was minimal: they were made of paper, they didn't have photos, and so on. In the U.S., at least, it was only when they started being used for a completely different purpose—age verification as a condition of buying alcohol—that forgeries started being a problem. In response, state governments changed their licenses to include a variety of anti-forgery features: photographs, watermarks and holograms, microprinting, and the like. Recently, their use has changed again. Since

9/11, they have been increasingly used as proof that a person isn't on a terrorist watch list. And now the government wants even more security features associated with them, like computer chips and enhanced security around their issuance.

We saw this with pair-bonding. Informal pair-bonding was enough to deal with Deacon's Paradox with respect to infidelity, but when inheritance became an issue, more formal mechanisms were required. Another example is joyriding; because joyriders never intended to keep the cars they stole, they couldn't be charged with theft—so before specific joyriding laws were enacted, they got off relatively lightly.

The market can also increase the incentive to defect. When the price of glass eels—immature eels that are a delicacy in Japan and Europe—started rising, more people began to fish for them. The result was a Tragedy of the Commons: illegal overfishing and poaching in England, France, and the northeastern U.S. resulted in reduced yields, which resulted in higher prices. This resulted in [even more overfishing](#), even further reduced yields, and even higher prices that rose from \$25 to \$950 per pound. Enforcement just couldn't keep up, and poachers have devastated the eel population. A technological advance might solve this societal dilemma; researchers are trying to breed and [farm these eels](#), which will increase supply and reduce the incentive to overfish.

Technological advances can magnify societal dilemmas as well. We'll talk about this in the next chapter, but for now, think of the difference between banking in person and banking online, manual door locks and electronic locks, or paper ballots and touch-screen voting machines. In all cases, the addition of technology makes some attacks easier.

A final way the incentive to defect can increase is when the scale of the societal dilemma changes. We saw this in the difference between a single sandwich seller in a market and a large sandwich-producing corporation, and between Fisherman Bob and the Robert Fish Corporation. Large organizations can gain more, and inflict more damage on the group, by defecting. As organizations grow in size and power, societal pressures that might have worked in the past won't necessarily work as well any longer.

Misunderstanding how different societal dilemmas interact. Societal dilemmas don't exist in isolation, and societal pressures designed to decrease the scope of defection in one societal dilemma can, as a side effect, increase the scope of defection in another.

For example, we recognize that the police force is both a solution and a problem. It is our agent in institutional pressures against criminals in general, but as an institution with its own self-interests, it has to be dissuaded from defecting. So we have all sorts of societal pressures protecting society from the police: rules

limiting search and seizure, rules against self-incrimination, rules about interrogation, rules about evidence, and so on. These necessarily affect the defection rate of criminals by making the police's job harder and more onerous, but we have them because—on balance—the result is a better police force and a better society. Recently, this has been changing. In our efforts to protect ourselves against terrorism, we have been dismantling many of the societal pressures we've put in place to protect ourselves from abuse by the police.

Similarly, over the past couple of decades we have dismantled a variety of financial regulations that limited the behavior of banks and other financial institutions.¹² Yes, those regulations made it harder for institutions to make money, but they also served to protect society from the effects of widespread bank defection.

Ignoring changing social norms. Sometimes societal norms change, and societal dilemmas start shifting to reflect the change. This often results in conflicting societal dilemmas as the new norms work their way through society, and in conflicts between subgroups within society who are either clinging to the old norms or embracing the new ones.

My favorite example is historical. **In ancient Rome**, it was important to worship the gods. It was also important that everyone in the community worship the gods. The gods were angered if some people shirked their religious responsibilities, like participating in festivals. This is one reason the Romans didn't like the early Christians. It's not that they worshipped their Christian god, it's that they didn't *also* worship the Roman gods. This was not simply a disagreement with Christians' personal choice; it was seen as a danger to the whole community.

Societal Dilemma: Worshipping Roman gods.	
Society: Society as a whole.	
Group interest: Making the Roman gods happy.	Competing interest: Making your own god happy.
Group norm: Worshipping the Roman gods.	Corresponding defection: Not worshipping the Roman gods as well.
<p>To encourage people to act in the group interest, the society implements a variety of societal pressures.</p> <p>Moral: From birth, Romans were taught their religion.</p> <p>Reputational: Romans who didn't participate in public religious ceremonies were penalized by the community.</p> <p>Institutional: Serious offenders were thrown to the lions.</p> <p>Security: Lions.</p>	

Eventually, social norms changed. Christians became a larger and larger minority. They were increasingly tolerated. Sometime in the early 300s AD, Emperor Constantine converted to Christianity. And slowly, what had been defection became cooperation.

Whether and when societal pressure failed depends on your point of view. If you believed in the Roman gods, then societal pressure failed when it didn't prevent Christians from offending the Roman gods. If you were an early Christian, then societal pressure failed when it didn't protect freedom of religion.

Another example is sexual harassment in the workplace. As long as those in power in the organization didn't enforce prohibitions against men harassing subordinate women, unwanted advances were relatively common and taken for granted. It wasn't until a larger society started enforcing sexual harassment rules that occurrences began to decline.

A similar dynamic is playing out with respect to gay marriage. It's a fundamentalist Christian belief that gay marriage isn't just a bad individual choice, but that its very existence threatens the traditional family: just like the Romans talking about Christianity. As such, it's a societal dilemma.

Societal Dilemma: Gay marriage.	
Society: Society as a whole.	
Group interest: Protecting the institution of marriage. Group norm: Only recognizing "approved" marriages.	Competing interest: Allowing everyone free choice in whom they can marry. Corresponding defection: Allowing gay couples to marry.
To encourage people to act in the group interest, the society implements a variety of societal pressures. Moral: Teach gay marriage is wrong. Reputational: Ostracize same-sex couples. Institutional: Refuse to give same-sex couples the same legal rights as different-sex couples. Pass laws making life especially difficult for same-sex couples. Security: None.	

Other people, though, don't see the dilemma. They don't accept that group defection would result in the social calamity the fundamentalists do. Not only do they defect, they don't even accept the dilemma as real.¹³

Norms can change quickly due to external threats. People are more willing to implement societal pressures—both the kinds that reward cooperators and the kinds that punish defectors—in times of war.

Most of the time, though, social norms change slowly. We’ve repeatedly talked about Deacon’s Paradox, and how pair-bonding is a societal pressure. Enforcement of that has changed. There was a time when you could be stoned to death for adultery, or for fornication out of wedlock. Now, in most of the world, that doesn’t happen. There are even parts of the world where it isn’t even frowned upon very heavily. And on the technological side, defecting from pair-bonds has become safer. The “wages of sin” used to include pregnancy, which came with it significant health and financial risks, and venereal disease. Cheap and effective birth control changed that, so much so that the current societal dilemma for women is a very different risk trade-off. More recently, unsafe sex practices brought with them a different set of health risks, ones that could be effectively mitigated with technological security measures like condoms.

Our evolving definitions of “society” show how societal norms evolve. As Barbara Jordan famously noted, the original definition of “we the people” in the U.S. didn’t include **women or slaves**. Over the centuries, our definitions of who is within the bounds of society have gradually become more inclusive.

You can see this evolution in the societal dilemma surrounding the current tone and integrity of **political debates** in the United States. The goal of politics—elections, policy debates, laws—is to govern the country by enacting the best policies for society and implementing the best laws to solve societal dilemmas. But there’s a competing interest of getting laws passed that benefit us in particular. We’re all better off if national policy debates are factual, honest, and civil, but it’s easy to resort to spin, distortions, smears, and lies. But if enough people do that, you get the circus that characterizes far too much of current American politics.¹⁴

Societal Dilemma: Policy debates.	
Society: Society as a whole.	
Group interest: Make the best policy decisions. Group norm: Debate public policy fairly, whatever that might mean.	Competing interest: For your side to win. Corresponding defection: Debate by whatever means necessary.
To encourage people to act in the group interest, the society implements a variety of societal pressures. <ul style="list-style-type: none"> Moral: Shame, honesty, honorability, and so on. Reputational: Shame and ridicule heaped on dishonest politicians. Reputation for statesmanship bestowed on honest ones. Institutional: For particularly egregious lies, libel laws. Anti-gerrymandering laws. Security: The proper use of rhetoric. Fact checking. 	

It's not clear that the level of dishonesty is new, but it seems to be carried out on a much broader scale today. Moral and reputational pressures used to work, but they are failing as the country bifurcates into two different groups with completely separate systems of values. Legal controls that impinge on free speech are a dangerous option. One solution is to stop gerrymandering safe legislative seats. By forcing these seats to be decided in the general election, as opposed to party-specific primaries or caucuses, candidates would have to appeal to swing centrist voters rather than their base. But potential legal societal pressures would be viewed as partisan, and untenable for that reason.

What's going on here is that the definition of "society" is changing. "Society as a whole" has less meaning in a polarized political climate such as the one in the U.S. in the early 21st century. People are defining their society as those who agree with them politically, and the other political side as "traitors," people who "hate America," or people who "want the terrorists to win." It's no surprise that there's widespread defection: with regard to the new, more restrictive, definition of "society," it's not defection at all.¹⁵

16

Technological Advances

Scale is one of the critical concepts necessary to understand societal pressures. The increasing scale of society is what forces us to shift from trust and trustworthiness based on personal relationships to impersonal trust—predictability and compliance—in both people and systems. Increasing scale is what forces us to augment our social pressures of morals and reputation with institutional pressure and security systems. Increasing scale is what's requiring more—and more complicated—security systems, making overall societal pressures more expensive and less effective. Increasing scale makes the failures more expensive and more onerous. And it makes our whole societal pressure system less flexible and adaptable.

This is all because increasing scale affects societal pressures from a number of different directions.

- *More people.* Having more people in society changes the effectiveness of different reputational pressures. It also increases the number of defectors, even if the percentage remains unchanged, giving them more opportunities to organize and grow stronger. Finally, more defectors makes it more likely that the defecting behavior is perceived as normal, which can result in a Bad Apple Effect.
- *Increased complexity.* More people means more interactions among people: more interactions, more often, over longer distances, about more things. This both causes new societal dilemmas to arise and causes interdependencies among dilemmas. Complex systems need to rely on technology more. This means that they **have more flaws** and can **fail in surprising** and catastrophic ways.
- *New systems.* As more and different technology permeates our lives and our societies, we find new areas of concern that need to be addressed, new

societal dilemmas, and new opportunities for defection. Airplane terrorism simply wasn't a problem before airplanes were invented; Internet fraud requires the Internet. The job of the defenders keeps getting bigger.

- *New security systems.* Technology gives certain societal pressure systems—specifically, reputational and institutional—the ability to scale. Those systems themselves require security, and that security can be attacked directly. So online reputation systems can be poisoned with fake data, or the computers that maintain them can be hacked and the data modified. Our webmail accounts can be hacked, and scammers can post messages asking for money in our name. Or our identities can be stolen from information taken from our home computers or centralized databases.
- *Increased technological intensity.* As society gets more technological, the amount of damage defectors can do grows. This means that even a very small defection rate can be as bad as a greater defection rate would have been when society was less technologically intense. This holds true for the sociopath intent on killing as many people as possible, and for a company intent on making as much profit as possible, regardless of the environmental damage. In both cases, technology increases the actor's potential harm. Think of how much damage a terrorist can do today versus what he could have done fifty years ago, and then try to extrapolate to what upcoming technologies might enable him to do fifty years from now.¹ Technology also allows defectors to better organize, potentially making their groups larger, more powerful, and more potent.
- *Increased frequency.* Frequency scales with technology as well. Think of the difference between someone robbing a bank with a gun and a getaway car versus someone stealing from a bank remotely over the Internet. The latter is much more efficient. If the hacker can automate his attack, he can steal from thousands of banks a day—even while he sleeps. [This aspect of scale](#) is becoming much more important as more aspects of our society are controlled not by people but by automatic systems.
- *Increased distance.* Defectors can act over both longer physical distances and greater time intervals. This matters because greater distances create the potential for more people, with weaker social ties, to be involved; this weakens moral and reputational pressure. And when physical distances cross national boundaries, institutional pressure becomes less effective as well.
- *Increased inertia and resistance to change.* Larger groups make slower decisions; and once made, those decisions persist and may be

very difficult to reverse or revise. This can cause societal pressures to stagnate.

In prehistoric times, the scale was smaller, and our emergent social pressures—moral and reputational—worked well because they evolved for the small-scale societies of the day. As civilization emerged and technology advanced, we invented institutions to help deal with societal dilemmas on the larger scales of our growing societies. We also invented security technologies to further enhance societal pressures. We needed to trust both these institutions and the security systems that increasingly affected our lives.

We also developed less tolerance for risk. For much of our species' history, life was dangerous. I'm not just talking about losing 15–25% of males to warfare in primitive societies, but infant mortality, childhood diseases, adult diseases, natural and man-made accidents, and violence from both man and beast. As technology, especially medical technology, improved, life became safer and longer. Our tolerance for risk diminished because there were fewer hazards in our lives. (Large, long-term risks like nuclear weapons, genetic engineering, and global warming are much harder for us to comprehend, and we [tend to minimize](#) them as a result.)

Today, societal scale continues to grow as global trade increases, the world's economies link up, global interdependencies multiply, and international legal bodies gain more power. On a more personal level, the Internet continues to bring distant people closer. Our risk tolerance has become so low that we have a fetish for eliminating—or at least [pretending to eliminate](#)—as much risk as possible from our lives.

Let's get back to societal pressures as a series of knobs. Technology is continuously improving, making new things possible and existing things easier, cheaper, better, or more reliable. But these same technological advances result in the [knobs being twiddled](#) in unpredictable ways. Also, as scale increases, new knobs get created, more people have their hands on the knobs, and knobs regulating different dilemmas get interlinked.

New technologies, new innovations, and new ideas increase the scope of defection in several dimensions. Defectors innovate. Attacks become easier, cheaper, more reliable. New attacks become possible. More people may defect because it's easier to do so, or their defections become more frequent or more intense.

This results in a security imbalance; the knob settings that society had deemed acceptable no longer are. In response, society innovates. It implements

new societal pressures. Perhaps they're based on new laws or new technology, perhaps there is some new group norm that gets reflected in society's reputational pressure, or perhaps it's more of what used to work. It's hard to get right at first, because of all the feedback loops we discussed, but eventually society settles on some new knob settings, and the scope of defection is reduced to whatever new level society deems tolerable. And then society is stable until the next technological innovation.




Figure 14: Societal Pressure Red Queen Effect

If Figure 14 looks familiar, it's because it's almost the same as Figure 3 from Chapter 2. This is a Red Queen Effect, fueled not just by natural selection but also by technological innovation. Think of airport security, counterfeiting, or software systems. The attackers improve, so the defenders improve, so the attackers improve, and so on. Both sides must continuously improve just to keep pace.

But it's not a normal Red Queen Effect; this one isn't fair. Defectors have a natural advantage, because they can make use of innovations to attack systems faster than society can use those innovations to defend itself. One of society's disadvantages is the delay between new societal pressures, and a corresponding change in the

scope of defection, which we talked about in the previous chapter. In fact, the right half of Figure 14 is the same as the main feedback loop of Figure 13, but with less detail.

More generally, defectors are quicker to use technological innovations. Society has to implement any new security technology as a group, which implies agreement and coordination and—in some instances—a lengthy bureaucratic procurement process. Unfamiliarity is also an issue. Meanwhile, a defector can just use the new technology. For example, it's easier for a bank robber to use his new motorcar as a getaway vehicle than it is for the police department to decide it needs one, get the budget to buy one, choose which one to buy, buy it, and then develop training and policies for it. And if only one police department does this, the bank robber can just move to another town. Corporations can make use of new technologies of influence and persuasion faster than society can develop resistance to them. It's easier for **hackers to find** security flaws in phone switches than it is for the phone companies to upgrade them. **Criminals can form** international partnerships faster than governments can. Defectors are more agile and more adaptable, making them much better at being early adopters of new technology.

We saw it in law enforcement's initial inability to deal with Internet crime. **Criminals were simply** more flexible. Traditional criminal organizations like the Mafia didn't move immediately onto the Internet; instead, new Internet-savvy criminals sprung up. They established websites like CardersMarket and DarkMarket, and established new crime organizations within a decade or so of the Internet's commercialization. Meanwhile, law enforcement simply didn't have the organizational fluidity to adapt as quickly. They couldn't fire their old-school detectives and replace them with people who understood the Internet. Their natural inertia and their tendency to sweep problems under the rug slowed things even more. They had to spend the better part of a decade playing catch-up.

There's one more problem. Defenders are in what the 19th-century military strategist Carl von Clausewitz called "the **position of the interior.**" They have to defend against every possible attack, while the defector just has to find one flaw and one way through the defenses. As systems get more

complicated due to technology, more attacks become possible. This means defectors have a first-mover advantage; they get to try the new attack first. As a result, society is constantly responding: shoe scanners in response to the shoe bomber, harder-to-counterfeit money in response to better counterfeiting technologies, better anti-virus software to combat the new computer viruses, and so on. The attacker's clear advantage increases the scope of defection further.

Of course, there are exceptions. Sometimes societal pressures improve without it being a reaction to an increase in the scope of defection. There are technologies that immediately benefit the defender and are of no use at all to the attacker. Fingerprint technology allowed police to identify suspects after they left the scene of the crime, and didn't provide any corresponding benefit to criminals, for example. The same thing happened with immobilizing technology for cars, alarm systems for houses, and computer authentication technologies. Some technologies benefit both, but still give more advantage to the defenders. The radio allowed street policemen to communicate remotely, which makes us safer than criminals communicating remotely endangers us.

Still, we tend to be reactive in security, and only implement new measures in response to an increased scope of defection.

Because the attackers generally innovate faster than the defenders, society needs time to get societal pressures right. The result of this is a security gap: the difference between the scope of defection that society is willing to tolerate and the scope of defection that exists. Generally, this gap hasn't been an insurmountable problem. Sure, some defectors are able to get away with whatever it is they're doing—sometimes for years or even decades—but society generally figures it out in the end. Technology has progressed slowly enough for the Red Queen Effect to work properly. And the slowness has even helped in some situations by minimizing overreactions.

The problem gets worse as technology improves, though. Look at Figure 15. On the top, you can see the difference between the defectors' use of technological innovation to attack systems and the defenders' use of technological innovations in security systems and other types of societal pressures. The security gap arising from the fact that the attackers are faster than the defenders is represented by the area under the technology curve between the two lines.

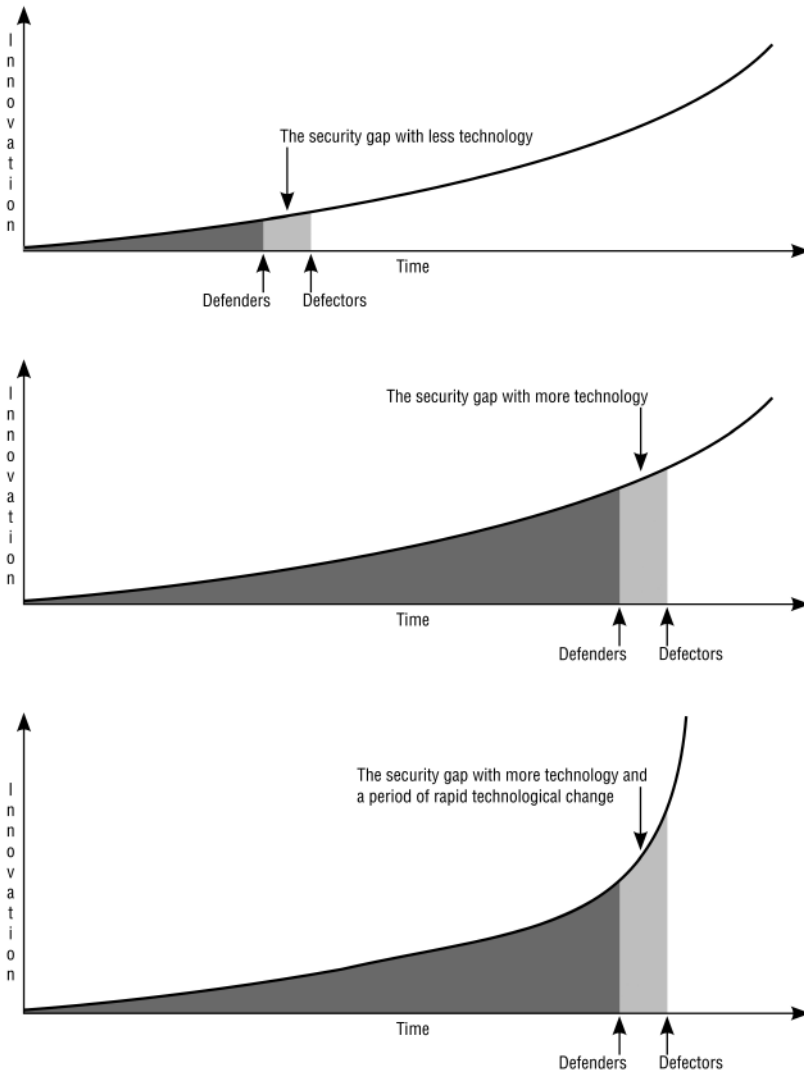


Figure 15: The Security Gap

Comparing the top picture with the middle one shows the difference between less and more technology. In the middle, the gap between attacker and defender is the same width, but because there's more technology, the area is greater. There are actually two dimensions to innovation: technological advancement and technological prevalence. In either dimension, the more technology there

is, the greater the security gap. In other words, if there are more innovations to exploit, there will be more damage resulting from society's inability to keep up with exploiters of all of them.

Think about it this way. Technology is available to both the attackers and the defenders, and it's pretty much all there is until moral, reputational, and institutional pressures catch up. When there's more technology out there, the attackers have more opportunity to increase the scope of defection before the defenders catch up. Technology can affect the scope of defection in many ways, but in general, it gives the attackers more leverage. So the more technological a society is, the greater the security gap is.

This is an intrinsic condition of the problem, for all the reasons we just talked about. The security gap cannot be eliminated.

The security gap is also greater in periods of rapid technological change, as society struggles to manage the broader social changes as well as quickly adapting defectors do. In 1970, futurist [Alvin Toffler](#) wrote about *future shock*, the psychological and social problems that result from people being forced to absorb too much technological change too quickly. His estimates about how much technological change people could deal with were way too low—the rate of technological change in the second decade of the 21st century is much faster than the seventh decade of the 20th—but his basic ideas are sound. [People learn how](#) to cope with new technologies at their own pace, some more easily than others. And groups of people move more slowly than some of their members. Defectors are not inherently less susceptible to future shock than society at large, but the more successful ones are. Successful defectors are always going to be able to outpace the average capability of society.

Again, look at Figure 15, the bottom this time. In a period of rapid change, technology increases faster, so the curve climbs higher in the same period of time than in the earlier figures. This faster growth rate makes for a larger area under the curve in the same period of time—a greater security gap.

This has happened before, notably in the 19th century. That's when we got railroads, steamships, the widespread use of paper mail, the telegraph, and then the telephone—all allowing people to communicate at greater distances and with greater speed. But perhaps even more important than any of that, there were significant changes in attitudes about people and the world. Society came to expect economic growth, along with universal education and universal betterment. The world changed, and that affected security.

The ease of rapid travel meant more people traveled. On one hand, this meant that you could no longer distrust people just because they came from “out of

town.” On the other, this allowed for a [new type of grifter](#), conning people out of their money and moving on before he could be caught. At the same time, cities got larger. Policing in 18th-century London was a hodge-podge of unpaid and unorganized constables and a draconian court system (160 different crimes carried the death penalty). This sort of community policing didn’t scale to a large modern city, so Sir Robert Peel organized the first modern police force and criminal justice system. [Other cities followed](#) suit.

Technology directly changed society as well. The telegraph meant that money could be transferred instantaneously, but the open nature of the system meant conversations could be eavesdropped on and spoofed. So operators developed codes to prevent that. Other examples were the mass production of timepieces, making it easier to manage employees; the rise of unions, giving employees more power with respect to their employers; and the telegraph and then the telephone, an enormous change in communication that affected everyone. It was an age where defectors adapted to a changing society, and society had to adapt to changing defectors.

Today, we’re seeing the effects of both more technology than ever before *and* a faster rate of technological change than ever before.² In particular, the revolutionary social and political changes brought about by information technology are causing security and trust problems to a whole new degree. We’ve already seen several manifestations of this: the global financial crisis, international terrorism, and cyberspace fraud. We’ve seen music and movie piracy grow from a minor annoyance to an international problem due to the ease of distributing pirated content on the Internet. We’ve seen Internet worms progress from minor annoyances to criminal tools to military-grade weapons that cause real-world damage, like the [Internet worm Stuxnet](#), the first military-grade cyberweapon the public has seen. All this has come about because information technology increases the scope of defection in several ways:

- *Migration of all data onto the Internet.* As data moves onto computer networks, there are more—and, more importantly, different—risks. The security that worked when the systems were manual, or housed on computers not attached to a global network, no longer works.³
- *Technological mediation of social systems.* Similarly, social systems—including systems of reputational pressure—are vulnerable to technological attacks as they become technologically enabled. For example, e-mail has security risks that paper mail does not. Electronic voting has security

risks that paper voting does not. Internet telephony has security risks that conventional telephony does not.

- *Migration of evolved social systems into deliberately created socio-technical systems.* In Chapter 14, we discussed the problem of delegating societal pressures to institutions, specifically government institutions. More and more, we are delegating societal pressures to corporations: the security of our conversations, our photographs, and our data. This trend of corporations acting as institutions gives those corporations more ability and incentive to defect.
- *Class breaks.* A product, or line of products, may have common vulnerabilities that impact every copy of the product that has ever been made. As globalization allows a single product to be used worldwide, the discovery of such a vulnerability can have a global impact. This is not new, but information systems are particularly prone to this type of problem. Information systems have common vulnerabilities that can be exploited *en masse*. Someone who finds, for example, a vulnerability in an operating system that allows him to steal data can steal data from the entire class of computers using that operating system.
- *Automation.* Information system attacks can be automated. Instead of manually having to break into computer systems, an attacker can write a program to do it automatically. This not only drastically increases the frequency of defection, it also has two other effects. One, it makes attacks whose probability of success is very small viable. And two, it makes attacks whose profitability is very small—so-called [salami attacks](#) because of how thinly sliced each instance of fraud is—viable.
- *Action at a distance.* Attacks that used to require the attacker to get up close and personal to his victims can now be done remotely, from anywhere on the planet. It used to be that a store in Los Angeles didn't have to worry about burglars living in London or Lagos; those places were simply too far away for it to be worth the burglar's time or expense to fly to Los Angeles. But on the Internet, every web store has to worry about every cyber burglar in the world. There are no natural defenses against distance. Similarly, 20 years ago, few Americans had to worry about encountering Ukrainian or Nigerian criminals. On the Internet, it happens constantly.
- *Technique propagation.* Because information system attacks can be automated and encapsulated in software, the capability to launch these attacks can propagate. No longer does a criminal have to learn how to attack a security system: pick a lock, defraud a bank, or whatever. On the Internet,

only the first attacker has to be skilled. Everyone else can [just use software](#).

- *Technique iteration and improvement.* Because attacks can be so efficient, it's easier for attackers to learn from their mistakes and improve their attacks. They can create ten varieties of a computer worm and test which one works best, or a hundred varieties of spam to see which one fools recipients best. Because so many Internet attack tools become public, it's easy for one attacker to learn from another's work.
- *Defector aggregation.* One thing that makes it easier to defect from society is finding a subgroup of defectors. This both makes it easier to overcome moral and reputational pressures, and allows defectors to trade tips on overcoming the legal pressure and security systems. The Internet itself lets defectors easily find and communicate with like-minded individuals. There's a whole online community of people who think childhood immunization is evil. There are terrorist-sympathetic websites, which might—it's hard to separate reality from media hype—also act as terrorist-recruiting websites. There are a gazillion places on the Internet where you can learn to hack computer systems and commit fraud.

There are two more changes that belong on this, too, but they won't fit neatly into bullet points: changes in organizational structure and changes in organizational behavior.

Let's start with organizational structure. The Internet reduces the cost of organization dramatically, enabling ad hoc and loosely connected organizations of individuals who contribute tiny amounts of effort towards a large goal.⁴ Linux and Wikipedia are both informally produced and freely available "products" created by legions of unpaid volunteers; and both are viable competition to corporate, traditionally created, alternatives. Crowdsourcing can produce results superior to more traditional mechanisms of delegating work.

From a societal pressure perspective, the normal competing interests we've come to expect from traditional organizations don't apply in the same way to these ad hoc organizations. For example, Microsoft can be—and in the past has been—pressured by the U.S. government to deliberately weaken encryption software in its products, so the government could better spy on people. This works because Microsoft is an American corporation, and in at least some ways beholden to American interests. Its operating system competitor, Linux, is not. Linux is an open-source operating system, not controlled by a business. The Linux team, even the few individuals at the core, are not motivated by

profit. They're not in any one country. They are probably unlikely to agree to a confidential meeting with government officials of any nationality. They are a different sort of actor. On the other hand, Microsoft probably has better systems in place to prevent infiltration by rogue programmers.

WikiLeaks is another stateless organization. WikiLeaks sits somewhere between a loose organization of activists and the personal mission of a single individual named Julian Assange. It exposes information that governments and powerful corporations would rather keep secret. In this way it is very much like an organization of journalists. But because it is not a commercial enterprise, and because it is not moored within a country, it's much more difficult to corral. And this scares countries like the United States.

Compare WikiLeaks to a traditional newspaper. That newspaper is in a societal dilemma with all the other newspapers in that country.

Societal Dilemma: Newspapers publishing government secrets.	
Society: All the newspapers in the country and the government.	
Group interest: Government not clamping down on freedom of the press.	Competing interest: Increase market share.
Group norm: Self-censor.	Corresponding defection: Publish any juicy secrets you discover.
To encourage people to act in the group interest, the society implements a variety of societal pressures.	
Moral: It's unpatriotic, or otherwise wrong, to publish government secrets.	
Reputational: Newspapers want good reputations because it keeps their readers, advertisers, and sources all happy.	
Institutional: Often, none. In fact, the U.S. Supreme Court has held that it is legal to publish secrets, even though it is illegal to leak them.	
Security: Potentially, espionage that lets the government know when a story is about to leak.	

This doesn't look like effective societal pressure, but it largely works. It works because, even in the absence of any laws, the pressure to cooperate—to self-censor—is surprisingly powerful. No press organization wants to be labeled as unpatriotic or traitorous, or jeopardize its advertisers.

The result is that newspapers sometimes publish embarrassing government secrets, and sometimes they don't. In 1971, the *New York Times* published the Pentagon Papers, a secret and damning history of U.S. military involvement in Vietnam. In mid-2004, the *New York Times* learned about the NSA's illegal

wiretapping of American citizens without a warrant, but [delayed publishing](#) the information for over a year—until well after the presidential election. Presumably there are things the *New York Times* has learned about and decided not to publish, period.

WikiLeaks changes that dynamic. It's not an American company. It's not even a for-profit company. It's not a company at all. And it's not really located in any legal jurisdiction. It simply isn't subject to the same pressures that the *New York Times* is. This means the government can't rely on the partial cooperation of WikiLeaks in the same way it can rely on that of traditional newspapers.⁵

In a blog post about the topic, [Clay Shirky referred](#) to the Supreme Court ruling in the Pentagon Papers case that said it's illegal to leak secrets but not illegal to publish leaks:

The legal bargain from 1971 simply does not and cannot produce the outcome it used to. This is one of the things freaking people in the US government out about the long-term change in the media environment—that the law has changed, but that the world has. Industrial era law, applied to internet-era publishing, might allow for media outlets which exhibit no self-restraint around national sensitivities, because they are run by people without any loyalty to—or, more importantly, need of—national affiliation to do their jobs.

Foreign journalists pose a similar problem. The U.S. government has much less leverage to pressure *El Pais* or Al Jazeera to change its coverage than it does with the *New York Times*. That mattered less before the Internet could bring all those news sources to everyone so easily.

This unmooring of institutions from nationality is upending many societal pressures; things that used to work no longer do. We saw the same dynamic in international corporations, which can more easily skirt national laws by moving between different countries.

Now to the final change, which is organization behavior. In addition to allowing organizations to grow in size, and therefore power, and facilitating new types of organizational structures, information technology is also changing how organizations act.

There have been many books and articles discussing how corporations today are putting short-term stock prices above all [other business considerations](#), including company health and long-term shareholder value. I've read lots of explanations for this change. That executives' bonuses are based on short-term numbers. That stocks are used more for short-term "bets" than for long-term

investments. That mutual funds and complex index options further remove investors from the companies they invest in. And that [investors have access](#) to more information faster—and can act on that information faster.

You get what you measure,⁶ and things like short-term profitability are much easier to measure than abstract concepts like long-term viability, or intangibles like customer satisfaction or reputation. An important facilitator for this dynamic—I don't know whether it's a cause or not—is information technology. Improved information technology makes the short-term numbers easier to monitor, so investors monitor them much more closely than ever before. This continuous monitoring makes them easier to optimize. We are better able to predict what a company's balance sheet will look like next week, and because we're so quick to trade one company for another, we care much less what it will look like in five years. This necessarily changes how investing works and how organizations behave: and the two are locked in a positive-feedback loop.

All these effects of ever-faster information technology affect other organizations at every scale, from the smallest groups to the entire world.



Modern large and technological trade-offs between group interest and competing interest are what social planners call [wicked problems](#). These are problems that are difficult (or impossible) to solve because of incomplete, poorly understood, contradictory, or changing requirements; because of complex interdependencies; and because of their uniqueness and novelty. Examples include global climate change, AIDS and pandemics in general, nuclear waste, terrorism and homeland security, drug trafficking and other international smuggling, and national healthcare. All of those problems involve societal pressures, and all of their solutions involve coercing people into following group norms ahead of other competing interests.

But—and this is important—all of the big societal pressure problems are about more than trust and security. They're interdependent with other societal dilemmas. They're interdependent with other societal systems. They have moral, social, economic, and political dimensions. Their solutions involve answering questions about how society organizes itself, the role of national and international government, the extent of individual liberties, and what sort of outcomes are optimal and desirable. And these aspects of the problems are far more important, and difficult, than the trust aspects. It's not simply a matter of implementing the best societal pressures to induce broad cooperation; everything else

matters more. The geopolitics that results in terrorism matter much more than any particular security measure against terrorists. The politics in which multinational corporations thrive matter much more than the societal pressures to ensure those corporations cooperate. The politics surrounding drug laws, tax laws, laws protecting civil liberties, and our social safety net matter much more than the societal pressures to ensure that those laws are followed. Look back to the figure in Chapter 15; the “constraints” and the “other considerations” are more important than the primary loop.

Here’s one example. In 2011, science fiction author Charles Stross gave a talk on the [ubiquity of data](#) that’s coming in the near future, from technologies like genetic mapping, “lifeblogging”—the audio and video recording of everything that happens to you—sensors on everyone and everything. Nothing he said required anything more than mild extrapolation. And then he talked about the issues that society is going to have to wrestle with once this data exists:

Is losing your genomic privacy an excessive price to pay for surviving cancer and evading plagues? (Broad analysis of everyone’s genetic data will result in significant new understanding about disease, and a flurry of medical results that will significantly benefit everyone. At the same time, an individual’s genetic data is both personal and private—even more so when companies start using it to prejudge people.)

Is compromising your sensory privacy through lifeblogging a reasonable price to pay for preventing malicious impersonation and apprehending criminals? (Lifeblogs have the potential to be a valuable police tool, not just by allowing victims to record crimes, but in the incidental recording of events in the background that later could be instrumental in identifying criminals.)

Is letting your insurance company know exactly how you steer and hit the gas and brake pedals, and where you drive, an acceptable price to pay for cheaper insurance? (Once insurance companies have all of this data, they could more easily offer differing insurance policy to different types of drivers.)

These are all societal dilemmas about how to balance group interest with self-interest. But before figuring out what kind of societal pressures to deploy to solve the problem, society first has to agree what the group interest is. We can’t start talking about what kind of societal pressures to set up to prevent people from keeping their genome secret, or protecting the privacy of their lifeblog, or limiting access to their car’s “black box” data, until we agree on what it means to

cooperate and what it means to defect in these situations. It's difficult to solve societal dilemmas while society itself is changing so quickly.

This isn't the first time technological change has caused social changes that forced us to rethink society, and it won't be the last. The trick will be getting societal pressure right in a society that's moving so fast that getting it wrong is an increasingly dangerous option. This means getting faster and better at setting societal pressure knobs. It means setting them right the first time, and then correcting them quickly in response to feedback, delays, and technological changes. To that end, here is a list of principles for designing effective societal pressures:

- *Understand the societal dilemma.* Not just what the group interest is, but what the group norm is, what the competing norms are, how the societal dilemma relates to other societal dilemmas, what the acceptable scope of defection is, and so on. A lot of ineffective societal pressures come from not understanding the true problem.
- *Consider all four societal pressures.* It's common to believe that one is enough: that reputation obviates the need for laws, or that a good security system is sufficient to enforce compliance. It's rare that this is true, and effective societal pressure usually involves all four categories, though not necessarily in equal measure. Considering all four will indicate how resources might be most effectively spent.
- *Pay attention to scale.* The scale of the societal dilemma influences how effective each of the four societal pressures will be. Noticing the scale, and noticing when the scale changes, is vital.
- *Foster empathy and community, increasing moral and reputational pressures.* In our large, anonymous society, it's easy to forget moral and reputational pressures and concentrate on legal pressure and security systems. This is a mistake; even though our informal social pressures fade into the background, they're still responsible for most of the cooperation in society.
- *Use security systems to scale moral and reputational pressures.* The two social pressures work best on the small scale, but security systems can enhance them to work at much larger scales. They don't work the same way, and the security systems are themselves open to attack. Still, we can't simply replace moral and reputational pressures with institutional pressures, so it is important to use technology in this way.
- *Harmonize institutional pressures across related technologies.* There shouldn't be one law for paper mail and another for e-mail, or one law for telephone conversations and another for Internet telephony. This sort of

thing used to work back when technology changed slowly. Now, by the time the legal system grinds through the process of creating a law, it may already be technologically obsolete. We need to make laws technologically invariant. This won't be easy, but we need to try.

- *Ensure that financial penalties account for the likelihood that a defection will be detected.* As I discussed in Chapter 13, a financial penalty that is too low can easily become a cost of doing business. If we expect a fine to be an effective societal pressure, it needs to be more expensive than the risk of defecting and paying it.
- *Choose general and reactive security systems.* Just as we need to make laws technologically invariant, we need to make security systems defector-invariant. That is, we need to concentrate on the broad motivations for defection, rather than on blocking specific tactics, to prevent defectors from working around security systems. One example is counterterrorism, where society is much **better off spending** money on intelligence, investigation, and emergency response than on preventing specific terrorist threats, like bombs hidden in shoes or underwear.
- *Reduce concentrations of power.* Power, whether it's concentrated in government, corporations, or non-government organizations, brings with it the ability to defect. The greater the power, the greater the scope of defection.⁷ One of the most important things society can do to reduce the risk of catastrophic defection is to reduce the amount of power held by individual actors in key positions.
- *Require transparency—especially in corporations and government institutions.* Transparency minimizes the principal-agent problem and ensures the maximum effect of reputational pressures. In our complex society, we can't monitor most societal dilemmas directly. We need to rely on others—proxies—to do the work for us. Checks and balances are the most powerful tool we have to facilitate this, and transparency is the best way to ensure that checks and balances work. A corollary of this is that society should not suppress information about defectors, their tactics, and the overall scope of defection.

We're currently in a period of history where technology is changing faster than it ever has. The worry is that if technology changes too fast, the defectors will be able to innovate so much faster than society can that the imbalances become even greater—increased scope of defection leading to an even more increased scope of defection—which can cause large societal failures. Think of

what would happen to the Red Queen Effect if the stoats evolved faster than the rabbits: they would become significantly faster than the rabbits, then eat all the rabbits, and then all starve (assume there's no other prey). Defectors in societal dilemmas can have the same effect if they evolve too quickly: they overwhelm the cooperators, which means there are no more cooperators, and the defectors themselves lose. Remember, parasites need society to be there in order to benefit from defecting; and being a parasite is a successful strategy only if you don't take too many resources from your host.

On the other hand, we're also in a period of history where the ability for large-scale cooperation is greater than it ever has been before. In 2011, law professor [Yochai Benkler](#) published a book that is in many ways a companion volume to this one: *The Penguin and The Leviathan: How Cooperation Triumphs Over Self-Interest*. Benkler writes that the Internet can and has enabled cooperation on a scale never seen before, and that politics—backed by science—is ready to embrace this new cooperation:

I am optimistic in thinking that we are not ripe to take on the task of using human cooperation to its fullest potential—to make our businesses more profitable, our economy more efficient, our scientific breakthroughs more radical, and our society safer, happier and more stable....

For decades we have been designing systems tailored to harness selfish tendencies, without regard to potential negative effects on the enormous potential for cooperation that pervades society. We can do better. We can design systems—be they legal or technical; corporate or civic; administrative or commercial—that let our humanity find a fuller expression; systems that tap into a far greater promise and potential of human endeavor than we have generally allowed in the past.

The lesson of this book isn't that defectors will inevitably ruin everything for everyone, but that we need to manage societal pressures to ensure they don't. We've seen how our prehistoric toolbox of social pressures—moral and reputational systems—does that on a small scale, how institutions enhance that on a larger scale, and how technology helps all three systems scale even more.

Over a decade ago, I wrote that "[security is a process](#), not a product." That's true for all societal pressures. The interplay of all the feedback loops means that both the scope of defection and the scope of defection society is willing to tolerate are constantly moving targets. There is no "getting it right"; this process never ends.



17

The Future

Society can't function without trust, and our complex, interconnected, and global society needs a lot of it. We need to be able to trust the people we interact with directly: as we sit next to them on airplanes, eat the food they serve us in the cabin, and get into their taxis when we land. We need to be able to trust the organizations and institutions that make modern society possible: that the airplanes we fly and the cars we ride in are well-made and well-maintained, that the food we buy is safe and their labels truthful, that the laws in the places we live and the places we travel will be enforced fairly. We need to be able to trust all sorts of technological systems: that the ATM network, the phone system, and the Internet will work wherever we are. We need to be able to trust strangers, singly and in organizations, all over the world all the time. We also need to be able to trust indirectly; we need to trust the trust people we don't already know and systems we don't yet understand. We need to trust trust.

Making this all work ourselves is impossible. We can't even begin to personally verify, and then deliberately decide whether or not to trust, the hundreds—thousands?—of people we interact with directly, and the millions of others we interact with indirectly, as we go about our daily lives. That's just too many, and we'll never meet them all. And even if we could magically decide to trust the people, we don't have the expertise to make technical and scientific decisions about trusting things like airplane safety, modern banking, and pharmacology.

Writing about trust, economist Bart Nooteboom said: “[Trust in things](#) or people entails the willingness to submit to the risk that they may fail us, with the expectation that they will not, or the neglect of lack of awareness of that possibility that they might.” Those three are all intertwined: we aren't willing to risk unless we're sure in our expectation that the risk is minor, so minor that most of the time we don't even have to think about it.

That's the value of societal pressures. They induce compliance with the group norms—that is, cooperation—so we're able to approximate the intimate trust we have in our friends on a much larger scale. It's not perfect, of course. The trust we have in actions and systems isn't as broad or deep as personal trust, but it's good enough. Societal pressures reduce the scope of defection. In a sense, by trusting societal pressures, we don't have to do the work of figuring out whether or not to trust individuals.

By inducing cooperation throughout society, societal pressures allow us to relax our guard a little bit. It's less stressful to live in a world where you trust people. Once you assume people can, in general and with qualifications, be trusted to be fair, nice, altruistic, cooperative, and trustworthy, you can stop expending energy constantly worrying about security. Then, even though you get burned by the occasional exception, your life is still more comfortable if you continue to believe.¹

We intuitively know this, even if we've never analyzed the mechanisms before. But the mechanisms of societal pressure are important. Societal pressures enable society's doves to thrive, even though there's a minority of hawks. Societal pressures enable society.

And despite the largest trust gap in our history, it largely works. It's easy to focus on defection—the crime, the rudeness, the complete mess of the political system in several countries around the world—but the evidence is all around you. Society is still here, alive and ticking. Trust is common, as is fairness, altruism, cooperation, and kindness. People don't automatically attack strangers or cheat each other. Murders, burglaries, fraud, and so on are rare.

We have a plethora of security systems to deal with the risks that remain. We know how to walk through the streets of our communities. We know how to shop on the Internet. We know how to interact with friends and strangers, whether—and how—to lock our doors at night, and what precautions to take against crime. The very fact that I was able to write and publish this book, and you were able to buy and read it, is a testament to all of our societal pressure systems. We might get it wrong sometimes, but we largely get it right.

At the same time, defection abounds. Defectors in our society have become more powerful, and they've learned to evade and sometimes manipulate societal pressures to enable their continued defection. They've used the rapid pace of technological change to increase their scope of defection, while society remains unable to implement new societal pressures fast enough in response. Societal pressures fail regularly.

The important thing to remember is this: no security system is perfect. It's hard to admit in our technologically advanced society that we can't do something, but in security there are a lot of things we can't do. This isn't a reason to live in fear, or even necessarily a cause for concern. This is the normal state of life. It might even be a good thing. Being alive entails risk, and there always will be outliers. Even if you reduced the murder rate to one in a million, three hundred unlucky people in the U.S. would be murdered every year.

These are not technical problems, though societal pressures are filled with those. No, the biggest and most important problems are at the policy level: global climate change, regulation and governance, political process, civil liberties, the social safety net. Historically, group interests either coalesced organically around the people concerned, or were dictated by a government. Today, understanding group interests increasingly involves scientific expertise, or new social constructs stemming from new technologies, or different problems resulting from yet another increase in scale.

Philosopher Sissela Bok wrote: "...trust is a social good to be protected just as much as the air we breathe or the water we drink. When it is damaged the community as a whole suffers; and when it is destroyed, societies falter and collapse." More generally, trust is the key component of social capital, and high-trust societies are better off in many dimensions than low-trust societies. And in the world today, levels of trust vary all over the map—although never down to the level of baboons.²

We're now at a critical juncture in society: we need to implement new societal systems to deal with the new world created by today's globalizing technologies. It is critical that we understand what societal pressures do and don't do, why they work and fail, and how scale affects them. If we do, we can continue building trust into our society. If we don't, the parasites will kill the host.

In closing, there are several points I want to make.

No matter how much societal pressure you deploy, there always will be defectors. All complex ecosystems contain parasites, and all human systems of cooperation and trust include people who try to take advantage of them. This will not change as long as societies are made up of humans. The possibility of perfect trust, or unbreakable security, is a science-fiction future that won't happen in the lifetime of anyone we know.

Increasing societal pressure isn't always worth it. It's not just the problem of diminishing returns discussed in Chapter 10. Looking back through history, the societies that enforce cooperation and conformance to the group norm, that

ruthlessly clamp down and punish defectors, and that monitor every aspect of their citizens' lives are not societies we think of as free. This is true whether the norms accurately reflect the desires of the group or are imposed from the top down.³ Security always has side effects and unwanted consequences.

This is okay. We've repeatedly talked about societal pressures as being necessary to sustain trust.⁴ This doesn't mean absolute trust, and it doesn't imply 100% cooperation. As long as the murder rate is low enough, speeders are few enough, and policemen on the take are rare enough, society flourishes.

Societal pressures can prevent cooperation, too. Not only do we sometimes fail to punish the guilty, we sometimes punish the innocent. People get reputations they don't deserve; people get convicted of crimes they didn't commit. And if the scope of defection is low enough, these false positives can be greater than the defection attempts thwarted. That's when you know it's time to dial back the knob.

We all defect at some times regarding some things. Sometimes we're simply being selfish. Sometimes we have another, stronger, self-interest. Sometimes we're just not paying attention. Sometimes our morality just doesn't permit us to cooperate with the group norm. And sometimes we feel a stronger attachment to another group, and its associated interests and norms. This is also okay.

Sometimes we defect honestly and innocently. Group norms can be too rigid for the way we live our lives. The white lies of our normal social interactions make relationships better, not worse. Sometimes assistants need to sign documents for their bosses, and sometimes attorneys and accountants need to innocently back-date documents. Sometimes defecting is a form of social lubricant: small social dishonesties that make life easier for everyone.

There are good defectors and there are bad defectors, and we can't always tell the difference—even though we think we can. We know that murderers are always bad and that pro-democracy demonstrators are always good, but even those truisms fray at the edges. Was the U.S.'s assassination of Osama bin Ladin good or bad? Is it okay that pro-democracy protesters in Egypt and other countries are anti-U.S. and anti-Israel? U.S. troops in Iraq may be either good or bad, depending on whether you're safely in the U.S., whether your daughter was just killed by one of them, or whether you own an oil company. Many defectors believe they are morally right: animal-rights activists who free animals from testing laboratories, the Sandinistas in Nicaragua, and the Nazis in Germany, just to name a few. And so did the Tiananmen Square protesters in China, and the United States' founding fathers.

I stumbled on [this parable](#) on the Internet as I was writing this book:

There was this kid who came from a poor family. He had no good options in life so he signed up for the military. After a few years he was deployed to a conflict infested, god-forsaken desert outpost. It was the worst tour of duty he could have been assigned. It was going to be hot and dangerous. Every day he had to live with a hostile populace who hated his presence and the very sight of his uniform. Plus, the place was swarming with insurgents and terrorists.

Anyhow, one morning the soldier goes to work and finds that he's been assigned that day to a detail that is supposed to oversee the execution of three convicted insurgents. The soldier shakes his head. He didn't sign up for this. His life just totally sucks. "They don't pay me enough," he thinks, "for the shit I have to do."

He doesn't know he's going to be executing the Son of God that day. He's just going to work, punching the time clock, keeping his head down. He's just trying to stay alive, get through the day, and send some money back home to Rome.

Systems of societal pressure can't tell the difference between good or bad defectors. Societal pressures are the mechanism by which societies impose rules upon themselves, even as the societies overlap and conflict. Those rules could be good, like a respect for human rights or a system for enforcing contracts. Those rules could be bad, like slavery, totalitarianism, persecution, or ritual murder. Or those rules could be perceived as good by some societies and bad by others: arranged marriages; heavy taxation; and prohibitions against drinking, dancing, pot smoking, or sharing music files via BitTorrent. Societal pressures simply enforce cooperation, without much consideration as to why the defector chose some competing interest. This is a good thing when it protects individuals from harm, loss, or social injustice, and a bad thing when it protects a regime that is not good to its people or prevents positive social change.

Society needs defectors. Groups benefit from the fact that some members [do not follow](#) the group norms. These are the outliers: the people who resist popular opinion for moral or other reasons. These are the people who invent new business models by copying and distributing music, movies, and books on the Internet. These are people like Copernicus and Galileo, who challenged official Church dogma on astronomy. These are the people who—to take a recent

example—disrupt energy auctions to protest government responsibility for climate change. They're also people living on the edge of society: squatters, survivalists, artists, cults, communes, hermits, and those who live off the grid or off the land. In 2011, U.S. Marine Dakota Meyer received the [Medal of Honor](#) for saving three dozen of his comrades who were under enemy fire. The thing is, he disobeyed orders in order to do so.

Defection represents an engine for innovation, an immunological challenge to ensure the health of the majority, a defense against the risk of monoculture, a reservoir of diversity, and a catalyst for social change. It's through defection from bad or merely outdated social norms that our society improves. In the stoat vs. rabbit Red Queen Effect from Chapter 2, it's the stoats that drive the change. Left to themselves, the rabbits will not improve.

This is important. The societies that societal pressures protect are not necessarily moral or desirable. In fact, they can protect some pretty awful ones. And because societal pressures necessarily become institutionalized—in police forces, in government agencies, in corporate security departments—they can be co-opted to justify and maintain those awful societies' awful institutions.

Sometimes a whistle-blower needs to publish documents proving his government has been waging an illegal bombing campaign in Laos and Cambodia. Sometimes a plutonium processing plant worker needs to [contact a reporter](#) to discuss her employer's inadequate safety practices. And sometimes a black woman needs to sit down at the front of a bus and not get up. Without defectors, social change would be impossible; stagnation would set in.

It's a tough balancing act, but I think we're up to it. Maybe not in the near term, but in the long term. History teaches how often we get it right. As Martin Luther King, Jr., said: "The arc of history is long, but it bends toward justice."⁵



Acknowledgments

I consider writing a book to be a social as well as a solitary activity. As such, I have a great number of people to thank.

First, thank you to all the participants of the four Security and Human Behavior Workshops, for sparking many of the ideas that went into this book. Thank you to my primary researcher, Kathleen Seidel, for her ability to find just what I need whether I know it exists or not, and also to Deborah Pursch for additional research. Thank you to the Hennepin County Library, which already had many of the books I needed, and got me absolutely everything else I asked for via interlibrary loan...even if I could only keep any one thing for three weeks.

Thank you to all the people who read this book in varying stages of completion. Steve Bass, Charles Faulkner, Greg Guerin, Victoria Gaile Laidler, Stephen Leigh, and David Prentiss read an early draft of the book, and were especially helpful in organizing my initial ideas. Dorothy Denning and Jay Walker read a later draft, and both suggested a reframing—on the same day, coincidentally—that turned out to be particularly useful. Rebecca Kessler gave the book a badly needed final edit. Andrew Acquisti, Andrew Adams, Michael Albaugh, Ross Anderson, Marcia Ballinger, Jason Becker, David Brown, Steve Brust, Tyler Burns, Jon Callas, David Campbell, Raphael Carter, Cody Charette, Dave Clark, Ron Clarke, Chris Cocking, Karen Cooper, David Cowan, Tammy Coxen, Cory Doctorow, John Douceur, Kevin Drum, Nicole Emery, Oisín Feeley, Eric Forste, Amy Forsyth, Peter Fraser-Mackenzie, J. Carl Ganter, Edward Goldstick, Sarah Green, Rachael Greenstadt, Jim Harper, Bill Herdle, Cormac Herley, Chris Hoofnagle, Leif Huhn, Owen Imholte, David Kahn, Jerry Kang, Arlene Katz, John Kelsey, Lori Kingerly, David Leach, David Mandel, Chris Manning, Petréa Mitchell, David Modic, Josh More, Doug Morgenstern, John Mueller, Peter Neumann, Andrew Odlyzko, Evan Oslick, Gerrit Padgham, Cirsten Paine,

Ross Patty, David Perry, Daniele Raffo, Coe Roberts, Peter Robinson, Dave Romm, David Ropeik, Marc Rotenberg, Stuart Schechter, Jeff Schmidt, Martin Schneier, David Schroth, Eric Seppanen, Susan Shapiro, Adam Shostack, Daniel Solove, Thomas Sprinkmeier, Nisheeth Srivastava, Frank Stajano, Mark Stewart, Steven L. Victor, Alan Wagner-Krankel, Paul Wallich, Chris Walsh, Rick Wash, Skyler White, and Jeff Yan also commented on all or part of the book.

Thank you to all the readers of my blog, who read my intermittent posts on the book's progress and regularly left insightful comments. Thank you to Carol Kennedy, who suggested the title in one of those comments. Thank you to Luke Fretwell, who designed a cover because he didn't like any of the candidates I posted. His idea ended up being the inspiration for the final cover.

Thank you to Josh More, for creating all the diagrams and putting up with my endless changes and tweaks. Thank you to Kee Nethery for proofreading the galleys. Thank you to Beth Friedman, for copyediting the book before it even got to my publisher, for dealing with all the MS-Word style issues, and for inserting herself between me and the copyeditor.

Thank you to Carol Long, Ashley Zurcher, Ellen Gerstein, and Tom Dinse at Wiley for helping turn the manuscript into an actual book.

Finally, thank you to the Suiboku Museum of Toyama, Japan, for the ink painting of a bulbul that adorned the manuscript cover during the writing process.



Notes

Chapter 1

Numbers preceding the notes refer to endnote numbers.

- (1) In his book, *The Speed of Trust*, Stephen Covey talks about **five levels of trust**, which he calls “waves”: self-trust, relationship trust, organizational trust, market trust, and societal trust.
- (2) Piero Ferrucci wrote:

To trust is to bet. Each time we trust, we put ourselves on the line. If we confide in a friend, we can be betrayed. If we put faith in a partner, we can be abandoned. If we trust in the world, we can be crushed. Far too often it ends that way. But the alternative is worse still, because if we do not put ourselves on the line, nothing will happen.
- (3) Diego Gambetta: “**When we say we trust** someone or that someone is trustworthy, we implicitly mean that the probability that he will perform an action that is beneficial or at least not detrimental to us is high enough for us to consider engaging in some form of cooperation with him.”
- (4) David Messick and Roderick Kramer: “**We will define trust** in these situations as making the decision *as if* the other person or persons will abide by ordinary ethical rules that are involved in the situation.”
- (5) Sociologist Anthony Giddens proposed a similar **three-level progression** of trust:

Trust in persons...is built upon mutuality of response and involvement: faith in the integrity of another is a prime source of a feeling of integrity and authenticity of the self. Trust in abstract systems provides for the security of day-to-day reliability, but by its very nature cannot supply either the mutuality or intimacy which personal trust relations offer....

In pre-modern settings, basic trust is slotted into personalised trust relations in the community, kinship ties, and friendships. Although any of these social connections can involve emotional intimacy, this is not a condition of the maintaining of personal trust. Institutionalised personal ties and informal or informalised codes of sincerity and honour provide (potential, by no means always actual) frameworks of trust....

With the development of abstract systems, trust in impersonal principles, as well as in anonymous others, becomes indispensable to social existence.

- (6) [Piotr Cofta](#) covered similar ground in his book *Trust, Complexity, and Control*.
- (7) Not coincidentally, I, along with colleagues Ross Anderson and Alessandro Acquisti, founded the annual Interdisciplinary Workshop on Security and Human Behavior in 2008.
- (8) Coming from mathematical security—cryptography—where research results are facts, it can be unsettling to research fields where there are theories, competing theories, overturned theories, and long-standing debates about theories. It sometimes seems that nothing is ever settled in the social sciences, and that for every explanation, there's a counter-explanation. Even worse, a reasonable case can be made that most research [findings are false](#) and there is [sloppy methodology](#) in the social sciences, primarily because of the pressure to produce newsworthy results. Also, that many results are based on experiments on a [narrow and unrepresentative](#) slice of humanity. The only way I can see to navigate this is to look at both the individual research results and the broader directions and meta-results.
- (9) [Adam Smith](#) wrote:

If there is any society among robbers and murderers, they must at least, according to the trite observation, abstain from robbing and murdering one another. Beneficence, therefore, is less essential to the existence of society than justice. Society may subsist, tho' not in the most comfortable state, without beneficence; but the prevalence of injustice must utterly destroy it.

Chapter 2

- (1) Chimpanzees have been observed using [sticks as weapons](#), and wrasses have been observed [using rocks](#) to open up shells.
- (2) Some of this can be pretty complex; a single [Brants's whistling rat](#) builds a burrow with dozens or hundreds of entrances, so there's always one close by to retreat to. There's even an [African rat](#) that applies a tree poison to its fur to make itself deadly.

- (3) Just recently, an [entirely separate](#), probably older, immune system was discovered in bacteria and archaea, called *Clustered Regularly Interspaced Short Palindromic Repeats* or CRISPRs.
- (4) In an [earlier book](#), I mistakenly called this the “establishing reflex.”
- (5) In one experiment, children were faster at picking out a [picture of a snake](#) than pictures of more benign objects.
- (6) Stephen Jay Gould used to call these “Just So Stories” because they rarely have any proof other than plausibility (and the fact that they make a good story). So while these seem like possible evolutionary explanations, there is still controversy in evolutionary biology over the levels of selection at work in any given instance. Certainly not all evolutionary biologists would accept these necessarily simple descriptions, although they would concur with the general outline that there was some evolutionary advantage to the possession of certain genes manifesting certain phenotypes in certain populations.
- (7) Among other things, human intelligence is unique in the complexity of its expression, and its ability to comprehend the passage of time. More related to security, humans are vastly ahead of even chimpanzees in their ability to understand [cause and effect](#) in the physical world.
- (8) No other creature on the planet does this. To use the words of philosopher Alfred Korzybski, humans are the only [time binding species](#): we are the only species that can pass information and knowledge between generations at an accelerating rate. Other animals can pass knowledge between generations, but we’re the only animal that does it at observable rates.
- (9) All 5,600 or so species of mammals are at least minimally social, if only in mating and child-rearing.
- (10) To use the words of philosopher Daniel Dennett, we need to adopt an [intentional stance](#) in order to understand each other. That is, instead of looking at people as physical objects or even biological systems, we have to look at them in terms of beliefs, intents, and thoughts.
- (11) There’s evidence from rodents that [social group size](#) is directly correlated with individuality.
- (12) There’s even a theory that reasoning evolved not because we needed to make better decisions, but because we needed to [win arguments](#) and convince other humans.
- (13) Of course, this does not necessarily mean that the sole purpose of the neocortex is to deceive.
- (14) It’s actually a range between 100 and 230; 150 is the most common value. Dunbar has often said “150, plus or minus 50.” [Others posit](#) the number is 200-ish. Groups that are more focused on survival tend to be larger, because “there’s safety in numbers.”

- (15) Larger group sizes aren't as stable because their members don't know each other well enough. We interact with people outside this circle more as categories or roles: the mailman, the emergency room nurse, that guy in the accounting department. We might recognize them as individuals, but we tend not to know a lot about them. A modern human might have a virtual network of 2,000 Facebook friends, but it's unlikely that he'll have more than a casual acquaintance with even a tenth of them.
- (16) Modern data from [primitive peoples](#) validates this number. In the primitive tribes of the New Guinea highlands, who lived apart from the rest of the world until the 1930s, about 25% of men and 5% of women died in warfare. The [Yanomamō](#) live in the upper reaches of the Orinoco River in Venezuela and Brazil. While they once had only sporadic contact with other cultures, they still lived apart in their traditional manner. They lost 24% of men and 7% of women to warfare.
- (17) [Big-game hunting](#) is inefficient because: 1) big game's low density means fewer encounters, 2) it's harder to catch, 3) it can hurt you when you hunt it, 4) it requires a lot of people to catch, 5) it takes a lot of work to butcher and preserve, and 6) it's perishable, and must be eaten quickly or preserved before it spoils.
- (18) [Chimpanzees' aggression rates](#) are two to three orders of magnitude higher than humans', although their lethal aggression rates are about the same as those of human subsistence societies.

Chapter 3

- (1) There is evidence that [increased specialization](#) is a function of group size. To be fair, there are researchers who maintain that division of labor is [not what makes](#) leafcutter ants so successful.
- (2) To some extent, this is also true of other social insects that don't have polymorphism. Bees, for instance, tend to change specializations as they age, but they can change early if some task is going undone. Leafcutter ants can't do this; they're physiologically distinct according to role.
- (3) This [startling statistic](#) comes from the fact that there are a lot of other organisms in our digestive tract: "The adult human organism is said to be composed of approximately 10^{13} eukaryotic animal cells. That statement is only an expression of a particular point of view. The various body surfaces and the gastrointestinal canals of humans may be colonized by as many as 10^{14} indigenous prokaryotic and eukaryotic microbial cells." Note that the percentage is by number, not by volume or weight. All those digestive organisms are much, much smaller than our own cells.
- (4) The initial paper is actually more complicated. In addition to hawks and doves, there are bullies who only pick on doves, retaliators who respond as hawks against hawks

and as doves against doves, and so on. And many other game theorists wrote papers analyzing this or that variant of the Hawk-Dove game, looking at other strategies or more complications in the simulation. But this simple representation is sufficient for our needs. Adding some sort of “fighting skill” parameter is a complexity that doesn’t add to our understanding, either.

- (5) Researchers have also conducted Hawk-Dove games with more fluid strategies. Instead of being 100% hawk or 100% dove, individuals could be a combination of both. That is, one individual might behave as a dove 80% of the time and as a hawk 20% of the time. What that individual does in any given situation might be random, or depend on circumstance. This complication better mirrors the behavior of real people.

Another way to make strategies more fluid is to allow individuals to use [some mixture](#) of hawk and dove strategies in a *single* encounter. So instead of either being all hawk or all dove, an individual might be 20% hawk/80% dove. That is, she might cooperate a lot but not fully or exclusively. This is definitely a more realistic model; we cooperate to different degrees with different people at different times.

In this more complicated model, it’s much harder for cooperative behavior to appear. If everyone is constantly switching from the dove camp to the hawk camp and vice versa—as happens in most species—a genetic mutation that enables a small amount of cooperation doesn’t confer enough benefit to take hold in the broader population before it gets stamped out by the defectors taking advantage of it.

- (6) Of course this is simplistic. The effects of laws on crime isn’t nearly as direct and linear as this example. We’ll talk about this more in Chapter 9. But the basic idea is correct.
- (7) The costs and benefits of being a hawk also depend on [population density](#). In simulations, dense populations have more doves, and sparse populations more hawks.
- (8) [South African meerkats](#) raise their young communally; even distantly related non-breeders will pitch in to protect newborn pups in their burrows, deliver them beetles, scorpions, and lizards to eat, and even pass along new foods mouth-to-mouth to help them become accustomed to unfamiliar flavors. [Red ruffed lemurs](#) engage in extensive alloparenting.
- (9) It’s [much less common](#) in the wild. It’s [also slow](#); there is evidence that mutualism appears to evolve more slowly than other traits.
- (10) Between species, mutualism is more commonly known as symbiosis. [Wrasse cleaner fish](#) are the canonical example; they eat parasites and dead skin off larger fish. This feeds the wrasses and provides a health benefit to the larger fish. Similarly, clownfish tend to stay within the tentacles of Ritteri sea anemones; each protects the other from predators. Pollination, too: the bees get food, and the plants get pollinated. It is easy for mutualistic relationships to evolve, which is why they are common throughout the natural world.
- (11) Sometimes, the benefit of fighting and winning is so great that most individuals will be hawks. Male [elephant seals](#) are an example; the winner gets to mate with all the

females on the beach. Sometimes, the risk of injury is so low that most individuals will be hawks: [some bullfrogs](#) function that way, because they can't really injure each other. Sometimes, the risk of injury is so high that almost [everyone is a dove](#): oryx and other hoofed animals with nasty horns, rattlesnakes, and so on. Most often, though, [there is a mixture](#) of hawks and doves within a population. Sometimes it's the more aggressive individuals that are hawks. Sometimes an animal is a hawk within its own territory and a dove outside it.

- (12) Economist [Kaushik Basu](#) described the problem in the introduction to his book *The Less Developed Economy*. Paraphrasing: Imagine that you are in a strange city, and you've hired a taxi to take you from the airport to your hotel. You and the taxi driver have never previously met, and you'll never meet again. Why do you pay him at the end? If you were just calculating, you might not bother. After all, the taxi driver has already driven you to your destination. Still, you might realize that if you didn't pay, the driver would make a huge fuss, embarrass you in public, perhaps resort to violence, and perhaps call the police. It's just not worth the risk for such a small amount of money. But here's the problem: even if you do pay, the taxi driver could still do all of that. If the taxi driver is just as calculating as you are, why doesn't he accuse you of nonpayment regardless? Double money for him, and he'll never see you again. So if he were going to do that, you might as well not pay. You can take the analysis even further. Maybe you both calculate that if the police got involved, the courts would figure out who wronged the other—maybe there was a camera in the taxi that recorded the whole thing—so it makes sense to be honest. But that doesn't help, either. If the police and the judges are just as calculating as you and the taxi driver, why should they attempt to resolve the dispute fairly, rather than in favor of the side that gave them the biggest bribe? They might fear they would get caught and punished, but that fear assumes those doing the catching and punishing aren't calculating and will attempt to be fair and honest.
- (13) [Neuroscience is starting](#) to make inroads into that question, too.
- (14) The [Ultimatum game](#) was first developed in 1982, and has been replicated repeatedly by different researchers using different variants in different cultures; there are hundreds of academic papers about the Ultimatum game.

Here's [how the game works](#). Two strangers are put in separate rooms and told they will divide a pot of money between them. They can't meet each other, and they can't communicate in any way. Instead, one of the subjects gets to divide the money any way he wants. That division is shown to the second subject, who gets to either accept or reject the division. If he accepts it, both subjects get their shares. If he rejects the division, neither subject gets a share. After this single division task, the experiment ends, and the two subjects leave via separate doors, never to meet.

Game theory predicts, and a rational economic analysis agrees, that the first player will make the most unfair division possible, and that the second player will accept

that unfair division. Here's the logic. The second player is smart to accept any division, even the most lopsided one, because some money is better than no money. And the first player, knowing that the second player will accept any division, is smart to offer him the most lopsided division possible. So if there's \$20 to divide, the first player will propose a \$19/\$1 split and the second player will accept it.

That makes sense on paper, but people aren't like that. Different experiments with this game found that first players generally offer between a third and a half of the money, and that the most frequent offer is a 50–50 split. That's right: they give money to strangers out of their own pocket, even though they are penalizing themselves economically for doing so, in an effort to be fair. Second players tend to reject divisions that are not at least reasonably fair; about half of the players **turn down offers** of less than 30%.

This experiment has been conducted with subjects from a wide variety of **cultural backgrounds**. It has been conducted with **large amounts of money**, and in places where small amounts of money make a big difference. Results are consistent.

- (15) The Dictator game is like the Ultimatum game, but with one critical difference: the second player is completely passive. The first player gets to divide the money, and both players receive their share. If the first player wants to keep all of it, he does. The second player has no say in the division or whether or not it is accepted.

In the Ultimatum game, the first player had to worry if the second player would penalize him. The Dictator game removes all of that second-guessing. The first player gets a pile of money, and hands the second player some, then keeps the rest. He is in complete control. Even in this game, people aren't as selfish as rational economic theory predicts. In one experiment, first players split the money evenly three-quarters of the time. Other experimental results are **more lopsided** than that, and the first player's division tends to be less fair than in the Ultimatum game, but not as unfair as it could be.

- (16) In the **Trust game**, the first player gets a pile of money. He can either keep it all or give a portion to the second player. Any money he gives to the second player is increased by some amount (generally 60%) by the researchers, then the second player can divide the increased result between the two players.

Assume \$10 is at stake here. If the first player is entirely selfish, he keeps his \$10. If he is entirely trusting, he gives it all to the second player, who ends up with \$16. If the second player is entirely selfish, he keeps the \$16. If he is completely fair, he gives the first player \$8 and keeps \$8.

Rational economic behavior predicts a very lopsided result. As in the Dictator game, the second player would be smart to give no money to the first player. And the first player, knowing this would be the second player's rational decision, would be smart to not give any money to the second player. Of course, that's **not what happens**. First players give, on average, 40% of the money to the second player. And second players, on average, give the first player back a third of the multiplied amount.

- (17) In a **Public Goods game**, each player gets a small pile of money. It's his to keep, but he can choose to pool some portion of it together with everyone else's. The researchers multiply this pool by a predetermined amount, then evenly divide it among all players.

A rational economic analysis of the game—that is, an analysis that assumes all players will be solely motivated by selfish interest or the bottom line—predicts that no one will contribute anything to the common pool; it's a smarter strategy to keep everything you have and get a portion of what everyone else contributes than it is to contribute to the common pool. But that's not what people do. Contrary to this prediction, people generally contribute 40–60% into the common pool. That is, people are generally not prepared to cooperate 100% and put themselves at the mercy of those who defect. But they're also generally not willing to be entirely selfish and not contribute anything. Stuck between those opposing poles, they more-or-less split the difference and contribute half.

- (18) One of the theories originally advanced to explain the first player's behavior in the Ultimatum game was **fear of rejection**. According to that theory, he is motivated to offer the second player a decent percentage of the total because he doesn't want the second player to penalize him by rejecting the offer. There's no rational reason for the second player to do that, but we—and presumably the first player—know he will. That explanation was proven wrong by the Dictator game.

Some researchers claim these experiments show that humans are naturally altruistic: they seek not only to maximize their own personal benefit but also the benefit of others, even strangers. **Others claim** that the human tendency at work in the different games is an aversion to being seen as greedy, which implies that reputation is the primary motivator.

Still other researchers try to explain results in terms of evolutionary psychology: individuals who cooperate with each other have a better chance of survival than those who don't. Today, we regularly interact with people we will never see again: fellow passengers on an airplane, members of the audience at public events, everyone we meet on our vacations, almost everyone we interact with if we live in a large city. But that didn't hold true in our evolutionary history. So while the Ultimatum, Dictator, and Trust games are one-time-only, our brains function as if we have a social network of not much more than 150 people, whom we are certain to meet again and again, often enough that the quality of our interactions matters in the long run.

- (19) We naturally gravitate toward **fair solutions**, and we naturally implement them: even when dealing with strangers, and even when being fair penalizes us financially. As one paper put it, “concerns for a fair distribution originate from personal and social rules that effectively constrain self-interested behavior.”

Joseph Henrich interviewed his subjects after Ultimatum game experiments and found that they thought a lot about fairness. First players wanted to **do what was fair**. Second

players accepted offers they thought were fair, and rejected offers they thought were unfair. They would rather receive no money at all than reward unfairness.

In variants of the Ultimatum and Dictator games where the first player won his position by an **act of skill**—doing better on a quiz, for example—he tended to offer less to the second player. It worked the **other way, too**; if the second player won his position in an act of skill, the first player tended to give him more.

- (20) There's a variant of the Public Goods game where subjects are allowed to spend their own money to **punish other players**; typically, it's something like \$3 deducted from the punished for every \$1 spent by the punisher. In one experiment, two-thirds of the subjects punished someone at least once, with the severity of the punishment rising with the severity of the non-cooperation. They did this even if they would never interact with the punished player again.

What's interesting is that **the punishment works**. Stingy players who have been punished are less stingy in future rounds of a Public Goods game—even if the punishers themselves aren't involved in those future rounds—and that behavior cascades to other players as well.

There's other research, with rewards as well as punishment, but the **results are mixed**; rewards seem to be less effective than punishment in modifying players' behavior.

- (21) A variant of the Dictator game illustrates this. Instead of giving, the first player can take money from the second player. And in many cases, he does. The rationalization goes along the following lines. In the standard version of the Dictator game, first players understand that the game is about giving, so they figure out how much to give. In this variant, the **game is about taking**, so they think about how much to take. A variant of the Trust game, called the **Distrust game**, illustrates a similar result.
- (22) Lots of fraud is based on feigning group identity.
- (23) About three-quarters of people give **half of the money** away in the Ultimatum game, but a few keep as much as possible for themselves. The majority of us might be altruistic and cooperative, but the minority is definitely selfish and uncooperative.
- (24) To be fair, there is a minority of researchers **who are skeptical** that mirror neurons are all that big a deal.
- (25) This is called the **prototype effect**, and has ramifications far greater than this one example.
- (26) **In many societies**, sharing when you have plenty obligates others to share with you when you're in need.
- (27) Notice that the four work best in increasingly larger group sizes. Direct reciprocity works best in very small groups. Indirect reciprocity works well in slightly larger groups. Network reciprocity works well in even larger groups. Group reciprocity

works well in even larger groups: groups of groups. I don't know of any research that has tried to establish the different human group sizes in which these operate, and how those sizes compare to Dunbar's numbers.

- (28) The [majority belief](#) is that it was primarily kin selection that sparked the evolution of altruistic behavior in humans, although Martin Nowak and Edward O. Wilson have recently caused quite a stir in the evolutionary biology community by proposing group selection as the driving mechanism. [One rebuttal](#) to this hypothesis was signed by 137 scientists. I have no idea how this debate will turn out, but it is likely that all mechanisms have operated throughout human evolutionary history, and [reinforced each other](#).
- (29) There's a lot here, and there have been [many books published](#) in the last few years on this general topic of neuropsychology: Michael Shermer's *The Science of Good and Evil*, Nigel Barber's *Kindness in a Cruel World*, Donald Pfaff's *The Neuroscience of Fair Play*, Martin Nowak's *SuperCooperators*, and Patricia Churchland's *Braintrust*. The last two are the best. There's also an [older book](#) on the topic by Matt Ridley.

Chapter 4

- (1) Very often, understanding how societal pressures work involves understanding human—and other animal—psychology in evolutionary terms, just as you might understand the function of the pelvis, the spleen, or male pattern baldness. This is [evolutionary psychology](#), first proposed by Edward O. Wilson in 1975, and which has really taken off in the last couple of decades. This is a new way of looking at psychology: not as a collection of behaviors, but as a manifestation of our species' development. It has the very real potential to revolutionize psychology by providing a meta-theoretical framework by which to integrate the entire field, just as evolution did for biology over 150 years ago.

To be fair, the validity of evolutionary psychology research is not universally accepted. Geneticist Anne Innis Dagg argues both that the [genetic science is flawed](#), and that the inability to perform experiments or collect prehistoric data render the conclusions nothing more than Gould's "Just So Stories."

However, evolutionary psychology is not only about genetic determinism. An evolutionary explanation for behavior does not equate to or imply the existence of a genetic explanation. Behaviors, especially human behaviors, are much more multifaceted than that. Certainly genes are involved in many of our psychological processes, especially those as deep-rooted as making security and trust trade-offs, but natural selection is possible with any characteristic that can be passed from parent to child. Learned characteristics, cultural characteristics, stories that illustrate model behavior, technical knowledge—all can be passed on. Evolutionary psychology is a mix of genetic and non-genetic inheritance.

- (2) It's called *ecological validity*. We are built for conditions of the past, when—for example—humans were worried about attack from large predators, not from small lead slugs from a gun 100 yards away in the dark. So the forehead protects us against blows from blunt objects, but is much less effective against bullets. Similarly, the skull is great protection for falls, but less effective against IEDs. The loss of ecological validity has meant the end of many species that could no longer adapt to changing conditions.
- (3) Of course, the cost of not paying that tax would be even more expensive. To take just one example, [Douglass North](#) wrote: “The inability of societies to develop effective, low-cost enforcement of contracts is the most important source of both historical stagnation and contemporary underdevelopment in the Third World.”
- (4) There's a reasonable argument that [no money would](#) be necessary, either. Reciprocal altruism would be enough for angels. Money is only required when debt becomes formal.
- (5) This is named after anthropologist [Terrence Deacon](#), who first described it.
- (6) Conservative estimates are that between 20% and 25% of all Americans have had sex with someone who is not their spouse while they are married.
- (7) A gaggle of recent animal studies across a variety of species demonstrate that there's [far more philandering](#) going on in the animal world than we previously thought. Of about 4,000 mammalian species, only a few are monogamous. Even birds, once regarded as the poster children of monogamy, aren't all that faithful to their mates. Once DNA fingerprinting became cheap in the 1990s, study after study showed that anything from 10% to 40% of chicks are not raised by their biological father.
- (8) [There's a balance](#) here. Archaeological evidence indicates that Neanderthals, while violent like any other primate, were more compassionate than early humans. Yet they died out while our ancestors survived. There is [preliminary evidence](#) that Neanderthals engaged in cannibalism.
- (9) These numbers are reflected in [military organization](#) throughout history: squads of 10 to 15, organized into platoons of three to four squads, organized into companies of three to four platoons, organized into battalions of three to four companies, organized into regiments of three to four battalions, organized into divisions of two to three regiments, and organized into corps of two to three divisions.
- (10) There are several theories on the evolutionary [origins of religion](#). While all talk about the ways it induces societal cohesion, they differ as to whether that's an essential aspect of religion or just a side effect.
- (11) The combination of these three are what sociologists call [social controls](#). I am [not using that term](#) because 1) it traditionally does not include coercive measures, and I need a term that encompasses both coercive and non-coercive measures, and 2) its definition has changed over the years and now is limited to crime and deviance. Also, the sociological term has never included physical security measures. Finally, I am

avoiding it for the same reason I am avoiding the game-theoretic term “social dilemmas”; I want to emphasize the societal aspect of these systems.

- (12) The research is by no means conclusive, but data from Facebook, Twitter, and elsewhere indicates that Dunbar’s numbers are **not growing** due to information technology. Facebook claims the average user has 130 friends; if you ignore people who don’t actually use their accounts, my guess is that the median is around 150. (<http://www.facebook.com/press/info.php?statistics>.) There’s even evidence that links the number of Facebook friends to the size of **certain brain regions**. Such social networks are changing the definition of “friend.” How else can you explain that so many of our Facebook pages include people we would never have even considered talking to in high school, and yet we help water their imaginary plants?

Chapter 5

- (1) The **Prisoner’s Dilemma** was originally framed in the 1950s by Merrill Flood and Melvin Dresher at the RAND Corporation, and was named several years later by Albert Tucker. **Many researchers** have informed and analyzed this game, most famously John Nash and then Robert Axelrod, who used it to help explain the evolution of cooperation.
- (2) I should probably explain about Alice and Bob. Cryptographers—and I started as a cryptographer—name the two actors in any security discussion Alice and Bob. To us, anyone we don’t know is either Alice or Bob. If you meet me, don’t be surprised if I call you Alice or Bob.
- (3) As stylized as the story is, this sort of thing is not uncommon. It’s basic plea bargaining.
- (4) I heard the story of someone who never stops at four-way stop signs, because he figures that the other person will stop. This hawkish strategy works great, as long as he only meets doves at intersections.
- (5) One database search yielded 73,000 academic papers with the phrase “Prisoner’s Dilemma” in the title.
- (6) Hardin used an **open grazing pasture** as an example. From the paper:

Picture a pasture open to all. It is to be expected that each herdsman will try to keep as many cattle as possible on the commons. Such an arrangement may work reasonably satisfactorily for centuries because tribal wars, poaching, and disease keep the numbers of both man and beast well below the carrying capacity of the land. Finally, however, comes the day of reckoning, that is, the day when the long-desired goal of social stability becomes a reality. At this point, the inherent logic of the commons remorselessly generates tragedy.

As a rational being, each herdsman seeks to maximize his gain. Explicitly or implicitly, more or less consciously, he asks, “What is the utility to me of adding one more animal to my herd?” This utility has one negative and one positive component.

1. The positive component is a function of the increment of one animal. Since the herdsman receives all the proceeds from the sale of the additional animal, the positive utility is nearly +1.
2. The negative component is a function of the additional overgrazing created by one more animal. Since, however, the effects of overgrazing are shared by all the herdsmen, the negative utility for any particular decision-making herdsman is only a fraction of -1.

Adding together the component partial utilities, the rational herdsman concludes that the only sensible course for him to pursue is to add another animal to his herd. And another.... But this is the conclusion reached by each and every rational herdsman sharing a commons. Therein is the tragedy. Each man is locked into a system that compels him to increase his herd without limit—in a world that is limited. Ruin is the destination toward which all men rush, each pursuing his own best interest in a society that believes in the freedom of the commons.

- (7) Normal games are zero-sum: someone wins, and someone else loses. The sum of the win (+1) and the loss (-1) totals zero.
- (8) Yes, these rules are sometimes made by autocratic rulers for their own benefit. We'll talk about this in Chapter 11.

Chapter 6

- (1) One way to think about defectors is that they are less risk-averse than cooperators. As a result, the cooperators tend to obtain moderate benefits with few severe costs, whereas defectors might get much larger benefits, but in the long run tend to pay more severe costs.
- (2) Dan Ariely's term, “[predictably irrational](#),” describes us pretty well.
- (3) The name comes from the movie *Rebel Without a Cause*, in which the antihero, Jim Stark, and the local bully race stolen cars toward a cliff; the first to jump out earns the shame of being called “chicken.” Of course, if no one defects, both cars fly over the edge and both players die. (If you don't have a convenient cliff, you can play the game by racing two cars directly at each other; the first person to swerve to avoid the oncoming car is the chicken.) In this game, cooperate–cooperate is the best solution, but cooperate–defect or defect–cooperate is much better than defect–defect. In foreign policy, this is known as brinkmanship, a strategy that almost led to disastrous

consequences during the [Cuban Missile Crisis](#) of 1962. There have been some fascinating experiments with Chicken that really seem to have brought out the [worst in people](#).

- (4) For many interactions, the Snowdrift Dilemma is a [better model](#) of the real world than the Prisoner's Dilemma.
- (5) There's also the unfortunately named Battle of the Sexes. He wants to do a stereotypically male thing on Saturday night. She wants to do a stereotypically female thing. The dilemma comes from the fact that each would rather do either of the two things with the other than do the stereotypical thing alone.
- (6) In behavioral economics, [Prospect Theory](#) has tried to capture these complexities. Daniel Kahneman is the only psychologist to ever win a Nobel Prize, and he won it in economics.
- (7) Many of the criticisms of Hardin's original paper on the Tragedy of the Commons pointed out that, in the real world, [systems of regulation](#) were commonly established by users of commons.
- (8) Douglas Hofstadter calls this "[superrationality](#)." He assumes that smart people will behave this way, regardless of culture. In his construction, a superrational player assumes he is playing against another superrational player, someone who will think like he does and make the same decisions he does. By that analysis, cooperate-cooperate is much better than defect-defect. In so doing, players are being [collectively rational](#), rather than individually rational. Collectively, cooperating is better.
- (9) In societies that prescribed a particular hand for eating and the other hand for wiping, this also made it impossible for the thief to eat in public without shaming himself.
- (10) Law professor Lawrence Lessig proposed a theory of regulation that identified [four different modalities](#) by which society can modify individual behavior: norms, markets, laws, and architecture. To use one of his examples, society could reduce smoking through a public ad campaign, a tax, smoking bans, or regulations on what quantity of addictive chemicals cigarettes can contain. According to Lessig, a smart regulator uses them all—or, at least, is aware of them all.

My model is similar. I've broken Lessig's "norms" into moral and reputational because, from the point of view of societal pressure, they're very different. Lessig's "markets" can either be informal or formal; in my model, that corresponds to reputational and institutional. And I've combined institutional markets with laws because, from a security perspective, they're similar enough to be treated together. My security is roughly analogous to Lessig's "architecture."

In *Freakonomics*, Steven Levitt and Stephen Dubner write that "there are [three basic flavors](#) of incentive: economic, social, and moral." These correspond to my institutional, reputational, and moral pressures.

Chapter 7

- (1) Voting by mail is much easier, which is why it is becoming increasingly common in jurisdictions that offer the option to everyone.
- (2) Unfortunately, [the same analysis](#) shows that it's not worth people's trouble to be informed voters; their most logical course of action is to vote but remain politically ignorant.
- (3) It's [related to other](#), more general, moral rules. For example, altruism is a major factor in predicting whether someone will vote or not.
- (4) I am [not distinguishing](#) between the terms "morals" and "ethics." Although many philosophers make distinctions between the two concepts, a debate about moral theory is far beyond the scope of this book. And my definition of "morals" is pretty inclusive.
- (5) "[National Voter Turnout](#) in Federal Elections 1960–2008," Infoplease.com, 2008. Note that this isn't the same as registered voters. In the U.S., voting is generally a two-step process. First you have to register. Then you have to vote. In most states, you can't even do both on the same day.
- (6) Here is how it's expressed in a variety of religions:

Judaism: "What is hateful to you, do not do to your fellow man. This is the entire Law; all the rest is commentary." —Talmud, Shabbat 31d.

Christianity: "So in everything, do to others what you would have them do to you, for this sums up the Law and the Prophets." —Matthew 7:12. Also "Do to others as you would have them do to you." —Luke 6:31.

Islam: "No one of you is a believer until he desires for his brother that which he desires for himself." —Forty Hadith of an-Nawawi 13.

Hinduism: "This is the sum of duty; do naught unto others what you would not have them do unto you." —Mahabharata 5,1517.

Confucianism: "Do not do to others what you would not like yourself. Then there will be no resentment against you, either in the family or in the state." —Analects 12:2.

Buddhism: "Hurt not others in ways that you yourself would find hurtful." —Udana-Varga 5,1.

Taoism: "Regard your neighbor's gain as your gain, and your neighbor's loss as your own loss." —Tai Shang Kan Yin P'ien, Chapter 49.

Jainism: "A man should wander about treating all creatures as he himself would be treated." —Sutrakritanga 1.11.33.

Zoroastrianism: “That nature alone is good which refrains from doing unto another whatsoever is not good for itself.” —Dadisten-I-dinik, 94,5.

Bahá’í: “And if thine eyes be turned towards justice, choose thou for thy neighbour that which thou choosest for thyself.” —Epistle to the Son of the Wolf, 30.

Philosophers and theologians see a [significant difference](#) in the positive and negative phrasing of this rule—“do unto others what you want...” versus “don’t do unto others what you don’t want...”—but that’s too far into the details for our purposes. As a societal pressure system, the altruistic and reciprocal nature of the rule is enough. Treat others well, because we will all be better off if everyone does the same.

(7) Here’s a random sampling:

From the [Chácobo of Bolivia](#): “If you are a human being, then you will share what you have with those who are in need.”

From the [Maori of New Zealand](#): “By many, by thousands is the object attained.”

From the [Yeyi of Botswana](#): “When staying in a happy community be happy: when staying in a sad community be sad,” and “It’s the termites which cause the tree to fall down”—basically, minor disputes undermine the strength of the community.

- (8) It’s also been demonstrated that people who [believe in free will](#) are less likely to cheat on tests or slack off on the job than those who [believe in predestination](#). No one is sure why: perhaps believing that you don’t have a choice in what you do undermines a person’s sense of integrity, or perhaps it just provides a convenient excuse for giving in to selfish temptations, as if they were an unavoidable destiny. Predictably, individuals who embrace the [concept of free will](#) are also more likely to hold other people responsible for their own actions, which in turn makes them more likely to punish defectors. I’m not saying that the concept of free will is innate, or that it evolved as a societal pressure system, but it seems to function as one.
- (9) Hauser is a discredited academic. Harvard recently found him guilty of scientific misconduct; a paper has been retracted, and he’s currently on leave and is no longer allowed to teach. Even so, his book has a lot of interesting insights into human moral systems.
- (10) [Inbreeding is likely](#) to result in recessive genetic disorders, making individuals less viable. This is why cheetahs, being so inbred because of how close to extinction they came at some point in their history, have such a high disease rate: there’s just not enough variety in the gene pool: Amish, too.
- (11) The game was the Ultimatum game (see note 14 in Chapter 3 for a full description). The goal was to find people [isolated from modern](#) society, and the Machiguenga tribe fit the bill. What Henrich found was that the first player tended to make what we would consider unfair divisions—85%/15% or so—and the second player would accept them. By contrast, people from modern societies playing the same game tend to reject such unbalanced

divisions. In post-game interviews, Machiguenga subjects told him they would accept any offer. It's not that they were more unwilling than their more urbane counterparts to either be unfair or to accept unfairness, but that they considered the unfairness to have occurred at the point where the first and second player were chosen. That is, first players considered themselves lucky to have been chosen as first players, and second players thought it bad luck to have been chosen as second players. Once they accepted their positions, the division wasn't tainted with notions of fairness. The minority of tribesmen who responded to the game in a manner more similar to players from industrialized societies were those who had spent the most time interacting with people beyond their tribe.

- (12) Believe it or not, there are security systems to help ensure that employees [wash their hands](#) before leaving the restroom, mostly involving hand stains that don't come out without vigorous washing.
- (13) The phrase “bad apple” has been misused recently. More and more, it's used to mean isolated wrongdoers whose actions don't affect anyone else in the group. The entire phrase is “one bad apple spoils the entire bunch,” and is intended to explicitly highlight how the reputation of one person can taint the reputation of all people in the group. Incidentally, this is actually true for apples stored in a root cellar. A spoiled apple will cause the rest of the apples to spoil.
- (14) The logical extreme of this idea is the “[broken windows theory](#)” of John Q. Wilson and George Kelling, that visible signs of criminal activity like broken windows and abandoned cars actually incite people to commit crimes. Wilson and Kelling believed that if you clean up these visible signs of lawlessness, a neighborhood will become safer overall; societal pressures against petty crime will cause a reduction in violent crime.

It sounds good, and Kelling used the theory to explain the dramatic drop in crime in New York City in the 1990s, but it turns out there's not much actual evidence that it's true. [Researchers compared](#) New York City and other cities, and found that New York's punitive measures against low-level visible lawlessness—a lot of which might be considered punitive measures against homelessness—didn't make much of a difference. It's not that this effect doesn't exist at all—there is [evidence that it does](#). It's that other causes of crime are more important, and focusing societal pressure on low-level criminal activities in the expectation that it will prevent other crimes is much less effective than directly preventing those other crimes.

Economist [Steven Levitt](#) looked at the reduction of crime across the U.S. in the 1990s and concluded: “Most of the supposed explanations...actually played little direct role in the crime decline, including the strong economy of the 1990s, changing demographics, better policing strategies, gun control laws, concealed weapons laws and increased use of the death penalty. Four factors, however, can account for virtually all of the observed decline in crime: increases in the number of police, the rising prison population, the waning crack epidemic and the legalization of abortion.”

- (15) A recent study of 75,000 households served by the Sacramento Municipal Utility District and Puget Sound Energy found that customers who received peer comparison

charts [reduced their energy](#) usage by an average of 1.2% to 2.1%, a change that was sustained over time. Of course, this isn't absolute. There are people who don't care, or don't care enough to make changes in their behavior—and there is evidence that [this system backfires](#) with some conservatives. Even so, enough [people are swayed](#) into cooperation by the comparison charts to make them an effective societal pressure system.

- (16) [In Rwanda](#), marriages between members of the Hutu and Tutsi ethnic groups are common. But when extremist Hutus came to power in the early 1990s, they pushed an increased stigmatization of Rwandese in mixed marriages. In the new Hutu morality, Tutsi women were vilified as immoral temptresses, and the men who succumbed to their charms were viewed as traitors.
- (17) We tend to empathize more with people suffering from acute problems than those with chronic need. Witness the outpouring of aid to the Indian Ocean tsunami victims of 2004 versus the aid given annually for things like malnutrition.
- (18) The [cash box was made](#) of wood, with a slot for money. Initially Feldman used an open basket of money, but some people took the money. He then tried a coffee can with a lid, but people stole from that, too. A locked wooden box is enough of a deterrent. The only way to take the money is to steal the box itself, which only happened about once a year.

There are [a host of unknowns](#) in these data. Did everyone pay 90%, or did nine in ten pay full price and one in ten pay nothing? This sort of honor system offers many ways to partially defect. Still, it offers interesting insights into how moral pressure works. As prices rose, the payment rate fell. This makes sense: as the financial benefit of non-payment increased, some people who were just barely on the side of cooperation were willing to overcome the moral prohibition against theft. Data from the number of bagels eaten showed that price-sensitive customers were more likely to defect than more consistent consumers. This also makes sense. People who purchased donuts—he started bringing them in, too—were more likely to underpay than people who purchased bagels. Maybe this meant that donut eaters were less cooperative than bagel eaters, although it might have had something to do with the perceived price versus value of the two items, or the fact that donuts are considered junk food whereas bagels are not. And there was a sharp and persistent increase in payment following the 9/11 terrorist attacks, in line with the in-group loyalty effects I talked about earlier.

Chapter 8

- (1) Researchers have used the Prisoner's Dilemma to study this. [People who defect](#) predict a 76% defection rate from other players, and people who cooperate predict a 68% cooperation rate. Put in layman's terms, people reflexively think others are like themselves. More interestingly, in one experiment, people were asked to predict the behavior of

other players after chatting with them for half an hour. Then, people were **better at predicting** who would cooperate and who would defect. In another experiment, players were asked to evaluate the intentions of their opponents at various points during a multi-round Prisoner's Dilemma game. Cooperative players were **better at recognizing** other cooperative players; defecting players regularly mischaracterized cooperative players as defecting. This isn't surprising since people tend to see themselves in others.

- (2) Reputation mattered in the various “game” experiments mentioned in Chapter 3: the Ultimatum game, the Dictator game, the Public Goods game, and so on. Subjects were more altruistic, more fair, and more cooperative when their actions were known to the researchers or when they met the other players, and less so when they were anonymous and alone in a room.
- (3) In 1984, political scientist **Robert Axelrod** studied an iterated Prisoner's Dilemma. He set up a computer tournament and invited academic colleagues from all over to compete against each other. What he found was interesting, and in hindsight fairly obvious. Successful strategies had four basic characteristics:

They were altruistic—Axelrod used the word “nice”—in that they did not defect before their opponent did.

They were retaliatory, and responded to defection with defection.

They were forgiving, and would cooperate again at some later point.

They were non-envious; their goal wasn't to outscore their opponent.

The most successful strategy—called “tit-for-tat”—was extremely simple. A tit-for-tat player would first cooperate, then mirror his opponent's previous move. If his counterpart cooperated in a round, then tit-for-tat would cooperate in the next. If his counterpart defected in a round, then tit-for-tat would defect in the next. If two tit-for-tats competed, they would both cooperate forever. Essentially, Axelrod discovered reputation.

- (4) The **oft-quoted line** is that the average dissatisfied customer will tell 9–10 of his friends, and that 13% will tell 20 or more people. On Facebook, they'll tell everyone they know; and on Yelp, they'll tell everyone they don't know. Of course, there's a difference between reputation learned firsthand and reputation learned secondhand, similar to the personal and impersonal trust discussed in Chapter 1.
- (5) Target stores used to go so far as to accept returns of items they knew weren't purchased at Target. They calculated it was better to accept the return than argue with the customer about where the item was purchased. They no longer do this; presumably too many defectors took advantage of the system.
- (6) Prisoner's Dilemma experiments confirm that when players know each other's reputations—instead of being anonymous—**cooperation jumps** from around 50% to around 80%.

- (7) **Dueling isn't always** irrational; an economic analysis of the practice demonstrates that it made sense, given the reputational realities of the day. Similarly, the deadly defense of reputations that occurs in the criminal underworld also makes economic sense.
- (8) **Chimpanzees are able** to learn about the reputation of others by eavesdropping on third-party interactions, but they do not directly communicate with each other about the reputation of other chimpanzees.
- (9) The Islamic notion of *ihsan*—that people should do right because God is always watching their thoughts and deeds—is relevant here. Pascal's Wager takes this view to a somewhat cynical conclusion: it's better to cooperate (believe in God, follow God's rules, and so on) than to defect, because the potential downside of defecting is so great.
- (10) Better yet, do good and let someone find out about it surreptitiously, as British essayist **Charles Lamb** commented: "The greatest pleasure I know, is to do a good action by stealth, and to have it found out by accident."
- (11) There is **counter-evidence** as well. In some circumstances, diversity seems to enhance cooperation. Eric Uslaner disputes Putnam's thesis, and argues that **diverse communities** can be more cooperative because people living in them are more likely to accept strangers into their "moral community." Clearly more research is required.
- (12) Two people living on opposite sides of the same Norwegian fjord would have spoken **different dialects**. Until recently, and possibly even still, it has been possible to identify the birthplace of native Britons to within 30 miles solely by their **English dialects**.
- (13) Anthropologist **David Nettle** ran an interesting simulation, along similar lines to the Hawk-Dove game. He set up an artificial world where cooperation was necessary for survival, and individuals could choose whom they wished to cooperate with. When he allowed individuals to cooperate only with others who spoke the same dialect, hawks were kept down to a much smaller percentage of the total population than when dialect wasn't a factor. None of this is very surprising; we already know that reciprocity based on proximity is one of the ways cooperation can evolve in a species. Most interestingly, Nettle found that this system of using dialects as societal pressure worked best when they changed rapidly from generation to generation. The simulation mirrored the manner in which these changes occur in life; historically, there are clear differences in human dialects over only a few generations.
- (14) We also try to adopt other cultural norms, to seem less foreign to others. We hand our business cards carefully with two hands to Japanese colleagues, and drink beer with German colleagues even if we prefer wine.
- (15) There's an **alternate analysis** of the Prisoner's Dilemma that bears this out. So far, we've been doing a fairly straightforward analysis of Alice's and Bob's options to determine which ones are better. One can extend that analysis by taking into account the probabilities that Alice and Bob will choose various options. If Alice and Bob are not

complete strangers, they may know something about each other and how the other is likely to proceed. If Alice is in a Prisoner's Dilemma with Bob, and knows that Bob is from the same culture, shares the same religion, and is a member of the same social class, Alice may reasonably anticipate that he will evaluate the situation the same way she does and—in the end—choose whatever option she does. Although she doesn't know Bob's decision beforehand, she knows that she and Bob are enough alike that they will probably choose the same option. Given that assumption, Alice is only choosing between cooperate–cooperate and defect–defect. That's no dilemma at all: cooperate–cooperate is better.

- (16) In general, “is” does not imply “ought.”
- (17) The system isn't entirely symmetrical. Once the john tears the bill in half, it's sunk cost. But the prostitute isn't yet at risk. If she doesn't keep her appointment, she doesn't gain but the john still loses his money.
- (18) There is a whole theory that **costly religious rituals**, such as expensive funerals or Bar Mitzvah parties, are a signaling mechanism to demonstrate a variety of prosocial behaviors.
- (19) The more costly and hard-to-fake the signals are, the more likely they are to be trustworthy. Similarly, the higher the stakes, the more likely signals are to be verified. If you're applying for a job as a surgeon, your résumé is likely to be checked more carefully than if you're applying for a job as a waiter.
- (20) As a side note, Maine lobstermen have a system where they notch a “V” into the tails of **breeding females**. Other lobstermen who catch those notched females are supposed to throw them back in the water. This is a societal dilemma that's primarily solved through morals and reputation; the “V” makes cooperation easier by making the females easier to spot and harder to sell.
- (21) Some examples of proverbs that illustrate this:
- “The gods visit the sins of the fathers upon the children.” —Euripides (c. 485–406 B.C.), Phrixus, fragment 970.
- “For the sins of your fathers you, though guiltless, must suffer.” —Horace, “Odes,” III, 6, l. 1.
- “The Lord is long-suffering, and of great mercy, forgiving iniquity and transgression, and by no means clearing the guilty, visiting the iniquity of the fathers upon the sons to the third and fourth generation.” —Exodus 34:6–7.
- “The sins of the father are to be laid upon the children.” —Shakespeare, *The Merchant of Venice*, Act III, Scene V, l. 1.
- (22) I tend not to trust ticket scalpers outside of stadiums. I'll never see them again, so they have little incentive not to rip me off. It was better when tickets were

hard-to-forge pieces of paper; that was a security system. But now that they're mostly printed receipts of online transactions and verified by bar codes, what's to stop a scalper from reprinting and reselling the same ticket over and over again? I'm essentially buying a sealed bag, and won't know if it's a real ticket or a box of rocks until I get to the gate.

- (23) Reputation doesn't scale down, either. If you're having dinner with your family, no one probably cares how much food you take when. As long as there's trust in this intimate setting, people already know who eats how much and how quickly, and trust that they will get their share eventually. Sometimes this sort of thing happens with close friends or in an intimate business setting, but there's more potential for defection.
- (24) **Edney listed several** reasons why a small group size is more effective: there's better communication within the group, it's easier to see how individuals react to scarcity, it's more difficult for individuals to avoid their responsibilities, there's less alienation, and the role of money is reduced. Edney wrote: "The improved focus on the group itself, the greater ease of monitoring exploitive power, and the opportunities for trust to develop among individuals with face-to-face contact are also enhanced." He doesn't use the terms, but he's talking about moral and reputational pressure.
- (25) **Michel Foucault said** something similar, when he was asked why he participated in student demonstrations when—as a tenured professor—he didn't need to get arrested and beaten up in order to show that he agreed with the student movement. He said: "I consider that it is a cop's job to use physical force. Anyone who opposes cops must not, therefore, let them maintain the hypocrisy of disguising this force behind orders that have to be immediately obeyed. They must carry out what they represent, see it through to the end."

Chapter 9

- (1) Historically, some countries, like England, France, the Netherlands, and the United States, have even sponsored pirates, giving them the designation of "privateers."
- (2) There's a **similar system in Sweden** to combat prostitution: the purchase of sex remains illegal, but the sale of sex has been decriminalized.
- (3) There's a lot more here that I am not going to get into. American prisons are nowhere near the forefront of penological science, and what penologists believe prisons are about isn't the same thing as what corrections officers believe prisons are about; and neither of these two things is what the public thinks prisons are about.
- (4) This is why I am using the word "sanction" instead of "punishment." Punishment implies an expectancy of felt guilt, an emotional satisfaction on the part of the punisher, some sort of existential balance restored. A sanction is a simple *quid pro quo* between the justice system and the accused.

- (5) Moreover, like most drivers, Alice is probably sure her [driving skill is better](#) than average, so she underestimates the risk that her speeding imposes on others.
- (6) I've always thought that the process of getting pulled over, and the wait while the policeman writes the ticket up, is a bigger incentive to obey the speed limit than the fine for a lot of people. The fine is only money, but getting pulled over directly counterbalances the incentive to speed: it results in you getting to your destination more slowly. The inequity of the same fines being assessed to people of all income levels is partly addressed through a points system, whereby states revoke a driver's license if he gets caught speeding too many times.
- (7) Technically, some taxes operate before (airline tickets), some during (road tolls), and some after (capital gains tax). But for our purposes, what matters is not when the money is collected, but that the tax only applies when someone does a particular thing.
- (8) [Electronic filing](#) makes it easier for the IRS to detect some types of fraud because all the data arrives digitally and can be automatically cross-checked.
- (9) This sort of thing has been observed many times. [Students perform better](#) on tests when they're told to try their best than when they're paid for each correct answer. Friends are more likely to help you move if you ask as a favor than if you offer them money. Pizza and beer at the end of the move don't count; that's reciprocal altruism. And [salary bonuses](#) in altruistic jobs can decrease performance. In general, the altruistic portion of a person's brain only works when the thrill center isn't stimulated by the possibility of financial compensation. If you try to stimulate both simultaneously, the thrill center wins.
- (10) [Ostrom's original rules](#) are:
1. The commons must be clearly defined, as must the list of individuals who can use it.
 2. What can be taken out of the commons, and what sort of resources are needed to maintain it, must be suited to local conditions.
 3. Those affected by the rules of the commons need to have a say in how those rules can be modified.
 4. The group charged with monitoring or auditing use of the commons must be accountable to the individuals being monitored.
 5. Individuals who overuse the commons must be assessed graduated penalties, in line with the seriousness of their offense.
 6. Individuals must have access to quick and cheap mechanisms to resolve the inevitable conflicts that come up.
 7. Individuals who use the commons must be able to come up with their own rules for managing it, without those rules being overruled by outside powers.
 8. If the commons is part of a larger system, all of this needs to be nested in multiple layers operating along the same lines.

- (11) [Jeremy Bentham](#) believed that crime could be abolished by using two knobs: making crimes harder to commit, and making punishments more draconian. However, he rightly pointed out that the punishment has to fit the crime. If, for example, both rape and murder are punishable by death, a calculating rapist will kill his victim so as to reduce the chance of his arrest. Similarly, if the fine is the same for driving three miles over the limit as it is for driving thirty miles over, you might as well drive faster—you'll get to your destination sooner, and the punishment for being caught is the same. [Gary Becker](#) expanded on this idea considerably.
- (12) Also note that [increasing the probability](#) of punishment is often cheaper—and more humane—than increasing the severity of punishment.
- (13) There's also [conflicting evidence](#) as to whether or not the probability of getting caught has a strong effect on breaking rules. One study measured how much people cheat on tests, given three different scenarios that changed their likelihood of getting caught. The rate of cheating did not increase with the probability that their cheating would remain undetected.
- (14) The trick with this pair of loopholes is to establish [two Irish subsidiaries](#): one based in a tax haven that holds the rights to its intellectual property outside the U.S., and another based in Ireland that receives the income gained from that property. In order to avoid Irish taxes, a third subsidiary—a Dutch corporation—serves as a transfer for royalties flowing from the subsidiary in Ireland to the tax haven. This byzantine arrangement is legal, even if those three corporations exist on paper only, and allows the parent company to avoid the IRS, even if it is entirely located in the United States.
- (15) That loophole closed after a year, but a [bigger one opened up](#)—and it's retroactive.

Chapter 10

- (1) This might be different in Third World countries. In 2010, [someone was sentenced](#) to three months in jail for stealing two towels from a Nigerian hotel.
- (2) It's not just physical sports. There's [doping in professional Scrabble](#). Some players take “smart drugs” like piracetam and modafinil.
- (3) The reality is much more complicated. While I'm sure that all doctors realize that doping is not in the group interest, as do most athletes, the general public is primarily interested in the spectacle and doesn't really care one way or the other.
- (4) In the 1970s, cyclists used corticosteroids and psychostimulants such as Ritalin, and newly developed norepinephrine-dopamine reuptake inhibitors such as Pemoline. They were banned, and by the end of the decade [assays were developed](#) to detect those substances. In the 1980s, athletes turned to newly developed analogues of endogenous

substances made possible through recombinant DNA technology, including human growth hormone, testosterone, anabolic steroids, and synthetic human erythropoietin (EPO). EPO, a glycoprotein hormone that controls red blood cell production, acts to increase oxygenation, an effect valued as highly by endurance athletes as it was by people suffering from anemia. **EPO use became** rampant in cycling and other sports, and continues to be rampant in spite of bans since the early 1990s and the development in the late 1990s of carbon-isotope ratio tests capable of determining whether substances are made naturally by the body or come from performance-enhancing drugs.

Next came **analogues of analogues**, such as darbepoetin alfa (Aranesp), a variation on the theme of EPO that became commercially available in 2001. It swiftly gained a following among bike racers and other endurance athletes; a test to detect it followed in 2003. A **new EPO replacement**, Mircera, found its way to both the medical and sports markets in 2007; assays to detect it were developed by 2008.

Norbolethone, first developed in 1966, was resurrected in the late 1990s and marketed as the **first designer steroid** by an entrepreneurial bodybuilder-turned-chemist intent on evading detection by the doping police. **Its fingerprint** was traceable by 2002. This **scenario was replayed** with tetrahydrogestrinone and madol, with assays developed within two years of their introduction into sports. The mid-to-late 2000s have seen an increase in blood doping: the use of blood transfusions to increase blood oxygen concentrations. This was soon followed by the development of **flow cytofluorometry** tests to detect it.

The as-yet-unrealized prospect of **gene doping** has led regulatory bodies to preemptively ban any non-therapeutic uses of genetic technology in sports. Presumably tests to detect athletes using them will follow.

- (5) In at least two instances, positive tests for norandrosterone, a steroid of which traces are found naturally in human urine, have been traced to **adulterated supplements** consumed by unsuspecting bicycle racers. Another athlete tested positive for benzodiazepine after consuming a **Chinese herbal product**. The most widely used **urine test for EPO** has been found to result in false positives in urine collected after strenuous physical exercise, though this conclusion has been **hotly contested** by the test's developer and others. **Rapid-screen immunoassays**—the most widely used tests—all too frequently yield false positives in individuals taking routine over-the-counter and prescription pain relievers and allergy, and acid reflux medications. Alpine skier Alain Baxter won the first British medal in Alpine skiing at the 2002 Winter Games in Salt Lake City. Two days after his victory, he was forced to return the bronze medal due to a positive test for methamphetamine resulting from a **Vicks Vapor Inhaler**.
- (6) Counterfeiting is a particularly hard problem, simply because of the economics. Anti-counterfeiting technologies must be cheap to copy in bulk, yet very expensive to copy

individually. To put it in concrete terms, it is certainly worth \$80 for a counterfeiter to make a passable forgery of a \$100 bill. But the government can't spend more than a few dollars on printing the real bills, so any anti-counterfeiting technology has to be inexpensive.

- (7) Studies show that despite knowing how easy it is for a criminal to create or clone a legitimate-looking website, people often use the [appearance of a website](#) as a gauge of credibility. A better way to judge legitimacy is the URL.
- (8) For example, a study on reducing terrorism risks at shopping centers found that the [least costly measure](#) suspicious package reporting, reduced risk by 60%, but the costly and inconvenient searching of bags at entrances achieved only a 15% additional risk reduction. Overall, in fact, the cheapest six security measures reduced risk by 70%, and the remaining 12 more costly security measures reduced risks by only another 25%.

Chapter 11

- (1) On the other hand, he might not steal because of pride. This dialogue appears in [Robert A. Heinlein's *To Sail Beyond the Sunset*](#):

“Thou shalt not steal. I couldn't improve that one, Father.”

“Would you steal to feed a baby?”

“Uh, yes.”

“Think about other exceptions; we'll discuss it in a year or two. But it is a good general rule. But why won't you steal? You're smart; you can probably get away with stealing all your life. Why won't you do it?”

“Uh—”

“Don't grunt.”

“Father, you're infuriating. I don't steal because I'm too stinkin' proud!”

“Exactly! Perfect. For the same reason you don't cheat in school, or cheat in games. Pride. Your own concept of yourself. ‘To thine own self be true, and it must follow, as the night the day—”

“—thou canst not then be false to any man.’ Yes, sir.”

“But you dropped the ‘g’ from the participle. Repeat it and this time pronounce it correctly: You don't steal because—”

“I am too...*stinking*...proud!”

“Good. A proud self-image is the strongest incentive you can have towards correct behavior. Too proud to steal, too proud to cheat, too proud to take candy from babies or to push little ducks into water. Maureen, a moral code for the tribe must be based on survival for the tribe...but for the individual correct behavior in the tightest pinch is based on pride, nor on personal survival. This is why a captain goes down with his ship; this is why ‘The Guard dies but does not surrender.’ A person who has nothing to die for has nothing to live for.”

- (2) Moral philosophers cover similar territory using a different vocabulary. Theologians talk about three [levels of moral meaning](#): the first is personal desire, the second is commitment to social order, and the third is “about the relations among extant order and the relations to past and future orders.” I’m making a gross generalization here, but someone at the first level will choose his self-interest and defect, someone at the second level will choose the long-term group interest and cooperate, and someone at the third level will either cooperate or defect depending on some higher moral principles.
- (3) William C. Crain provides a [good summary](#) of Kohlberg’s six stages:

At stage 1 children think of what is right as that which authority says is right. Doing the right thing is obeying authority and avoiding punishment. At stage 2, children are no longer so impressed by any single authority; they see that there are different sides to any issue. Since everything is relative, one is free to pursue one’s own interests, although it is often useful to make deals and exchange favors with others.

At stages 3 and 4, young people think as members of the conventional society with its values, norms, and expectations. At stage 3, they emphasize being a good person, which basically means having helpful motives toward people close to one. At stage 4, the concern shifts toward obeying laws to maintain society as a whole.

At stages 5 and 6 people are less concerned with maintaining society for its own sake, and more concerned with the principles and values that make for a good society. At stage 5 they emphasize basic rights and the democratic processes that give everyone a say, and at stage 6 they define the principles by which agreement will be most just.

- (4) [Social identity theory](#) has a lot to say about the relative strength of different groups.
- (5) Between 800 and 3,000 people worldwide [immolated themselves](#) in the 40 years between 1963 and 2002 in support of various political and social causes.
- (6) Author and poet [Brian Christian](#) writes this about relative morals:

Thomas Jefferson owned slaves; Aristotle was sexist. Yet we consider them wise? Honorable? Enlightened? But to own slaves in a slave-owning society and to

be sexist in a sexist society are low-entropy personality traits. In a compressed biography of people, we leave those out. But we also tend on the whole to pass *less judgment* on the low-entropy aspects of someone's personality compared to the high-entropy aspects. The *diffs* between them and their society are, one could argue, by and large wise and honorable. Does this suggest, then, a *moral* dimension to compression?

- (7) If you think back to the Prisoner's Dilemma, the police deliberately put the prisoners in that artificial and difficult situation to induce their cooperation. It turns out this is a [useful mechanism](#) for social control.
- (8) The [Stop Snitching campaign](#) can also be explained as a pair of societal dilemmas. The trade-off is between cooperating with society as a whole, and cooperating with the people in the local neighborhood.
- (9) On the other hand, there's a lot less cod in the stores now than there was in the 1970s. And what there is is a lot more expensive.
- (10) [Nepotism is making](#) a comeback in the United States, especially in politics. George W. Bush and Dick Cheney both brought relatives into the federal government while they were in the White House, as did many in their administration. When Republican [Senator Frank Murkowski](#) became governor of Alaska, he appointed his daughter as his Senate replacement. Republican [Representative Richard Pombo](#) might be the worst recent offender in the country; he used his office to funnel money to all sorts of family and friends. Not to pick only on Republicans, Democratic [Representative Eddie Bernice Johnson](#) awarded thousands of dollars in college scholarships to four of her relatives and two of her top aide's children. [Even Bernie Sanders](#) has paid family from campaign donations, and he's a socialist.

It's not all big government, either. One study of Detroit libraries found that [one in six staffers](#) had a relative who also worked in the library system. And [Rupert Murdoch's News Corp.](#) was sued in 2011 by shareholders for nepotism when it bought his daughter's company.

- (11) [Many states have policies](#) about this.

Chapter 12

- (1) [One.Tel in Australia](#) was an example of this. CEO compensation was based on the number of subscribers. As a result, CEOs initiated new-customer campaigns with very cheap contracts—so cheap that the company was losing money on each new subscriber. As a result, the CEOs got their bonuses and One.Tel went bankrupt.

- (2) I am not trying to imply that organizations encourage employee loyalty in order to make them more likely to defect from society as a whole, only that it's one effect of employee loyalty.
- (3) There's another complication. A bishop is not just an employer or supervisor of a priest. In the theological understanding of the church, a bishop is considered to have something of a paternal relationship to a priest. Therefore, the bishop has a responsibility to his priests that a bank supervisor would not have to one of his subordinates. The bishop legitimately is supposed to look out for his priests, especially since his priests have given up all their normal family social connections, and dedicated their lives to the church.
- (4) There was no evidence of a conspiracy, and the Bush Justice Department never followed through with prosecution. Although President Barack Obama had previously praised whistle-blowers as "often the best source of information about waste, fraud, and abuse in government," in April 2010—two and a half years after the original raid—the Obama Justice Department indicted Drake under the Espionage Act, putting him at risk of 35 years' imprisonment on charges of "wilfully retaining" copies of documents he had provided to Congressional investigators. The case was halted on the eve of trial; the government dropped all of the major charges, the financially devastated Drake pleaded guilty to a single misdemeanor, and he was sentenced to community service and a year of probation.
- (5) At the time of writing, Manning has not been convicted of being the source of the WikiLeaks cables, nor has he confessed to the crime.
- (6) Substandard safety by [Massey Energy](#) is a similar example. In 2010, its Upper Big Branch mine exploded and killed 25 people. Sacrificing safety to save money was one of the causes.
- (7) Here's one example, from investment banker Jonathan Knee:

The bankers who pressed these [questionable telecom credits](#) at Morgan in their quest for market share, fees, and internal status coined an acronym that could well be a rallying cry for what the entire investment banking industry had become more broadly. "IBG YBG" stood for "I'll Be Gone, You'll Be Gone." When a particularly troubling fact came up in due diligence on one of these companies, a whispered "IBG YBG" among the banking team members would ensure that a way would be found to do the business, even if investors, or Morgan Stanley itself, would pay the price down the road. Don't sweat it, was the implication, we'll all be long gone by then.

- (8) Famously, Henry Blodget of Merrill-Lynch [described dot.coms](#) as "crap" while at the same time talking them up to investors.

Chapter 13

- (1) There's a theory about which industries will attempt to fix prices in a free-market economy: mature industries where there are only a few major companies that have been lobbying together for a long time. Those companies are likely to have executives who have worked for all the other companies during their careers, and are personally friendly with all the other executives. They are also likely to have former regulators working for them, and former employees as regulators. At this point, there's enough trust amongst them for them to **band together into** a cartel. Another researcher wrote that the two features that are necessary for successful cartels are high seller market sales concentration and product homogeneity. **High barriers to entry** help ensure that a cartel is long-lived.
- (2) The only markets where we have routinely allowed for monopolies are utilities: power, gas, telephone, etc. The idea is that the cost of infrastructure is so high, and the potential for profit is so slim, that market economics will simply drive sellers out of business. Given that, society has given companies monopolies and then heavily regulated them. If technology changes the cost of infrastructure, it makes sense to deregulate those industries.
- (3) I am ignoring any effects from the garment going out of season, or out of style, as it hangs unsold on the rack.
- (4) The same societal dilemma exists in the labor market. Individual sellers—potential employees—are competing for buyers: jobs. And just as competition in the sandwich market results in the cheapest possible sandwiches, competition in the labor market results in the lowest possible wages. But in this case, society recognizes there is an inherent value to higher labor prices. So we allow sellers to organize themselves into cooperative groups: unions.
- (5) Of course, by this I mean the average customer. There will be customers who notice that the sandwiches are worse, and they'll either find it impossible to buy better sandwiches or they'll have to go to special "high quality" sandwich shops for their now-more-expensive sandwiches. Today, we now have to buy organic food, at higher prices, sometimes in high-end grocery stores, to get the same quality of food that was commonly available 50 years ago.
- (6) Calling it "medicine" allowed the company to exploit a loophole in the Prohibition laws.
- (7) Two examples: Rupert Murdoch and his News Corp. founded Fox News; and David and Charles Koch and their immense manufacturing and investment company Koch Industries were among the founders of the Tea Party.

- (8) **There are exceptions.** The Patagonia clothing company is an example of socialist capitalism at its finest.
- (9) There was a big debate in the UK in the 19th century about whether limited companies should be easy to set up, or if an Act of Parliament should be required for each separate company. Much of the debate focused on the fact that companies don't have souls and thus cannot be guilty of treason. It's the same "immortal sociopath" argument.
- (10) Advertising can actually **implant false memories**.
- (11) **The economic term** for this is *lock-in*. Think about your cell phone and cell plan, your computer and operating system, your game console, and so on. It's hard to switch to a competitor, because it involves things like losing months on a subscription service, buying new applications and having to learn how to use them, giving up your already-purchased stock of peripherals, and so on. Industries with low switching costs are very susceptible to changes in reputation. If you drink a Coke today and don't like it, you can easily switch to Pepsi tomorrow. Industries with high switching costs are more robust; if your cell phone company provides lousy service, you're much less likely to switch, because switching is hard and expensive. Raising switching costs is one of the ways corporations artificially limit the effects of a bad reputation on their sales—and another way a modern corporate economy tries to break the fundamental societal dilemmas of a market economy.
- (12) The company, Innovative Marketing, and its CEO James M. Reno, were eventually **able to bargain down** their \$1.8 million judgment to a measly \$17,000 in back taxes and \$100,000 in forfeitures. Given that their scam was alleged to be in the vicinity of \$100 million, they definitely came out ahead.
- (13) In April 2011, a Congressional committee report revealed that between 2005 and 2009, the 14 leading **hydraulic fracturing** companies in the United States used over 2,500 hydraulic fracturing products containing 750 compounds, more than 650 of which were known or possible human carcinogens, substances regulated under the Safe Drinking Water Act, or hazardous air pollutants.
- (14) The **company's arguments** were basically 1) we think it's safe, and 2) those chemicals are trade secrets.
- (15) **The same dynamic** explains why many large projects fail when management adds more people to them.
- (16) There are two basic ways to increase Coase's ceiling. The first is to decrease the cost of internal organizational tasks. The second is to decrease the cost of building a hierarchical organization of organizations. Technology aids in both of those: travel technology to allow people to move around, communications technology to allow better coordination and cooperation, and information technology to allow information

to move around the organization. The fact that all of these technologies have vastly improved in the past few decades is why organizations are growing in size.

- (17) [Senator Bernie Sanders](#) actually had a reasonable point when he said that any company that is too big to fail is also too big to exist.
- (18) The people who use sites like Google and Facebook are not those companies' customers. They are the products that those companies sell to their customers. In general: if you're not paying for it, then you're the product. Sometimes you're the product even if you are paying for it. This isn't new with the Internet. Radio and television programs were traditionally distributed for free, and the audience was the product sold to advertisers. Newspapers are priced far below production costs, with the difference made up by readers being sold to advertisers.
- (19) For example, many large chemical companies use hazardous substances like phosgene, methyl isocyanate, and ethylene oxide in their plants, but don't ship them between locations. They minimize the amounts that are stored as process intermediates. In rare cases of extremely hazardous materials, no significant amounts are stored; instead, they are only present in pipes connecting the reactors that make them with the reactors that consume them.
- (20) For individuals, this is called being judgment-proof, and generally involves minimizing assets. Corporations can achieve the same thing with subsidiaries, so that liability falls on a corporate shell with no assets.

Chapter 14

- (1) And by those no longer in power. Some systems of societal pressures can be hard to get rid of once they're in place.
- (2) This quote, attributed to Louis XIV of France, translates as "The state, it's me." More colloquially, "I am the state." Or in the terms of this book: "As ruler of this country, what is in my interest is necessarily in society's interest."
- (3) In general, terrorism is an [ineffective tactic](#) to advance a political agenda. Political scientist Max Abrams analyzed the political motivations of 28 terrorist groups—the complete list of "foreign terrorist organizations" designated by the U.S. Department of State since 2001. He listed 42 policy objectives of those groups, and found that they only achieved them 7% of the time.
- (4) This isn't to say that we have a good intuition about what level of security is reasonable. A strict cost/benefit analysis of most airline security measures demonstrates that they don't make much sense. But of course, security trade-offs are subjective and have a strong psychological component. There are several aspects of terrorism that cause us

to exaggerate the threat. I'll talk about them in Chapter 15, but basically, we feel less secure than we actually are. So we want more societal pressure than would make strict economic sense.

- (5) If you do the math, more people have died because they chose to drive instead of fly than the terrorists killed on 9/11.
- (6) This **isn't just theoretical**. There is evidence that these considerations affect policy.

Chapter 15

- (1) Of course, there's a lot more to the trade-off of paying taxes than free riding. The tax rates might be so high that it is impossible for someone to survive if he pays his taxes. The taxes might be used to fund an immoral government. And it's possible for the system to collapse even if everyone pays their taxes; the government might allocate the money badly. The former Soviet Union serves as a nice example of this.
- (2) Those of you who have **studied systems dynamics** will recognize this diagram as a combination of two systems archetypes: Fixes that Fail, and Limits to Success.
- (3) Traditional examples of **experiential goods** include vacations, college educations, therapists, and management consulting. This is opposed to something like a desk chair or a can of Coke, where you pretty much know what you're getting before you buy it. Other experiential goods are restaurant dinners, fine art, home improvements, and a move to a new city. Even things that are pretty much commoditized have aspects of experience: a new car, a big-screen television, or a pet gerbil. We know from psychology that people tend to overestimate how much happier they expect a big purchase to make them. Security systems suffer from this same psychological problem; even if people knew exactly how much security a system would give them, they couldn't predict how much safer that additional security would make them feel.
- (4) **Ben Franklin said**: "Those who would give up essential liberty to purchase a little safety deserve neither liberty nor safety."
- (5) It's **also human nature** to not consider, or at least not consider with sufficient weight, the possibility of unintended consequences.
- (6) To take one example, criminals can threaten store owners and steal money from them. Lone criminals generally use guns for this purpose, although they have other ways. Criminal organizations are far more efficient. They can run protection rackets, where they extort money from store owners by threat of violence. They can make far more money this way, often without ever brandishing weapons or even making overt threats. "Nice store you have here" can go a long way if you have a good enough reputation.
- (7) There was a **major political backlash** in the UK against trash monitoring technologies.

- (8) In Europe, [life-cycle management](#) laws are beginning to reduce the amount of trash generated by forcing manufacturers of automobiles to pay for disposal of their products when they are eventually junked.
- (9) [The Innocence Project](#), which works to exonerate convicted felons using DNA evidence, has found that approximately 25% of the 273 people they exonerated in the past 20 years confessed to crimes they didn't commit.
- (10) [Cheating on test scores](#) in response to the No Child Left Behind Act also happened in Chicago, Atlanta, across Pennsylvania, and probably elsewhere in the U.S. as well. [One teacher described](#) the societal pressure to ensure cooperation with the group of teachers: "It's easy to lose your moral compass when you are constantly being bullied."
- (11) In *The Dilbert Principle*, [Scott Adams wrote](#):
- A manager wants to find and fix software bugs more quickly. He offers an incentive plan: \$20 for each bug the Quality Assurance people find and \$20 for each bug the programmers fix. (These are the same programmers who create the bugs.) Result: An underground economy in "bugs" springs up instantly. The plan is rethought after one employee nets \$1,700 the first week.
- (12) It's [18 years](#) if you count from 1994, when banks were first allowed to engage in interstate banking (yes, no banks operated in multiple states before then); 15, if you count from the Fed's relaxation of Glass-Steagall restrictions; 12, if you count from the repeal of Glass-Steagall.
- (13) Not accepting the dilemma as claimed is common among many defectors, including pot smokers, music pirates, and people who count cards at casinos.
- (14) The [potential failure](#) from widespread defection is great. Alexis de Tocqueville said: "The American Republic will endure until the day Congress discovers that it can bribe the public with the public's money."
- (15) I believe that the modern representative democracy is outdated as a political institution. I like to say that it's the best form of government that the mid-18th century could produce. Think about it: because both travel and communications were hard, local groups had to pick one of their own to go all the way to the capital and help make laws in the group's name. Now that travel and communications are easy, there's probably a better system.

Chapter 16

- (1) It would be interesting to chart, as a function of historical time, how much damage an armed group of ten men could do in society before they were subdued. The amount would be pretty stable until the invention of gunpowder, and then would grow continuously

until today. Future advances in chemical, nuclear, and biological weapon capabilities will increase that number even more in the future.

- (2) I don't mean to compare now with ten years ago, or even thirty years ago. I mean to compare it with 100 years ago, 500 years ago, and 1,000 years ago. If you drew a graph, it would be jagged, but over the long term, the rate of technological change has been steadily increasing.

What might be different today is that the rate of change might never again slow down. Not only is the rate of change increasing, but the rate of the rate of change is accelerating as well. Future shock is affecting more of us and more aspects of our lives. The endgame may be [the singularity](#)—which plenty of other people have written and spoken about—but what do we do between now and then? The singularity does answer the question of what comes next for societal pressure. After moral, reputational, institutional, and security pressures comes group mind—technologically-enhanced moral pressure—à la the Borg on *Star Trek*. I don't advocate this as a research direction, but it would give us a huge advantage over the leafcutter ants.

- (3) I don't just mean security against criminals and spies, I also mean security [against the government](#). Over the decades, countries have developed social security systems that prevent law enforcement from abusing the power society delegates to them. In the U.S., these include the warrant process, rules of evidence, search and seizure rules, rules of interrogation, rules prohibiting self-incrimination, and so on. When our communications and writings were on paper, the police would need to demonstrate probable cause and receive a warrant from a judge. Today, our communications and writings are on commercial networks: Facebook, Google Docs, our e-mail providers, and so on. In many cases, the police can simply ask the companies for that data: with no probable cause, without a warrant, and without you even knowing.
- (4) [Clay Shirky writes](#) extensively about these types of organizations.
- (5) The difference is obvious when you look at SafeHouse, a copycat version of WikiLeaks run by the *Wall Street Journal*. Its [terms of service](#) state that SafeHouse “reserve[s] the right to disclose any information about you to law enforcement authorities or to a requesting third party, without notice, in order to comply with any applicable laws and/or requests under legal process....” The *Wall Street Journal* can't do otherwise; the costs of defecting are just too great.
- (6) This is a simplification of something [Lord Kelvin said](#):
- I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of Science, whatever the matter may be.
- (7) Or, as Lord Acton said over 100 years ago: “Power tends to corrupt, and absolute power corrupts absolutely.”

Chapter 17

- (1) In some ways, this is similar to Kierkegaard's [leap of faith](#), the non-logical acceptance of belief required for most religions.
- (2) The [World Values Survey](#) measures impersonal trust in about 70 different countries by asking the question: "Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?" The Scandinavian countries reported the highest level of trust (60% in Norway, Sweden, and Denmark believe most people can be trusted), while countries like Peru, Turkey, Rwanda, and Trinidad and Tobago reported the lowest. The United States ranked towards the higher end. The Gallup World Poll also measures trust by asking three questions: whether it was likely that a neighbor, a stranger, or the police would return to the owner a lost wallet with the money and valuables intact. Again, the [results differ widely](#) by country, and the perceived trustworthiness of neighbors, strangers, and the police differ as well.

All of these surveys collect data on what people say, not what they do. I have not found any study that actually tested these wallet numbers, but *Reader's Digest* [tried something similar](#) with cell phones. Researchers left cell phones unattended in conspicuous places in cities around the world. They then called the phones to see if anyone would answer and return them to their owners. Return rates varied wildly in different cities: Ljubljana won with a 97% return rate, while New York had an 80% return rate, Sydney a 60% return rate, and Singapore a 53% return rate. Hong Kong placed last with a 42% return rate. In a more controlled [laboratory experiment](#) with people from six different world cultures, researchers found significant differences in the level of trust displayed, especially when there was the potential for punishment.

- (3) It's commonly asserted that countries with authoritarian governments have low crime rates: that in these countries, both good and bad defectors are stamped out. And if we want to live in a free country where dissent is tolerated, we necessarily need to tolerate some level of crime. It's a good story, and it may be true, but there's not much in the way of supporting data. The problems are twofold. One: in authoritarian regimes, government-generated data pertaining to crime rates is vulnerable to distortion and manipulation, especially since the regime is motivated to flatter and defend itself. And two: crime statistics provided by authoritarian regimes are likely to be skewed by the absence of figures for crimes condoned or carried out by the state or against marginalized groups. So while rates of reported street crime like muggings, burglaries, and murders are often said to be lower under authoritarian regimes such as the former USSR, former East Germany, and Nazi Germany than in democratic countries, it might be that stamping out dissent doesn't actually make the streets safer. Mussolini didn't make the trains run on time; he just made it illegal to complain about them.

- (4) There's an interesting analogy between protecting against defectors and vaccinating to achieve herd immunity. Society doesn't have to completely fix the problem of defections; it just has to fix it well enough that individuals are not likely to run into the problem. Doing so is much more cost-effective than trying to bring the scope of defection down to zero.
- (5) This quote, widely [attributed to King](#), is actually his paraphrase of an older quote by the abolitionist Theodore Parker from 1853: "I do not pretend to understand the moral universe. The arc is a long one. My eye reaches but little ways. I cannot calculate the curve and complete the figure by experience of sight. I can divine it by conscience. And from what I see I am sure it bends toward justice."



References

Chapter 1

Numbers preceding the references are page numbers.

- 2 *contain parasites* Cory Doctorow (2005), “All Complex Ecosystems Have Parasites,” O’Reilly Emerging Technology Conference, San Diego, California. Christopher Langton, ed. (1994), *Artificial Life III*, Westview Press.
- 3 *their own privacy* Bruce Schneier (15 Jul 2009), “Facebook Should Compete On Privacy, Not Hide It Away,” *The Guardian*.
- 4 *effectively looted* Jeff Gottlieb (22 Jul 2010), “Bell Council Found Loophole in Law to Allow Big Salaries,” *Los Angeles Times*. Jeff Gottlieb and Ruben Vives (8 Aug 2010), “Bell Councilman ‘Ashamed,’ ‘Disgusted’ that Rizzo Earned \$1.5 Million,” *Los Angeles Times*.
- 4 *internal cheating* Greg Hoggund and Gary McGraw (2007), *Exploiting Online Games; Cheating Massively Distributed Systems*, Addison-Wesley Professional.
- 5 *trust is a bet* Piotr Sztompka (1999), *Trust: A Sociological Theory*, Cambridge University Press, 25.
- 5 *Trust involves* Russell Hardin (1992), “The Street-Level Epistemology of Trust,” *Analyse & Kritik*, 14:152–76.
- 5 *impersonal trust* Susan P. Shapiro (1987), “The Social Control of Impersonal Trust,” *American Journal of Sociology*, 93:623–58.
- 5 *described trust* Don Tapscott and David Ticoll (2003), *The Naked Corporation: How the Age of Transparency Will Revolutionize Business*, Free Press.
- 6 *Trust is the expectation* Francis Fukuyama (1995), *Trust: The Social Virtues and the Creation of Prosperity*, Simon & Schuster, 26.
- 6 *three critical functions* Barbara Misztal (1996), *Trust in Modern Societies: The Search for the Bases of Social Order*, Cambridge Polity Press.

- 6 **recent example** David Remnick (1993), *Lenin's Tomb: The Last Days of the Soviet Empire*, Random House.
- 7 **value of trust** Francis Fukuyama (1995), *Trust: The Social Virtues and the Creation of Prosperity*, Simon & Schuster. Barbara Misztal (1996), *Trust in Modern Societies: The Search for the Bases of Social Order*, Cambridge Polity Press. Adam B. Seligman (1997), *The Problem of Trust*, Princeton University Press. Piotr Sztompka (1999), *Trust: A Sociological Theory*, Cambridge University Press. Steven M.R. Covey (2006), *The Speed of Trust: The One Thing That Changes Everything*, Free Press.
- 7 **trust is the atmosphere** Sissela Bok (1978), *Lying: Moral Choice in Private and Public Life*, Pantheon Books, 31.

Chapter 1 Notes

- 249 **five levels of trust** Steven M.R. Covey (2006), *The Speed of Trust: The One Thing That Changes Everything*, Free Press.
- 249 **To trust is to bet** Piero Ferrucci (2006), *The Power of Kindness: The Unexpected Benefits of Leading a Compassionate Life*, Tarcher, 80–1.
- 249 **When we say we trust** Diego Gambetta (1988), “Can We Trust Trust?” in Diego Gambetta, ed., *Trust: Making and Breaking Cooperative Relations*, Blackwell Publishers, 213–37.
- 249 **We will define trust** David M. Messick and Roderick M. Kramer (2001), “Trust as a Form of Shallow Morality,” in Karen S. Cook, ed., *Trust in Society*, Russell Sage Foundation, 89–117.
- 249 **three-level progression** Anthony Giddens (1990), *The Consequences of Modernity*, Stanford University Press, 114–24.
- 250 **Piotr Cofta covered** Piotr Cofta (2007), *Trust, Complexity, and Control: Confidence in a Convergent World*, John Wiley & Sons.
- 250 **findings are false** John P.A. Ioannidis (2005), “Why Most Published Research Findings Are False,” *PLoS Medicine*, 2(8): e124.
- 250 **sloppy methodology** Eric-Jan Wagenmakers, Rudd Wetzels, Denny Borsboom, and Han van der Maas (2011), “Why Psychologists Must Change the Way They Analyze Their Data: The Case of Psi,” *Journal of Personality and Social Psychology*, 100:426–32. Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn (2011), “False Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant,” *Psychological Science*, 22:1359–66.
- 250 **narrow and unrepresentative** Joseph Henrich, Steven J. Heine, and Ara Norenzayan (2010), “The Weirdest People in the World?” *Behavioral & Brain Sciences*, 33:61–83.
- 250 **Adam Smith** Adam Smith (1759), *The Theory of Moral Sentiments*.

Chapter 2

- 17 *territorial chorus* Tim Clutton-Brock (2009), "Cooperation between Non-Kin in Animal Societies," *Nature*, 462:51–7.
- 17 *hornworms* André Kessler, Rayko Halitschke, Celia Diezel, and Ian T. Baldwin (2006), "Priming of Plant Defense Responses in Nature by Airborne Signaling between *Artemisia tridentata* and *Nicotiana attenuate*," *Oecologia*, 148:280–92.
- 17 *plasmids secrete* Michael B. Yarmolinsky (1995), "Programmed Cell Death in Bacterial Populations," *Science*, 267:836–7.
- 17 *Heat and light* Henry Lutz Ehrlich and Dianne K. Newman (2009), *Geomicrobiology*, CRC Press. Smithsonian Institution (2009), "The Archean: The First Life on Earth," in *Geologic Time: The Story of a Changing Earth, Department of Paleobiology*, National Museum of Natural History, Smithsonian Institution.
- 18 *first animal predator* Ben Harder (9 Apr 2002), "Was This the Earth's First Predator?" *National Geographic*.
- 18 *The Selfish Gene* Richard Dawkins (1976), *The Selfish Gene*, Oxford University Press.
- 19 *defend against* William D. Hamilton, Robert Axelrod, and Reiko Tanese (1990), "Sexual Reproduction as an Adaptation to Resist Parasites (A Review)," *Proceedings of the National Academy of Sciences of the United States of America*, 97:3566–73. Levi T. Morran, Olivia G. Schmidt, Ian A. Gelarden, Raymond C. Parrish II, and Curtis M. Lively (2011), "Running with the Red Queen: Host-Parasite Coevolution Selects for Biparental Sex," *Science*, 333:216–18.
- 19 *jawed fish* Charles A. Janeway (2006), *Immunobiology: The Immune System in Health and Disease*, Sixth Edition, Taylor & Francis Group.
- 19 *orienting response* Ivan P. Pavlov (1927), *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*, tr. G.V. Anrep, Oxford University Press. Evgeni Nikolaevich Sokolov (1963), "Higher Nervous Functions: The Orienting Reflex," *Annual Review of Physiology*, 25:545–80.
- 19 *particularly good* Joshua New, Leda Cosmides, and John Tooby (2007), "Category-Specific Attention for Animals Reflects Ancestral Priorities, Not Expertise," *Proceedings of the National Academy of Sciences of the United States of America*, 104:16598–603. Lynn Isabell (2009), *The Fruit, the Tree, and the Serpent: Why We See So Well*, Harvard University Press.
- 19 *throw things* Justin N. Wood, David D. Glynn, and Marc D. Hauser (2007), "The Uniquely Human Capacity to Throw Evolved from a Non-Throwing Primate: An Evolutionary Dissociation between Action and Perception," *Biology Letters*, 3:360–5.
- 19 *size-weight misperception* Qin Zhu and Geoffrey P. Bingham (2011), "Human Readiness to Throw: The Size-Weight Illusion Is Not an Illusion When Picking the Best Objects to Throw," *Evolution & Human Behavior*, 32:288–93.
- 19 *Similar stories* Stephen Jay Gould (Oct 1985), "Not Necessarily a Wing," *Natural History*, 94:12–25.

- 20 *weird security* Randolph M. Nesse (2001), "The Smoke Detector Principle: Natural Selection and the Regulation of Defensive Responses," *Annals of the New York Academy of Sciences*, 935:75–85.
- 21 *Red Queen Effect* Leigh Van Valen (1973), "A New Evolutionary Law," *Evolutionary Theory*, 1:1–30. Leigh Van Valen (1977), "The Red Queen," *The American Naturalist*, 111:809–10. Matt Ridley (1994), *The Red Queen: Sex and the Evolution of Human Nature*, MacMillan Publishing Co.
- 21 *continuously improve* Seth Finnegan, Jonathan L. Payne, and Steve C. Wang (2008), "The Red Queen Revisited: Reevaluating the Age Selectivity of Phanerozoic Marine Genus Extinctions," *Paleobiology*, 34:318–41. Michael J. Benton (2009), "The Red Queen and the Court Jester: Species Diversity and the Role of Biotic and Abiotic Factors Through Time," *Science*, 323:728–32. Michael J. Benton (2010), "Evolutionary Biology: New Take On the Red Queen," *Nature*, 463:306–7.
- 22 *other animals* Edward A. Wasserman, Thomas R. Zentall (2006), *Comparative Cognition: Experimental Explorations of Animal Intelligence*, Oxford University Press. Zhanna Reznikova (2007), *Animal Intelligence: From Individual to Social Cognition*, Cambridge University Press. Jeremy Taylor (2009), *Not a Chimpanzee: The Hunt to Find the Genes That Make Us Human*, Oxford University Press.
- 22 *supply our brains* Adam D. Pfefferle, Lisa R. Warner, Catrina W. Wang, William J. Nielsen, Courtney C. Babbitt, Olivier Fedrigo, and Gregory A. Wray (2011), "Comparative Expression Analysis of the Phosphocreatine Circuit in Extant Primates: Implications for Human Brain Evolution," *Journal of Human Evolution*, 60:205–12.
- 22 *cultural evolution* Gregory Cochran and Henry Harpending (2009), *The 10,000 Year Explosion: How Civilization Accelerated Human Evolution*, Basic Books.
- 22 *Nicholas Humphrey* Nicholas Humphrey (1976), "The Social Function of Intellect," in Paul Patrick Gordon Bateson and Robert A. Hinde, eds., *Growing Points in Ethology*, Cambridge University Press, 303–17.
- 23 *Daniel Gilbert* Daniel Gilbert (2 Jul 2006), "If Only Gay Sex Caused Global Warming," *Los Angeles Times*.
- 23 *capacity for deception* Richard W. Byrne and Nadia Corp (2004), "Neocortex Size Predicts Deception Rate in Primates," *Proceedings of the Royal Society B: Biological Sciences*, 271:1693–9.
- 23 *non-primate mammals* James K. Rilling and Thomas R. Insel (1999), "The Primate Neocortex in Comparative Perspective Using Magnetic Resonance Imaging," *Journal of Human Evolution*, 37:191–223.
- 24 *neocortex correlates* Robin I.M. Dunbar (1992), "Neocortex Size as a Constraint on Group Size in Primates," *Journal of Human Evolution*, 20:469–93.
- 24 *mean human group* Robin I.M. Dunbar (2003), "The Social Brain: Mind, Language, and Society in Evolutionary Perspective," *Annual Review of Anthropology*, 32:163–81. Alberto Hernando, Diego Villuendas, Cristina Vesperinas, Marta Abad, and Ángel Plastino (2010), "Unravelling the Size Distribution of Social Groups with Information Theory on Complex Networks," *The European Physical Journal B: Condensed Matter & Complex Systems*, 76:87–97.

- 24 *number appears regularly* R.A. Hill and Robin I.M. Dunbar (2003), "Social Network Size in Humans," *Human Nature*, 14:53–72.
- 24 *mean group size* Robin I.M. Dunbar (2003), "The Social Brain: Mind, Language, and Society in Evolutionary Perspective," *Annual Review of Anthropology*, 32:163–81.
- 24 *died in warfare* Steven A. LeBlanc and Katherine E. Register (2003), *Constant Battles: Why We Fight*, St. Martin's Press. Samuel Bowles (2009), "Did Warfare Among Ancestral Hunter-Gatherers Affect the Evolution of Human Social Behaviors?" *Science*, 324:1293–8.
- 24 *Paul Seabright* Paul Seabright (2004), *The Company of Strangers: A Natural History of Economic Life*, Princeton University Press.
- 24 *other species* J. Maynard Smith and George R. Price (1973), "The Logic of Animal Conflict," *Nature*, 246:15–8.
- 24 *should primarily hunt* Craig B. Stanford (2001), "A Comparison of Social Meat-Foraging by Chimpanzees and Human Foragers," in Craig B. Stanford and H. Bunn, eds., *Meat-Eating and Human Evolution*, Oxford University Press. Gottfried Hohmann (2009), "The Diets of Non-Human Primates: Frugivory, Food Processing, and Food Sharing," in Jean-Jacques Hublin, Michael P. Richards, eds. (2009), *The Evolution of Hominin Diets: Integrating Approaches to the Study of Palaeolithic Subsistence*, Springer.
- 24 *all primitive societies* Steven A. LeBlanc and Katherine E. Register (2003), *Constant Battles: Why We Fight*, St. Martin's Press.
- 25 *David Buss* David M. Buss (2006), *The Murderer Next Door: Why the Mind Is Designed to Kill*, Penguin, 40.
- 25 *quite violent* Steven A. LeBlanc and Katherine E. Register (2003), *Constant Battles: Why We Fight*, St. Martin's Press. David M. Buss (2006), *The Murderer Next Door: Why the Mind Is Designed to Kill*, Penguin. Bureau of Justice Statistics (1994), "Violent Crime," U.S. Department of Justice.
- 25 *some argue* Steven Pinker (2011), *The Better Angels of Our Nature: Why Violence Has Declined*, Viking.
- 25 *kill in war* Dave Grossman (1995), *On Killing: The Psychological Cost of Learning to Kill in War and Society*, Little, Brown & Co.
- 25 *with other primates* Lars Rodseth, Richard W. Wrangham, Alisa M. Harrigan, and Barbara B. Smuts (1991), "The Human Community as a Primate Society," *Current Anthropology*, 32:221–54. Bruce M. Knauft (1991), "Violence and Sociality in Human Evolution," *Current Anthropology*, 32:391–428. Christoph P.E. Zollikofer, Marcia S. Ponce de Leon, Bernard Vandermeersch, and Francois Leveque (2002), "Evidence for Interpersonal Violence in the St. Cesaire Neanderthal," *Proceedings of the National Academy of Sciences of the United States of America*, 99:6444–8. Margaret C. Crofoot and Richard W. Wrangham (2010), "Intergroup Aggression in Primates and Humans: The Case for a Unified Theory," in Peter M. Kappeler and Joan M. Silk, eds., *Mind the Gap: Tracing the Origins of Human Universals*, Springer.

Chapter 2 Notes

- 250 *sticks as weapons* Jill D. Pruetz and Paco Bertolani (2007), "Savanna Chimpanzees, *Pan troglodytes verus*, Hunt with Tools," *Current Biology*, 17:412–17.
- 250 *using rocks* Giacomo Bernardi (2011), "The Use of Tools by Wrasses (*Labridae*)," *Coral Reefs*, published online before print.
- 250 *Brants' whistling rat* Tim P. Jackson (2000), "Adaptation to Living in an Open Arid Environment: Lessons from the Burrow Structure of the Two Southern African Whistling Rats, *Parotomys brantsii* and *P. littedalei*," *Journal of Arid Environments*, 46:345–55.
- 250 *African rat* Jonathan Kingdon, Bernard Agwanda, Margaret Kinnaird, Timothy O'Brien, Christopher Holland, Thomas Gheysens, Maxime Boulet-Audet, and Fritz Vollrath (2011), "A Poisonous Surprise under the Coat of the African Crested Rat," *Proceedings of the Royal Society B: Biological Sciences*, published online before print.
- 251 *entirely separate* R. Jansen, J.D. Embden, W. Gaastra, L.M. Schouls (2002), "Identification of Genes That Are Associated with DNA Repeats in Prokaryotes," *Molecular Microbiology*, 43:1565–75.
- 251 *earlier book* Bruce Schneier (2006), *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus.
- 251 *picture of a snake* Vanessa LoBue and Judy DeLoache (2008), "Detecting the Snake in the Grass: Attention to Fear-Relevant Stimuli by Adults and Young Children," *Psychological Science*, 19:284–9.
- 251 *cause and effect* Gema Martin-Ordas, Josep Call, and Fernando Colmenares (2007), "Tubes, Tables, and Traps: Great Apes Solve Two Functionally Equivalent Trap Tasks but Show No Evidence of Transfer Across Tasks," *Animal Cognition*, 11:423–30.
- 251 *time binding species* Alfred Korzybski (1933), *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*, International Non-Aristotelian Library Publishing Corp.
- 251 *intentional stance* Daniel C. Dennett (1987), *The Intentional Stance*, MIT Press.
- 251 *social group size* Kimberly A. Pollard and Daniel T. Blumstein (2011), "Social Group Size Predicts the Evolution of Individuality," *Current Biology*, 21:413–7.
- 251 *win arguments* Hugo Mercier and Dan Sperber (2011), "Why Do Humans Reason? Arguments for an Argumentative Theory," *Behavioral & Brain Sciences*, 34:54–74.
- 251 *Others posit* Christopher McCarty, Peter D. Killworth, H. Russell Bernard, Eugene C. Johnsen, and Gene A. Shelley (2000), "Comparing Two Methods for Estimating Network Size," *Human Organization*, 60:28–39. H. Russell Bernard, Gene Ann Shelley, and Peter Killworth (1987), "How Much of a Network Does the GSS and RSW Dredge Up?" *Social Networks*, 9:49–63.
- 252 *primitive peoples* Mervyn J. Meggitt (1977), *Blood Is Their Argument*, Mayfield.
- 252 *Yanomamö* Napoleon A. Chagnon (1974), *Studying the Yanomamo*, Holt, Rinehart & Winston.

- 252 *Big-game hunting* David A. Byers and Andrew Ugan (2005), "Should We Expect Large Game Specialization in the Late Pleistocene? An Optimal Foraging Perspective on Early Paleoindian Prey Choice," *Journal of Archaeological Science*, 32:1624–40. Kristen Hawkes, James F. O'Connell, and Nicholas G. Blurton-Jones (1991), "Hunting Income Patterns among the Hadza: Big Game, Common Goods, Foraging Goals and the Evolution of the Human Diet," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 33:243–50. John D. Speth, Khori Newlander, Andrew A. White, Ashley K. Lemke, and Lars E. Anderson (2010), "Early Paleoindian Big-Game Hunting in North America: Provisioning or Politics?" *Quaternary International*, published online before print. Mary C. Stiner, Natalie D. Munro, and Todd A. Surovell (2000), "The Tortoise and the Hare: Small-Game Use, the Broad-Spectrum Revolution, and Paleolithic Demography," *Current Anthropology*, 41: 39–73.
- 252 *Chimpanzees' aggression rates* Richard W. Wrangham, Michael L. Wilson, and Martin N. Muller (2006), "Comparative Rates of Violence in Chimpanzees and Humans," *Primates*, 47:14–26.

Chapter 3

- 27 *division of labor* Edward O. Wilson (1987), "Causes of Ecological Success: The Case of the Ants," *Journal of Animal Ecology*, 56:1–9. Bert Holldobler and Edward O. Wilson (2009), *The Superorganism: The Beauty, Elegance, and Strangeness of Insect Societies*, W.W. Norton & Co. Bert Holldobler and Edward O. Wilson (2010), *The Leafcutter Ants: Civilization by Instinct*, W.W. Norton & Co.
- 28 *Hawk-Dove game* John Maynard Smith and George R. Price (1973), "The Logic of Animal Conflict," *Nature*, 246:15–8.
- 31 *war of all against all* Thomas Hobbes (1651), *Leviathan*, Printed for Andrew Crooke, at the Green Dragon in St. Paul's Churchyard.
- 33 *compassion extended* Penny Spikins, Holly Rutherford, and Andy Needham (2010), "From Homininity to Humanity: Compassion from the Earliest Archaic to Modern Humans," *Time and Mind*, 3:303–25. Priyali Rajagopal and Nicole Votolato Montgomery (2011), "I Imagine, I Experience, I Like: The False Experience Effect," *The Journal of Consumer Research*, 38:578–94.
- 33 *invention of agriculture* Robert Boyd and Peter Richerson (2004), *Not by Genes Alone: How Culture Transformed Human Evolution*, University of Chicago Press.
- 33 *Unrelated elephants* Joshua M. Plotnik, Richard Lair, Wirot Suphachoksahakun, and Frans de Waal (2011), "Elephants Know When They Need a Helping Trunk in a Cooperative Task," *Proceedings of the National Academy of Sciences of the United States of America*, published online before print.
- 33 *ritualized battles* Mark Briffa (2010), "Territoriality and Aggression," *Nature Education Knowledge*, 1:19. Craig Packer and Anne E. Pusey (1985), "Asymmetric Contests in Social Mammals Respect, Manipulation and Age-Specific Aspects," in Paul J.

- Greenwood, Paul H. Harvey, and Montgomery Slatkin, eds., *Evolution: Essays in Honour of John Maynard Smith*, Cambridge University Press, 173–86. Tabitha M. Innocent and Stuart A. West (2006), “Social Evolution: Cooperation by Conflict,” *Current Biology*, 16:365–67. Hanna Kokko, Andrés López-Sepulcre, and Lesley J. Morrell (2006), “From Hawks and Doves to Self-Consistent Games of Territorial Behavior,” *The American Naturalist*, 167:901–12.
- 34 **Vampire bats** Gerald S. Wilkinson (1984), “Reciprocal Food Sharing in the Vampire Bat,” *Nature*, 308:181–4.
- 34 **large frontal lobes** Kunwar P. Bhatnagar (2008), “The Brain of the Common Vampire Bat, *Desmodus rotundus murinus* (Wagner, 1840): A Cytoarchitectural Atlas,” *Brazilian Journal of Biology*, 68:583–99.
- 34 **help each other** Robert L. Trivers (1971), “The Evolution of Reciprocal Altruism,” *Quarterly Review of Biology*, 46:35–7.
- 35 **transactional analysis** Eric Berne (1996), “Principles of Transactional Analysis,” *Indian Journal of Psychiatry*, 38:154–9.
- 35 **potential mate** Amotz Zahavi (1975), “Mate Selection: A Selection for a Handicap,” *Journal of Theoretical Biology*, 53:205–14. Amotz Zahavi (1977), “The Cost of Honesty (Further Remarks on the Handicap Principle),” *Journal of Theoretical Biology*, 67:603–5. Amotz Zahavi (1977), “Reliability in Communication Systems and the Evolution of Altruism,” in Bernard Stonehouse and Christopher M. Perrins, eds., *Evolutionary Ecology*, Macmillan, 253–9. Rufus A. Johnstone (1995), “Sexual Selection, Honest Advertisement and the Handicap Principle: Reviewing the Evidence,” *Biological Reviews*, 70:1–65. Amotz Zahavi and Avishag Zahavi (1997), *The Handicap Principle: A Missing Piece of Darwin’s Puzzle*, Oxford University Press. Amotz Zahavi (2003), “Indirect Selection and Individual Selection in Sociobiology: My Personal Views on Theories of Social Behaviour,” *Animal Behaviour*, 65:859–63.
- 35 **altruistic acts** Randolph M. Nesse (2007), “Runaway Social Selection for Displays of Partner Value and Altruism,” *Biological Theory*, 2:143–55.
- 35 **kind people** Charlie L. Hardy and Mark Van Vugt (2006), “Nice Guys Finish First: The Competitive Altruism Hypothesis,” *Personality & Social Psychology Bulletin*, 32:1402–13. Lauri A. Jensen-Campbell, Jennifer M. Knack, and Haylie L. Gomez (2010), “The Psychology of Nice People,” *Social & Personality Psychology Compass*, 4:1042–56. Pat Barclay (2010), “Altruism as a Courtship Display: Some Effects of Third-Party Generosity on Audience Perceptions,” *British Journal of Psychology*, 101:123–35. Timothy A. Judge, Beth A. Livingston, and Charlice Hurst (2011), “Do Nice Guys and Gals Really Finish Last? The Joint Effects of Sex and Agreeableness on Income,” *Journal of Personality & Social Psychology*, in press.
- 35 **George Price** Oren S. Harman (2010), *The Price of Altruism: George Price and the Search for the Origins of Kindness*, W.W. Norton & Co.
- 35 **bargaining games** Gary E. Bolton (1998), “Bargaining and Dilemma Games: From Laboratory Data Towards Theoretical Synthesis,” *Experimental Economics*, 1:257–81.

- 36 *found a coin* Paula F. Levin and Alice M. Isen (1972), "The Effect of Feeling Good on Helping: Cookies and Kindness," *Journal of Personality & Social Psychology*, 21:384–8.
- 36 *flying through clouds* Lawrence J. Sanna, Edward C. Chang, Paul M. Miceli, and Kristjen B. Lundberg (2011), "Rising Up to Higher Virtues: Experiencing Elevated Physical Heights Uplifts Prosocial Actions," *Journal of Experimental & Social Psychology*, 47:472–6.
- 36 *mirror neurons* Giuseppe di Pellegrino, Luciano Fadiga, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzolatti (1992), "Understanding Motor Events: A Neurophysiological Study," *Experimental Brain Research*, 91:176–80. Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti (1996), "Action Recognition in the Premotor Cortex," *Brain*, 119:593–609. Vittorio Gallese and Alvin Goldman (1998), "Mirror Neurons and the Simulation Theory of Mind-Reading," *Trends in Cognitive Sciences*, 2:493–501.
- 36 *altruistic innately* Ernst Fehr and Simon Gächter (2002), "Altruistic Punishment in Humans," *Nature*, 415:137–40. Alan G. Sanfey, James K. Rilling, Jessica A. Aronson, Leigh E. Nystrom, and Jonathan D. Cohen (2003), "The Neural Basis of Economic Decision-Making in the Ultimatum Game," *Science*, 300:1755–8. Tania Singer, Ben Seymour, John P. O'Doherty, Klass E. Stephan, Raymond J. Dolan, and Chris D. Frith (2006), "Empathic Neural Responses Are Modulated by the Perceived Fairness of Others," *Nature*, 439:466–9. Molly J. Crockett, Luke Clark, Golnaz Trabibnia, Matthew D. Lieberman, and Trevor W. Robbins (2008), "Serotonin Modulates Behavioral Reactions to Unfairness," *Science*, 320:1739. Molly J. Crockett, Luke Clark, Marc D. Hauser, and Trevor W. Robbins (2010), "Serotonin Selectively Influences Moral Judgment and Behavior Through Effects on Harm Aversion," *Proceedings of the National Academy of Sciences of the United States of America*, 107:17433–8.
- 37 *fairness and justice* Katarina Gospic, Erik Mohlin, Peter Fransso, Predrag Petrovic, Magnus Johannesson, and Martin Ingvar (2011), "Limbic Justice: Amygdala Involvement in Immediate Rejection in the Ultimatum Game," *PLoS Biology*, 9 (5): e1001054.
- 37 *and attractive* Tania Singer, Ben Seymour, John P. O'Doherty, Klass E. Stephan, Raymond J. Dolan, and Chris D. Frith (2006), "Empathic Neural Responses Are Modulated by the Perceived Fairness of Others," *Nature*, 439:466–9.
- 37 *found that oxytocin* Paul J. Zak, Robert Kurzban, and William T. Matzner (2003), "Oxytocin Is Associated with Interpersonal Trust in Humans," *PLoS ONE*, 2:11, e1128. Zoe R. Donaldson and Larry J. Young (2008), "Oxytocin, Vasopressin, and the Neurogenetics of Sociality," *Science*, 322:900–4.
- 37 *someone's oxytocin* Michael Kosfield, Marcus Heinefichs, Paul J. Zak, Urs Fischbacher, and Ernst Fehr (2003), "Oxytocin Increases Trust in Humans," *Nature*, 435:673–6.
- 37 *Adam Smith* Adam Smith (1759), *The Theory of Moral Sentiments*.
- 38 *non-kin groups* Martin A. Nowak (2006), "Five Rules for the Evolution of Cooperation," *Science*, 314:1560–3. Martin A. Nowak and Roger Highfield (2011),

SuperCooperators: Altruism, Evolution, and Why We Need Each Other to Succeed, Free Press.

- 38 *Indirect reciprocity* Martin A. Nowak and Karl Sigmund (1998), “Evolution of Indirect Reciprocity by Image Scoring,” *Nature*, 393:573–7. Martin A. Nowak and Karl Sigmund (1998), “The Dynamics of Indirect Reciprocity,” *Journal of Theoretical Biology*, 194:561–74. Manfred Milinski, Dirk Semmann, Theo C.M. Bakker, and Hans-Jürgen Krambeck (2001), “Cooperation Through Indirect Reciprocity: Image Scoring or Standing Strategy?” *Proceedings of the Royal Society B: Biological Sciences*, 268:2495–501.
- 38 *biological models* Herbert Gintis (2003), “The Hitch-Hikers Guide to Altruism: Genes, Culture, and the Internalization of Norms,” *Journal of Theoretical Biology*, 220:407–18.
- 38 *Patricia Churchland* Patricia S. Churchland (2011), *Braintrust: What Neuroscience Tells Us About Morality*, Princeton University Press, 93.
- 38 *our murderousness* Samuel Bowles (2006), “Group Competition, Reproductive Leveling, and the Evolution of Human Altruism,” *Science*, 314:1569–72. Jung-Kyoo Choi and Samuel Bowles (2007), “The Coevolution of Parochial Altruism and War,” *Science*, 318:636–40.
- 39 *And these processes* Robert Boyd and Peter J. Richerson (1985), *Culture and the Evolutionary Process*, University of Chicago Press. Robert Bettinger, Peter J. Richerson, and Robert Boyd (1995), “Can Group Functional Behaviors Evolve by Cultural Group Selection?” *Current Anthropology*, 36:473–94.

Chapter 3 Notes

- 252 *increased specialization* Raphaël Jeanson, Jennifer H. Fewell, Root Gorelick, and Susan M. Bertram (2007), “Emergence of Increased Division of Labor as a Function of Group Size,” *Behavioral Ecology & Sociobiology*, 62:289–98.
- 252 *not what makes* Anna Dornhaus (2008), “Specialization Does Not Predict Individual Efficiency in an Ant,” *PLOS Biology*, 6:e285.
- 252 *startling statistic* Dwayne C. Savage (1977), “Microbial Ecology of the Gastrointestinal Tract,” *Annual Review of Microbiology*, 31:107–33.
- 253 *some mixture* Stephen Le and Robert Boyd (2007), “Evolutionary Dynamics of the Continuous Iterated Prisoner’s Dilemma,” *Journal of Theoretical Biology*, 245:258–67.
- 253 *population density* Pierre Auger and Dominique Pontier (1998), “Fast Game Theory Coupled to Slow Population Dynamics: The Case of Domestic Cat Populations,” *Mathematical Biosciences*, 148:65–82.
- 253 *South African meerkats* Tim H. Clutton-Brock, Peter N.M. Brotherton, M. Justin O’Riain, Ashleigh S. Griffin, David Gaynor, Ruth Kansky, Lynda Sharpe, and Grant M. McIlrath (2001), “Contributions to Cooperative Rearing in Meerkats,” *Animal Behaviour*, 61:705–10. Alex Thornton (2008), “Social Learning about Novel Foods in Young Meerkats,” *Animal Behaviour*, 76:1411–21.

- 253 *Red ruffed lemurs* Natalie Vasey (2007), "The Breeding System of Wild Red Ruffed Lemurs (*Varecia rubra*): A Preliminary Report," *Primates*, 48:41–54.
- 253 *much less common* Lee Alan Dugatkin (1997), *Cooperation Among Animals: An Evolutionary Perspective*, Oxford University Press.
- 253 *also slow* Carl T. Bergstrom and Michael Lachmann (2003), "The Red King Effect: Evolutionary Rates and the Division of Surpluses in Mutualisms," in Peter Hammerstein, ed., *Genetic and Cultural Evolution of Cooperation*, MIT Press, 223–8.
- 253 *Wrasse cleaner fish* Frans de Waal (Apr 2005), "How Animals Do Business," *Scientific American*, 72–79.
- 253 *Elephant seals* Burney J. Le Boeuf (1974), "Male-Male Competition and Reproductive Success in Elephant Seals," *American Zoologist*, 14:163–76.
- 254 *some bullfrogs* Richard D. Howard (1978), "The Evolution of Mating Strategies in Bullfrogs, *Rana catesbeiana*," *Evolution*, 32:850–71.
- 254 *everyone is a dove* Valerius Geist (1974), "On the Relationship of Social Evolution and Ecology in Ungulates," *American Zoologist*, 14:205–20.
- 254 *there is a mixture* John Maynard Smith and Geoffrey A. Parker (1976), "The Logic of Asymmetric Contests," *Animal Behaviour*, 24:159–75. Peter Hammerstein (1981), "The Role of Asymmetries in Animal Contests," *Animal Behaviour*, 29:193–205.
- 254 *Kaushik Basu* Kaushik Basu (1984), *The Less-Developed Economy: A Critique of Contemporary Theory*, Oxford University Press, 5–6.
- 254 *Neuroscience is starting* Kerri Smith (2011), "Neuroscience vs Philosophy: Taking Aim at Free Will," *Nature*, 477:23–5.
- 254 *Ultimatum game* Charles A. Holt (2000), "Y2K Bibliography of Experimental Economics and Social Science: Ultimatum Game Experiments," University of Virginia. Hessel Oosterbeek, Randolph Sloof, and Gijs van de Kuilen (2004), "Cultural Differences in Ultimatum Game Experiments: Evidence From a Meta-Analysis," *Experimental Economics*, 7:171–88.
- 254 *how the game works* Werner Güth, Rolf Schmittberger, and Bernd Schwarze (1982), "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior & Organization*, 3:267–88.
- 255 *turn down offers* Hessel Oosterbeek, Randolph Sloof, and Gijs van de Kuilen (2004), "Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis," *Experimental Economics*, 7:171–88.
- 255 *cultural backgrounds* Donna L. Bahry (2004), "Trust in Transitional Societies: Experimental Results from Russia," Paper presented at the American Political Science Association Meeting, Chicago. Michael Gurven and Jeffrey Winking (2008), "Collective Action in Action: Prosocial Behavior In and Out of the Laboratory," 110:179–90. Joseph Henrich (2000), "Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon," *American Economic Review*, 90:973–9. Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ensminger, Kim Hill, Francisco Gil-White, Michael Gurven, Frank

- Marlowe, John Q. Patton, Natalie Smith, and David Tracer (2005), “‘Economic Man’ in Cross-Cultural Perspective: Behavioral Experiments in 15 Small-Scale Societies,” *Behavioral & Brain Sciences*, 28:795–855. Joseph Henrich, Richard McElreath, Abigail Barr, Jean Ensminger Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwina Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer, and John Ziker (2006), “Costly Punishment Across Human Societies,” *Science*, 312:1767–70.
- 255 *large amounts of money* Juan Camilo Cardenas and Jeffrey P. Carpenter (2005), “Experiments and Economic Development: Lessons from Field Labs in the Developing World,” Middlebury College Economics Discussion Paper No. 0505.
- 255 *more lopsided* Daniel Kahneman, John L. Knetsch, and Richard H. Thaler (1986), “Fairness and the Assumptions of Economics,” *Journal of Business*, 59:S285–S300. Christoph Engel (2011), “Dictator Games: A Meta Study,” *Experimental Economics*, 14:584–610.
- 255 *Trust game* Joyce Berg, John Dickhaut, and Kevin McCabe (1995), “Trust, Reciprocity, and Social History,” *Games & Economic Behavior*, 10:122–42.
- 255 *not what happens* Colin Cramer (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Russell Sage Foundation.
- 256 *Public Goods game* John O. Ledyard (1995), “Public Goods: A Survey of Experimental Research,” in Alvin E. Roth and John H. Kagel, eds., *Handbook of Experimental Economics*, Princeton University Press.
- 256 *fear of rejection* Daniel Kahneman, John L. Knetsch, and Richard H. Thaler (1986), “Fairness and the Assumptions of Economics,” *Journal of Business*, 59:S285–S300. Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ensminger, Kim Hill, Francisco Gil-White, Michael Gurven, Frank Marlowe, John Q. Patton, Natalie Smith, and David Tracer (2005), “‘Economic Man’ in Cross-Cultural Perspective: Behavioral Experiments in 15 Small-Scale Societies,” *Behavioral & Brain Sciences*, 28:795–855.
- 256 *Some researchers claim* Ernst Fehr and Urs Fischbacher (2005), “Human Altruism: Proximate Patterns and Evolutionary Origins,” *Analyse & Kritik*, 27:6–47.
- 256 *Others claim* Roland Bénabou and Jean Tirole (2006), “Incentives and Prosocial Behavior,” *American Economic Review*, 96:1652–78. Amihai Glazer and Kai A. Konrad (1996) “A Signaling Explanation for Charity,” *American Economic Review*, 86:1019–28. Dan Ariely, Anat Bracha, and Stephan Meier (2008), “Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially,” Federal Reserve Bank of Boston Working Paper No. 07–9.
- 256 *fair solutions* Gary E. Bolton, Elena Katok, and Rami Zwick (1998), “Dictator Game Giving: Rules of Fairness Versus Acts of Kindness,” *International Journal of Game Theory*, 27:269–99.
- 256 *do what was fair* Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean

- Ensminger, Kim Hill, Francisco Gil-White, Michael Gurven, Frank Marlowe, John Q. Patton, Natalie Smith, and David Tracer (2005), "'Economic Man' in Cross-Cultural Perspective: Behavioral Experiments in 15 Small-Scale Societies," *Behavioral & Brain Sciences*, 28:795–855.
- 257 *act of skill* Elizabeth Hoffman, Kevin McCabe, Keith Shachat, and Vernon L. Smith (1994), "Preferences, Property Rights, and Anonymity in Bargaining Games," *Games & Economic Behavior*, 7:346–80.
- 257 *other way, too* Bradley J. Ruffle (1998), "More Is Better, But Fair Is Fair: Tipping in Dictator and Ultimatum Games," *Games & Economic Behavior*, 23:247–76.
- 257 *punish other players* Christopher T. Dawes, James H. Fowler, Tim Johnson, Richard McElreath, and Oleg Smirnov (2007), "Egalitarian Motivations in Humans," *Nature*, 446:794–6.
- 257 *the punishment works* Ernst Fehr and Simon Gächter (2002), "Altruistic Punishment in Humans," *Nature*, 415:137–40. James H. Fowler and Nicholas A. Christakis (2010), "Cooperative Social Behavior Cascades in Human Social Networks," *PNAS Early Edition*.
- 257 *results are mixed* David Rand, Anna Dreber, Tore Ellingsen, Drew Fudenberg, and Martin Nowak (2009), "Positive Interactions Promote Public Cooperation," *Science*, 325:1272–5. Christopher T. Dawes, James H. Fowler, Tim Johnson, Richard McElreath, and Oleg Smirnov (2007), "Egalitarian Motivations in Humans," *Nature*, 446:794–6. James Andreoni, William Harbaugh, and Lise Vesterlund (2003), "The Carrot or the Stick: Rewards, Punishments, and Cooperation," *The American Economic Review*, 93:893–902.
- 257 *game is about taking* Nicholas Bardsley (2008), "Dictator Game Giving: Altruism or Artifact?" *Experimental Economics*, 11:122–33.
- 257 *Distrust game* Iris Bohnet and Stephan Meier (2005), "Deciding to Distrust," Federal Reserve Bank of Boston Public Policy Discussion Paper 05–4.
- 257 *half of the money* Ernst Fehr and Klaus M. Schmidt (1999), "A Theory of Fairness, Competition, and Cooperation," *The Quarterly Journal of Economics*, 114:817–68. Daniel Kahneman, John L. Knetsch, and Richard H. Thaler (1986), "Fairness and the Assumptions of Economics," *Journal of Business*, 59:S285–S300.
- 257 *who are skeptical* Ilan Dinstein, Cibu Thomas, Marlene Behrmann, and David J. Heeger (2008), "A Mirror Up to Nature," *Current Biology*, 18:R13–8. Kaspar Meyer and Antonio Damasio (2008), "Mirror Neurons: Behind the Looking Glass," *Nature*, 454:167–8. Gregory Hickok (2009), "Eight Problems for the Mirror Neuron Theory of Action Understanding in Monkeys and Humans," *Cognitive Neurosciences*, 21:1229–43.
- 257 *prototype effect* Eleanor Rosch (1975), "Cognitive Representation of Semantic Categories," *Journal of Experimental Psychology*, 104:192–233. George Lakoff (1987), *Women, Fire and Dangerous Things: What Categories Reveal About the Mind*, University of Chicago Press.
- 257 *In many societies* Rebecca L. Bliege Bird and Douglas W. Bird (1997), "Delayed Reciprocity and Tolerated Theft: The Behavioral Ecology of Food-Sharing Strategies,"

- Current Anthropology*, 38:49–78. Serge Bahuchet (1990), “Food Sharing Among the Pygmies of Central Africa,” *African Study Monographs*, 11:27–53. Catherline Feeley (2011), “The Host’s Dilemma: Game Theory and Homeric Hospitality,” Classical Association of Canada Annual Meeting, Vancouver, Canada, 11 May 2002, unpublished conference paper.
- 258 *majority belief* Martin A. Nowak, Corina E. Tarnita, and Edward O. Wilson (2010), “The Evolution of Eusociality,” *Nature*, 466:1057–62.
- 258 *One rebuttal* Patrick Abbot, Jun Abe, John Alcock, et al. (2011), “Inclusive Fitness Theory and Eusociality,” *Nature*, 471:E1–4.
- 258 *reinforced each other* Robert Wright (1994), *The Moral Animal: The New Science of Evolutionary Psychology*, Pantheon Books, 201–2.
- 258 *many books published* Michael Shermer (2004), *The Science of Good and Evil: Why People Cheat, Gossip, Care, Share, and Follow the Golden Rule*, Times Books. Nigel Barber (2004), *Kindness in a Cruel World: The Evolution of Altruism*, Prometheus Books. Donald W. Pfaff (2007), *The Neuroscience of Fair Play: Why We (Usually) Follow the Golden Rule*, Dana Press. Martin A. Nowak and Roger Highfield (2011), *SuperCooperators: Altruism, Evolution, and Why We Need Each Other to Succeed*, Free Press. Patricia S. Churchland (2011), *Braintrust: What Neuroscience Tells Us About Morality*, Princeton University Press.
- 258 *older book* Matt Ridley (1996), *The Origins of Virtue*, Viking.

Chapter 4

- 41 *Robert Sapolsky* Robert Sapolsky (2003), “A Bozo of a Baboon: A Talk with Robert Sapolsky,” *Edge*.
- 41 *Matt Ridley* Matt Ridley (1993), *The Red Queen: Sex and the Evolution of Human Nature*, HarperCollins Publishers, 9–10.
- 42 *hyperbolic discounting* Richard H. Thaler (1981), “Some Empirical Evidence on Dynamic Inconsistency,” *Economics Letters*, 8:201–7. Shane Frederick, George Loewenstein, and Ted O’Donoghue (2002), “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 40:351–401.
- 43 *Francis Fukuyama* Francis Fukuyama (1995), *Trust: The Social Virtues and the Creation of Prosperity*, Simon & Schuster, 27–8.
- 43 *James Madison* James Madison (1788), *The Federalist*, 51.
- 46 *many distinct levels* Richard J. Smith (1996), “Biology and Body Size in Human Evolution: Statistical Inference Misapplied,” *Current Anthropology*, 37:451–81.
- 46 *begin to fail* Bruce Schneier (Jul 2009), “Security, Group Size, and the Human Brain,” *IEEE Security & Privacy*, 7:88.
- 50 *Code of Hammurabi* Martha T. Roth (1997), *Law Collections from Mesopotamia and Asia Minor*, Scholars Press.

Chapter 4 Notes

- 258 *evolutionary psychology* Edward O. Wilson (1975), *Sociobiology: The New Synthesis*, Harvard University Press. Edward O. Wilson (1978), *On Human Nature*, Harvard University Press.
- 258 *genetic science is flawed* Anne Innis Dagg (2004), “Love of Shopping” Is Not a Gene: Problems with Darwinian Psychology, Black Rose Books.
- 259 *Douglass North* Douglass C. North (1990), *Institutions, Institutional Change, and Economic Performance*, Cambridge University Press, 54.
- 259 *no money would* David Graeber (2011), *Debt: The First 5,000 Years*, Melville House.
- 259 *Terrence Deacon* Terrence W. Deacon (1997), *The Symbolic Species: The Co-Evolution of Language and the Human Brain*, W.W. Norton & Co., 384–401.
- 259 *far more philandering* Simon C. Griffith, Ian P. Owens, and Katherine A. Thuman (2002), “Extra Pair Paternity in Birds: A Review of Interspecific Variation and Adaptive Function,” *Molecular Ecology*, 11:2195–212. Anna Dubiec (2009), “Avian Infidelity,” *Research in Progress: Behavioral Ecology*, 3:32–4.
- 259 *There’s a balance* Thomas D. Berger and Erik Trinkaus (1995), “Patterns of Trauma Among the Neanderthals,” *Journal of Archaeological Science*, 22:841–52. Daniel S. Adler, Guy Bar-Oz, Anna Belfer-Cohen, and Ofer Bar-Yosef (2006), “Ahead of the Game: Middle and Upper Palaeolithic Hunting Behaviors in the Southern Caucasus,” *Current Anthropology*, 47:89–118. Penny Spikins, Holly Rutherford, and Andy Needham (2010), “From Homininity to Humanity: Compassion from the Earliest Archaic to Modern Humans,” *Time & Mind*, 3:303–25.
- 259 *preliminary evidence* Carles Lalueza-Fox, Antonio Rosas, Almudena Estalrich, Elena Gigli, Paula F. Campos, Antonio García-Taberner, Samuel García-Vargas, Federico Sánchez-Quinto, Oscar Ramírez, Sergi Civit, Markus Bastir, Rosa Huguet, David Santamaría, M. Thomas P. Gilbert, Eske Willerslev, and Marco de la Rasilla (2011), “Genetic Evidence for Patrilocal Mating Behavior Among Neanderthal Groups,” *Proceedings of the National Academy of Sciences of the United States of America*, 108:250–3.
- 259 *military organization* Wei-Xing Zhou, Didier Sornette, Russell A. Hill, and Robin I.M. Dunbar (2005), “Discrete Hierarchical Organization of Social Group Sizes,” *Proceedings of the Royal Society B: Biological Sciences*, 272:439–44.
- 259 *origins of religion* Scott Atran (2004), *In Gods We Trust: The Evolutionary Landscape of Religion*, Oxford. Pascal Boyer (2002), *Religion Explained: The Evolutionary Origins of Religious Thought*, Basic Books. David Sloan Wilson (2003), *Darwin’s Cathedral: Evolution, Religion, and the Nature of Society*, University of Chicago Press. Lewis Wolpert (2007), *Six Impossible Things Before Breakfast: The Evolutionary Origins of Belief*, Norton.
- 259 *social controls* Edward Alsworth Ross (1896), “Social Control,” *The American Journal of Sociology*, 1:513–35.

- 259 *not using that term* Mathieu Deflem (2007), “The Concept of Social Control: Theories and Applications,” paper presented at the International Conference on Charities as Instruments of Social Control in Nineteenth-Century Britain, Université de Haute Bretagne (Rennes 2), Rennes, France, November 22–23.
- 260 *not growing* Zeynep Tufekci (2008), “Grooming, Gossip, Facebook and Myspace,” *Information, Communication & Society*, 11:544–64. Bruno Goncalves, Nicola Perra, and Alessandro Vespignani (2011), “Validation of Dunbar’s Number in Twitter Conversations,” *PLoS ONE*, 6:e22656. *The Economist* (26 Feb 2009), “Primates on Facebook: Even Online, the Neocortex Is the Limit.”
- 260 *certain brain regions* Ryota Kanai, Bahador Bahrami, Rebecca Roylance, and Geraint Rees (2011), “Online Social Network Size Is Reflected in Human Brain Structure,” *Proceedings of the Royal Society B: Biological Sciences*, published online before print.

Chapter 5

- 53 *Douglas Hofstadter* Douglas Hofstadter (1985), *Metamagical Themas*, Bantam Dell Publishing Group.
- 57 *free-rider problem* Robert Albanese and David D. van Fleet (1985), “Rational Behavior in Groups: The Free-Riding Tendency,” *The Academy of Management Review*, 10:244–55.
- 57 *Whooping cough* Paul Offit (2011), *Deadly Choices: How the Anti-Vaccine Movement Threatens Us All*, Basic Books.

Chapter 5 Notes

- 260 *Prisoner’s Dilemma* Merrill M. Flood (1952), “Some Experimental Games,” Research Memorandum RM 789–1, The RAND Corporation. Republished as: Merrill M. Flood (1958), “Some Experimental Games,” *Management Science*, 5:5–26. Albert W. Tucker (1980), “A Two-Person Dilemma,” *UMAP Journal*, 1:101–3. Albert W. Tucker (1983), “The Mathematics of Tucker: A Sampler,” *The Two-Year College Mathematics Journal*, 14:228–32.
- 260 *Many researchers* Sylvia Nasar (2001), *A Beautiful Mind: The Life of Mathematical Genius and Nobel Laureate John Nash*, Simon & Schuster. John Nash (2008), “The Agencies Method for Modeling Coalitions & Cooperations in Games,” *International Game Theory Review*, 10:539–64. Robert Axelrod and William D. Hamilton (1981), “The Evolution of Cooperation,” *Science*, 211:1390–6. Robert Axelrod (1984), *The Evolution of Cooperation*, Basic Books.
- 260 *open grazing pasture* Garrett Hardin (1968), “The Tragedy of the Commons,” *Science*, 162:1243–8.

Chapter 6

- 63 *don't overfish* Boris Worm, Ray Hilborn, Julia K. Baum, Trevor A. Branch, Jeremy S. Collie, Christopher Costello, Michael J. Fogarty, Elizabeth A. Fulton, Jeffrey A. Hutchings, Simon Jennings, Olaf P. Jensen, Heike K. Lotze, Pamela M. Mace, Tim R. McClanahan, C  il  n Minto, Stephen R. Palumbi, Ana M. Parma, Daniel Ricard, Andrew A. Rosenberg, Reg Watson, and Dirk Zeller (2009), "Rebuilding Global Fisheries," *Science*, 325:578–85. Ed Pilkington (17 May 2010), "Saving Global Fish Stocks Would Cost 20 Million Jobs, Says UN," *The Guardian*.
- 64 *Jean-Jacques Rousseau* Jean-Jacques Rousseau (1754), *A Discourse on a Subject Proposed by the Academy of Dijon: What Is the Origin of Inequality of Men, and Is it Authorized by Natural Law?*
- 64 *defections happen* Brian Skyrms (2004), *The Stag Hunt and the Evolution of Social Structure*, Cambridge University Press.
- 65 *Snowdrift Dilemma* Michael Doebeli and Christoph Hauert (2005), "Models of Cooperation Based on the Prisoner's Dilemma and the Snowdrift Game," *Ecology Letters*, 8:748–66.
- 65 *social dilemmas* Robyn M. Dawes (1980), "Social Dilemmas," *Annual Review of Psychology*, 31:169–93. Samuel S. Komorita and Craig D. Parks (1994), *Social Dilemmas*, Westview Press.
- 65 *much more complicated* Richard H. McAdams (2008), "Beyond the Prisoners' Dilemma: Coordination, Game Theory and the Law," John M. Olin Law and Economics Working Paper No. 437, The Law School, University of Chicago.
- 68 *need to conform* Solomon Asch (Nov 1955), "Opinions and Social Pressure," *Scientific American*, 193:31–5.

Chapter 6 Notes

- 261 *predictably irrational* Dan Ariely (2008), *Predictably Irrational: The Hidden Forces That Shape our Decisions*, Harper Perennial.
- 261 *Cuban Missile Crisis* Steven J. Brams (24 Jan 2001), "Game Theory and the Cuban Missile Crisis," *Plus Magazine*.
- 262 *worst in people* Morton Deutsch and Robert M. Krauss (1960), "The Effect of Threat upon Interpersonal Bargaining," *Journal of Abnormal & Normal Social Psychology*, 61:181–9.
- 262 *better model* Rolf K  mmerli, Caroline Colliard, Nicolas Fiechter, Blaise Petitpierre, Flavien Russier, and Laurent Keller (2007), "Human Cooperation in Social Dilemmas: Comparing the Snowdrift Game with the Prisoner's Dilemma," *Proceedings of the Royal Society B: Biological Sciences*, 274:2965–70.

- 262 *Prospect Theory* Daniel Kahneman and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47:263–92.
- 262 *systems of regulation* Susan Jane Buck Cox (1985), "No Tragedy on the Commons," *Environmental Ethics*, 7:49–62.
- 262 *superrationality* Douglas Hofstadter (1985), *Metamagical Themas*, Bantam Dell Publishing Group.
- 262 *collectively rational* Anatol Rapoport, ed. (1974), *Game Theory as a Theory of Conflict Resolution*, Reidel Publishing Co.
- 262 *four different modalities* Lawrence Lessig (1998), "The New Chicago School," *The Journal of Legal Studies*, 27:661–91. Lawrence Lessig (2000), *Code and Other Laws of Cyberspace*, Basic Books.
- 262 *three basic flavors* Steven D. Levitt and Stephen J. Dubner (2006), *Freakonomics: A Rogue Economist Explores the Hidden Side of Everything*, Morrow, 17.

Chapter 7

- 75 *not to vote* William Riker and Peter Ordeshook (1968), "A Theory of the Calculus of Voting," *American Political Science Review*, 62:25–42.
- 75 *people vote* Derek Parfit (1984), *Reasons and Persons*, Oxford University Press.
- 75 *actual voters* Aaron Edlin, Andrew Gelman, and Noah Kaplan (2007), "Voting as a Rational Choice: Why and How People Vote to Improve the Well-Being of Others," *Rationality and Society*, 19:293–314.
- 76 *modified our brains* Patricia S. Churchland (2011), *Braintrust: What Neuroscience Tells Us About Morality*, Princeton University Press.
- 77 *the murder rate* Criminal Justice Information Services Division (2010), "Crime in the United States by Volume and Rate per 100,000 Inhabitants, 1990–2009," U.S. Department of Justice, Federal Bureau of Investigation.
- 77 *Elsewhere in the* United Nations Office on Drugs & Crime (2008), "The Tenth United Nations Survey of Crime Trends and Operations of Criminal Justice Systems (2005–2006)."
- 77 *Golden Rule* Jeffrey Wattles (1996), *The Golden Rule*, Oxford University Press.
- 78 *associated with religion* Azim F. Shariff and Ara Norenzayan, "God is Watching You: Priming God Concepts Increases Prosocial Behavior in an Anonymous Economic Game," *Psychological Science*, 18:803–9.
- 78 *Ten Commandments* Nina Mazar and Dan Ariely (2006), "Dishonesty in Everyday Life and Its Policy Implications," *Journal of Public Policy & Marketing*, 25:117–26.
- 78 *harsh, punitive, vengeful* Azim F. Shariff and Ara Norenzayan (2011), "Mean Gods Make Good People: Different Views of God Predict Cheating Behavior," *International Journal for the Psychology of Religion*, 21:85–96.
- 78 *pretty much useless* Bernard Gert (1989), "Morality Versus Slogans," in Daryl Close and Nicholas Meier, eds. (1995), *Morality in Criminal Justice*, Wadsworth, 51–60.

- 78 *kill animals* Jeff McMahan (2002), *The Ethics of Killing: Problems at the Margins of Life*, Oxford University Press.
- 78 *This is the stuff* Alasdair MacIntyre (1981), *After Virtue: A Study in Moral Theology*, University of Notre Dame Press.
- 79 *moral instinct* Stephen Pinker (13 Jan 2008), "The Moral Instinct," *New York Times Magazine*.
- 79 *Others relevant* Donald E. Brown (1991), *Human Universals*, Temple University Press.
- 79 *specific brain functions* Marc D. Hauser (2006), *Moral Minds: How Nature Designed our Universal Sense of Right and Wrong*, Ecco. John Mikhail (2007), "Universal Moral Grammar: Trends, Evidence, and the Future," *TRENDS in Cognitive Sciences*, 11:143–52.
- 79 *five fundamental systems* Jonathan Haidt and Craig Joseph (2004), "Intuitive Ethics: How Innately Prepared Intuitions Generate Culturally Variable Virtues," *Daedalus*, 55–66. Jonathan Haidt and Jesse Graham (2009), "Planet of the Durkheimians, Where Community, Authority, and Sacredness Are Foundations of Morality," in John T. Jost, Aaron C. Kay, and Hulda Thorisdottir, eds., *Social and Psychological Bases of Ideology and System Justification*, Oxford University Press.
- 79 *strong tendency* David Berreby (2005), *Us and Them: Understanding Your Tribal Mind*, Little, Brown and Co.
- 79 *defer to authority* Stanley Milgram (1963). "Behavioral Study of Obedience," *Journal of Abnormal & Social Psychology*, 67:371–8. Neil Lutsky (1995), "When Is 'Obedience' Obedience? Conceptual and Historical Commentary," *Journal of Social Issues*, 51: 55–65. Thomas Blass (1999), "The Milgram Paradigm After 35 Years: Some Things We Now Know About Obedience to Authority," *Journal of Applied Social Psychology*, 29:955–78.
- 79 *purity and sanctity* Mary Douglas (1978), *Purity and Danger: An Analysis of the Concepts of Pollution and Taboo*, Routledge & Kegan Paul.
- 79 *Spontaneous cooperation* Rebecca Solnit (2009), *A Paradise Built in Hell: The Extraordinary Communities that Arise in Disaster*, Viking Adult.
- 79 *increase in solidarity* Randall Collins (2004), "Rituals of Solidarity and Security in the Wake of Terrorist Attack," *Sociological Theory*, 22:58–87.
- 80 *increase in prosocial* Kathryn S. Steinberg and Patrick M. Rooney (2005), "America Gives: A Survey of Americans' Generosity After September 11," Center on Philanthropy at Indiana University.
- 80 *Crime in New York* Jason Bram (2003), "New York City's Economy Before and After September 11," *Federal Reserve Bank of New York Second District Highlights, Current Issues in Economics & Finance*, 9:1–6.
- 80 *in-group/out-group* Colin Holbrook, Paolo Sousa, and Jennifer Hahn-Holbrook (2011), "Unconscious Vigilance: Worldview Defense Without Adaptations for Terror, Coalition, or Uncertainty Management," *Journal of Personality & Social Psychology*, 101:451–66

- 80 *increase in hate crimes* Tanya Schevitz (26 Nov 2002), “FBI Sees Leap in Anti-Muslim Hate Crimes, 9/11 Attacks Blamed for Bias—Blacks Still Most Frequent Victims,” *San Francisco Chronicle*.
- 80 *Kin aversion* Mark Erickson (1989), “Incest Avoidance and Familial Bonding,” *Journal of Anthropological Research*, 45:267–91. Debra Lieberman, John Tooby, and Leda Cosmides (2003), “Does Morality Have a Biological Basis? An Empirical Test of the Factors Governing Moral Sentiments Relating to Incest,” *Proceedings of the Royal Society B: Biological Sciences*, 270:819–26. John M. Ingham and David H. Spain (2005), “Sensual Attachment and Incest Avoidance in Human Evolution and Child Development,” *Journal of the Royal Anthropological Institute*, 11:677–701. Robin Fox (1984), *The Red Lamp of Incest: An Enquiry into the Origin of Mind and Society*, University of Notre Dame Press.
- 80 *protective impulses* Sarah Hall Sternglanz, James L. Gray, and Melvin Murakami (1977), “Adult Preferences for Infantile Facial Features: An Ethological Approach,” *Animal Behaviour*, 25:108–15. Katherine A. Hildebrandt and Hiram E. Fitzgerald (1979), “Facial Feature Determinants of Perceived Infant Attractiveness,” *Infant Behavior & Development*, 2:329–39. Morten L. Kringelbach, Annukka Lehtonen, Sarah Squire, Allison G. Harvey, Michelle G. Craske, Ian E. Holliday, Alexander L. Green, Tipu Z. Aziz, Peter C. Hansen, Piers L. Cornelissen, and Alan Stein (2008), “A Specific and Rapid Neural Signature for Parental Instinct,” *PLoS ONE*, 3:e1664. Melanie L. Glocker, Daniel D. Langleben, Kosha Ruparel, James W. Loughhead, Ruben C. Gur, and Norbert Sachser (2009), “Baby Schema in Infant Faces Induces Cuteness Perception and Motivation for Caretaking in Adults,” *Ethology*, 115:257–63.
- 80 *small animals* Wakako Sanefuji, Hidehiro Ohgami, and Kazuhide Hashiya (2007), “Development of Preference for Baby Faces Across Species in Humans (*Homo sapiens*),” *Journal of Ethology*, 25:249–54.
- 80 *even dolls* Robert A. Hinde and L.A. Barden (1985), “The Evolution of the Teddy Bear,” *Animal Behaviour*, 33:1371–3. Paul H. Morris, Vasudevi Reddy, and R.C. Bunting (1995), “The Survival of the Cutest: Who’s Responsible for the Evolution of the Teddy Bear?” *Animal Behaviour*, 50:1697–700.
- 80 *notions of fairness* Joseph Henrich, Jean Ensminger, Richard McElreath, Abigail Barr, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwina Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer, and John Ziker (2010), “Markets, Religion, Community Size, and the Evolution of Fairness and Punishment,” *Science*, 327:1480–4.
- 81 *Don’t mess with Texas* Lilli Rockwell (29 Sep 2006), “Don’t Mess With Texas’ Named Top Ad,” *Austin American-Statesman*.
- 81 *effective in changing* Robert B. Cialdini (2003), “Crafting Normative Messages to Protect the Environment,” *Current Directions in Psychological Science*, 12:105–9.
- 81 *taking advantage* Jelena Grujić, Constanza Fosco, Lourdes Araujo, José A. Cuesta, and Angel Sánchez (2010), “Social Experiments in the Mesoscale: Humans Playing a Spatial Prisoner’s Dilemma,” *PLoS ONE*, 5:11, e13749.

- 82 **Andrew Colman** Andrew M. Colman (1995), *Game Theory and Its Applications in the Social and Biological Sciences*, Second Edition, Routledge.
- 82 **how people evaluate** Mark Grenovetter (1978), "Threshold Models of Collective Behavior," *American Journal of Sociology*, 83:1420–43.
- 82 **more likely to litter** Robert B. Cialdini, Carl A. Kallgren, and Raymond R. Reno (1991), "A Focus Theory of Normative Conduct: A Theoretical Refinement and Reevaluation of the Role of Norms in Human Behavior," *Advances in Experimental Social Psychology*, 21:201–34. Susan M. Reiter and William Samuel (1980), "Littering as a Function of Prior Litter and the Presence or Absence of Prohibitive Signs," *Journal of Applied Social Psychology*, 10:45–55. E. Scott Geller, Jill F. Witmer, and Margaret A. Tuso (1977), "Environmental Interventions for Litter Control," *Journal of Applied Psychology*, 62:344–51.
- 82 **seeing someone** Harold Grasmick, Robert Bursik, and Karyl Kinsey (1991), "Shame and Embarrassment as Deterrents to Noncompliance with the Law: The Case of an Antilittering Campaign," *Environment & Behavior*, 23:233–51. Carl A. Kallgren, Raymond R. Reno, and Robert B. Cialdini (2000), "A Focus Theory of Normative Conduct: When Norms Do and Do Not Affect Behavior," *Personality & Social Psychology Bulletin*, 26:1002–12.
- 82 **general breakdown** James B. Stewart (2011), *Tangled Webs: How False Statements Are Undermining America: From Martha Stewart to Bernie Madoff*, Penguin Press.
- 82 **unpunished free rider** Robert O. Kurzban and Daniel Houser (2001), "Individual Differences in Cooperation in a Circular Public Goods Game," *European Journal of Personality*, 15:537–552. David P. Myatt and Chris Wallace (2008), "When Does One Bad Apple Spoil the Barrel? An Evolutionary Analysis of Collective Action," *Review of Economic Studies*, 75:499–527.
- 82 **In Islam** Sahih Muslim, Book 037, Hadith Number 6658.
- 83 **propaganda campaigns** David Livingstone Smith (2011), *Less than Human: Why We Demean, Enslave, and Exterminate Others*, St. Martin's Press.
- 83 **failure or absence** Simon Baron-Cohen (2011), *Zero Degrees of Empathy: A New Theory of Human Cruelty*, Penguin/Allen Lane; published in the U.S. as *The Science of Evil: On Empathy and the Origins of Human Cruelty*, Basic Books.
- 84 **food on the honor** Stephen J. Dubner and Steven D. Levitt (6 Jun 2004), "What the Bagel Man Saw: An Accidental Glimpse at Human Nature," *New York Times Magazine*, 62–5.
- 84 **follow social norms** Daniel Kahneman and Dale T. Miller (1986), "Norm Theory: Comparing Reality to Its Alternatives," *Psychological Review*, 93:136–53. Cass R. Sunstein (1996), "Social Norms and Social Roles," *Columbia Law Review*, 96:903–68. Helen Bernhard, Ernst Fehr, and Urs Fischbacher (2006), "Group Affiliation and Altruistic Norm Enforcement," *American Economic Review*, 96:217–21.
- 84 **Emmanuel Levinas** Michael L. Morgan (2011), *The Cambridge Introduction to Emmanuel Levinas*, Cambridge University Press.
- 84 **people turned in** Akiko Fujita (17 Aug 2011), "Honest Japanese Return \$78 Million in Cash Found in Quake Rubble," *ABC News*.

- 85 *practice of religion* David Kowalewski (1980), "Protest for Religious Rights in the USSR: Characteristics and Consequences," *Russian Review*, 39:426–41. Sabrina Petra Ramet, ed. (1993), *Religious Policy in the Soviet Union*, Cambridge University Press.
- 85 *Confidence tricksters* Frank Stajano and Paul Wilson (Mar 2011), "Understanding Scam Victims: Seven Principles for Systems Security," *Communications of the ACM*, 54:70–5.
- 85 *we have trouble* David McCullough (2003), *Truman*, Simon & Schuster, 510.

Chapter 7 Notes

- 263 *the same analysis* Ilya Somin (2006), "Knowledge about Ignorance: New Directions in the Study of Political Information," *Critical Review*, 18:255–78.
- 263 *related to other* James H. Fowler (2006), "Altruism and Turnout," *Journal of Politics*, 68:674–83. James H. Fowler and Cindy D. Kam (2007), "Beyond the Self: Altruism, Social Identity, and Political Participation," *Journal of Politics*, 69:811–25. James H. Fowler and Christopher T. Dawes (2008), "Two Genes Predict Voter Turnout," *Journal of Politics*, 70:579–94.
- 263 *not distinguishing* Mark Johnson (1993), *Moral Imagination: Implications of Cognitive Science for Ethics*, University of Chicago Press.
- 263 *National Voter Turnout* Demos (2009), "Demos Fact Sheet: Election Day Registration."
- 264 *significant difference* Stephen Anderson (2009), "The Golden Rule: Not So Golden Anymore," *Philosophy Now*, 74:26–9.
- 264 *Chácobo of Bolivia* Gilbert R. Prost (1983) "Chácobo: Society of Equality," unpublished manuscript, Department of Anthropology, University of Florida; quoted in Michael Gurven (2004), "To Give and to Give Not: The Behavioral Ecology of Human Food Transfers," *Behavioral & Brain Sciences*, 27:543–59.
- 264 *Maori of New Zealand* Raymond Firth (1926), "Proverbs in Native Life, with Special Reference to Those of the Maori," *Folklore*, 37:134–53.
- 264 *Yeyi of Botswana* Alistair J. Sutherland (1981), "Local Level Dispute Processes in Botswana: The Yeyi Moot Encapsulated," *Journal of African Law*, 25:94–114.
- 264 *believe in free will* Kathleen D. Vohs and Jonathan W. Schooler (2008), "The Value of Believing in Free Will: Encouraging a Belief in Determinism Increases Cheating," *Psychological Science*, 19:49–54
- 264 *believe in predestination* Tyler F. Stillman, Roy F. Baumeister, Kathleen D. Vohs, Nathaniel M. Lambert, Frank D. Fincham, and Lauren E. Brewer (2010), "Personal Philosophy and Personnel Achievement: Belief in Free Will Predicts Better Job Performance," *Social Psychological & Personality Science*, 1:43–50.
- 264 *concept of free will* Shaun Nichols (2011), "Experimental Philosophy and the Problem of Free Will," *Science*, 331:1401–3.
- 264 *Inbreeding is likely* M. Menotti-Raymond and S.J. O'Brien (1993), "Dating the Genetic Bottleneck of the African Cheetah," *Proceedings of the National Academy of Sciences of the United States of America*, 90:3172–6. William Amos and John Harwood (1998),

- “Factors Affecting Levels of Genetic Diversity in Natural Populations,” *Philosophical Transactions of the Royal Society B*, 353:177–86. Mary Jane McKay (11 Feb 2009), “Genetic Disorders Hit Amish Hard,” *CBS News*.
- 264 *isolated from modern* Joseph Henrich (2000), “Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon,” *American Economic Review*, 90:973–9.
- 265 *wash their hands* John M. Lynn (2000), “Method and Apparatus for Helping to Ensure the Washing of Hands,” U.S. Patent #6,031,461. John M. Lynn (2000), “Method and Apparatus for Helping to Ensure the Washing of Hands,” U.S. Patent #6,147,607. Fiona A. Lynn and John M. Lynn (2001), “Method and Apparatus for Helping to Ensure the Washing of Hands,” U.S. Patent #6,211,788.
- 265 *broken windows theory* James Q. Wilson and George L. Kelling (Mar 1982), “Broken Windows,” *Atlantic Monthly*, 127:29–38. George L. Kelling (17 Jul 2009), “How New York Became Safe: The Full Story,” *City Journal*.
- 265 *Researchers compared* Khaled Taqi-Eddin and Dan Macallair (1999), “Shattering Broken Windows: An Analysis of San Francisco’s Alternative Crime Policies,” Justice Policy Institute. Randall G. Sheldon (2003), “Assessing ‘Broken Windows’: A Brief Critique,” Center on Juvenile & Criminal Justice. Bernard E. Harcourt and Jens Ludwig (2006), “Broken Windows: New Evidence from New York City and a Five-City Social Experiment,” *University of Chicago Law Review*, 73:271–320.
- 265 *evidence that it does* Kees Keizer, Siegwart Lindenberg, and Linda Steg (2008), “The Spreading of Disorder,” *Science*, 322:1681–5.
- 265 *Steven Levitt* Steven D. Levitt (2004), “Understanding Why Crime Fell in 1990s: Four Factors that Explain the Decline and Six that Do Not,” *Journal of Economic Perspectives*, 18:163–90.
- 266 *reduced their energy* Ian Ayres, Sophie Raseman, and Alice Shih (2009), “Evidence from Two Large Field Experiments That Peer Comparison Feedback Can Reduce Residential Energy Usage,” NBER Working Paper 15386, National Bureau of Economic Research.
- 266 *this system backfires* Dora L. Costa and Matthew E. Kahn (2010), “Energy Conservation ‘Nudges’ and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment,” NBER Working Paper No. 15939, National Bureau of Economic Research.
- 266 *people are swayed* Matt Davis (2011), “Behavior and Energy Savings: Evidence from a Series of Experimental Interventions,” Environmental Defense Fund.
- 266 *In Rwanda* Heather B. Hamilton (2000), “Rwanda’s Women: The Key to Reconstruction,” *Journal of Humanitarian Assistance*.
- 266 *cash box was made* Stephen J. Dubner and Steven D. Levitt (6 Jun 2004), “What the Bagel Man Saw: An Accidental Glimpse at Human Nature,” *New York Times Magazine*, 62–5.
- 266 *a host of unknowns* Steven D. Levitt (2006) “White-Collar Crime Writ Small: A Case Study of Bagels, Donuts, and the Honor System,” *The American Economic Review*, 96:290–4.

Chapter 8

- 87 *ability to deceive* David Livingstone Smith (2004), *Why We Lie: The Evolutionary Roots of Deception and the Unconscious Mind*, St. Martin's Press.
- 87 *Wason Selection Task* Peter C. Wason (1966), "Reasoning," in B.M. Foss, *New Horizons in Psychology*, Penguin. Peter C. Wason and Diana Shapiro (1971), "Natural and Contrived Experience in a Reasoning Problem," *Quarterly Journal of Experimental Psychology*, 23:63–71.
- 87 *better at solving* Leda Cosmides (1989), "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task," *Cognition*, 31:187–276.
- 87 *scans of the brains* Nicola Canessa, Alessandra Gorini, Stefano F. Cappa, Massimo Piattelli-Palmarini, Massimo Danna, Ferruccio Fazio, and Daniela Perani (2005), "The Effect of Social Context on Deductive Reasoning: An fMRI Study," *Human Brain Mapping*, 26:30–43.
- 88 *we cooperate primarily* Natalie Henrich and Joseph Henrich (2007), "Evolutionary Theory and the Social Psychology of Human Cooperation," in Joseph Henrich and Natalie Henrich, eds., *Why Humans Cooperate: A Cultural and Evolutionary Explanation*, Oxford University Press: 35–74. Robert Boyd and Peter J. Richerson (1992), "Punishment Allows the Evolution of Cooperation (or Anything Else) in Sizable Groups," *Ethology & Sociobiology*, 13:171–95.
- 89 *reputation information* Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman (2000), "Reputation Systems," *Communications of the ACM*, 43:45–8. Daniel Houser and John Wooders (2006), "Reputation in Auctions: Theory, and Evidence from eBay," *Journal of Economics & Management Strategy*, 15:353–69. Paul Resnick, Richard Zeckhauser, John Zwanson, and Kate Lockwood (2006), "The Value of Reputation on eBay: A Controlled Experiment," *Experimental Economics*, 9:79–101. Jian Yang, Xiaorhi Hu, and Han Zhang (2007), "Effects of a Reputation Feedback System on an Online Consumer-to-Consumer Auction Market," *Decision Support Systems*, 44:93–105.
- 89 *eBay changed this* Max Sherry (30 Oct 2008), "eBay Feedback System Negative for Sellers," suite101.com. Greg Kusch (8 Jan 2010), "Leaving Buyer Feedback on eBay (I'll Scratch Your Back. Will You Scratch Mine?)," Ezinearticles.com.
- 89 *remember negative* Felicia Pratto and Oliver P. John (1991), "Automatic Vigilance: The Attention-Grabbing Power of Negative Social Information," *Journal of Personality & Social Psychology*, 61:380–91. Roy F. Baumeister, Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs (2001), "Bad Is Stronger Than Good," *Review of General Psychology*, 5:323–70. Tiffany A. Ito, Jeff T. Larsen, N. Kyle Smith, and John T. Cacioppo (1998), "Negative Information Weighs More Heavily on the Brain: The Negativity Bias in Evaluative Categorizations," *Journal of Personality & Social Psychology*, 75:887–900. Elizabeth A. Kensinger (2007), "Negative Emotion Enhances

- Memory Accuracy: Behavioral and Neuroimaging Evidence,” *Current Directions in Psychological Science*, 16:213–18.
- 89 *theory that gossip* Peter J. Wilson (1974), “Filcher of Good Names: An Enquiry Into Anthropology and Gossip,” *Man, New Series*, 9:93–102. Ronald S. Burt and Mark Knez (1995), “Kinds of Third-Party Effects on Trust,” *Rationality & Society*, 7:255–92. Ralf D. Sommerfeld, Hans-Jurgen Krambeck, Dirk Semmann, and Manfred Milinski (2007), “Gossip as an Alternative for Direct Observation in Games of Indirect Reciprocity,” *Proceedings of the National Academy of Sciences of the United States of America*, 104:17435–40. Robin I.M. Dunbar (22 Nov 1992), “Why Gossip Is Good for You,” *New Scientist*, 28–31. Robin I.M. Dunbar (1996), *Grooming, Gossip and the Evolution of Language*, Harvard University Press.
- 89 *whom to interact* Robin I.M. Dunbar (2004), “Gossip in Evolutionary Perspective,” *Review of General Psychology*, 8:100–10. Ralf D. Sommerfeld, Hans-Jurgen Krambeck, Dirk Semmann, and Manfred Milinski (2007), “Gossip as an Alternative for Direct Observation in Games of Indirect Reciprocity,” *Proceedings of the National Academy of Sciences of the United States of America*, 104:17435–40.
- 89 *establish group interests* Richard H. McAdams (1997), “The Origin, Development, and Regulation of Norms,” *Michigan Law Review*, 96:338–433. Roy F. Baumeister, Liqing Zhang, and Kathleen D. Vohs (2004), “Gossip as Cultural Learning,” *Review of General Psychology*, 8:111–21.
- 89 *gossip helps keep* Richard McAdams (1996), “Group Norms, Gossip, and Blackmail,” *University of Pennsylvania Law Review*, 144:2237–92. Dirk Semmann, Hans-Jürgen Krambeck, and Manfred Milinski (2005), “Reputation Is Valuable Within and Outside One’s Own Social Group,” *Behavioral Ecology & Sociobiology*, 57:611–6. Bianca Beersma and Gerben A. Van Kleef (2011), “How the Grapevine Keeps You in Line: Gossip Increases Contributions to the Group,” *Social Psychological & Personality Science*, 2:642–9.
- 89 *Diamond merchants* Barak D. Richman (2002), “Community Enforcement of Informal Contracts: Jewish Diamond Merchants in New York,” Harvard Law School John M. Olin Center for Law, Economics and Business Discussion Paper Series, Paper 384.
- 90 *a pair of eyes* Melissa Bateson, Daniel Nettle, and Gilbert Roberts (2006), “Cues of Being Watched Enhance Cooperation in a Real-World Setting,” *Biology Letters*, 2:212–4.
- 90 *names and addresses* Arthur L. Beaman, Bonnel Kentz, Edward Diener, and Soren Svanum (1979), “Self-Awareness and Transgression in Children: Two Field Studies,” *Journal of Personality & Social Psychology*, 37:1835–46.
- 90 *There’s a great word* Oliver Conway (22 Jun 2004), “Congo Word ‘Most Untranslatable,’” *BBC News*.
- 91 *Quakers would cooperate* Adrian Cadbury (2003), “Beliefs and Business: The Experience of Quaker Companies,” The Foundation of Lady Katherine Leveson. James Surowiecki (2004), *The Wisdom of Crowds*, Anchor. Steven Davison (2011), “The Double-Culture Period: Factors in Quaker Success,” unpublished manuscript.

- 92 *Maghribi traders* Avner Greif (2008), "Contract Enforcement and Institutions among the Maghribi Traders: Refuting Edwards and Ogilvie," SIEPR Discussion Paper 08–018, Stanford Institute for Economic Policy Research.
- 92 *Robert Putnam* Robert Putnam (2007), "E Pluribus Unum: Diversity and Community in the Twenty-First Century," *Scandinavian Political Studies*, 30:137–74. Robert Putnam (2001), *Bowling Alone: The Collapse and Revival of American Community*, Touchstone/Simon & Schuster.
- 92 *vestigial remnants* David Sally (2000), "A General Theory of Sympathy, Mind-Reading, and Social Interaction, with an Application to the Prisoners' Dilemma," *Social Science Information*, 39:567–643.
- 92 *people who look like* Lisa M. DeBruine (2004), "Facial Resemblance Increases the Attractiveness of Same-Sex Faces More Than Other-Sex Faces," *Proceedings of the Royal Society B: Biological Sciences*, 271:2085–90. Lisa M. DeBruine (2005), "Trustworthy but Not Lust-Worthy: Context-Specific Effects of Facial Resemblance," *Proceedings of the Royal Society B: Biological Sciences*, 272:919–22. Lisa M. DeBruine, Benedict C. Jones, Anthony C. Little, and David I. Perrett (2008), "Social Perception of Facial Resemblance in Humans," *Archives of Sexual Behavior*, 27:64–77. Alexandra Alvergne, Elise Huchard, Damien Caillaud, Marie J. E. Charpentier, Joanna M. Setchell, Charlène Ruppli, Delphine Féjan, Laura Martinez, Guy Cowlshaw, and Michel Raymond (2009), "Human Ability to Recognize Kin Visually Within Primates," *International Journal of Primatology*, 30:199–210.
- 92 *same accent* Shirley Lev-Ari and Boaz Keysar (2010), "Why Don't We Believe Non-Native Speakers? The Influence of Accent on Credibility," *Journal of Experimental Social Psychology*, 46:1093–6.
- 92 *change our patterns* Tanya L. Chartrand and John A. Bargh (1999), "The Chameleon Effect: The Perception-Behavior Link and Social Interaction," *Journal of Personality & Social Psychology*, 76:893–910. Jessica L. Lakin, Valerie E. Jefferis, Clara Michelle Cheng, and Tanya L. Chartrand (2003), "The Chameleon Effect As Social Glue: Evidence for the Evolutionary Significance of Nonconscious Mimicry," *Journal of Nonverbal Behavior*, 27:145–62. Tanya L. Chartrand and Rick van Baaren (2009), "Human Mimicry," *Advances in Experimental Social Psychology*, 41:219–74.
- 92 *don't sound like* John Baugh (2010), "Linguistic Profiling," in Sinfree Makoni, Geneva Smitherman, and Arnetha Ball, eds., *Black Linguistics*, Routledge. Shiri Lev-Ari and Boaz Keysar (2010), "Why Don't We Believe Non-Native Speakers? The Influence of Accent on Credibility," *Journal of Experimental Social Psychology*, 46:1093–96.
- 93 *how to cooperate* Ernst Fehr and Urst Fischbacher (2004), "Social Norms and Human Cooperation," *Trends in Cognitive Sciences*, 8:185–90.
- 93 *Platinum Rule* Tony Alessandra and Michael J. O'Connor (1996), *The Platinum Rule: Discover the Four Basic Business Personalities—And How They Can Lead You to Success*, Warner Books.
- 93 *reputation scale* Avinash Dixit and Barry Nalebuff (1991), *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*, Norton, 144–61.

- 95 *Criminals have signals* Diego Gambetta (2009), *Codes of the Underworld: How Criminals Communicate*, Princeton University Press.
- 95 *Shared brand names* Jean Tirole (1996), "A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality)," *Review of Economic Studies*, 63:1–22.
- 95 *Saudi Binladin Group* Holly Williams (20 Nov 2001), "Bin Laden Group Reputation Brief on Global Scale," *PR Week UK*.
- 95 *pay a premium* Stefano Castriota and Marco Delmastro (2009), "The Economics of Collective Reputation: Minimum Quality Standards, Vertical Differentiation, and Optimal Group Size," American Association of Wine Economists Working Paper 50.
- 95 *more reputable producer* Stuart Landon and Constance E. Smith (1998), "Quality Expectations, Reputation, and Price," *Southern Economic Journal*, 64:628–47. Olivier Gergaud and Florine Livat (2004), "Team versus Individual Reputations: A Model of Interaction and Some Empirical Evidence," *Cahiers de la Maison des Sciences Economiques*, No. bla04015.
- 95 *cultivate their brand* Smooch Reynolds (Dec 2002), "Career Branding: Is There Really Such a Concept?" *Public Relations Tactics*, 9:7,22.
- 96 *A small complaint* Daniel J. Solove (2007), *The Future of Reputation: Gossip, Rumor, and Privacy on the Internet*, Yale University Press.
- 96 *isn't an effective* Daniel Balliet, Laetitia B. Mulder, and Paul A.M. Van Lange (2011), "Reward, Punishment, and Cooperation: A Meta-Analysis," *Psychological Bulletin*, 137:594–615.
- 97 *Maine lobstermen* James M. Acheson (Apr 1972), "Territories of the Lobstermen," *Natural History Magazine*, 81 (4): 60–9.
- 97 *Shame is a common* Thomas J. Scheff (2000), "Shame and the Social Bond: A Sociological Theory," *Sociological Theory*, 18:84–99.
- 97 *excessive CEO pay* Sandeep Gopalan (2007), "Shame Sanctions and Excessive CEO Pay," *Delaware Journal of Corporate Law*, 32:757–97.
- 98 *Informal punishments* Marc Galanter and David Luban (1993), "Poetic Justice: Punitive Damages and Legal Pluralism," *The American University Law Review*, 42:1393–1463.
- 98 *bad apples diminishes* Jaap W. Ouwerkerk, Paul A.M. van Lange, Marcello Gallucci, and Norbert L. Kerr (2004), "Avoiding the Social Death Penalty: Ostracism and Cooperation in Social Dilemmas," Presentation at the 7th Annual Sydney Symposium of Social Psychology: "The Social Outcast: Ostracism, Social Exclusion, Rejection, and Bullying," 15–18 Mar, 2004. Norbert L. Kerr, Ann C. Rumble, Ernest S. Park, Jaap W. Ourwerkerk, Craig D. Parks, Marcello Gallucci, and Paul A.M. van Lange (2009), "How Many Bad Apples Does It Take to Spoil the Whole Barrel?" Social Exclusion and Toleration for Bad Apples," *Journal of Experimental Social Psychology*, 45:603–13.
- 98 *if Alice kills Bob* Napoleon A. Chagnon (1988), "Life Histories, Blood Revenge, and Warfare in a Tribal Population," *Science, New Series*, 239:985–92. Jon Elster (1990), "Norms of Revenge," *Ethics*, 100:862–85.

- 98 *Nazis instituted this* Robert Loeffel (2007), "Sippenhaft, Terror and Fear in Nazi Germany: Examining One Facet of Terror in the Aftermath of the Plot of 20 July 1944," *Contemporary European History*, 16:51–69.
- 98 *various cultures* Jon Elster (1990), "Norms of Revenge," *Ethics*, 100:862–85. Stephen D. White (1986), "Feuding and Peace-Making in the Touraine Around the Year 1100," *Traditio*, 42:195–263. Trevor Dean (1997), "Marriage and Mutilation: Vendetta in Late Medieval Italy," *Past & Present*, 157:3–36. Joel T. Rosenthal (1966), "Marriage and the Blood Feud in 'Heroic' Europe," *The British Journal of Sociology*, 17:133–44.
- 98 *Hatfields versus the McCoys* Otis K. Rice (1982), *The Hatfields and McCoys*, University of Kentucky Press.
- 98 *reputation can solve* Brian Skyrms (1996), *The Evolution of the Social Contract*, Cambridge University Press, 56–8.
- 99 *John Wayne Gacy* Clifford L. Linedecker (1980), *The Man Who Killed Boys: A True Story of Mass Murder in a Chicago Suburb*, St. Martin's Press.
- 99 *Dr. Harold Shipman* Sarah Ramsay (2001), "Audit Further Exposes UK's Worst Serial Killer," *Lancet*, 357:123–4.
- 101 *settling disputes* Robert C. Ellickson (1991), *Order without Law: How Neighbors Settle Disputes*, Harvard University Press.
- 102 *Dunbar wrote* Robin I.M. Dunbar (Nov 1992), "Why Gossip Is Good For You," *New Scientist*, 28–31.
- 102 *Hardin: "Perhaps"* Garrett Hardin (1994), "The Tragedy of the Unmanaged Commons," *Trends in Ecology & Evolution*, 9:199.

Chapter 8 Notes

- 266 *People who defect* James E. Alcock and Diana Mansell (1977), "Predisposition and Behaviour in a Collective Dilemma," *Journal of Conflict Resolution*, 21:443–57.
- 267 *better at predicting* Robert H. Frank (1988), *Passions within Reason: The Strategic Role of the Emotions*, W.W. Norton & Co., 137–43.
- 267 *better at recognizing* Harold H. Kelly and Anthony J. Stahelski (1970), "Errors in Perception of Intentions in a Mixed-Motive Game," *Journal of Experimental Social Psychology*, 6:279–400.
- 267 *Robert Axelrod* Robert Axelrod (1984), *The Evolution of Cooperation*, Basic Books.
- 267 *oft-quoted line* Pete Blackshaw (2008), *Satisfied Customers Tell Three Friends, Angry Customers Tell 3,000*, Crown Business.
- 267 *cooperation jumps* Claus Wedekind and Manfred Milinski (1996), "Human Cooperation in the Simultaneous and the Alternating Prisoner's Dilemma: Pavlov Versus Generous Tit-for-Tat," *Proceedings of the National Academy of Sciences of the United States of America*, 93:2686–9.
- 268 *Dueling isn't always* Christopher Kingston and Robert E. Wright (2009), "The Deadliest of Games: The Institution of Dueling," *Southern Economic Journal*, 76:1094–1106.

- Diego Gambetta (2009), *Codes of the Underworld: How Criminals Communicate*, Princeton University Press.
- 268 *Chimpanzees are able* Francys Subaiul, Jennifer Vonk, Sanae Okamoto-Barth, and Jochen Barth (2008), "Do Chimpanzees Learn Reputation by Observation? Evidence from Direct and Indirect Experience with Generous and Selfish Strangers," *Animal Cognition*, 11:611–23.
- 268 *Charles Lamb* Charles Lamb (1833), "Table-Talk by the Late Elia," in Edward Verrall Lucas, ed. (1913), *The Works of Charles and Mary Lamb: Miscellaneous Prose, 1798–1834*, Macmillan, 400.
- 268 *counter-evidence* Francisco C. Santos, Marta D. Santos, Jorge M. Pacheco (2008), "Social Diversity Promotes the Emergence of Cooperation in Public Goods Games," *Nature*, 454:213–6. Sven Van Segbroeck, Francisco C. Santos, Tom Lenaerts, Jorge M. Pacheco (2009), "Reacting Differently to Adverse Ties Promotes Cooperation in Social Networks," *Physical Review Letters*, 102:058105–9.
- 268 *diverse communities* Eric M. Uslaner (2002), *The Moral Foundations of Trust*, Cambridge University Press. Eric M. Uslaner (2004), "Trust and Social Bonds: Faith in Others and Policy Outcomes Reconsidered," *Political Research Quarterly*, 57:501–7. Eric M. Uslaner (2011), "Trust, Diversity, and Segregation in the United States and the United Kingdom," *Comparative Sociology*, 10:221–47.
- 268 *different dialects* Ottar Tinglum (1929), "Landsmaal-Riksmaal: The Feud Between the Dialects and the Official Language of Norway," *American Journal of Sociology*, 34:686–92.
- 268 *English dialects* Sally Johnson, Tommaso M. Milani, and Clive Upton (2008), "Whose Voices? A Hypermodal Approach to Language Ideological Debates on the BBC 'Voices' Website," Working Paper No. 127, Lancaster University Department of Linguistics and English Language Centre for Language in Social Life. Clive Upton and John David Allison Widdowson (2006), *An Atlas of English Dialects*, Oxford University Press.
- 268 *David Nettle ran* David Nettle (1999), *Linguistic Diversity*, Oxford University Press.
- 268 *alternate analysis* Brian Skyrms (1996), *The Evolution of the Social Contract*, Cambridge University Press, 57–61.
- 269 *"is" does not imply* David Hume (1739), *A Treatise on Human Nature*.
- 269 *costly religious rituals* William Irons (1996), "Morality as an Evolved Adaptation," in J.P. Hurd, ed., *Investigating the Biological Foundations of Morality*, Edwin Mellen Press, 1–34. William Irons (2001), "Religion as a Hard-to-Fake Sign of Commitment," in Randolph Nesse, ed., *The Evolution of Commitment*, Russell Sage Foundation, 155–78. Richard Sosis (2000), "Religion and Intra-Group Cooperation: Preliminary Results of a Comparative Analysis of Utopian Communities," *Cross-Cultural Research*, 37:211–39.
- 269 *breeding females* James Acheson and Roy Gardner (2011), "The Evolution of the Maine Lobster V-Notch Practice: Cooperation in a Prisoner's Dilemma Game," *Ecology & Society*, 16:41.

- 270 *Edney listed several* Julian J. Edney (1981), "Paradoxes on the Commons: Scarcity and the Problem of Inequality," *Journal of Community Psychology*, 9:3–34.
- 270 *Michel Foucault said* Didier Eribon (1991), *Michel Foucault*, translated from French by Betsy Wing, Harvard University Press, 265.

Chapter 9

- 103 *continuous operation clause* Marc C. Singer (2005), "'Going Dark' Provisions: Making Clients See the Light," *New Jersey Law Journal*, 179:1–3.
- 104 *solitary, poor, nasty* Thomas Hobbes (1651), *Leviathan*, Printed for Andrew Crooke, at the Green Dragon in St. Paul's Churchyard.
- 104 *Martin Luther said* Martin Luther (1524), *Von Kaufshandlung und Wucher*.
- 104 *Immanuel Kant* Immanuel Kant (1795), *First Supplement on the Guarantee for Perpetual Peace*.
- 104 *John Locke* John Locke (1689), *Second Treatise of Government*.
- 104 *Jean-Jacques Rousseau* Jean-Jacques Rousseau (1762), *Du Contrat Social (The Social Contract)*.
- 104 *John Rawls* John Rawls (1971), *A Theory of Justice*, Belknap Press of Harvard University Press.
- 104 *In Plato's Republic* Plato (c. 427–347 BC), *The Republic*.
- 105 *Niccolò Machiavelli* Niccolò Machiavelli (1517), *Discourses Upon the First Ten Books of Titus Livy*.
- 105 *Garrett Hardin* Garrett Hardin (1994). "The Tragedy of the Unmanaged Commons," *Trends in Ecology & Evolution*, 9:199.
- 106 *voting is required* Elliot Frankal (4 Jul 2005), "Compulsory Voting Around the World," *The Guardian*.
- 107 *marriage rites were informal* Stephanie Coontz (2005), *Marriage, a History: From Obedience to Intimacy, or How Love Conquered Marriage*, Viking.
- 107 *accomplices at worst* Marla Dickerson and Cecilia Sanchez (5 Aug 2008), "Mexican Police Linked to Rising Kidnappings," *Los Angeles Times*.
- 107 *occasionally jails* British Broadcasting Corporation (20 June 2011), "Somalia: Six Jailed for 'Pirate Ransom' Case," *BBC News*.
- 108 *offering a bribe* Kaushik Basu (2011), "Why, for a Class of Bribes, the Act of Giving a Bribe Should Be Treated as Legal," Working Paper 1-2011-DEA, Ministry of Finance, Government of India.
- 108 *easy to make exceptions* Ed Moloney (2002), *A Secret History of the IRA*, Viking Penguin.
- 110 *Alexander Hamilton* Alexander Hamilton (1787), "The Federalist #15," in Alexander Hamilton, John Jay and James Madison; John C. Hamilton, ed. (1866), *The Federalist: A Commentary On the Constitution of the United States*, Lippincott, 143.
- 110 *immoral, ineffective* Toni M. Massaro (1991), "Shame, Culture, and American Criminal Law," *Michigan Law Review*, 89:1880–1944. James Q. Whitman (1998), "What

- Is Wrong with Inflicting Shame Sanctions?" *The Yale Law Journal*, 107:1055–92.
- Brian Netter (2005), "Avoiding the Shameful Backlash: Social Repercussions for the Increased Use of Alternative Sanctions," *The Journal of Criminal Law & Criminology*, 96:187–216. Dan M. Kahan (2006), "What's Really Wrong with Shaming Sanctions," *Texas Law Review*, 84:2075–95.
- 111 *finer reduce speeding* Robert Beattie Crawford (1989), "Toward a Theory on Speeding and Penalties: Have the Increased Fines Reduced Speeding in Massachusetts?" Master's thesis, Massachusetts Institute of Technology. D.C. Webster and P.A. Wells (2000), "The Characteristics of Speeders," Road Safety Division, Department of the Environment, Transport and the Regions, TRL Report. Saranath Lawpoolsri, Jingyi Li, and Elisa R. Braver (2007), "Do Speeding Tickets Reduce the Likelihood of Receiving Subsequent Speeding Tickets? A Longitudinal Study of Speeding Violators in Maryland," *Traffic Injury Prevention*, 8:26–34.
- 111 *uninsured drivers* Ray Massey (23 Nov 2010), "Uninsured Drivers Kill 160 People a Year but Face Inadequate Fines as Low as £50," *Daily Mail*.
- 111 *basing traffic fines* British Broadcasting Corporation (14 Jan 2002), "Nokia Boss Gets Record Speeding Fine," *BBC News*. British Broadcasting Corporation (10 Feb 2004), "Finn's Speed Fine Is a Bit Rich," *BBC News*. Jason Paur (8 Jan 2010), "Swiss Slap Speeder With \$290K Fine," *Wired*. British Broadcasting Corporation (12 Aug 2010), "Swede Faces World-Record \$1m Speeding Penalty," *BBC News*.
- 113 *tax the use of antibiotics* Infectious Diseases Society of America (2011), "Combating Antimicrobial Resistance: Policy Recommendations to Save Lives," *Clinical Infectious Diseases*, 52 (suppl. 5): S397–428.
- 114 *In one experiment* Ann E. Tenbrunsel and David M. Messick (1999), "Sanctioning Systems: Decision Frames and Cooperation," *Administrative Science Quarterly*, 44:684–707.
- 114 *instituted a fine* Uri Gneezy and Aldo Rustichini (2000), "A Fine is a Price," *Journal of Legal Studies*, 29:1–17.
- 114 *towns are experimenting* Ben Hamilton-Baillie (2004), "Urban Design: Why Don't We Do It In The Road? Modifying Traffic Behavior Through Legible Urban Design," *Journal of Urban Technology*, 11:43–6. Peter Peters (2005), "Exchanging Travel Speed: Time Politics in Mobility Practices," *Configurations*, 13:395–419,436. Tom Vanderbilt (Summer 2008), "The Traffic Guru," *The Wilson Quarterly*, 26–32.
- 114 *incident in Switzerland* Bruno S. Fray and Felix Oberholzer-Gee (1997), "The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out," *American Economic Review*, 87:746–55.
- 115 *gave different weights* Don A. More, Lloyd Tantu, and Max H. Bazerman (2010), "Conflict of Interest and the Intrusion of Bias," *Judgment & Decision Making*, 5:37–53.
- 115 *require advisors* Daylian M. Cain, George Loewenstein, and Don A. More (2005), "The Dirt on Coming Clean: Perverse Effects of Disclosing Conflict of Interest," *The Journal of Legal Studies*, 34:1–25.

- 116 *able to self-regulate* Fikret Berkes, David Feeny, Bonnie J. McCay, and James M. Acheson (1989), "The Benefits of the Commons," *Nature*, 340:91–3. David Feeny, Fikret Berkes, Bonnie J. McCay, and James M. Acheson (1990), "The Tragedy of the Commons: Twenty-Two Years Later," *Human Ecology*, 18:1–19. Carl J. Dahlman (1992), "The Tragedy of the Commons that Wasn't: On Technical Solutions to the Institutions Game," *Population & Environment*, 12:285–96. Elinor Ostrom, Joanna Burger, Christopher B. Field, Richard B. Norgaard, and David Policansky (1999), "Revisiting the Commons: Local Lessons, Global Challenges," *Science*, 284:278–82.
- 116 *outsiders come in* Mansel G. Blackford (2009), "Fishers, Fishing, and Overfishing: American Experiences in Global Perspective, 1976–2006," *Business History Review*, 83:239–66.
- 117 *The scarlet letter* George Elliott Howard (1904), *A History of Matrimonial Institutions: Chiefly in England and the United States, with an Introductory Analysis of the Literature and the Theories of Primitive Marriage and the Family*, Fred B. Rothman & Co., 169–78.
- 118 *existential threat* John Mueller and Mark G. Stewart (2 Apr 2010), "Hardly Existential: Thinking Rationally About Terrorism," *Foreign Affairs*.
- 118 *death penalty reduces* Hashem Dezhbakhsh, Paul H. Rubin, and Joanna M. Shepherd (2003), "Does Capital Punishment Have a Deterrent Effect? New Evidence from Post-moratorium Panel Data," *American Law & Economics Review*, 5:344–76. Paul R. Zimmerman (2004), "State Executions, Deterrence, and the Incidence of Murder," *Journal of Applied Economics*, 7:163–93. Naci Mocan and Kaj Gittings (2010), "The Impact of Incentives on Human Behavior: Can We Make it Disappear? The Case of the Death Penalty," in Rafael Di Tella, Sebastian Edwards, and Ernesto Schargrotsky, eds., *The Economics of Crime: Lessons for and from Latin America*, University of Chicago Press. Steven S. Cuellar and Kyle Buehring (2010), "Does Capital Punishment Have a Deterrence Effect on the Murder Rate? Issues and Evidence," Sonoma State University.
- 118 *demonstrate it doesn't* Jeffrey Fagan (2005), "Deterrence and the Death Penalty: A Critical Review of New Evidence," Testimony to the New York State Assembly Standing Committee on Codes, Assembly Standing Committee on Judiciary and Assembly Standing Committee on Correction, Hearings on the Future of Capital Punishment in the State of New York, Jan 21, 2005. John J. Donahue and Justin Wolfers (2005), "Uses and Abuses of Empirical Evidence in the Death Penalty Debate," *Stanford Law Review*, 58:791–845. John J. Donahue and Justin Wolfers (2009), "Estimating the Impact of the Death Penalty on Murder," *American Law & Economics Review*, 11:249–309. Michael Radelet and Traci Lacock (2009), "Do Executions Lower Homicide Rates? The Views of Leading Criminologists," *The Journal of Criminal Law & Criminology*, 99:489–508.
- 118 *probability of punishment* Gary S. Becker (1968), "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76:169–217. Jeffrey Grogger (1991), "Certainty vs. Severity of Punishment," *Economic Inquiry*, 29:297–309. Lixuan

- Zhang, Wayne W. Smith, and William C. McDowell (2009), "Examining Digital Piracy: Self-Control, Punishment, and Self-Efficacy," *Information Resources Management Journal*, 22:24–44.
- 118 *societal causes of crime* Peter W. English and Bruce D. Sales (2005), *More Than the Law: Behavioral and Social Facts in Legal Decision Making*, American Psychological Association.
- 119 *Laws can have loopholes* Saul Levmore (2010), "Ambiguous Statutes," *The University of Chicago Law Review*, 77:1073–89.
- 119 *Google pays* Jesse Drucker (21 Oct 2010), "Google 2.4% Rate Shows how \$60 Billion Lost to Tax Loopholes," Bloomberg.
- 119 *large paper mills* Steven Mufson (28 Mar 2009), "Papermakers Dig Deep in Highway Bill To Hit Gold," *The Washington Post*. Jad Mouawad and Clifford Krauss (19 Apr 2009), "Lawmakers May Limit Paper Mills' Windfall," *New York Times. Wall Street Journal* (30 Jun 2009), "The Black Liquor War."
- 119 *video games* David Kochieniewski (10 Sep 2011), "Rich Tax Breaks Bolster Makers of Video Games," *New York Times*.
- 119 *Japan exploits* Eldon V.C. Greenberg, Paul S. Hoff, and Michael I. Goulding (2002), "Japan's Whale Research Program and International Law," *California Western International Law Journal*, 32:151–209. Reuben B. Ackerman (2002), "Japanese Whaling in the Pacific Ocean: Defiance of International Whaling Norms in the Name of 'Scientific Research,' Culture, and Tradition," *Boston College International & Comparative Law Review*, 323–42.
- 120 *Law of the Sea* Marcos A. Orellana (2004), "The Law on Highly Migratory Fish Stocks: ITLOS Jurisprudence in Context," *Golden Gate University Law Review*, 34:459–95.
- 120 *Security Council Resolution 1441* Michael Byers (2004), "Agreeing to Disagree: Security Council Resolution 1441 and Intentional Ambiguity," *Global Governance*, 10:165–86.
- 120 *Judge Gordon Hewart* Gordon Hewart (1923), *Rex v. Sussex Justices, Ex parte McCarthy*, King's Bench Division, *All England Law Reports*: 233.
- 120 *Invasive species* Jason Van Driesche and Roy Van Driesche (2000), *Nature Out of Place: Biological Invasions in the Global Age*, Island Press.
- 120 *ecological disaster* Daniel Simberloff (2001) "Biological Invasions: How Are They Affecting Us, and What Can We Do About Them?" *Western North American Naturalist*, 61:308–15.
- 120 *silver carp* M. Freeze and S. Henderson (1982) "Distribution and Status of the Big-head Carp and Silver Carp in Arkansas," *North American Journal of Fisheries Management*, 2:197–200.
- 120 *pregnant brown tree snake* Thomas H. Fritts and Gordon H. Rodda (1998), "The Role of Introduced Species in the Degradation of Island Ecosystems: A Cade History of Guam," *Annual Review of Ecology & Systematics*, 29:113–40.

- 120 *zebra mussel larvae* David L. Strayer (2010), “Alien Species in Fresh Waters: Ecological Effects, Interactions With Other Stressors, and Prospects for the Future,” *Freshwater Biology*, 55:152–74.
- 120 *Monk parakeets* Michael A. Russello, Michael L. Avery, and Timothy F. Wright (2008), “Genetic Evidence Links Invasive Monk Parakeet Populations in the United States to the International Pet Trade,” *BMC Evolutionary Biology*, 8:217.
- 120 *crop damage* Eric A. Tillman (2000), “Bird Damage to Tropical Fruit in South Florida,” *The Ninth Wildlife Damage Management Conference Proceedings*.
- 120 *also cause fires* Michael L. Avery (2002), “Monk Parakeet Management at Electric Utility Facilities in South Florida,” Internet Center for Wildlife Damage Management, USDA National Wildlife Research Center Staff Publications. Dan Mangan and Erin Calabrese (21 Apr 2009), “Fowl Play in Qns.: Blackout Parrots’ Nests Cut Electricity,” *New York Post*.
- 121 *state laws vary wildly* Michael A. Russello, Michael L. Avery, and Timothy F. Wright (2008), “Genetic Evidence Links Invasive Monk Parakeet Populations in the United States to the International Pet Trade,” *BMC Evolutionary Biology*, 8:217.

Chapter 9 Notes

- 270 *similar system in Sweden* Heather Monasky (2011), “On Comprehensive Prostitution Reform: Criminalizing the Trafficker and the Trick, But Not the Victim—Sweden’s Sexköpslagen in America,” *William Mitchell Law Review*, 37:101–56.
- 271 *driving skill is better* Justin Kruger and David Dunning (1999), “Unskilled and Unaware of It: How Difficulties in Recognizing One’s Own Incompetence Lead to Inflated Self-Assessments,” *Journal of Personality & Social Psychology*, 77:1121–34. Corey L. Guenther and Mark D. Alicke (2010), “Deconstructing the Better-Than-Average Effect,” *Journal of Personality & Social Psychology*, 99:755–70.
- 271 *Electronic filing* U.S. Internal Revenue Service (8 Jul 2009), “Update on Reducing the Federal Tax Gap and Improving Voluntary Compliance,” Department of the Treasury. Barbara Weltman (24 Jan 2011), “Ten Ways to Avoid a Tax Audit,” *Wall Street Journal*.
- 271 *Students perform better* Uri Gneezy and Aldo Rustichini (2000), “Pay Enough or Don’t Pay at All,” *Quarterly Journal of Economics*, 115:791–810.
- 271 *salary bonuses* Randall Eberts, Kevin Hollenbeck, and Joe Stone (2002), “Teacher Performance Incentives and Student Outcomes,” *Journal of Human Resources*, 37:913–27.
- 271 *Ostrom’s original rules are* Elinor Ostrom (1980), *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press.
- 272 *Jeremy Bentham* Jeremy Bentham (1789), *An Introduction to the Principles of Morals and Legislation*.
- 272 *Gary Becker* Gary Becker (1996), “The Economic Way of Looking at Human Behavior: The Nobel Lecture,” *Journal of Political Economy*, 101:385–409.

- 272 *increasing the probability* Mark Kleinman (2009), *When Brute Force Fails: How to Have Less Crime and Less Punishment*, Princeton University Press.
- 272 *conflicting evidence* Nina Mazar and Dan Ariely (2006), "Dishonesty in Everyday Life and Its Policy Implications," *Journal of Public Policy & Marketing*, 25:117–26.
- 272 *two Irish subsidiaries* Joseph B. Darby III and Kelsey Lemaster (15 May 2007), "Double Irish More than Doubles the Tax Savings," *Practical US/International Tax Strategies*, 11:2,11–16.
- 272 *bigger one opened up* Todd Neeley (28 Mar 2011), "Pulp, Paper Companies Amend Tax Returns Actions Could Cost Taxpayers Billions of Dollars," *DTN: The Progressive Farmer*.

Chapter 10

- 123 *making towels souvenirs* Lynn Yaeger (2 Jun 2009), "Stealing Hotel Amenities: Right or Wrong?" MSNBC.
- 124 *reduced towel theft* Sara J. Welch (2011), "Gee, How Did That Towel End Up in My Suitcase," *New York Times* blog.
- 125 *patterns of fraud* Yufeng Kou, Chang-Tien Lu, Sirirat Sirwongwattana, and Yo-Ping Huang (2004), "Survey of Fraud Detection Techniques," *IEEE International Conference on Networking, Sensing and Control*, vol. 2, 749–54. Jon T.S. Quah and M. Sriganesh (2008), "Real-Time Credit Card Fraud Detection Using Computational Intelligence," *Expert Systems with Applications*, 35:1721–32. Edwin Raj Benson and A. Annie Portia (2011), "Analysis on Credit Card Fraud Detection Methods," *2011 International Conference on Computer, Communication and Electrical Technology*, 152–6. Arturo Elias, Alberto Ochoa-Zezzatti, Alejandro Padilla, and Julio Ponce (2011), "Outlier Analysis for Plastic Card Fraud Detection: A Hybridized and Multi-Objective Approach," *Hybrid Artificial Intelligence Systems, Lecture Notes in Computer Science*, Springer-Verlag, 6679:1–9. Amlan Kundu, Shamik Sural, and Arun K. Majumdar (2006), "Two-Stage Credit Card Fraud Detection Using Sequence Alignment," *Information Systems Security, Lecture Notes in Computer Science*, Springer-Verlag, 4332:260–75.
- 126 *predictive policing programs* Martin B. Short, Maria R. D'Orsogna, Virginia B. Pasour, George E. Tita, P. Jeffrey Brantingham, Andrea L. Bertozzi, and Lincoln B. Chayes (2008), "A Statistical Model of Criminal Behavior," *Mathematical Models and Methods in Applied Sciences*, 18 (Supplement):1249–67. Beth Pearsall (2010), "Predictive Policing: The Future of Law Enforcement?" *NIJ Journal*, 266:16–9. Nancy Murray (2011), "Profiling in the Age of Total Information Awareness," *Race & Class*, 51:3–24.
- 127 *Timothy McVeigh's van* Associated Press (28 Sep 2009), "Attorney: Oklahoma City Bombing Tapes Appear Edited," *Oklahoman*.
- 127 *reduce car theft* Ian Ayres and Steven Levitt (1998), "Measuring Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack," *Quarterly*

- Journal of Economics*, 113:43–77. Marco Gonzalez-Navarro (2008), “Deterrence and Displacement in Auto Theft,” Center for Economic Policy Studies Working Paper.
- 128 *confidence in vaccines* Heidi Larson, Louis Z. Cooper, Juhani Eskola, Samuel L. Katz, and Scott Ratzan (6 Aug 2011), “Addressing the Vaccine Confidence Gap,” *The Lancet*, 378:526–35.
- 129 *new version of Monopoly* Stephanie Clifford (16 Feb 2011), “No Dice, No Money, No Cheating. Are You Sure This is Monopoly?” *New York Times*.
- 130 *Crime Prevention Techniques* Ronald V. Clarke and John Eck (2003), *Become a Problem Solving Crime Analyst*, Willan Publishing. Derek B. Cornish and Ronald V. Clarke (2003) “Opportunities, Precipitators and Criminal Decisions,” in *Theory for Practice in Situational Crime Prevention, Crime Prevention Studies, Vol. 16*, Criminal Justice Press.
- 131 *enforcing anti-doping* David R. Mottram, ed. (2011), *Drugs in Sport, Fifth Edition*, Taylor & Francis.
- 131 *Alex Zülle* John Hoberman (20 Sep 2007), “Dopers on Wheels: The Tour’s Sorry History,” MSNBC.
- 132 *almost a third* Patrick Mignon (2003), “The Tour de France and the Doping Issue,” *International Journal of the History of Sport*, 20:227–45.
- 132 *never intended* Jamey Keaten (20 Jul 2007), “At Tour de France, Asthma Inhalers Can Be a Rider’s Best Friend,” *USA Today*.
- 132 *stripped of her gold medal* *Sports Illustrated* (28 Sep 2000), “Arbitrators Uphold Decision to Strip Raducan of Gold,” CNN.
- 132 *install security cameras* Bruce Schneier (26 Jun 2008), “CCTV Doesn’t Keep Us Safe, Yet the Cameras Are Everywhere,” *The Guardian*.
- 132 *full-body scanners* Bruce Schneier (2 Dec 2010), “Why the TSA Can’t Back Down,” *The Atlantic*.
- 132 *military spends billions* *New York Times* (2 Oct 1984), “A Lemon the Size of the Pentagon,” *New York Times*. Michael Ditton (Aug 1988), “The DIVAD Procurement: A Weapon System Case Study,” *The Army Lawyer*, 3–9. Irene Willhite (Mar 2002), “40-mm Division Air Defense Gun: DIVAD (Sgt. York),” *Cold War Times*, 15–22.
- 132 *better slugs* Ronald V. Clarke, Ronald P. Cody, and Mangai Natarajan (1994), “Subway Slugs: Tracking Displacement on the London Underground,” *British Journal of Criminology*, 34:122–38.
- 132 *door locks and safes* Mark Weber Thomas (2000), *Locks, Safes and Security*, Second Edition, Charles Thomas Publisher, Ltd. Matt Blaze (2003), “Cryptography and Physical Security: Rights Amplification in Master-Keyed Mechanical Locks,” IEEE Security and Privacy. Matt Blaze (2004), “Safecracking for the Computer Scientist,” *U. Penn CIS Department Technical Report*.
- 133 *fax machines* Bruce Schneier (29 May 2008), “Why Do We Accept Signatures by Fax?” *Wired News*.
- 133 *Electronic voting machines* Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten (2006), “Security Analysis of the Diebold AccuVote-TS Voting Machine,” Center for

- Information Technology Policy, Princeton University. Adam Aviv, Pavol Cerny, Sandy Clark, Eric Cronin, Gaurav Shah, Micah Sherr, and Matt Blaze (2008), “Security Evaluation of ES&S Voting Machines and Election Management System,” USENIX/ACCURATE Electronic Voting Technology Workshop (EVT ’08), paper presented at “USENIX/ACCURATE. Hari K. Prasad, J. Alex Halderman, Rop Gonggrijp, Scott Wolchok, Eric Wustrow, Arun Kankipati, Sai Krishna Sakhamuri, and Vasavya Yagati (2010), “Security Analysis of India’s Electronic Voting Machines,” *Proceedings of the 17th ACM Conference on Computer and Communications Security*, Association of Computing Machinery. Doug Jones and Barbara Simons (2012), *Broken Ballots: Will Your Vote Count in the Electronic Age?* CSLI Publications.
- 133 *electronics in cars* Stephen Checkoway, Damon McCoy, Brian Kantor, Danny Anderson, Hovav Shacham, Stefan Savage, Karl Koscher, Alexei Czeskis, Franziska Roesner, and Tadayoshi Kohno (2011), “Comprehensive Experimental Analyses of Automotive Attack Surfaces,” USENIX Security.
- 133 *Networked medical devices* William H. Maisel and Tadayoshi Kohno (2010), “Improving the Security and Privacy of Implantable Medical Devices,” *New England Journal of Medicine*, 362:1164–6.
- 133 *cursive is not being taught* Katie Zezima (28 Apr 2011), “The Case for Cursive,” *New York Times*.
- 134 *read trust signals* Jens Riegelsberger, M. Angela Sasse, and John D. McCarthy (2005), “The Mechanics of Trust: A Framework for Research and Design,” *International Journal of Human-Computer Studies*, 62:381–422.
- 134 *harder-to-forge ID* Bruce Schneier (2007), “Will REAL ID Actually Make Us Safer? An Examination of Privacy and Civil Liberties Concerns,” testimony to the Senate Judiciary Committee.
- 135 *spent about \$1 trillion* John Mueller and Mark Stewart (2011), *Terrorism, Security, and Money: Balancing the Risks, Benefits, and Costs of Homeland Security*, Oxford University Press.
- 135 *doesn’t count the wars* Linda J. Bilmes and Joseph E. Stiglitz (2008), *The Three Trillion Dollar War: The True Cost of the Iraq Conflict*, W.W. Norton & Co.
- 136 *cost-benefit analysis* Gary Becker (1968), “Crime and Punishment: An Economic Approach,” *The Journal of Political Economy*, 76:169–217.
- 136 *passion overrides rationality* Howard Engel (2001), *Crimes of Passion: An Unblinking Look at Murderous Love*, Firefly Books.

Chapter 10 Notes

- 272 *someone was sentenced* Barbara de Lollis (15 Sep 2010), “Woman Faces Three Months in Jail for Stealing Hotel Towels,” *USA Today*.
- 272 *doping in professional Scrabble* Stefan Fatiss (2001), *Word Freak: Heartbreak, Triumph, Genius, and Obsession in the World of Competitive Scrabble Players*, Houghton Mifflin Harcourt.

- 272 *assays were developed* John H. Lewis (1979), "A Routine Screening Method for the Major Metabolite of Methyl Phenidate in the Urine of Greyhounds," *Forensic Science International*, 14:221–7. Orisa J. Igwe and Jerry W. Blake (1981), "Gas/Liquid Chromatographic Analysis of Pemoline in Biological Fluids Using Electron Capture Detection," *Journal of Chromatographic Science*, 19:617–24.
- 273 *EPO use became* Robin Parisotto, Christopher J. Gore, Kerry R. Emslie, Michael J. Ashenden, Carlo Brugnara. Chris Howe, David T. Martin, Graham J. Trout, and Allan G. Hahn (2000), "A Novel Method Utilising Markers of Altered Erythropoiesis for the Detection of Recombinant Human Erythropoietin Abuse in Athletes," *Haematologica*, 85:564–72. Caroline K. Hatton (2007), "Beyond Sports Doping Headlines: The Science of Laboratory Tests for Performance-Enhancing Drugs," *Pediatric Clinics of North America*, 54:713–33. Rafael Maia de Almeida Bento, Lucia Menezes Pinto Damasceno, and Francisco Radler de Aquino Neto (2003), "Recombinant Human Erythropoietin in Sports: A Review," *Revista Brasileira de Medicina do Esporte*, 9:181–90.
- 273 *analogues of analogues* Andreas Breidbach, Don H. Catlin, Gary A. Green, Inna Tregub, Henry Truong, and Jeffrey Gorzek (2003), "Detection of Recombinant Human Erythropoietin in Urine by Isoelectric Focusing," *Clinical Chemistry*, 49:901–7.
- 273 *new EPO replacement* Don H. Catlin, George A. Maylin, Sabrina Benchaar, Sandra M. Neades, Miranda S. Timmons, and Kenneth H. McKeever (2009), "Detection of Pegylated Epoetin in Horse Plasma after Intravenous Administration," *Comparative Exercise Physiology*, 6:105–8.
- 273 *first designer steroid* Henry K. Lee (4 Aug 2006), "Inventor of Clear Steroid Gets Three Months in Prison," *San Francisco Chronicle*.
- 273 *Its fingerprint* Don H. Catlin, Brian D. Ahrens, and Yulia Kucherova (2002), "Detection of Norbolethone, an Anabolic Steroid Never Marketed, in Athletes' Urine," *Rapid Communications in Mass Spectrometry*, 16:1273–5.
- 273 *scenario was replayed* Don H. Catlin, Michael H. Sekera, Brian D. Ahrens, Borislav Starcevic, Yu-Chen Chang, and Caroline K. Hatton (2004), "Tetrahydrogestrinone: Discovery, Synthesis, and Detection in Urine," *Rapid Communications in Mass Spectrometry*, 18:1245–9. Michael H. Sekera, Brian D. Ahrens, Yu-Chen Chang, Borislav Starcevic, Costas Georgakopoulos, and Don H. Catlin (2005), "Another Designer Steroid: Discovery, Synthesis, and Detection of Madol in Urine," *Rapid Communications in Mass Spectrometry*, 19:781–4.
- 273 *flow cytofluorometry* Margaret Nelson, Michael Ashenden, Mark Langshaw, and Hazel Popp (2002), "Detection of Homologous Blood Transfusion by Flow Cytometry: A Deterrent against Blood Doping," *Haematologica*, 87:881. Patricia A. Arndt and Belinda M. Kumpel (2008), "Blood Doping in Athletes: Detection of Allogeneic Blood Transfusions by Flow Cytofluorometry," *American Journal of Hematology*, 83:657–67.
- 273 *gene doping* H. Lee Sweeney (Jul 2004), "Gene Doping," *Scientific American*, 291:63–9. Gary R. Gaffney and Robin Parisotto (2007), "Gene Doping: A Review of Performance-Enhancing Genetics," *Pediatric Clinics of North America*, 54:807–22.

- 273 *adulterated supplements* Jeff Jones (4 Dec 2002), “Tainted Supplements: Positive or Not?” *Cycling News*.
- 273 *Chinese herbal product* Patricia L. Eachus (1996), “Positive Drug Screen for Benzodiazepine Due to a Chinese Herbal Product,” *Journal of Athletic Training*, 31:165–6.
- 273 *urine test for EPO* Monique Beullens, Joris R. Delanghe, and Mathieu Bollen (2006), “False-Positive Detection of Recombinant Human Erythropoietin in Urine Following Strenuous Physical Exercise,” *Blood*, 107:4711–13. Joris R. Delanghe, Mathieu Bollen, and Monique Beullens (2007), “Testing for Recombinant Erythropoietin,” *American Journal of Hematology*, 83:237–41.
- 273 *hotly contested* Don Catlin, Gary Green, Michael Sekera, Paul Scott, and Borislav Starcevic (2006), “False-Positive EPO Test Concerns Unfounded,” *Blood*, 108:1778. Françoise Lasne (2006), “No Doubt About the Validity of the Urine Test for Detection of Recombinant Human Erythropoietin,” *Blood*, 108:1778–9.
- 273 *Rapid-screen immunoassays* Nancy C. Brahm, Lynn L. Yeager, Mark D. Fox, Kevin C. Farmer, and Tony A. Palmer (2010), “Commonly Prescribed Medications and Potential False-Positive Urine Drug Screens,” *American Journal of Health Systems & Pharmacy*, 67:1344–50.
- 273 *Vicks Vapor Inhaler* Telegraph (21 Mar 2002), “British Skier Stripped of Bronze.”
- 274 *appearance of a website* Rachna Dhamija, J.D. Tygar, and Marti Hearst (2006), “Why Phishing Works,” *CHI '06 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- 274 *least costly measures* Tom LaTourette, David R. Howell, David E. Mosher, and John MacDonald (2006), *Reducing Terrorism Risk at Shopping Centers: An Analysis of Potential Security Options*, RAND Technical Report.

Chapter 11

- 139 *Rhonda Breard embezzled* Kevin Opsahl (11 Aug 2010), “Kirkland Broker Rhonda Breard to Spend Nearly Seven Years in Prison for Fraud,” *Seattle Post-Intelligencer*.
- 139 *Envy can motivate* Simone Moran and Maurice E. Schweitzer (2008), “When Better is Worse: Envy and the Use of Deception,” *Negotiation & Conflict Management Research*, 1:3–29.
- 139 *things out of anger* Yang Wang, Saranga Komanduri, Pedro Giovanni Leon, Gregory Norcie, Alessandro Acquisti, and Lorrie Faith Cranor (2011), “‘I Regretted the Minute I Pressed Share’: A Qualitative Study of Regrets on Facebook,” *Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS2011)*.
- 140 *heinous behavior* Ervin Staub (1992), *The Roots of Evil: The Origins of Genocide and Other Group Violence*, Cambridge University Press.
- 140 *extreme situations* David DeSteno and Piercarlo Valdesolo (2011), *Out of Character: Surprising Truths About the Liar, Cheat, Sinner (and Saint) Lurking in All of Us*, Crown Archetype.
- 140 *lethal altruists* Adolf Tobeña (2009), “Lethal Altruists: Itineraries along the Dark Outskirts of Moralistic Prosociality,” *Annals of the New York Academy of Sciences*, 1167:5–15.

- 140 *Abraham Maslow* Abraham H. Maslow (1943), "A Theory of Human Motivation," *Psychological Review*, 50:370–96.
- 140 *Rational Choice Theory* Gary S. Becker (1976), *The Economic Approach to Human Behavior*, University of Chicago Press.
- 141 *Bounded Rationality* Bryan D. Jones (1999), "Bounded Rationality," *Annual Review of Political Science*, 2:297–321. Daniel Kahneman (2003), "Maps of Bounded Rationality: Psychology for Behavioral Economics," *The American Economic Review*, 93:1449–75. Gerd Gigerenzer (2007), *Gut Feelings: The Intelligence of the Unconscious*, Viking Adult. Dan Ariely (2008), *Predictably Irrational: The Hidden Forces that Shape our Decisions*, Harper Perennial. Ori Brafman and Rom Brafman (2008), *Sway: The Irresistible Pull of Irrational Behavior*, Crown Business. Shankar Vedantam (2010), *The Hidden Brain: How Our Unconscious Minds Elect Presidents, Control Markets, Wage Wars, and Save our Lives*, Spiegel & Grau. Daniel Kahneman (2011), *Thinking Fast and Slow*, Farrar, Straus and Giroux. Duncan J. Watts (2011), *Everything is Obvious: *Once You Know the Answer*, Crown Business.
- 141 *trade-off process* Bruce Schneier (2008), "The Psychology of Security," *AFRICACRYPT 2008*, Springer-Verlag, 50–79. Daniel Gardner (2008), *The Science of Fear: Why We Fear the Things We Shouldn't—and Put Ourselves in Greater Danger*, Dutton Adult. Paul Slovic (2000), *The Perception of Risk*, Earthscan Publications Ltd. Barry Glassner (1999), *The Culture of Fear: Why Americans Are Afraid of the Wrong Things*, Basic Books.
- 141 *different type of trade-off* Lawrence Kohlberg (1981), *Essays on Moral Development, Vol. I: The Philosophy of Moral Development*, Harper & Row.
- 142 *Kohlberg's Stages* University of Canberra Faculty of Education (2011), "University of Canberra Graduate Diploma of Education E-reserve Lecture Notes: Ed Foundations, Module B, Part 3."
- 143 *Emily Dickinson* Emily Dickinson (1862), "The soul selects her own society," poem 303.
- 144 *morality to be central* William Damon (1984), "Self-Understanding and Moral Development from Childhood to Adolescence," in William M. Kurtines and Jacob L. Gewirtz, eds., *Morality, Moral Behavior, and Moral Development*, John Wiley & Sons.
- 144 *spiritual geniuses* René Girard (1999), *Je Vois Satan Tomber Comme l'Éclair*, Grasset English translation (2001), *I See Satan Fall Like Lightning*, Orbis Books.
- 144 *Ralph Waldo Emerson* Ralph Waldo Emerson (1841), "Self-Reliance," in *Essays: First Series*.
- 144 *Henry David Thoreau* Henry David Thoreau (1852), *Walden*.
- 144 *Laurel Thatcher Ulrich* Laurel T. Ulrich (2007), *Well-Behaved Women Seldom Make History*, Alfred A. Knopf.
- 144 *Socrates's morals* Plato (360 BCE), *Crito*, tr. Benjamin Jowett, MIT Classics Archive. Plato (360 BCE), *Phaedo*, tr. Benjamin Jowett, MIT Classics Archive. Anthony D'Amato (1976), "Obligation to Obey the Law: A Study of the Death of Socrates," *California Law Review*, 49:1079–1108.

- 144 *Misunderstanding the defector* Bruce Schneier (2003), *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus Books, Chapter 5.
- 145 *there's no guarantee* Laurie P. Cohen (29 Nov 2004), "Split Decisions: Federal Cases Show Big Gap in Reward For Cooperation," *Wall Street Journal*.
- 145 *"Stop Snitching" campaign* Rick Hampson (28 Mar 2006), "Anti-Snitch Campaign Riles Police, Prosecutors," *USA Today*. Rick Frei (2010), "Witness Intimidation and the Snitching Project," written testimony submitted to the Subcommittee on Drugs and Crime, U.S. Senate Committee on the Judiciary.
- 147 *Con artists try* David Maurer (1940), *The Big Con: The Story of the Confidence Man*, Bobbs Merrill.
- 147 *Fake anti-virus software* Brian Krebs (3 Aug 2011), "Fake Antivirus Industry Down, But Not Out," *Krebs on Security*.
- 147 *Internet money laundering* Mitchell Zuckoff (15 May 2005), "Annals of Crime: The Perfect Mark," *The New Yorker*, 36–42.
- 148 *doctrine of necessity* Leslie Wolf-Phillips (1979), "Constitutional Legitimacy: A Study of the Doctrine of Necessity," *Third World Quarterly*, 1:99–133.
- 148 *competing interests* H.E. Mason, ed. (1996), *Moral Dilemmas and Moral Theory*, Oxford University Press.
- 149 *publicly confessed* Sean O'Callaghan (27 Jan 1997), "The Lying: A Former Terrorist Describes His Life in the IRA," *National Review*.
- 149 *Defecting from the IRA* Sean O'Callaghan (1999), *The Informer: The Real Life Story of One Man's War Against Terrorism*, Corgi.
- 150 *Ulysses S. Grant* Louis A. Coolidge (1917), *Ulysses S. Grant*, Houghton Mifflin. Lawrence M. Salinger (2005), *Encyclopedia of White-Collar & Corporate Crime*, Sage Publications.
- 152 *interlocking directorate* Mark S. Mizruchi (1996), "What Do Interlocks Do? An Analysis, Critique, and Assessment of Research on Interlocking Directorates," *Annual Review of Sociology*, 22:271–98.
- 153 *Larry Froistad confessed* Debra J. Saunders (10 May 1998), "Spilled Milk," *San Francisco Chronicle*. Lisa DiCarlo (11 May 1998), "Murder, She Read," *New York Magazine*.
- 153 *Police solidarity* Jerome H. Skolnick (2002), "Corruption and the Blue Code of Silence," *Police Practice & Research*, 3:7–19. Louise Westmarland (2005), "Police Ethics and Integrity: Breaking the Blue Code of Silence," *Policing & Society*, 15:145–65. Gary R. Rothwell and J. Norman Baldwin, "Whistle-Blowing and the Code of Conduct in Police Agencies," *Crime & Delinquency*, 53:8–10. Barry Wright (2010), "Civilianising the 'Blue Code'? An Examination of Attitudes to Misconduct in the Police Extended Family," *International Journal of Police Science & Management*, 12:339–56.
- 153 *policeman beating* Jayme Poisson (10 Jun 2011), "Cop Unmasked: Meet Const. Glenn Weddell, the Officer ID'd in G20 Assault," *The Toronto Star*.

Chapter 11 Notes

- 274 [Robert A. Heinlein](#) Robert A. Heinlein (1988), *To Sail Beyond the Sunset*, G.P. Putnam's Sons, 38–9.
- 275 [levels of moral meaning](#) Kenneth R. Melchin (1998), *Living with Other People: An Introduction to the Christian Ethics Based on Bernard Lonergan*, The Liturgical Press, 56–60.
- 275 [good summary](#) William C. Crain (1985), “Kohlberg’s Stages of Moral Development,” Chapter 7 of *Theories of Development: Concepts and Applications*, Prentice-Hall, 118–36.
- 275 [Social identity theory](#) Henri Tajfel (1974), “Social Identity and Intergroup Behaviour,” *Social Science Information*, 13:65–93. Naomi Ellemers, Russell Spears, and Bertjan Doosje (2001), “Self and Social Identity,” *Annual Review of Psychology*, 53:161–86.
- 275 [immolated themselves](#) Michael Biggs (2005), “Dying Without Killing: Self-Immolations, 1963–2002,” in Diego Gambetta, ed., *Making Sense of Suicide Missions*, Oxford University Press.
- 275 [Brian Christian](#) Brian Christian (2011), *The Most Human Human: What Talking with Computers Teaches Us About What it Means to Be Alive*, Doubleday.
- 276 [useful mechanism](#) Gordon Tullock (1967), “The Prisoner’s Dilemma and Mutual Trust,” *Ethics*, 77:229–30.
- 276 [Stop Snitching campaign](#) Rachael Woldoff and Karen G. Weiss (2010), “Stop Snitchin’: Exploring Definitions of the Snitch and Implications for Urban Black Communities,” *Journal of Criminal Justice & Popular Culture*, 17:184–223. Bret D. Asbury (2011), “Anti-Snitching Norms and Community Loyalty,” *Oregon Law Review*, 89:1257–312.
- 276 [Nepotism is making](#) Helen Thomas (16 Aug 2001), “Bush Keeps It All in the Family,” *Seattle Post-Intelligencer*.
- 276 [Senator Frank Murkowski](#) Jonathan Turley (13 Jan 2003), “Public Payroll: A Family Affair,” *Los Angeles Times*.
- 276 [Representative Richard Pombo](#) League of Conservation Voters (2005), “Rep. Richard Pombo’s Family & Friends Network.”
- 276 [Representative Eddie Bernice Johnson](#) Todd J. Gillman and Christy Hoppe (28 Aug 2010), “Rep. Eddie Bernice Johnson Violated Rules, Steered Scholarships to Relatives,” *Dallas Morning News*.
- 276 [Even Bernie Sanders](#) Vermont Guardian (21 Apr 2005), “Nepotism Crosses Party Lines,” *Vermont Guardian*.
- 276 [one in six staffers](#) Christine MacDonald (6 May 2011), “Nepotism Rampant at Detroit Libraries: 1 in 6 Staffers Have Relatives Who Work in Strapped Department,” *The Detroit News*.
- 276 [Rupert Murdoch’s News Corp](#) Mark Sweney (17 Mar 2011), “Rupert Murdoch’s News Corp Sued Over ‘Nepotism’ in Buying His Daughter’s Firm: Investors Allege Group Is Overpaying in \$675m Deal to Acquire Elisabeth Murdoch’s TV Production Business Shine,” *The Guardian*.

- 276 *Many states have policies* New York State Bar Association Committee on Professional Ethics (1975), "Topic: Conflict of Interests; Lawyer Spouses as Assistant District Attorney and Assistant Public Defender," Opinion #409—8/28/75 (69–75). State Bar of Montana (1988), "May Spouses Serve as Opposing Counsel, Actively Engaging in the Same Suit or Matter on Opposite Sides?" Ethics Opinion 950407. Sandra Carlin Guin (1985), "Marital Status as a Reason for Attorney-Spouse Disqualification," *Journal of the Legal Profession*, 9:181–8.

Chapter 12

- 158 *Carol Braun was* Jed Block (2004), *Betrayal*, Goodwill Industries of North Central Wisconsin, Inc.
- 159 *principal-agent problem* Kathleen M. Eisenhardt (1989), "Agency Theory: An Assessment and Review," *The Academy of Management Review*, 14:57–74. John M. Darley (2010), "Constructive and Destructive Obedience: A Taxonomy of Principal-Agent Relationships," *Journal of Social Issues*, 41:124–54.
- 160 *corporate looting* George A. Akerlof, Paul M. Romer, Robert E. Hall, and N. Gregory Mankiw (1993), "Looting: The Economic Underworld of Bankruptcy for Profit," *Brookings Papers on Economic Activity*, 1993(2): 1–73.
- 160 *Sambo's restaurants* Charles Bernstein (1984), *Sambo's: Only a Fraction of the Action: The Inside Story of a Restaurant Empire's Rise and Fall*, National Literary Guild.
- 163 *moral considerations* James A. Waters (1978), "Catch 205: Corporate Morality as an Organizational Phenomenon," *Organizational Dynamics*, 6:3–19. James Weber (1990), "Managers' Moral Reasoning: Assessing Their Responses to Three Moral Dilemmas," *Human Relations*, 43:687–702. Thomas M. Jones and Lori Versteegen Ryan (1998), "The Effect of Organizational Forces on Individual Morality: Judgment, Moral Approval, and Behavior," *Business Ethics Quarterly*, 8:431–45. J. Stuart Bunderson (2001), "Normal Injustices and Morality in Complex Organizations," *Journal of Business Ethics*, 33:181–90. Linda K. Treviño, Gary R. Weaver, and Scott J. Reynolds (2006), "Behavioral Ethics in Organizations: A Review," *Journal of Management*, 32:951–90. Nicole Andreoli and Joel Lefkowitz (2009), "Individual and Organizational Antecedents of Misconduct in Organizations," *Journal of Business Ethics*, 85:309–32. William Carbonaro and Jessica L. Collett (2010), "Pushing the Man: Moral Judgments in Group Settings," Paper presented at the Annual Meeting of the American Sociological Association, San Francisco, California.
- 163 *many reasons for this* James Weber (1990), "Managers' Moral Reasoning: Assessing Their Responses to Three Moral Dilemmas," *Human Relations*, 43:687–702. Maurice Punch (2000), "Suite Violence: Why Managers Murder and Corporations Kill," *Crime, Law & Social Change*, 33:243–80. Yoav Vardi (2001), "The Effects of Organizational and Ethical Climates on Misconduct at Work," *Journal of Business Ethics*, 29:325–37.

- 163 *Beech-Nut Corporation* Leonard Buder (14 Nov 1987), "Beech-Nut is Fined \$2 Million for Sale of Fake Apple Juice," *New York Times*. James Traub (24 Jul 1988), "Into the Mouths of Babes," *New York Times Magazine*.
- 163 *being victimized* BishopAccountability.org (2011), "What Percent of Priests Were Accused?"
- 163 *actual number of victims* BishopAccountability.org (2011), "Data on the Crisis: The Human Toll." BishopAccountability.org (2011), "A Documentary History of the Crisis."
- 163 *retained accused priests* Brooks Egerton and Reese Dunklin (12 Jun 2002), "Two-Thirds of Bishops Let Accused Priests Work," *Dallas Morning News*.
- 163 *Vatican's exhortations* Office of the Holy See (1962), "Instruction on the Manner of Proceeding in Cases of Solicitation: The Decree Crimen Sollicitationis," Vatican Press. Catholics for a Free Choice (2002), "The Holy See and the Convention on the Rights of the Child: A Shadow Report," Catholics for a Free Choice. Antony Barnett (17 Aug 2003), "Vatican Told Bishops to Cover Up Sex Abuse: Expulsion Threat in Secret Documents," *The Guardian*. Associated Press (18 Jan 2011), "Vatican Letter Told Ireland's Catholic Bishops Not to Report Child Abuse," *The Guardian*.
- 164 *forced the victims* Richard Nangle (1 Apr 2002), "Deal Mandated Silence: Accuser Not Allowed to Discuss Abuse Allegations," *Worcester Telegraph & Gazette*. Stephen Gillers (2002), "Speak No Evil: Settlement Agreements Conditioned on Noncooperation Are Illegal and Unethical," *Hofstra Law Review*, 31:1–22.
- 165 *top WorldCom executives* Susan Pulliam and Deborah Soloman (8 Nov 2008), "How Three Unlikely Sleuths Exposed Fraud at WorldCom: Firm's Own Employees Sniffed Out Cryptic Clues and Followed Hunches," *Wall Street Journal*. Amanda Ripley (30 Dec 2002), "Cynthia Cooper: The Night Detective," *Time*. WorldCom (2003), "Report of Investigation by the Special Investigative Committee of the Board of Directors of WorldCom, Inc." Cynthia Cooper (2008), *Extraordinary Circumstances: The Journey of a Corporate Whistleblower*, John Wiley & Sons, Inc.
- 166 *corporate board members* Peter J. Henning (2009), "Board Dysfunction: Dealing with the Threat of Corporate Criminal Liability," Wayne State University Law School Legal Studies Research Paper.
- 167 *analyst Thomas Drake* Jane Mayer (23 May 2011), "The Secret Sharer: Is Thomas Drake an Enemy of the State?" *The New Yorker*. Glenn Greenwald (30 Jul 2011), "Obama's Whistleblower War Suffers Two Defeats," *Salon*. Thomas Drake and Jessalyn Radack (1 Aug 2011), "A Surprising War on Leaks Under Obama," *Philadelphia Inquirer*.
- 167 *1962 Vatican Instruction* Office of the Holy See (1962), "Instruction on the Manner of Proceeding in Cases of Solicitation: The Decree Crimen Sollicitationis," Vatican Press.
- 167 *drug was ineffective* Adrian Viens and Julian Savulescu (2004), "The Olivieri Symposium: Introduction to The Olivieri Symposium," *Journal of Medical Ethics*, 30:1–7. Rosamond Rhodes and J.J. Strain (2004), "The Olivieri Symposium: Whistleblowing

- in Academic Medicine,” *Journal of Medical Ethics*, 30:35–9. Thomas Faunce, Steve Bolsin, and Wei-Ping Chan (2004), “The Olivieri Symposium: Supporting Whistleblowers in Academic Medicine: Training and Respecting the Courage of Professional Conscience,” *Journal of Medical Ethics*, 30:40–3.
- 167 *exposed misconduct* Jaime Adame (2 Jun 2004), “Police Whistleblowers,” *Gotham Gazette*.
- 168 *military pilots defected* John Hooper and Ian Black (21 Feb 2011), “Libya Defectors: Pilots Told to Bomb Protesters Flee to Malta,” *The Guardian*.
- 169 *U.S.’s use of torture* Ian Fishback (28 Sep 2005), “A Matter of Honor,” *Washington Post*.
- 170 *Deepwater Horizon* Joe Nocera (18 Jun 2010), “BP Ignored the Omens of Disaster,” *New York Times*.
- 170 *employees of BP* National Commission on the Deepwater Horizon Oil Spill and Offshore Drilling (2011), *Deep Water: The Gulf Oil Disaster and the Future of Offshore Drilling, Final Report*, Government Printing Office. National Commission on the Deepwater Horizon Oil Spill and Offshore Drilling (2011), *Macondo: The Gulf Oil Disaster, Chief Counsel’s Report*, Government Printing Office.
- 170 *The “Big Dig,”* Andrea Estes (24 Jan 2008), “Big Dig Settlement Will Take Quick Hit,” *Boston Globe*. Jonathan Saltzman (15 Nov 2008), “Companies to Settle for \$26m in Tunnel Collapse,” *Boston Globe*. Gregory W. Sullivan (2005), “A Big Dig Cost Recovery Referral: Waterproofing Mismanagement by Bechtel/Parsons Brinckerhoff in East Boston,” Office of the Inspector General, Commonwealth of Massachusetts. David W. Fowler (2010), “The Big Dig: Lessons in Engineering and Ethics,” CAEE Annual Alumni Ethics Seminar, Austin, Texas.
- 171 *conflict between the* Gregory M. Hebb (2002), “Conflict of Interest in Commercial Bank Equity Underwriting,” *Financial Review*, 37:185–205. Dan R. Dalton, S. Trevis Certo, and Catherine M. Daily (2003), “Initial Public Offerings as a Web of Conflicts of Interest: An Empirical Assessment,” *Business Ethics Quarterly*, 13:289–314. Lily Fang and Ayako Asuda (2009), “The Effectiveness of Reputation as a Disciplinary Mechanism in Sell-Side Research,” *Review of Financial Studies*, 22:3735–77.
- 171 *A full discussion* Erik Sirri (2004), “Investment Banks, Scope, and Unavoidable Conflicts of Interest,” *Federal Reserve Bank of Atlanta Economic Review*, Fourth Quarter 2004. Frederick S. Mishkin (2005), “Conflicts of Interest in the Financial Industry,” in *The Economics of Money, Banking and Financial Markets, Alternate Edition*, Pearson Education. Frederick S. Mishkin (2005), “Policy Remedies for Conflicts of Interest in the Financial System,” Bank of Canada. Cheolwoo Lee (2008), “Analyst Firm Parent-Subsidiary Relationship and Conflict of Interest: Evidence from IPO Recommendations,” 2008 Financial Management Association Annual Meeting. Roni Michaely and Kent L. Womack (1999), “Conflict of Interest and the Credibility of Underwriter Analyst Recommendations,” *Review of Financial Studies*, 12:653–86. Hamid Mehran and Rene M. Stulz (2007), “The Economics of Conflicts of Interest in Financial Institutions,” *Journal of Financial Economics*, 85 (2007): 267–96.

Chapter 12 Notes

- 276 *One.Tel in Australia* Elizabeth Sexton (27 Oct 2007), “Adler Settles In One.Tel Bonus Case,” *Sydney Morning Herald*.
- 277 *Massey Energy* U.S. Department of Labor Mine Safety and Health Administration (29 Jun 2011), “Upper Big Branch Mine, Mine ID: 46-08436, April 5, 2010 Accident, Public Briefing.” J. Davitt McAteer, Katie Beall, James A. Beck, Jr., Patrick C. McGinley, Celeste Monforton, Deborah C. Roberts, Beth Spence, and Suzanne Weise (2011), “Upper Big Branch: The April 5, 2010, Explosion: A Failure of Basic Coal Mine Safety Practices,” Governor’s Independent Investigation Panel.
- 277 *questionable telecom credits* Jonathan A. Knee (2006), *The Accidental Investment Banker: Inside the Decade that Transformed Wall Street*, Oxford University Press, xvii.
- 277 *described dot.coms* *In the matter of Merrill Lynch & Co.* (2002), Decision and Order, Supreme Court of the State of New York.

Chapter 13

- 174 *result is price-fixing* Robert Schuettinger and Eamonn F Butler (1979), *Forty Centuries of Wage and Price Controls: How Not to Fight Inflation*, The Heritage Institute.
- 174 *Adam Smith* Adam Smith (1776), *The Wealth of Nations*.
- 174 *global citric acid cartel* John M. Connor (1998), “The Global Citric Acid Conspiracy,” *Agribusiness*, 14:435–52.
- 175 *computer-memory cartel* U.S. Department of Justice (5 Oct 2005), “Samsung Agrees to Plead Guilty and to Pay \$300 Million Criminal Fine for Role in Price Fixing Conspiracy: Korean Company Pays Second Largest Criminal Antitrust Fine in U.S. History,” press release.
- 178 *Ginger Paralysis Association* John Parascandola (1995), “The Public Health Service and Jamaica Ginger Paralysis in the 1930s,” *Public Health Reports*, 110:361–3. Cecil Munsey (Winter 2006), “Paralysis in a Bottle: The Jake Walk Story,” *Bottles & Extras*, 7–12.
- 178 *this incident led* Paul M. Wax (1995), “Elixirs, Diluents, and the Passage of the 1938 Federal Food, Drug and Cosmetic Act,” *Annals of Internal Medicine*, 122:456–61.
- 179 *follow its charter* Joel Baken (2004), *The Corporation: The Pathological Pursuit of Profit and Power*, Free Press.
- 180 *Moral pressure is* Maurice Punch (2000), “Suite Violence: Why Managers Murder and Corporations Kill,” *Crime, Law & Social Change*, 33:243–80. Jenny White, Albert Bandura, and Lisa A. Bero (2009), “Moral Disengagement in the Corporate World,” *Accountability in Research*, 16:41–74. Jennifer J. Kish-Gephart, David A. Harrison, and Linda Klebe Trevino (2010), “Bad Apples, Bad Cases, and Bad Barrels: Meta-Analytic Evidence About Sources of Unethical Decisions at Work,” *Journal of Applied Psychology*, 95:1–3.

- 180 *overseas sweatshops* Dexter Roberts, Pete Engardio, Aaron Bernstein, Stanley Holmes, and Xiang Ji (27 Nov 2006), “Secrets, Lies, and Sweatshops,” *Bloomberg Business Week*. Eric V. Edmonds and Nina Pavcnik (2005), “Child Labor in the Global Economy,” *Journal of Economic Perspectives*, 19:199–220.
- 181 *replace management* Leemore Dafny and David Dranove (2009), “Regulatory Exploitation and Management Changes: Upcoding in the Hospital Industry,” *Journal of Law & Economics*, 52:223–50.
- 181 *immortal sociopaths* Joel Baken (2004), *The Corporation: The Pathological Pursuit of Profit and Power*, Free Press.
- 181 *good for their reputation* Charles J. Fombrun, Naomi A. Gardberg, and Michael L. Barnett (2000), “Opportunity Platforms and Safety Nets: Corporate Citizenship and Reputational Risk,” *Business & Society Review*, 105:85–106. David M. Kreps (2000), “Corporate Culture and Economic Theory,” in Nicolai J. Foss, ed., *The Theory of the Firm: Critical Perspectives on Business and Management*, Taylor & Francis. Einer Elhauge (2005), “Sacrificing Corporate Profits in the Public Interest,” *New York University Law Review*, 80:733–869.
- 181 *Baron Thurlow* John Poynder (1844), *Literary Extracts, Vol. 1*, J. Hatchard & Sons, 2.
- 181 *no soul to damn* John C. Coffee (1981), “‘No Soul to Damn, No Body to Kick’: An Unscandalized Inquiry Into the Problem of Corporate Punishment,” *Michigan Law Review*, 79:386–459.
- 182 *a lot of money* Kantar Media (17 Mar 2011), “Kantar Media Reports U.S. Advertising Expenditures Increased 6.5 Percent In 2010,” press release.
- 182 *psychological manipulation* Robert B. Cialdini (1998), *Influence: The Psychology of Persuasion*, Collins.
- 182 *dose of neuroscience* Samuel M. McClure, Jian Li, Damon Tomlin, Kim S. Cypert, Latané M. Montague, and P. Read Montague (2004). “Neural Correlates of Behavioral Preference for Culturally Familiar Drinks,” *Neuron*, 44:379–87. Brian Knutson, Scott Rick, G. Eliot Wimmer, Drazen Prelec, and George Loewenstein (2007), “Neural Predictors of Purchases,” *Neuron*, 53:147–56. Dan Ariely and Gregory S. Berns (2010), “Neuromarketing: The Hope and Hype of Neuroimaging in Business,” *Nature Reviews Neuroscience*, 11:284–92. Douglas Rushkoff (2000), “Advertising,” Chapter 5 of *Coercion: Why We Listen To What ‘They’ Say*, Putnam.
- 183 *150 people died* U.S. Food and Drug Administration (2009), “Information on Adverse Event Reports and Heparin.” Alan Coukell (2009), “Protecting Consumers from Adulterated Drugs,” presentation at U.S. Food and Drug Administration, Public Meeting on Economically Motivated Adulteration, College Park, Maryland.
- 183 *Cargill recalled* Tom Watkins (2011), “Cargill Recalls 36 Million Pounds of Ground Turkey,” *CNN News*.
- 183 *ChoicePoint allowed* Bob Sullivan (14 Feb 2005), “Database Giant Gives Access to Fake Firms,” *MSNBC*. Jaikumar Vijayan (26 Jan 2006), “FTC Imposes \$10M Fine Against ChoicePoint for Data Breach,” *Computer World*.

- 184 *There's recognition* Daniel G. Goldstein and Gerd Gigerenzer (2002), "Models of Ecological Rationality: The Recognition Heuristic," *Psychological Review*, 109:75–90.
- 184 *There's social proof* Herbert C. Kelman (1958), "Compliance, Identification, and Internalization: Three Processes of Attitude Change," *Journal of Conflict Resolution*, 2:51–60.
- 184 *attribute substitution* Daniel Kahneman and Shane Frederick (2002), "Representativeness Revisited: Attribute Substitution in Intuitive Judgment," in Thomas Gilovich, Dale Griffin, and Daniel Kahneman, eds., *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press, 49–81.
- 184 *a lemons market* George Akerlof (1970), "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 83:488–500. George E. Hoffer and Michael D. Pratt (1987), "Used Vehicles, Lemons Markets, and Used Car Rules: Some Empirical Evidence," *Journal of Consumer Policy*, 10:409–14. Steven E. Kaplan, Pamela B. Roush, and Linda Thorne (2007), "Andersen and the Market for Lemons in Audit Reports," *Journal of Business Ethics*, 70:363–73.
- 184 *Globalization is making* Rick Mullin (2011), "Shock to the System: Big Questions About Drug Safety Arise in the Wake of Rampant Supply-Chain Globalization," *Chemical & Engineering News*, 89:11–20.
- 184 *Glaxo Smith-Kline* Toni Clarke (26 Oct 2010), "Glaxo to Pay \$750 Million in Adulterated Drugs Case," Reuters.
- 184 *died of kidney failure* Walt Bogdanich and Jake Hooker (6 May 2007), "From China to Panama, a Trail of Poisoned Medicine," *New York Times*.
- 184 *prone to rollover* Joan Claybrook (2001), "Criminal Penalties for Corporate Misbehavior," *Tort Source*, 3:1, 4.
- 185 *salmonella outbreak* Mary Clare Jalonick (22 Aug 2010), "A Supplier in Egg Recall Has History of Violations, Including in Maine," *Portland Press Herald*. Brendan Flaherty (2010), "Salmonella Outbreak in Eggs: An Update," *Food & Drug Law Institute Update*, 6: 16–19.
- 185 *the drug Rebif* U.S. Department of Justice (4 May 2011), "Pharmaceutical Giant, Serono, Agrees to Pay \$44.3 Million to Settle False Claims Act Case," press release. *Wall Street Journal* (5 May 2011), "Settlement Reached In MS Drug Case."
- 185 *using child labor* United States International Child Labor Program (1998), "By the Sweat and Toil of Children, Vol. V: Efforts to Eliminate Child Labor," Federal Publications, paper 144.
- 186 *oil industry wanting* Jonathan Fahey (12 May 2011), "How the Oil Industry Saves \$4.4 Billion in Taxes," *USA Today*.
- 186 *law passed limiting* Fred J. Hellinger and Gary J. Young (2005), "Health Plan Liability and ERISA: The Expanding Scope of State Legislation," *American Journal of Public Health*, 95:217–23. Katherine L. Record (2010), "Wielding the Wand Without Facing the Music: Allowing Utilization Review Physicians to Trump Doctors' Orders, but Protecting Them from the Legal Risk Ordinarily Attached to the Medical Degree," *Duke Law Journal*, 59:955–1000.

- 186 *licensing interior designers* Janet Zink (1 Apr 2011), “Real Knock-Down, Drag-Out Fight in Florida Is over Commercial Interior Design,” *St. Petersburg Times*. George F. Will (22 Mar 2007), “Wallpapering with Red Tape,” *The Washington Post*. Dick M. Carpenter II (2007), “Designing Cartels: How Industry Insiders Cut Out Competition,” Institute for Justice.
- 187 *horizontal drilling technology* Carl T. Montgomery and Michael B. Smith (2010), “Hydraulic Fracturing: History of an Enduring Technology,” *Journal of Petroleum Technology*, 62:26–41.
- 187 *shoddy building practices* Heather Timmons (25 Apr 2003), “Shoddy Building in the Housing Boom?” *Bloomberg BusinessWeek*.
- 188 *used to be smaller* John Stopford (1998), “Multinational Corporations,” *Foreign Policy*, 113:12–24. Gardiner C. Means (1931), “The Growth in the Relative Importance of the Large Corporation in American Economic Life,” *The American Economic Review*, 21:10–42.
- 188 *Ronald Coase first* Ronald Coase (1937), “The Nature of the Firm,” *Economica*, 4:386–405.
- 189 *Nick Leeson’s* Richard W. Stevenson (28 Feb 1995), “The Collapse of Barings: The Overview: Young Trader’s \$29 Billion Bet Brings down a Venerable Firm,” *New York Times*. Erik Ipsen (19 Jul 1995), “Bank of England Cites Fraud in Barings Collapse,” *New York Times*. Peter Culshaw (8 Jan 2009), “Nick Leeson: How the Original Rogue Trader at Barings Bank Is Thriving in the Credit Crunch,” *The Telegraph*.
- 189 *Kweku Adoboli* Victoria Howley and Emma Thomasson (16 Sep 2011), “UBS \$2 Billion Rogue Trade Suspect Held in London,” Reuters. Haig Simonian (24 Sep 2011), “UBS Chief Resigns Over Rogue Trade Affair,” *Financial Times*.
- 189 *Different countries* John M. Connor and Darren Bush (2008), “How to Block Cartel Formation and Price Fixing: Using Extraterritorial Application of the Antitrust Laws as a Deterrence Mechanism,” *Penn State Law Review*, 112:813–57. John M. Connor (2002), “The Food and Agricultural Global Cartels of the 1990s: Overview and Update,” Purdue University Department of Agricultural Economics, Staff Paper 02–4.
- 189 *Amazon.com uses its* Jeanine Poggi (2011), “Amazon Sales Tax: The Battle, State by State,” *TheStreet.com*.
- 191 *free is very different* Chris Anderson (2009), *Free: The Future of a Radical Price*, Hyperion. Dan Ariely (2008), *Predictably Irrational: The Hidden Forces That Shape Our Decisions*, HarperCollins.
- 191 *Chemical plants* David Bollier, Frank Clemente, Laura MacCleery, Hugh Jackson, and Craig Aaron (2004), “Homeland Unsecured,” Public Citizen. Linda-Jo Schierow (2006), “Chemical Facility Security,” Congressional Research Service Report RL-31530.
- 191 *hundreds of plants* Margaret E. Kosal (2006), “Terrorism Targeting Industrial Chemical Facilities: Strategic Motivation and the Implications for U.S. Security,” *Studies in Conflict & Terrorism*, 29:719–51.

- 191 *The problem is* Jeffrey Hahn, Donna Post Guillen, and Thomas Anderson (2005), “Process Control Systems in the Chemical Industry: Safety vs. Security,” Paper presented at the 20th Annual CCPS International Conference, Idaho National Laboratory, April 2005.

Chapter 13 Notes

- 278 *band together into* Tim Wu (2010), *The Master Switch: The Rise and Fall of Information Empires*, Alfred A. Knopf.
- 278 *High barriers to entry* John M. Connor (2002), “The Food and Agricultural Global Cartels of the 1990s: Overview and Update,” Purdue University Department of Agricultural Economics, Staff Paper 02-4.
- 279 *There are exceptions* Yvonne Chouinard (2005), *Let My People Go Surfing: The Education of a Reluctant Businessman*, Penguin Press.
- 279 *implant false memories* Priyali Rajagopal and Nicole Votolato Montgomery (2011), “I Imagine, I Experience, I Like: The False Experience Effect,” *The Journal of Consumer Research*, 38:578–94.
- 279 *The economic term* Paul A. David (1985), “Clio and the Economics of QWERTY,” *The American Economic Review*, 75:332–7. W. Brian Arthur (1989), “Competing Technologies, Increasing Returns, and Lock-In by Historical Events,” *The Economic Journal*, 99:116–31. Stan J. Liebowitz and Stephen E. Margolis (1995), “Path Dependence, Lock-In, and History,” *Journal of Law, Economics & Organization*, 11:205–26. Begoña Garcia Mariñoso (2001), “Technological Incompatibility, Endogenous Switching Costs and Lock-In,” *The Journal of Industrial Economics*, 49:281–98. Gal Zauberman (2003), “The Intertemporal Dynamics of Consumer Lock-In,” *Journal of Consumer Research*, 30:405–19. William Barnes, Myles Gartland, and Martin Stack (2004), “Old Habits Die Hard: Path Dependency and Behavioral Lock-In,” *Journal of Economic Issues*, 38:371–7. Joseph Farrell and Paul Klemperer (2007), “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization, Volume 3*, North-Holland, 1967–2072.
- 279 *able to bargain down* Tricia Bishop (13 Dec 2008), “Court Orders ‘Scareware’ Shut Down,” *Baltimore Sun*. Lucian Constantin (16 Jun 2009), “ByteHosting Settles with the FTC in Scareware Advertising Lawsuit,” *Softpedia*.
- 279 *hydraulic fracturing* U.S. House of Representatives Committee on Energy and Commerce Minority Staff (2011), “Chemicals Used in Hydraulic Fracturing.”
- 279 *company’s arguments* *New York Times* (3 Nov 2009), “The Halliburton Loophole (Editorial),” *New York Times*.
- 279 *The same dynamic* Frederick P. Brooks, Jr. (1975), *The Mythical Man-Month*, Addison-Wesley.
- 280 *Senator Bernie Sanders* Bernard Sanders (6 Nov 2009), “Too Big To Fail—Too Big To Exist,” *Huffington Post*.

Chapter 14

- 196 *Full-body scanners* Leon Kaufman and Joseph W. Carlson (2010), "An Evaluation of Airport X-Ray Backscatter Units Based on Image Characteristics," *Journal of Transportation Security*, 4:73–94. Andrew Welch (2010), "Full-Body Scanners: Full Protection from Terrorist Attacks or Full-On Violation of the Constitution," *Transportation Law Journal*, 37:167–98. David J. Brenner (2011), "Are X-Ray Backscatter Scanners Safe for Airport Passenger Screening?" *Radiology*, 259:6–10.
- 196 *the underwear bomber* Steve Lord (17 Mar 2010), "Aviation Security: TSA Is Increasing Procurement and Deployment of the Advanced Imaging Technology, but Challenges to This Effort and Other Areas of Aviation Security Remain," Testimony of the Director, Homeland Security and Justice Issues, United States Government Accountability Office, before the Subcommittee on Transportation Security and Infrastructure Protection, Committee on Homeland Security, House of Representatives, Report GAO-10-484T.
- 197 *strip-search every* Electronic Privacy Information Center (10 Mar 2011), "DHS: We Have the Authority to Routinely Strip-Search Air Travelers," press release. Electronic Privacy Information Center (15 Jul 2011), "Federal Appeals Court: TSA Violated Federal Law, Must Take Public Comment on Body Scanners," press release. *EPIC v. DHS* (15 Jul 2011), Opinion, Case No. 10N1157. DC Circuit Court of Appeals, filed November 1, 2010.
- 197 *September 11 attacks* Nate Silver (4 Jan 2010), "The Skies Are as Friendly as Ever: 9/11, Al Qaeda Obscure Statistics on Airline Safety," *FiveThirtyEight.com*.
- 198 *scale is too large* Bruce Schneier (2008), "Seven Habits of Highly Unsuccessful Terrorists," *Wired News*. Max Abrams (2008), "What Terrorists Really Want," *International Security*, 32:78–105.
- 202 *regulatory capture* Jean J. Laffont and Jean Tirole (1991), "The Politics of Government Decision-Making: A Theory of Regulatory Capture," *The Quarterly Journal of Economics*, 106:1089–127. Mark Jickling (2004), "Barriers to Corporate Fraud: How They Work, Why They Fail," CRS Report for Congress RL32718, Congressional Research Service. Dieter Helm (2006), "Regulatory Reform, Capture, and the Regulatory Burden," *Oxford Review of Economic Policy*, 22:169–86. Ernesto Dal Bó (2006), "Regulatory Capture: A Review," *Oxford Review of Economic Policy*, 22:203–25.
- 202 *steering contracts to* Charlie Savage (11 Sep 2008), "Sex, Drug Use and Graft Cited in Interior Department," *New York Times*. Earl E. Devaney (9 Sep 2008), "OIG Investigations of MMS Employees," memorandum to Secretary Dirk Kempthorne from the Office of the Inspector General, United States Department of the Interior.
- 202 *Federal Aviation Administration* Mark C. Niles (2002), "On the Hijacking of Agencies (And Airplanes): The Federal Aviation Administration, 'Agency Capture,' and Airline Security," *Journal of Gender, Social Policy & the Law*, 10:381–442. John Downer

- (2010), "Trust and Technology: The Social Foundations of Aviation Regulation," *The British Journal of Sociology*, 61:83–106. Johanna Neumann (3 Apr 2008), "FAA's 'Culture of Coziness' Targeted in Airline Safety Hearing," *Los Angeles Times*.
- 202 *Securities and Exchange* Susan E. Woodward (1998), "Regulatory Capture at the Securities and Exchange Commission," prepared for the Milken Institute Conference on Capital Markets, March 16, 1998, Santa Monica, California. Mark Quintyn and Michael W. Taylor (2002), "Regulatory and Supervisory Independence and Financial Stability," International Monetary Fund Working Paper 02–46. Nicholas Dorn (2010), "The Governance of Securities: Ponzi Finance, Regulatory Convergence, Credit Crunch," *British Journal of Criminology*, 50:23–45.
- 204 *illegally spying* James Risen and Eric Lichtblau (16 Dec 2005), "Bush Lets U.S. Spy on Callers without Courts," *New York Times*. *American Civil Liberties Union et al. v. National Security Agency et al.* (18 Aug 2006), Order, United States District Court for the Eastern District of Michigan, Case No. 06-CV-10204. *American Civil Liberties Union et al. v. National Security Agency et al.* (4 Oct 2006), Order, United States Court of Appeals for the Sixth Circuit, Case Nos. 06–2095/2140. Charlie Savage and James Risen (31 Mar 2010), "Federal Judge Finds N.S.A. Wiretaps Were Illegal," *New York Times*. *Al-Haramain Islamic Foundation, Inc., et al., v. Barack H. Obama, et al.* (31 Mar 2010), Opinion, United States Court of Appeals for the Ninth Circuit, Case No. C 07–0109 VRW.
- 204 *authorizing the torture* Jane Mayer (13 Aug 2007), "The Black Sites," *The New Yorker*. Daniel Nasaw (24 Aug 2009), "Newly Released CIA Report Details Interrogation Techniques," *The Guardian*. Associated Press (14 Jun 2011), "New Scrutiny of CIA Torture, Botched Rendition," *CBS News*. Eric Lichtblau and Eric Schmitt (1 Jul 2011), "U.S. Widens Inquiries into 2 Jail Deaths," *New York Times*.
- 204 *right to assassinate* John Griffing (5 Oct 2011), "Obama, the Hitman: Killing Due Process," *The American Thinker*.
- 204 *court rejected EPIC's* Electronic Privacy Information Center (15 Jul 2011), "Federal Appeals Court: TSA Violated Federal Law, Must Take Public Comment on Body Scanners," press release. *EPIC v. DHS* (15 Jul 2011), Opinion, Case No. 10N1157 (filed Nov 1, 2010), DC Circuit Court of Appeals.

Chapter 14 Notes

- 280 *ineffective tactic* Max Abrams (2006), "Why Terrorism Does Not Work," *International Security*, 31:42–78. Bruce Schneier (12 Jul 2007), "The Evolutionary Brain Glitch that Makes Terrorism Fail," *Wired News*.
- 281 *isn't just theoretical* A. Peter McGraw, Alexander Todorov, and Howard Kunreuther (2011), "A Policy Makers Dilemma: Preventing Terrorism or Preventing Blame," *Organizational Behavior & Human Decision Processes*, 115:25–34.

Chapter 15

- 207 *people paying taxes* Luigi Bosco and Luigi Mittone (1997), "Tax Evasion and Moral Constraints: Some Experimental Evidence," *Kyklos*, 50:297–324.
- 207 *a lot of fraud* U.S. Department of the Treasury (2006), "IRS Updates Tax Gap Estimates," Informational release IR-2006–28.
- 207 *third-party estimate* Edgar L. Feige and Richard Cebula (2011), "America's Underground Economy: Measuring the Size, Growth and Determinants of Income Tax Evasion in the U.S.," Munich Personal RePEc Archive, MPRA Paper No. 29672.
- 207 *cheating on their taxes* James A. Tackett, Joe Antenucci, and Fran Wolf (2006), "A Criminological Perspective of Tax Evasion," *Tax Notes*, 110:654–8.
- 207 *reasons people cheat* Stuart P. Green (2009), "What Is Wrong With Tax Evasion?" *Houston Business & Tax Journal*, 9:220–33.
- 207 *very little enforcement* U.S. Department of the Treasury (2008), *IRS Oversight Board Annual Report 2007*.
- 207 *IRS examined less* Department of the Treasury (2008), Internal Revenue Service Data Book 2007.
- 207 *Corporate audits are down* Transactional Records Access Clearing House (2008), "Audits of Largest Corporations Slide to All Time Low."
- 207 *increasing the penalties* Mukhtar M. Ali, H. Wayne Cecil, and James A. Knoblett (2001), "The Effects of Tax Rates and Enforcement Policies on Taxpayer Compliance: A Study of Self-Employed Taxpayers," *Atlantic Economics Journal*, 29:186–202.
- 207 *increases compliance* Robert Halperin and Joseph Tzur (2002), "Tax Evasion and the Low Penalty, Low Audit Rate Phenomenon," *Journal of Accounting & Public Policy*, 9:179–96.
- 208 *Second, it's profitable* Patricia Sabatini (25 Mar 2007), "Tax Cheats Cost U.S. Hundreds of Billions," *Pittsburgh Post-Gazette*.
- 208 *failing to declare* Richard S. Dunham (15 Apr 1998), "Maybe Those Meanies at the IRS Aren't So Bad After All," *Business Week*.
- 208 *demonized taxes* Robert W. McGee (2004), *The Philosophy of Taxation and Public Finance*, Springer-Verlag.
- 208 *these reasons interact* Edgar L. Feige and Richard Cebula (2011), "America's Underground Economy: Measuring the Size, Growth and Determinants of Income Tax Evasion in the U.S.," Munich Personal RePEc Archive, MPRA Paper No. 29672 (Jan 2011).
- 208 *blamed income inequality* Kim M. Bloomquist (2003), "Tax Evasion, Income Inequality and Opportunity Costs of Compliance," Paper presented at the 96th Annual Conference of the National Tax Association.
- 208 *carried-interest tax* Nicholas Kristof (6 Jul 2011), "Taxes and Billionaires," *New York Times*. Laura Saunders (6 Aug 2011), "'Carried Interest' in the Cross Hairs," *Wall Street Journal*.

- 208 *investment tax credit* Don MacKenzie, Louise Bedsworth, and David Friedman (2005), "Fuel Economy Fraud: Closing the Loopholes That Increase U.S. Oil Dependence," Union of Concerned Scientists. Jim Hopkins (10 Feb 2004), "SUV Sales Climb on Tax Loophole," *USA Today*.
- 209 *lost federal revenue* Scott Horsley (30 Jun 2011), "What Tax 'Loopholes' Does Obama Want to Close?" National Public Radio.
- 209 *perceive tax evasion* Erich Kirchler, Boris Maciejovsky, and Friedrich Schneider (2003), "Everyday Representations of Tax Avoidance, Tax Evasion, and Tax Flight: Do Legal Differences Matter?" *Journal of Economic Psychology*, 24:535–53.
- 209 *just as immoral* Zoë Prebble and John M. Prebble (2010), "The Morality of Tax Avoidance," *Creighton Law Review*, 20:101–58.
- 210 *first people arrested* Bradley S. Davis (1994), "It's Virus Season Again, Has Your Computer Been Vaccinated? A Survey of Computer Crime Legislation as a Response to Malevolent Software," *Washington University Law Quarterly*, 72:410–40. Patrick J. Leahy (1992), "New Laws for New Technologies: Current Issues Facing the Subcommittee on Technology and the Law," *Harvard Journal of Law & Technology*, 5:1–23.
- 210 *laws regulating surveillance* Valerie Caproni (17 Feb 2011), "Going Dark: Lawful Electronic Surveillance in the Face of New Technologies," testimony before the House Judiciary Committee.
- 211 *legalization of abortion* John J. Donohue III and Steven D. Levitt, "The Impact of Legalized Abortion on Crime," *The Quarterly Journal of Economics*, 116:379–420.
- 211 *unintended consequences* Thomas P. Bonczar and Allen J. Beck (1997), "Lifetime Likelihood of Going to State or Federal Prison," Special Report NCJ-160092, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics. Heather C. West, William J. Sabol, and Sarah J. Greenman (2010), "Prisoners in 2009," Report NCJ 231675, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics. Bruce Western and Christopher Wildeman (2009), "The Black Family and Mass Incarceration," *The Annals of the American Academy of Political and Social Science*, 621:221–42.
- 211 *increasing shoplifting* Kerry Segrave (2001), *Shoplifting: A Social History*, Mcfarland & Co.
- 212 *fewer bicycle riders* Maxwell H. Cameron, A. Peter Vulcan, Caroline F. Finch, and Stuart V. Newstead (1994). "Mandatory Bicycle Helmet Use Following a Decade of Helmet Promotion in Victoria, Australia: An Evaluation," *Accident Analysis & Prevention*, 26:325–7. Dorothy L. Robinson (2007), "Bicycle Helmet Legislation: Can We Reach a Consensus?" *Accident Analysis & Prevention*, 39:86–93. Christopher S. Carpenter and Mark F. Stehr (2010), "Intended and Unintended Effects of Youth Bicycle Helmet Laws," NBER Working Paper No. 15658, National Bureau of Economic Research.
- 212 *chilling effect on* Electronic Frontier Foundation (2008), "Unintended Consequences: Ten Years Under DMCA."
- 212 *increased violence* Dan Werb, Greg Rowell, Gordon Guyatt, Thomas Kerr, Julio Montaner, and Evan Wood (2011), "Effect of Drug Law Enforcement on Drug Market Violence: A Systematic Review," *International Journal of Drug Policy*, 22:87–94.

- 212 *trade-off between* Bruce Schneier (2003), *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus Books, Chapter 1.
- 213 *different characteristics* Bruce Schneier (2003), *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus Books, Chapter 5.
- 214 *unit pricing was* Laetitia B. Mulder, Eric van Dijk, David De Cremer, and Henk A.M. Wilke (2006), "When Sanctions Fail to Increase Cooperation in Social Dilemmas: Considering the Presence of an Alternative Option to Defect," *Personality & Social Psychology Bulletin*, 32:1312–24. Marie Lynn Miranda and Joseph E. Aldy (1998), "Unit Pricing of Residential Municipal Waste: Lessons From Nine Case Studies," *Journal of Environmental Management*, 52:79–93.
- 214 *Stuffing more trash* Don Fullerton and Thomas C. Kinneman (1996), "Household Responses to Pricing Garbage by the Bag," *The American Economic Review*, 86:971–84.
- 215 *based on perceived* Bruce Schneier (3 Apr 2008), "The Difference Between Feeling and Reality in Security," *Wired News*.
- 215 *Natural Biases* Barry Glassner (1999), *The Culture of Fear: Why Americans Are Afraid of the Wrong Things*, Basic Books. Paul Slovic (2000), *The Perception of Risk*, Earthscan Publications. Daniel Gilbert (2 Jul 2006), "If Only Gay Sex Caused Global Warming," *Los Angeles Times*. Jeffrey Kluger (26 Nov 2006), "How Americans Are Living Dangerously," *Time*. Cass Sunstein and Richard Zeckhauser (2011), "Overreaction to Fearsome Risks," *Environmental & Resource Economics*, 48:435–49.
- 215 *John Mueller wrote* John Mueller (2004), "A False Sense of Insecurity?" *Regulation*, 27:42–6.
- 216 *exaggerate the risk* John Mueller (2006), *Overblown: How Politicians and the Terrorism Industry Inflate National Security Threats, and Why We Believe Them*, Free Press.
- 216 *tolerance for risk* Meir Statman (2010), "The Cultures of Risk Tolerance," *Social Sciences Research Network Behavioral & Experimental Finance eJournal*, 1–23.
- 216 *income inequality* Kate Pickett and Richard Wilkinson (2011), *The Spirit Level: Why Greater Equality Makes Societies Stronger*, Bloomsbury Press.
- 217 *false confessions* Saul M. Kassin, "False Confessions: Causes, Consequences, and Implications for Reform," *Current Directions in Psychological Science*, 17: 249–53. Jennifer T. Perillo and Saul M. Kassin (2011), "Inside Interrogation: The Lie, The Bluff, and False Confessions," *Law & Human Behavior*, 35:327–37.
- 217 *teachers cheated on* Jack Gillum and Marisol Bello (30 Mar 2011), "When Standardized Test Scores Soared in D.C., Were the Gains Real?" *USA Today*.
- 218 *analysis of test data* Brian A. Jacob and Steven D. Levitt (2003), "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating," *The Quarterly Journal of Economics*, 118:843–77. Brian A. Jacob and Steven D. Levitt (2003), "Catching Cheating Teachers: The Results of an Unusual Experiment in Implementing Theory," in William G. Gale and Janet Rothenberg Pack, eds., *Brookings-Wharton Papers on Urban Affairs*, Brookings Institution Press, 185–220.
- 218 *If you pay CEOs* Joel Spolsky (2001), "Measurement," *Joel on Software*.

- 219 *even more overfishing* British Broadcasting Corporation (5 Apr 2006), “River Patrol to Halt Eel Poachers,” *BBC News*. Steven Morris (28 May 2010), “Poachers Devastate Somerset’s Elver Fisheries,” *The Guardian*. Caroline Wyatt (11 Apr 2005), “Asia Demand Threatens French Eels,” *BBC News*. Adam Sage (13 Dec 2007), “Eel Poachers Put ‘White Gold’ on Critical List,” *Times*. Bob Ritchie and Dan McKiernan (1999), “Environmental Police Prepare for ‘Glass Eel’ Poaching Season,” *Massachusetts Department of Marine Fisheries DMF News*, 19:6. Atlantic States Marine Fisheries Commission, American Eel Plan Development Team (2000), “Interstate Fishery Management Plan for American Eel (*Anguilla rostrata*),” Fishery Management Report No. 36. Samuel J. Baldwin (11 May 2011), “Glass Eels Hit \$950 Per Pound,” *The Lincoln County News*.
- 219 *farm these eels* *Japan Times* (10 Apr 2010), “Japan Succeeds in World’s First Complete Cultivation of Eels.”
- 220 *In ancient Rome* Robert L. Wilken (1985), *The Christians as the Romans Saw Them*, Yale University Press. Peter Brown (1995), *Authority and the Sacred: Aspects of the Christianization of the Roman World*, Cambridge University Press.
- 221 *Norms can change quickly* Ayelet Gneezy and Daniel M.T. Fessler (2011), “Conflict, Sticks, and Carrots: War Increases Prosocial Punishments and Rewards,” *Proceedings of the Royal Society B: Biological Sciences*, published online before print, 8 Jun 2011:1–5.
- 222 *women or slaves* Barbara C. Jordan (24 Jul 1974), “Statement on House Judiciary Proceedings to Impeach President Richard Nixon,” House Judiciary Committee.
- 222 *political debates* Howard Fineman (2008), *The Thirteen American Arguments: Enduring Debates That Define and Inspire Our Country*, Random House.

Chapter 15 Notes

- 281 *studied systems dynamics* Daniel H. Kim and Virginia Anderson (1989), *Systems Archetype Basics*, Pegasus Communications. Peter Senge (1990), *The Fifth Discipline: The Art and Practice of the Learning Organization*, Currency.
- 281 *experiential goods* Daniel Gilbert (2006), *Stumbling on Happiness*, Knopf.
- 281 *Ben Franklin said* Richard Jackson (1759), *An Historical Review of the Constitution and Government of Pennsylvania*.
- 281 *also human nature* Joyce Ehrlinger and Richard P. Eibach (2011), “Focalism and the Failure to Foresee Unintended Consequences,” *Basic & Applied Social Psychology*, 33:1, 59–68.
- 281 *major political backlash* Sophie Goodchild (27 Aug 2006), “500,000 Wheelie Bins ‘Have a Spy In the Lid,’” *The Independent*. Louise Gray (21 Jul 2009), “Councils ‘Snooping’ In Bins To See Who Is Recycling Properly,” *The Telegraph*. Chris Gourlay and Maurice Chittenden (11 Oct 2009), “Bin Police Make You Save Every Scrap,” *London Times*. Philip Reynolds (30 Apr 2010), “Bin Sensor Technology Revolutionizes Waste Collection,” *Materials Recycling Week*. British Broadcasting Corporation (5 Mar 2011), “Warning of Rise in Microchips in Council Bins,” *BBC News*.

- 282 *life-cycle management* U.S. Environmental Protection Agency (2010), "Recycling and Reuse: End-of-Life Vehicles and Producer Responsibility."
- 282 *The Innocence Project* Innocence Project (2011), "Facts on Post-Conviction DNA Evidence."
- 282 *Cheating on test scores* Heather Vogell (6 Jul 2011), "Investigation Into APS Cheating Finds Unethical Behavior Across Every Level," *Atlanta Journal-Constitution*. Jaime Sarrio (10 Jul 2011), "Cheating Scandal Adds Fuel to Debate Over High-Stakes Tests," *Atlanta Journal-Constitution*. Michael Winerip (31 Jul 2011), "Pa. Joins States Facing a School Cheating Scandal," *New York Times*.
- 282 *One teacher described* Benjamin Herold (28 Jul 2011), "Confession of a Cheating Teacher," *The Notebook*.
- 282 *Scott Adams wrote* Scott Adams (1996), *The Dilbert Principle: A Cubicle's-Eye View of Bosses, Meetings, Management Fads & Other Workplace Afflictions*, HarperBusiness, 12.
- 282 *18 years if* Matthew Sherman (2009), "A Short History of Financial Deregulation in the United States," Center for Economic and Policy Research.
- 282 *potential failure* Alexis de Tocqueville (1835), *Democracy in America*, Saunders and Otley.

Chapter 16

- 225 *have more flaws* Bruce Schneier (Mar 2000), "Software Complexity and Security," *Crypto-Gram*.
- 225 *fail in surprising* Charles Perrow (1984), *Normal Accidents: Living with High-Risk Technologies*, Basic Books. Dietrich Dörner (1996), *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*, Perseus Books.
- 226 *This aspect of scale* Jane Wakefield (22 Aug 2011), "When Algorithms Control the World," *BBC News*.
- 227 *tend to minimize* Daniel Gilbert (2 Jul 2006), "If Only Gay Sex Caused Global Warming," *Los Angeles Times*.
- 227 *pretending to eliminate* Frank Furedi (1997), *The Culture of Fear*, Continuum International Publishing Group.
- 227 *knobs being twiddled* Bruce Schneier (2003), *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus Books, Chapter 7.
- 229 *hackers to find* Ron Rosenbaum (Oct 1971), "Secrets of the Little Blue Box," *Esquire*, 117–25, 222–36. Gary D. Robson (Spring 2004), "The Origins of Phreaking," *Blacklisted 411*, 6:17–23.
- 229 *Criminals can form* Allan Castle (1997), "Transnational Organized Crime and International Security," Institute of International Relations, The University of British Columbia Working Paper No. 19. Phil Williams (2001), "Transnational Criminal Networks," in John Arquilla and David F. Ronfeldt, eds., *Networks and Netwars: The Future of Terror, Crime, and Militancy*, RAND Corporation, 61–97. Oded Löwenheim

- (2002), "Transnational Criminal Organizations and Security: The Case Against Inflating the Threat," *International Journal*, 57:513–36.
- 229 *Criminals were simply* Warwick Ashford (6 Oct 2010), "ISSE 2010: Police Are Playing Catch-Up as Criminals Embrace IT," *Computer Weekly*. Stephen Pritchard (2 Jun 2011), "Vulnerabilities: Battle Is Joined on Two Fronts," *Financial Times*.
- 229 *position of the interior* Carl von Clausewitz (1832), *On War*.
- 232 *Alvin Toffler wrote* Alvin Toffler (1970), *Future Shock*, Random House.
- 232 *People learn how* Michael Anissimov (2004), "Future Shock Level Analysis," *AcceleratingFuture.com*.
- 233 *new type of grifter* David Maurer (1940), *The Big Con: The Story of the Confidence Man*, Bobbs Merrill.
- 233 *Other cities followed* J.L. Lyman (1964), "The Metropolitan Police Act of 1829: An Analysis of Certain Events Influencing the Passage and Character of the Metropolitan Police Act in England," *The Journal of Criminal Law, Criminology, & Police Science*, 55:141–54.
- 233 *Internet worm Stuxnet* Kim Zetter (2011), "How Digital Detectives Deciphered Stuxnet, the Most Menacing Malware in History," *Threat Level* blog, *Wired.com*.
- 234 *salami attacks* Michael E. Kabay (24 Jul 2002), "Salami Fraud," *Network World*.
- 235 *just use software* Bruce Schneier (2000), *Secrets & Lies: Digital Security in a Networked World*, John Wiley & Sons.
- 237 *delayed publishing* Eric Lichtblau (26 Mar 2008), "The Education of a 9/11 Reporter: The Inside Drama Behind the Times' Warrantless Wiretapping Story," *Slate*.
- 237 *Clay Shirky referred* Clay Shirky (31 Dec 2010), "Half-Formed Thought on WikiLeaks & Global Action," *Shirky.com*.
- 237 *other business considerations* Alfred Rappaport (2005), "The Economics of Short-Term Performance Obsession," *Financial Analysts Journal*, 61:65–79. Thomas P. Byrne (2010), "False Profits: Reviving the Corporation's Public Purpose," *UCLA Law Review Discourse*, 25:25–49. Patrick Bolton, Jose Scheinkman, and Wei Xiong (2006), "Executive Compensation and Short-Termist Behaviour in Speculative Markets," *Review of Economic Studies*, 73:577–610. Lynne Dallas (2011), "Short-Termism, the Financial Crisis and Corporate Governance," *San Diego Legal Studies Paper No. 11–052*. Lawrence Mitchell (2011), *Corporate Irresponsibility: America's Newest Export*, Yale University Press. Alfred Rappaport (2011), *Saving Capitalism from Short-Termism*, McGraw-Hill.
- 238 *investors have access* John C. Bogle (2005), *The Battle for the Soul of Capitalism*, Yale University Press.
- 238 *wicked problems* C. West Churchman (1967), "Wicked Problems," *Management Science*, 14:141–2. Horst Rittel and Melvin Webber (1973), "Dilemmas in a General Theory of Planning," *Policy Sciences*, 4:155–69. E. Jeffrey Conklin (2006), *Dialog Mapping: Building a Shared Understanding of Wicked Problems*, John Wiley & Sons.
- 239 *ubiquity of data* Charles Stross (2011), "Network Security in the Medium Term, 2061–2561 AD," paper presented at USenix Security.

- 241 *better off spending* Bruce Schneier (2009), “Beyond Security Theater,” *New Internationalist*, 427:10–13.
- 242 *Yochai Benkler* Yochai Benkler (2011), *The Penguin and the Leviathan: How Cooperation Triumphs Over Self-Interest*, Crown Business, 25–6.
- 242 *security is a process* Bruce Schneier (2000), *Secrets and Lies: Digital Security in a Networked World*, John Wiley & Sons, 273, 395.

Chapter 16 Notes

- 283 *the singularity* Vernor Vinge (1993), “Technological Singularity,” paper presented at the VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute, 30–31 March. Raymond Kurzweil (2005), *The Singularity is Near: When Humans Transcend Biology*, Penguin Press.
- 283 *against the government* Leonard Deutchman and Sean Morgan (2005), “The ECPA, ISPs and Obtaining E-mail: A Primer for Local Prosecutors,” American Prosecutors Research Institute. U.S. Department of Justice, Computer Crime & Intellectual Property Section (2009), *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*, Third Edition, Office of Legal Education Executive Office for United States Attorneys.
- 283 *Clay Shirky writes* Clay Shirky (2008), *Here Comes Everybody: The Power of Organizing Without Organizations*, Penguin Press.
- 283 *terms of service* *Wall Street Journal* (4 May 2011), “Terms of Use for SafeHouse,” *Wall Street Journal*.
- 283 *Lord Kelvin said* Silvanus Phillips Thompson (2011), *The Life of William Thomson, Baron Kelvin of Largs*, Cambridge University Press.

Chapter 17

- 243 *Trust in things* Bart Nooteboom (2002), *Trust: Forms, Foundations, Functions, Failures and Figures*, Edward Elgar, 45.
- 245 *trust is a social good* Sissela Bok (1978), *Lying: Moral Choice in Private and Public Life*, Pantheon Books, 26–7.
- 245 *trust is the key* James Colman (1988), *Foundations of Social Theory*, Harvard University Press. Francis Fukuyama (1995), “Social Capital and the Global Economy: A Redrawn Map of the World,” *Foreign Affairs*, 75:89–103. Luigi Guiso, Paola Sapienza, and Luigi Zingales (2004), “The Role of Social Capital in Financial Development,” *American Economic Review*, 94:426–56.
- 247 *this parable* Richard Beck (2011), “Tales of the Demonic,” *Experimental Theology*.
- 247 *do not follow* Tom Zeller, Jr. (26 Jul 2011), “Tim De Christopher: Is Civil Disobedience the Key to Climate Change Action?” *The Huffington Post*.
- 248 *Medal of Honor* Michael A. Memoli and Michael Muskal (15 Sep 2011), “Obama Awards Medal of Honor to Kentucky Marine,” *Los Angeles Times*.

- 248 *contact a reporter* Richard L. Rashke (2000), *The Killing of Karen Silkwood: The Story Behind the Kerr-McGee Plutonium Case*, Second Edition, Cornell University Press.

Chapter 17 Notes

- 284 *leap of faith* Søren Kierkegaard (1844), *The Concept of Anxiety*. Alistair Hannay and Gordon Marino, eds. (1997), *The Cambridge Companion to Kierkegaard*, Cambridge University Press.
- 284 *World Values Survey* Jaime Diez Medrano (2011), “Interpersonal Trust,” World Values Survey Archive and ASEP/JDS.
- 284 *results differ widely* Magali Rheault (5 Oct 2007), “Many World Citizens Trust Neighbors More Than Police: Trust in Neighbors and Police About Equal in 21 Countries,” Gallup.
- 284 *tried something similar* *Reader’s Digest* (Jul 2007), “The Reader’s Digest Global Phone Test.”
- 284 *laboratory experiment* Simon Gächter, Benedikt Herrmann, and Christian Thöni (2010), “Culture and Cooperation,” CESifo Working Paper No. 3070.
- 285 *attributed to King* Jamie Stiehm (3 Sep 2010), “Oval Office Rug Gets History Wrong,” *Washington Post*.

About the Author



Photo by Norbert Schiller

BRUCE SCHNEIER is an internationally renowned security technologist who studies the human side of security. A prolific author, he has written hundreds of articles, essays, and academic papers, as well as eleven books that together have sold more than 400,000 copies. He has testified before Congress, is a frequent guest on television and radio, and is regularly quoted in the press. His blog and monthly newsletter at www.schneier.com reach over 250,000 devoted readers worldwide.

“The closest thing the security industry has to a rock star.”

—*The Register*



Index

A

actors in society, 10
Adams, Scott, w *Dilbert Principle*, 284
advertising reputation, 90
airplane terrorism, 197–199
Alice and Bob
 cryptographers and, 262
 Prisoner's Dilemma and, 270–271
altruism
 corporations, 181
 direct reciprocity and, 38
 existence, 34–35
 group selection and, 38, 260
 indirect reciprocity and, 38
 innate, 36–37
 network reciprocity and, 38
 neuroscientific explanation, 36
 pleasure in, 37
 reciprocal, 34, 261, 273
 signals, 35
 transactional analysis, 35
angels *versus* men, 43
animal intelligence, weapons and
 tools, 252
appearance, trust and, 92

apple juice dilemma (Beech-Nut), 163
attacker, terminology, 13
auctions, 176
audit/forensic security system, 125
authoritarian government, 286
authority/respect system of
 morality, 79

B

Bad Apple Effect, 82
 bad apple usage, 267
 punishment threat, 98
 taxes and, 208
bagel honor system, 83–84, 268
Baird, Jeff, 167
banishment, 96–97
Battle of the Sexes, 264
Beech-Nut Corporation apple juice,
 163
behaviors
 individual, 84
 prosocial, 32, 35–36
Big Dig highway project, 170
Blodget, Henry, 279

Bok, Sissela, 7
 bottled water, 182
 Bounded Rationality, 141
 BP Deepwater Horizon explosion, 170
 brain, neuroplasticity, 42
 branding, reputation and, 95
 bribery, 108
 broken windows theory, 267
 burglary, costs, 105–106
 Buss, David, murderousness and, 25

C

cameras, societal pressure and, 126–127
 cartels, 174–175
 certification, reputation formalization and, 116
 change, societal pressures and, 50
 cheating
 No Child Left Behind Act and, 284
 students' tests, 218, 274
 chemical plants, 191–192
 Chicken game, 65, 263–264. *See also* Snowdrift Dilemma
 chimpanzees
 aggression rates, 254
 reputation and, 270
 weapons, 252
 civilization description, 1
 Clarke, Ronald V., Crime Prevention Techniques chart, 130
 cliques, 46
 Code of Hammurabi, 50
 coercion, formalizing, 109
 commensalistic relationships, 28
 commerce
 defrauding merchants, 59–60
 as Prisoner's Dilemma, 53
 reputation and, 53, 91
 societal dilemmas and, 53
 commitment, reputation and, 93–94
 community, symbolic contracts and, 45
 competing interests, 9, 139–140, 141
 defection and, 142–143, 145
 exemptions, 148
 families, 153–154
 groups, 142–144
 moral, not reporting and, 146
 multiple, 140–141
 nepotism and, 150–151
 not reporting, 145–146
 O'Callaghan, Sean, 149
 organizations, 156
 scale, 143
 terminology, 11
 trade-offs between group interest, 238–239
 competing norms, 11
 compliance, societal pressures and, 49
 confidence tricksters, 99
 contextualities
 intuition, 134–135
 morals, 78–79
 contracts, 117
 symbolic, sexual fidelity and, 45
 Cooper, Cynthia, 165–166
 cooperation, 2
 versus defection, 34, 68
 dialect and, 270
 diversity and, 270
 empathy and, 83
 evolutionary benefits, 34
 Golden Rule, 77
 non-kin, 33–34

- Nowak, Martin A., 38
- Pascal's Wager, 270
- Platinum Rule, 93
- police, 147
 - IRA and, 150
- Prisoner's Dilemma and, 262, 269
- reputation and, 89–90
- reward/punishment and, 88
- rewards and, 96–98
- social pressures and, 49
- societal pressures, 145
- spontaneous, 79–80
- terminology, 11
- trust and, 5
- coordination games, 65
- corporate security officers, 201
- corporations. *See also* organizations
 - altruistic behavior, 181
 - competition and moral pressures, 181
 - defection and, 179–180, 189
 - institutional pressures, 185
 - licensing requirements, 186–187
 - name changes, 100
 - punishment, 189
 - reputation and, 181–183
 - risk trade-off, 180
 - societal pressure, manipulation, 189
 - societal pressure failure, 185–186
 - Thurlow, Baron, on, 181
- counterfeiting, 275–276
- counterterrorism policy, 215–216
- crime
 - Internet fraud, 146
 - not reporting, 146–147
 - organizations, 283
- Crime Prevention Techniques
 - chart, 130
- criminals, 12–13
 - as defectors (*See also* defectors)
 - organization, societal dilemmas and, 60
 - society of, 252
 - testifying against each other, 109, 124
- crises, spontaneous cooperation, 79–80
- CRISPRs (Clustered Regularly Interspaced Short Palindromic Repeats), 253
- cultural evolution
 - intelligence and, 22
 - speed, trust and, 42
- cultural morals, 80–81
 - timing, 144
- customers, reputation and, 88–89
- D**
- day care late pick-up fees, 114
- De Beers, 174–175
- Deacon, Terrence, 261
- Deacon's Paradox, 45
 - pair bonds, 106–107
- deception
 - detecting, 87, 89
 - frequency in primates, 23–24
 - intelligence evolution and, 22–23
 - neocortex and, 253
- Deepwater Horizon explosion, 170
- defection, 2
 - benefits, reduction, 68
 - Big Dig highway project, 170
 - BP Deepwater Horizon explosion, 170

- competing interests and, 139–143, 145
- consequences, 68
- versus cooperation, 34, 68
- cooperation and, 68
- corporations, 179–180, 189
- costs, 111–112
- costs/benefits, 31
- damage limits, 68
- defector groups, 100
- detecting defectors, 89
- difficulty, 68
- failure from widespread, 284
- Hobbes, Thomas, on, 104
- incentive, 217–219
- institutions, 203
- laws as punishment, 109
- laws' effects, 121
- Luther, Martin, 104
- Machiavelli on, 105
- observer and, 90
- organizations, 159, 213
- parasitic relationships, 28
- Prisoner's Dilemma and, 268–269
- punishment, 96–98
- Red Queen Effect and, 228–229
- reputation and, 89–90
 - minimizing effects, 99–100
- scope of, 13
 - societal pressures and, 68–69
- society's decisions, 209–210
- technology and, 228–229
- terminology, 11
- Ulrich, Laurel Thatcher, 144
- value versus reputation damage, 100–101
- Wall Street securities, 170
- whistle-blowers, 165–169
- defectors. *See also* criminals
 - famous, 14
 - good and bad, 246
 - need for, 247–248
 - rise of power, 244
 - risk and, 263
 - societal pressures
 - strategies for dealing with, 43–44
 - terminology, 11–13
 - vaccinations, 58
- defender, terminology, 13
- defenses, as security system, 125
- defrauding merchants, 59–60
- delegation
 - enforcing institutions, 201–202
 - societal pressures, 198–199
- desertion from military, 167–168
- detecting deception, 87, 89
- detection/response security
 - system, 125
- dialect, 270
 - cooperation and, 270
 - trust and, 92
- diamond merchants, 174–175
- Dictator game, 36, 257, 259
- Dilbert comic strip, 159
- The Dilbert Principle* (Adams), 284
- dilemmas. *See* societal dilemmas
- direct reciprocity, altruism and, 38, 259–260
- distrust, tax and, 43–44
- diversity, cooperation and, 270
- division of labor, 27, 44–45
- doctrine of scandal, 164
- doping in professional sports, 131–132, 274
 - EPO, 275
 - norandrosterone, 275

norbolethone, 275
 World Anti-Doping Agency, 201
 Drake, Thomas, 167
 dueling, 270
 Dunbar, Robin
 Dunbar number, 23–24, 46, 47, 253
 Facebook and, 262
 primate groups and neocortex,
 23–24

E

ecological validity, 261
 economic behavior, societal pressures,
 173–178
 ecosystems, 1–2
 ego-preservation interest, 139
 not reporting and, 146
 organizations, 156
 elephant seals, 255–256
 embezzlement, 158
 Emerson, Ralph Waldo, group interest
 and, 144
 empathy
 acute need *versus* chronic, 268
 cooperation and, 83
 groups and, 85
 energy conservation, 82–83, 267–268
 EPO (erythropoietin), doping
 and, 275
 establishing reflex, 253
 ethics. *See* morality
 ethnic diversity, trust and, 92
 evolution
 cultural
 intelligence and, 22
 speed and trust, 42
 genetic, speed and trust, 42
 intuition and, 8
 religion origins, 261
 and security, 18–20
 of societal pressures, 44
 evolutionary psychology, 260
 exemptions, 148
 experiential goods, 283
 externalities, 53

F

FAA (Federal Aviation
 Administration), 202
 Facebook
 Dunbar number and, 262
 privacy, 203–204
 face-to-face interaction, morality
 and, 84
 failure in trust, 3–4
 failure of societal pressures, 10
 fairness/reciprocity system of
 morality, 79
 false memories, advertising and, 281
 fear of punishment, honesty and,
 104–105
 Ferrucci, Piero, 251
 financial network, 4
 institutionalized liabilities, 111–112
 financial penalties, 110–112
 financial rewards, 114–115
 fines, institutional penalties, 185
 Food, Drug, and Cosmetic Act, 179
 Ford Explorer, 184–185
 forgiveness, *ilunga*, 90–91
 four-way stop signs, 262
 fracking, 187, 281
 fraction of the action at Sambo's, 160
 fraud, patterns, 259

Freakonomics, 264
 free riders, 58–59
 lying and, 82
 taxes and, 207–209
 free will, *versus* predestination, 266
 full-body scanners, 196–197

G

Gacy, John Wayne, 99
 game theory
 coordination games, 65
 parasites and, 28–29
 Prisoner's Dilemma, 51–53
 games, security systems and, 128
 gang members, groups and, 100
 gay marriage societal dilemma,
 221–222
 genetic determinism, 135
 genetic evolution, speed, trust and, 42
 genocide, moral pressures and, 83
 Ginger Paralysis Association,
 medicine terminology, 280
 Glaucon (*Republic*), on fear of
 punishment, 104
 Glaxo Smith-Kline, 184
 globalization, 4
 Golden Rule, 77
 expressions, 265–266
 Platinum Rule, 93
 gossip, reputation and, 89
 government
 angels *versus* men, 43
 authoritarian, 286
 newspapers publishing secrets, 236
 Social Contract Theory, 104
 graffiti, 81
 gridlock, 107, 128
 group interest
 moral interest and, 152–153
 trade-offs between competing
 interests, 238–239
 group norms, 81–82
 group reciprocity, 259–260
 group selection, altruism and, 38
 groups. *See also* competing interests
 cliques, 46
 Emerson, Ralph Waldo, 144
 empathy and, 85
 gang members, 100
 group interest, 11
 of another group, 140
 not reporting and, 146
 hunting, 24
 Hutterites, 102
 individuality and, 253
 megabands, 46
 membership markers, 92–93
 norms, 196
 ostracization, 45
 Prisoner's Dilemma and, 51–53
 reputational pressures and, 184–185
 size
 effectiveness, 272
 organizations and, 187–193
 reputational pressure and, 101
 specialization, 254
 social identity and, 277
 societal dilemmas, 9
 societal pressures, 50
 sympathy groups, 46
 Thoreau, Henry David, 144
 Tragedy of the Commons, 55–58
 tribes, 46
 types, 149
 guilt, social pressures and, 44

H

- Hamilton, Alexander, punishment, 110
- hand washing, security systems, 267
- Hardin, Garrett
 institutional pressure, 105
 Tragedy of the Commons, 55
- harm/care system of morality, 79
- Hawk-Dove model, 28–31, 254–255
 four-way stop signs, 262
 Prisoner's Dilemma and, 54
 reputation solving societal dilemma, 98–99
- herd immunity, vaccinations and, 58
- Hobbes, Thomas, on defection, 104
- honor among thieves, 80
- honor system, 83–84
- hotel towel theft, 123–124, 274
- human process model, 33
- Humphrey, Nicholas, 22–23
- Hutterites' community size, 102
- hyperbolic discounting rate, 42
- identity theft, 3
- ignorance (competing interests), 140
- ihsan*, 270
- ilunga*, 90–91
- immolation, 277
- immune system, CRISPRs, 253
- impersonal trust, 5
- impersonation fraud. *See* identity theft
- inbreeding, genetic disorders, 266
- incentives, 113–115
- incest. *See* kin aversion
- indirect reciprocity, 259–260
- altruism and, 38
- individuality and group size, 253
- in-group/loyalty system of morality, 79
- Innocence Project, 284
- inspections, 114
- institutional penalties, fines, 185
- institutional pressures, 9, 69–74. *See also* laws
 burglary, 105–106
 corporations, 185
 failure, 118–121
 financial penalties, 110–111
 financial rewards, 114–115
- Hardin, Garrett, 105
- incentives, 113–115
- inspections, 114
- intensity levels, 118
- laws and, 109–110
- legal loopholes, 119–120
- organizations, 157
- Ostrom, Elinor, 115–116
- reputational pressures and, 116–117
- resistance, 107–108
- security-augmented, 128
- taxation, 112–113
- voting, 106
- institutionalized liabilities, 111–112
- institutions, 47–48, 203
- intelligence
 cultural evolution and, 22
 deception and, 22–23
 murderousness and, 24–25
 social groups and, 24
 socialization and, 22–23
 time passage and, 253
 trust of non-kin, 34
 trust's effects, 45

intentional stance, 253
 Interdisciplinary Workshop on
 Security and Human Behavior, 252
 Internet, 146–147
 interventions, 125
 introduction letters, 91–92
 intuition
 context, 134–135
 evolution and, 8
 technology and, 134–135
 invasive species, 120–121

J

Jamaica Ginger during Prohibition,
 178–179
 judgment-proof, 282
 Just So Stories, 253, 260

K

Kahneman, Daniel, 264
 kidnapping, 107
 kin aversion, 80, 92–93
 kin helping kin, 33
 Kohlberg, Lawrence, six stages of
 moral development, 141–142

L

Lamb, Charles, stealthy good actions,
 270
 language, dialect, trust and, 92
 language instinct and moral
 instinct, 79
 large-game hunting, 24
 law. *See also* punishment

application, 120
 calculating crimes, 119
 commitment formalization, 117
 computer viruses, 210
 enforcing, 110, 119
 formalized coercion, 109
 intended effect, 118
 interpretation problems, 119
 legitimate behavior, 120–121
 licenses and, 186–187
 loopholes, 119–120
 reputation formalization and,
 116–117
 rulers and, 195–196
 sanctions, 110
 law breakers. *See also* defection
 as defectors, 13
 legal contracts, 117
 legitimate behavior, laws and,
 120–121
 lemons market, 184
 Lenin, Vladimir, 14
 letters of introduction, 91–92
 levels of trust, 251
 licenses, 117
 manipulating laws and, 186–187
 lock-ins, 281
 loopholes in laws, 119–120, 274
 Luther, Martin, on defection, 104
 lying, ethics breakdown and, 82

M

Machiavelli, Niccolò, defectors
 and, 105
 Machiavellian Intelligence
 Theory, 22–23
 Maine lobster V-notch practice, 271

- mall shopping
 - hours dilemma, 103
 - terrorism risk reduction, 276
- Manning, Bradley, 169, 279
- marriage
 - extra-marital sex, 261
 - mixed in Rwanda, 268
- martyrs, 144
- Maslow's needs, 140
- Massey Energy, 279
- meerkats, 255
- megaband, 46
- membership, 92–93, 116
- memes, moral contextualities, 79
- men *versus* angels, 43
- merchants
 - auctions, 176
 - cartels, 174–175
 - price fixing, 173–176
 - sellers options, 177–178
- Merck Serono, 185
- meta rules, 78
- military desertion, 167–168
- military organization, 261
- Minerals Management Service, 202
- mirror neurons, 36, 259
- Misztal, Barbara, functions performed
 - by trust, 6
- monogamy, 261
- monopolies, 280
- moral instinct, language instinct
 - and, 79
- moral interest, 11
 - competing moral interest, 140
 - group interest and, 152–153
- moral pressures, 8–9, 66–67, 69–74
 - failures, 84–85
 - market competition and, 181
 - organizations, 157
 - security-augmented, 127–128
 - totalitarian regimes and, 85
 - voting and, 76
- morality, 76
 - face-to-face interaction and, 84
 - martyrs, 144
 - moral meaning, levels, 277
 - morals *versus* ethics, 265
 - relative morals, 277–278
 - Socrates, 144
 - spiritual geniuses, 144
 - stages of moral development,
 - 141–142, 277
- morals, 76
 - conflicting, 84–85
 - contextualities, 78–79
 - cultural, 80–81
 - general, 77–78
 - Golden Rule, 77
 - kin aversion, 80
 - manipulation, 85
 - meta rules, 78
 - moral reminders, 77–78
 - overreaching, 85
 - religion and, 78
 - scaling, 85
 - social pressures and, 44
 - societal dilemmas and, 77–85
 - societal pressures and, 76, 81–82, 85
 - systems underlying, 79
 - universal, 77
- murder
 - murderousness and intelligence,
 - 24–25
 - as societal dilemma, 76
- mutualism, 33, 255
- mutualistic relationships, 28

N

natural world and security, 17
 Neanderthals, 261
 negative reputation, 116–117
 neocortex, deception and, 253
 nepotism
 and competing interests, 150–151
 politics, 278
 nesting structure of organizations,
 171–172
 network reciprocity, altruism and, 38,
 259–260
 neurons, mirror, 36, 259
 neuroplasticity of the brain, 42
 neuroscience, 260
 altruism and, 36
 non-zero-sum property, 56
 Nooteboom, Bart, 243
 norandrosterone, doping and, 275
 norbolethone, doping and, 275
 North, Oliver, 14
 not reporting competing interests, 146
 Nowak, Martin A., cooperative
 behavior, 38

O

observer, defection and, 90
 O'Callaghan, Sean, 149
 Olivieri, Nancy Fern, 167
 omertà, 13
 One.Tel, 278
 organizations, 155. *See also*
 corporations
 Beech-Nut Corporation apple
 juice, 163
 competing interests, 156

 criminal, 283
 defection, 159, 213
 Dilbert comic strip, 159
 embezzlement, 158
 group size, 187–193
 nesting structure, 171–172
 principal-agent problem, 159
 Robert Fish Organization
 example, 161
 Sambo's restaurants, 160
 Sarbanes-Oxley, 160
 security measures, *versus* job
 type, 159
 security systems, 124–125
 societal dilemmas, 158
 societal pressures, 156–157
 whistle-blowers, 165–169
 ostracization, 45
 Ostrom, Elinor
 rules, 273
 Tragedies of the Commons, 115–116
 overfishing societal dilemma, 55–57,
 161–162

P

pair bonds, 44–45
 Deacon's Paradox, 106–107
 parasites
 defectors and, 28
 ecosystems, 2
 game theory and, 28–29
 human organism and, 254
 reproduction, 19
 societal, 28
 parasitic relationships, 28
 Parker, Theodore, 287
 participation, 96–98

- Pascal's Wager, 270
- pasture example, Tragedy of
 Commons and, 262–263
- pertussis (whooping cough)
 vaccination, 57–58
- piracy, 107
 privateers, 272
- Platinum Rule, 93
- plea bargaining, 262
- pleasure, altruism and, 37
- police
 cooperation and, 147
 IRA and, 150
 Foucault, Michel, 272
- policy debates, 222–223
- power
 institutions, 203
 societal pressures and, 195
- predators and prey, 20–21
- preemptive interventions security
 system, 125
- price fixing, 174–176, 280
- principal-agent problem, 159
- Prisoner's Dilemma, 8, 51–53
 Alice and Bob, 270–271
 commerce and, 53
 cooperation and, 262, 269
 defection and, 268–269
 Hawk-Dove model, 54
 reputation and, 269
 risk trade-off and, 63
 shopping malls, 104
 social control and, 278
 versus The Stag hunt, 64–65
 Tragedy of the Commons and, 56
- prisons, 272
- privacy
 Facebook, 203–204
 social networking sites, 3
- privateers, 272
- Prohibition, 178–179
 Ginger Paralysis, medicine
 terminology, 280
- property rules, 59–60
- prosocial behaviors, 32
 experimenters' games, 35–36
 religion and, costly rituals, 271
- Prospect Theory, 264
- prostitution, 272
- psychological motivations (competing
 interest), 139–140
 not reporting, 146
 organizations, 156
- psychology, evolutionary, 260
- Public Goods game, 36, 258
 player punishment, 259
- public money for projects, 186–187
- punishment. *See* sanctions
- purity/sanctity system of morality, 79
- ## R
- Rational Choice Theory, 140–141
- real estate collapse, 3
- reasoning, argumentative theory, 253
- reciprocal altruism, 34, 259–260,
 261, 273
- recovery security system, 125
- recycling, 283–284
- Red Queen Effect, 20, 228
 feedback loop, 21
 murder and, 25
 societal dilemmas and, 87
- regulatory capture, 202
- relational interest, 11, 140
 not reporting and, 146

- relative morals, 277–278
- religion
 - costly rituals and prosocial behaviors, 271
 - morals and, 78
 - origins, 261
 - Quakers, trust and, 91–92
 - reputation and observance, 94
 - Soviet prohibition, 85
- Republic* (Plato), fear of punishment, 104
- reputation
 - advertising, 90
 - altruism and, 181
 - branding and, 95
 - chimpanzees, 270
 - commerce and, 91
 - commitment and, 93–94
 - confidence tricksters, 99
 - cooperation and, 89–90
 - corporate name changes, 100
 - corporations and, 181–183
 - customers and, 88–89
 - defection and, 89–90
 - minimizing effects, 99–100
 - doctrine of scandal, 164
 - forgiveness, *ilunga*, 90–91
 - game experiments, 269
 - gossip and, 89
 - group membership and, 92–93
 - groups and, 184–185
 - laws formalizing, 116–117
 - merchants and, 53
 - negative, 116–117
 - observer and, 90
 - Prisoner's Dilemma and, 269
 - Quakers and, 91–92
 - religious observance, 94
 - scale, 91–92
 - social pressures and, 44, 269
 - societal pressures and, failure, 99–102
 - systemizing, 95–96
 - untrue stories, 100
- reputational pressures, 9, 69–74
 - customers and, 88–89
 - genetic determinism and, 135
 - group size, 101
 - institutional pressure and, 116–117
 - institutions and, 47
 - organizations, 157
 - security-augmented, 128
 - technological security addition, 133–134
- research
 - findings, 252
 - methodology, 252
- revenge, 97
- revenue streams, 190–191
- reward/punishment
 - cooperation and, 88, 96–98
 - institutional pressures, 113–115
- risk perception, 215
- risk trade-off, 51
 - arms race and, 54
 - Bad Apple Effect, 82
 - corporations, 180
 - Prisoner's Dilemmas and, 63
 - societal pressures and, 71–73
 - vaccinations, 57–58
- ritual, commitment and, 94
- Robert Fish Corporation example, 161
- Roman Catholic Church societal dilemma, 163–164
 - bishops' role, 279
 - Crimen Sollicitationis*, 167

Rousseau, Jean-Jacques, Stag Hunt, 64–65
 Rwanda, mixed marriages, 268

S

SafeHouse, 285
 sales techniques, 176–177
 Sambo's restaurants, 160
 sanctions, 110, 272. *See also*
 punishment; reward/punishment
 Bentham, Jeremy, 274
 coercion, formalizing, 109
 corporations, 189
 defection, 96–98
 fear of, honesty and, 104–105
 financial penalties, 110–111
 Hamilton, Alexander, on, 110
 laws and, 109
 Sarbanes-Oxley, 160
 scale
 competing interests, 143
 morals, 85
 reputation, 91–92
 societal pressures, 70–71,
 225–227
 technological security
 systems, 135
 society, technology and, 48
 trust needs and, 46
 Schindler, Oskar, 14
 science exploration, 8
 scope of defection, 13
 overfishing and, 56
 societal pressures and, 68–69
 SEC (Securities and Exchange
 Commission), 202
 security
 diminishing returns, 136
 genetic determinism, 135
 history of, 17–25
 natural world and, 17
 as societal pressure, 7
 too much, 136
 security gap, 231–232
 security systems
 audit/forensic systems, 125
 augmented opening for attack,
 133–134
 cheating at games, 129
 Clarke, Ronald V., chart, 129–130
 defenses, 125
 detection/response systems, 125
 doping in professional sports,
 131–132
 expense, 135–136
 failure, 132–133
 gridlock and, 128
 interventions, 125
 organizations, 124–125, 157
 preemptive interventions, 125–126
 recovery systems, 125
 societal pressures and, 9, 126–136
 vaccines and, 128
 security-augmented pressures,
 127–128
 self-immolation, 277
 self-interest, 11
 defecting and, 13
 Tragedy of the Commons and, 55–58
 selfish competing interest
 (organizations), 156
 selfish self-interest competing
 interest, 139
 not reporting, 145
 self-preservation interest, 139

- competing interest
 - (organizations), 156
 - not reporting and, 145–146
- sellers options, 177–178
- sexual fidelity, symbolic contracts
 - and, 45
- shame, 97
- Shipman, Dr. Harold, 99
- shopping malls, Prisoner's Dilemma
 - and, 104
- signals
 - of altruism, 35
 - verifying, 271
- sippenhaft*, 98
- Smith, Adam, trust, 37
- snakes, children identifying, 253
- Snowdrift Dilemma, 65
 - Prisoner's Dilemma and, 264
- Social Brain Hypothesis, 22–23
- Social Contract Theory, 104
- social controls, 261–262
 - Prisoner's Dilemma and, 278
- social dilemmas, 65
- social identity, group strength
 - and, 277
- social interaction
 - groups, intelligence and, 24
 - neocortex and, 23–24
- social networking, privacy, 3
- social order. *See also* Social Contract Theory
- social pressures
 - cooperation and, 49
 - morals and, 44
 - pair bonds and, 45
 - reputation and, 44
- social systems, socio-technical
 - systems and, 134
- socialist capitalism, 281
- socialization, human intelligence and, 22–23
- societal dilemmas, 8. *See also*
 - competing interests
 - airplane security
 - implementation, 200
 - airplane terrorism, 197–199
 - cheating customers, 88
 - cheating students, 218
 - commerce and, 53
 - competing, 151–152
 - competing interests, 139–140, 141
 - competing on to-buy prices, 176
 - contracts, 117
 - cooperating with police, 147
 - corporate expenses, 129
 - criminal organization, 60
 - criminals, testifying against each other, 109, 124
 - defrauding merchants, 59–60
 - different people and, 65
 - doping in professional sports, 131–132
 - fracking, 187
 - gay marriage, 221–222
 - groups as actors, 9
 - interaction misunderstanding, 219–220
 - mall hours, 103
 - military desertion, 168
 - moral pressure, 66–67
 - morals and, 77–85
 - murder, 76
 - newspapers publishing government secrets, 236
 - non-zero-sum property, 56
 - overfishing, 55–57, 162

- policy debates, 222–223
- price setting, 174–176, 280
- Prisoner's Dilemma, 51–53
- protecting pedophiles, 164
- public money for projects, 186–187
- Red Queen Effect and, 87
- Roman Catholic Church, 163–164
- sales techniques, 176–177
- security systems, 128–129
- society's solutions, 64
- speeding, 111
- stealing, 67
- symmetry, 53
- technological advances and, 219
- terminology, 11
- Tragedy of the Commons, 55–58
- vaccinations, 57–58
- voting, 75–76
- whistle-blowing, 165–169
- working within an organization, 158
- worshipping Roman gods, 220–221
- societal pressures
 - cameras and, 126–127
 - change and, 50
 - Code of Hammurabi and, 50
 - compliance and, 49
 - cooperation and, 145
 - corporations, 189
 - defectors, 245
 - delegation, 198–199
 - design tips, 240–241
 - economics and, 173–178
 - evolution, 44
 - failure, 10, 185–186
 - changing social norms, 220–221
 - dilemmas encouraging
 - deception, 216
 - incentive misunderstanding, 213
 - incentive to defect, 217–219
 - mistakes, 212–213
 - misunderstood actors, 213
 - risk misinterpretation, 215
 - societal dilemma interaction, 219–220
 - taxes, 207–209
 - feedback loops, 211–212
 - goals, 209
 - gossip and, 89
 - as immune system, 74
 - implementing delay, 210
 - inspections, 114
 - institutional, 9, 69–74
 - security-augmented, 128
 - moral, 8–9, 69–74
 - group norms, 81–82
 - security-augmented, 127–128
 - morals and, 76
 - weakest of pressures, 85
 - organizations, 156–157
 - personal trash, 214
 - power and, 195
 - propagation, 211
 - Red Queen Effect, 228
 - reputational, 9, 69–74
 - failure, 99–102
 - security-augmented, 128
 - risk trade-off and, 71–73
 - scale, 70–71, 225–227
 - technological security
 - systems, 135
 - scope of defection and, 68–69
 - security, 7
 - security systems, 9, 126–136
 - sippenhaft*, 98
 - technology and, 49
 - terminology, 13

trust and, 7, 243–244

society

- actors, 10
- of criminals, 252
- group interests, 11
- parasites, 28
- scale, 48–49
- terminology, 10
- trust and, 2
 - health of society, 6
 - inducing, 7

society decides, 195

socio-technical systems, 134

Socrates, morals, 144

soldier parable, 247

spam, 3

specialization, group size, 254

speeding, 111

spiritual geniuses, 144

spontaneous cooperation, 79–80

sports doping, 131–132, 274

- EPO, 275
- norandrosterone, 275
- norboletone, 275
- World Anti-Doping Agency, 201

Stag Hunt, 64–65

stages of moral development, 141–142, 277

stealthy good actions, 270

stereotyping, 93

Stross, Charles, 239

superrationality, 264

sybiosis, 255

symbolic contracts, sexual fidelity and, 45

sympathy group, 46

systems, reputation and, 95–96

T

Target return policy, 269

taxes

- avoidance, 208
- Bad Apple Effect and, 208
- distrust and, 43–44
- free-riders and, 207–209
- as institutional pressure, 112–113

technology

- advances, societal pressure and, 225–242
- defectors and, 228–229
- information technology
 - defection and, 233–235
- intuition and, 134–135
- law enforcement, 229
- scale and, 48
- security additions to reputational systems, 133–134
- security gap, 231–232
- societal complexity and, 50
- societal pressures and, 49
- socio-technical systems, 134
- specialization and, 49–50
- trade-offs between interests, 238–239

telegraph, 233

terrorism. *See also* TSA

- airplane terrorism, 197
- chemical plants and, 191–192
- counterterrorism policy, 215–216
- full-body scanners, 196–197
- mall shopping, 276
- politics and, 282
- Thatcher, Margaret and, 108

Thatcher, Margaret, terrorist

negotiations, 108
 theft of hotel towels, 123–124
 Thoreau, Henry David, group
 norm, 144
 Thurlow, Baron, 181
 time passage, 253
 totalitarian regimes, moral codes
 and, 85
 Tragedy of the Commons, 55–58
 criticisms, 264
 Ostrom, Elinor, 115–116
 pasture example, 262–263
 upper limit for commons, 101
 transactional analysis, 35
 transferring, 98
 trash burning, 214
 tribes, 46
 trust, 251–252
 cooperation and, 5
 definitions, 5
 division of labor and, 27, 44–45
 ethnic diversity and, 92
 evolutionary advantage, 32–33
 failure in, 3–4
 Ferrucci, Piero, 251
 healthy society and, 6
 history of, 41–50
 human need for, 37
 impersonal, 5
 intelligence and, 45
 kin helping kin, 32–33
 letters of introduction, 91–92
 levels, 251
 Misztal, Barbara, 6
 natural world examples, 33
 necessity, 8–9
 others' trustworthiness and, 35

pair bonds and, 44–45
 primates, 41
 Quakers and, 91–92
 scale and, 46
 Smith, Adam, 37
 societal pressures and, 243–244
 society and, 2
 systems *versus* people, 6
 Trust game, 36, 257
 TSA (Transportation Security
 Administration), 199–201
 full-body scanners, 197
 implementing security, 200

U

Ulrich, Laurel Thatcher, 144
 Ultimatum game, 36, 256–257
 fairness and, 258–259
 Machiguenga tribe and, 266–267
 universal moral, 77

V

vaccinations, 57–58
 security systems and, 128
 violence, intelligence and, 24–25
 vision, security and, 19
 voting
 institutional pressure and, 106
 by mail, 265
 as societal dilemma, 75–76

W–Z

Wall Street securities, 170
 Wason Selection Task, 87

water, bottled, 182

weapons, chimpanzees, 252

whistle-blowers, 165–169, 279

WikiLeaks, 237

 Manning, 169, 279

word usages. *See* terminology

World Anti-Doping Agency, 201

World Values Survey, 286

zero-sum games, 263

"The closest the security industry has to a rock star."
—*The Register*

BRUCE SCHNEIER

A portrait of Bruce Schneier, a man with a beard and a ponytail, wearing a dark suit jacket over a light blue shirt. He is looking directly at the camera with a neutral expression. The background is a plain, light grey color.

CARRY ON

SOUND ADVICE FROM
SCHNEIER ON SECURITY

WILEY

Carry On

Carry On

Sound Advice
from Schneier
on Security

Bruce Schneier

WILEY

Carry On: Sound Advice from Schneider on Security

Published by
John Wiley & Sons, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2014 by Bruce Schneier

Published by John Wiley & Sons, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-1-118-79081-6
ISBN: 978-1-118-79083-0 (ebk)
ISBN: 978-1-118-79082-3 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or website may provide or recommendations it may make. Further, readers should be aware that Internet websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2013954201

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Credits

Executive Editor

Carol Long

Project Editor

Tom Dinse

Senior Production Editor

Kathleen Wisor

Editorial Manager

Mary Beth Wakefield

Freelancer Editorial Manager

Rosemarie Graham

Associate Director of Marketing

David Mayhew

Marketing Manager

Ashley Zurcher

Business Manager

Amy Knies

Vice President and Executive Group

Publisher

Richard Swadley

Associate Publisher

Jim Minatel

Project Coordinator, Cover

Katie Crocker

Proofreader

Nancy Carrasco

Indexer

Johnna VanHoose Dinse

Cover Image

Steve Woit

Cover Designer

Ryan Sneed

Contents

Introduction	xv
1 The Business and Economics of Security	1
Consolidation: Plague or Progress.	1
Prediction: RSA Conference Will Shrink Like a Punctured Balloon	2
How to Sell Security	4
Why People Are Willing to Take Risks	4
How to Sell Security.	6
Why Do We Accept Signatures by Fax?	7
The Pros and Cons of LifeLock	9
The Problem Is Information Insecurity	12
Security ROI: Fact or Fiction?	14
The Data Imperative	15
Caveat Emptor.	16
Social Networking Risks	17
Do You Know Where Your Data Are?	18
Be Careful When You Come to Put Your Trust in the Clouds	21
Is Perfect Access Control Possible?	22
News Media Strategies for Survival for Journalists.	24
Security and Function Creep	26
Weighing the Risk of Hiring Hackers	27
Should Enterprises Give In to IT Consumerization at the Expense of Security?	29
The Vulnerabilities Market and the Future of Security	30
So You Want to Be a Security Expert.	33
When It Comes to Security, We're Back to Feudalism	34
I Pledge Allegiance to the United States of Convenience	35
The Good, the Bad, and the Ugly	36
You Have No Control Over Security on the Feudal Internet	37

2	Crime, Terrorism, Spying, and War	41
	America's Dilemma: Close Security Holes, or Exploit Them Ourselves	41
	Are Photographers Really a Threat?	43
	CCTV Doesn't Keep Us Safe, Yet the Cameras Are Everywhere	45
	Chinese Cyberattacks: Myth or Menace?	47
	How a Classic Man-in-the-Middle Attack Saved Colombian Hostages	48
	How to Create the Perfect Fake Identity	51
	A Fetishistic Approach to Security Is a Perverse Way to Keep Us Safe.	52
	The Seven Habits of Highly Ineffective Terrorists	54
	Why Society Should Pay the True Costs of Security.	56
	Why Technology Won't Prevent Identity Theft	58
	Terrorists May Use Google Earth, but Fear Is No Reason to Ban It.	60
	Thwarting an Internal Hacker	62
	An Enterprising Criminal Has Spotted a Gap in the Market.	65
	We Shouldn't Poison Our Minds with Fear of Bioterrorism	66
	Raising the Cost of Paperwork Errors Will Improve Accuracy.	68
	So-Called Cyberattack Was Overblown	70
	Why Framing Your Enemies Is Now Virtually Child's Play	72
	Beyond Security Theater.	73
	Feeling and Reality.	74
	Refuse to Be Terrorized	76
	Cold War Encryption Is Unrealistic in Today's Trenches.	77
	Profiling Makes Us Less Safe	80
	Fixing Intelligence Failures	81
	Spy Cameras Won't Make Us Safer	82
	Scanners, Sensors Are Wrong Way to Secure the Subway.	84
	Preventing Terrorist Attacks in Crowded Areas	86
	Where Are All the Terrorist Attacks?	87

Hard to Pull Off	88
Few Terrorists.	88
Small Attacks Aren't Enough	89
Worst-Case Thinking Makes Us Nuts, Not Safe	89
Threat of "Cyberwar" Has Been Hugely Hyped.	92
Cyberwar and the Future of Cyber Conflict.	94
Why Terror Alert Codes Never Made Sense	96
Debate Club: An International Cyberwar Treaty Is the Only Way to Stem the Threat	97
Overreaction and Overly Specific Reactions to Rare Risks	99
Militarizing Cyberspace Will Do More Harm Than Good.	101
Rhetoric of Cyber War Breeds Fear—and More Cyber War	103
Attacks from China	103
GhostNet.	104
Profitable.	105
The Boston Marathon Bombing: Keep Calm and Carry On	105
Why FBI and CIA Didn't Connect the Dots	107
The FBI's New Wiretapping Plan Is Great News for Criminals	109
US Offensive Cyberwar Policy.	112
3 Human Aspects of Security	117
Secret Questions Blow a Hole in Security	117
When You Lose a Piece of Kit, the Real Loss Is the Data It Contains.	118
The Kindness of Strangers	120
Blaming the User Is Easy—But It's Better to Bypass Them Altogether	122
The Value of Self-Enforcing Protocols	123
Reputation Is Everything in IT Security	125
When to Change Passwords	127
The Big Idea: Bruce Schneier	129
High-Tech Cheats in a World of Trust	131

Detecting Cheaters.	134
Lance Armstrong and the Prisoner's Dilemma of Doping in Professional Sports.	137
The Doping Arms Race as Prisoner's Dilemma	138
The Ever-Evolving Problem.	139
Testing and Enforcing	140
Trust and Society	141
How Secure Is the Papal Election?.	143
The Court of Public Opinion.	147
On Security Awareness Training	150
Our New Regimes of Trust	152
4 Privacy and Surveillance	155
The Myth of the "Transparent Society"	155
Our Data, Ourselves.	157
The Future of Ephemeral Conversation.	158
How to Prevent Digital Snooping	160
Architecture of Privacy	162
Privacy in the Age of Persistence.	164
Should We Have an Expectation of Online Privacy?	167
Offhand but On Record	168
Google's and Facebook's Privacy Illusion.	171
The Internet: Anonymous Forever.	173
A Taxonomy of Social Networking Data	175
The Difficulty of Surveillance Crowdsourcing	177
The Internet Is a Surveillance State	179
Surveillance and the Internet of Things.	181
Government Secrets and the Need for Whistleblowers.	184
Before Prosecuting, Investigate the Government	187

5 Psychology of Security 189

- The Security Mindset 189
- The Difference between Feeling and Reality in Security 191
- How the Human Brain Buys Security 194
- Does Risk Management Make Sense? 195
- How the Great Conficker Panic Hacked into Human Credulity 197
- How Science Fiction Writers Can Help, or Hurt, Homeland Security 198
- Privacy Salience and Social Networking Sites 201
- Security, Group Size, and the Human Brain 203
- People Understand Risks—But Do Security Staff Understand People? 205
- Nature’s Fears Extend to Online Behavior 206

6 Security and Technology 209

- The Ethics of Vulnerability Research 209
- I’ve Seen the Future, and It Has a Kill Switch. 211
- Software Makers Should Take Responsibility 212
- Lesson from the DNS Bug: Patching Isn’t Enough. 214
- Why Being Open about Security Makes Us All Safer in the Long Run 216
- Boston Court’s Meddling with “Full Disclosure” Is Unwelcome. 218
- Quantum Cryptography: As Awesome as It Is Pointless 220
- Passwords Are Not Broken, but How We Choose Them Sure Is. 222
- America’s Next Top Hash Function Begins 223
- Tigers Use Scent, Birds Use Calls—Biometrics Are Just Animal Instinct. 225
- The Secret Question Is: Why Do IT Systems Use Insecure Passwords? 227
- The Pros and Cons of Password Masking 229
- Technology Shouldn’t Give Big Brother a Head Start 231
- Lockpicking and the Internet 233

The Battle Is On against Facebook and Co. to Regain Control of Our Files.	235
The Difficulty of Un-Authentication	237
Is Antivirus Dead?	238
Virus and Protocol Scares Happen Every Day— but Don't Let Them Worry You.	240
The Failure of Cryptography to Secure Modern Networks	242
The Story behind the Stuxnet Virus	244
The Dangers of a Software Monoculture	247
How Changing Technology Affects Security.	249
The Importance of Security Engineering.	251
Technologies of Surveillance.	253
When Technology Overtakes Security	255
Rethinking Security	255
7 Travel and Security	259
Crossing Borders with Laptops and PDAs	259
The TSA's Useless Photo ID Rules	261
The Two Classes of Airport Contraband	262
Fixing Airport Security	264
Laptop Security while Crossing Borders	265
Breaching the Secure Area in Airports	268
Stop the Panic on Air Security.	269
A Waste of Money and Time.	271
Why the TSA Can't Back Down.	273
The Trouble with Airport Profiling	275
8 Security, Policy, Liberty, and Law	279
Memo to Next President: How to Get Cybersecurity Right.	279
CRB Checking	281
State Data Breach Notification Laws: Have They Helped?	283
How to Ensure Police Database Accuracy	285
How Perverse Incentives Drive Bad Security Decisions	287

It's Time to Drop the "Expectation of Privacy" Test 288

Who Should Be in Charge of Cybersecurity? 291

Coordinate, but Distribute Responsibility 294

"Zero Tolerance" Really Means Zero Discretion 295

US Enables Chinese Hacking of Google 297

Should the Government Stop Outsourcing
Code Development? 299

Punishing Security Breaches 300

Three Reasons to Kill the Internet Kill Switch Idea 302

 Internet without Borders 302

 Unpredictable Side Effects 303

 Security Flaws 303

Web Snooping Is a Dangerous Move 304

The Plan to Quarantine Infected Computers 307

Close the Washington Monument 310

Whitelisting and Blacklisting 312

Securing Medical Research: a Cybersecurity Point of View 313

Fear Pays the Bills, but Accounts Must Be Settled 317

Power and the Internet 319

Danger Lurks in Growing New Internet Nationalism 321

IT for Oppression 323

The Public/Private Surveillance Partnership 325

Transparency and Accountability Don't Hurt Security—
They're Crucial to It 327

It's Smart Politics to Exaggerate Terrorist Threats 329

References 333

Index 347

Introduction

I like writing essays. I like the length: 600 to 1,200 words is my personal sweet spot. I like the format: a tight argument designed to make a particular point. And I like the style: explaining complicated topics to a lay audience is something I do well. Books are long, both in actual words and in the time they take to write. But I can write an essay in a fit of inspiration in a morning, and get it published the next day if everything goes well.

Not that it always goes that well, of course. Some essays are harder to write than others, and some are *very* hard. I like to take a few days to consider an issue before I write about it, which means that mine is generally not the first essay on the Internet after a news event. Editors, of course, hate this. They want something that catches the current news cycle.

Still, writing is something I'm good at and something I do a lot of. Since 1992, I have written almost 500 essays, op-eds, and articles for a wide variety of publications. They're all on my website—www.schneier.com, if you don't already know—and a selection of them has been collected into two books. The first collection, *Schneier on Security*, covered essays from April 2002 to February 2008. This volume covers essays from March 2008 to June 2013.

Looking back at the entire body of work, I have some lessons, observations, and advice for others trying to get their own articles published. And while my writing is mostly about security, much of the advice is general.

- **Opinions are cheap.** Charles McCabe famously said, “Any clod can have the facts, but having opinions is an art.” He’s right, but it doesn’t follow that any clod lacks an opinion. On the Internet, opinions are a dime a dozen. I rarely get paid for my essays. Oh, there were a few fun years where *Wired* paid me to write a regular column, but they eventually realized that it was cheaper to not bother paying me for them, since I was going to keep writing in any case. I’m not saying that it’s impossible to get paid for writing opinions—of course it is—only that it’s increasingly rare and difficult.
- **Persuading someone is hard—and rare.** My goal is to write persuasive essays, but I doubt they do a lot of actual persuading. More often, I’m writing to people who already agree with me, giving them new ways to think about the issue, or new words to use when doing their own persuading.

- **It's hard not to repeat yourself.** I write for many different audiences, often on very similar topics. And I often repeat myself. If I find a turn of phrase I like, I reuse it. If I have a perfect paragraph on a topic, I'm likely to use it again. I used to write restaurant reviews semiprofessionally, and would regularly complain about how few ways there are to say "this tastes good." It's not that bad in my security writing, but sometimes it feels as if it comes close.
- **Stories repeat.** Again and again, essays I wrote five or ten years ago suddenly become relevant after some news event. The essay I wrote about data mining in 2001 was important after the Boston Marathon bombing. The essay I wrote about fingerprint scanners in 1998 was important when Apple released an iPhone with a fingerprint scanner. The essay I wrote about Chinese cyberattacks in 2008 has been pertinent every couple of years since then. Drug testing in sports, TSA security, the value of privacy, ubiquitous surveillance, security against lone shooters: it all becomes relevant again after a news event. Sometimes I dust off an old essay, tack on a new introduction, and republish it. But most of the time I try for a new perspective. I don't like resaying old things, even if they are new again.
- **Editors rewrite.** Sometimes they only rewrite a little, but sometimes they rewrite a lot. Sometimes their rewrites are improvements, and sometimes they're just different. It's okay to push back on rewrites that don't improve your work. Once I refused to let a publication publish an essay of mine because they changed too much and wouldn't change things back. And there were a few times I wish that I'd yanked essays where the editors cut too much.
- **The headline isn't your problem.** As an essay writer, you don't get any say in your headline. If you're lucky, you'll get to see it before it's published, but you probably won't. The headline is how the publication entices readers to your essay. As such, it'll be more sensationalist than you want. Or it'll be simpler than you want. Or it'll be less descriptive than you want. Let it go—you can't change it.
- **Links rot.** It's frustrating, but they do. Links you include in essays you write today are likely to return "Page not found" errors a few years from now. For my last volume of essays, I included links at the end. I was going to do the same for this book, but link-checking showed that almost a tenth of them were already dead. These aren't ancient essays; the oldest are six years old and the newest are current.

- **Mistakes happen.** Don't be afraid to admit your mistakes. If you're going to write in anything resembling real time, your writing will sometimes contain errors—of fact, of logic, of conclusion, of opinion—of pretty much everything. When you do, admit them. Don't hedge. Don't mumble. Just admit them. You'll feel better, and your audience will respect you for it.
- **Opinions can change.** Don't be afraid to change your mind. If you're going to write over anything resembling a reasonable length of time, you're going to change your mind about some things. Maybe you'll discover new facts that cause you to reach different conclusions. Maybe you'll just think about things in a new light and reach different conclusions. That's fine. Just explain it. John Maynard Keynes said, "When the facts change, I change my mind. What do you do, sir?" Exactly.
- **You need to write in order to be read.** The world is full of people with great ideas who never make them available to the wider world. My first rule of writing is that you can't improve it until it's written down. So write that first draft; it's really the only way you'll see the weaker parts of your argument. (The world is also filled with people with terrible ideas who make them available to everyone—but that's a separate problem.)
- **Beta readers are important.** Cultivate a stable of them. The more people you have reading your essays before publication, the better your writing will be. Don't be afraid of criticism. Divorce your ego from your writing. That's the key for accepting criticism, and being able to process and use it; you can't let your ego interfere with hearing what your beta readers are telling you. The way I think of it is that people will criticize my work regardless, but if they criticize a draft, I have the opportunity to fix it before publication. Almost all of my essays have been improved by someone else's comments on an early draft, and some of my essays would have been terrible without those improvements.

When I write a book, it's easy to thank the people who read and commented on it. It's impossible to do the same with essays. So here, in this collection of essays, I would like to thank all the people who have read and commented on essay drafts: David M. Perry, Greg Guerin, Steve Bass, Bill Herdle, David Prentiss, Vicki Laidler, Stephen Leigh, Moshe Yudkowsky, Jon Callas, Doug Whiting, Stefan Lucks, and Jesse Walker. I apologize for any names I inadvertently omitted. I haven't kept a list, and I know I'm not remembering everybody.

So, welcome to my second collection of essays. I think there's something in here for everyone's tastes, as long as their tastes include security: technology and security, economics and security, psychology and security, politics and security. I'm still writing, and will probably publish a third volume of these in five or so years. Thanks for reading.

Bruce Schneier

The essays in this book previously appeared in various publications and may follow the usage conventions of the original publishers.

Carry On

1

The Business and Economics of Security

Consolidation: Plague or Progress

Originally published in Information Security, March 2008

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

We know what we don't like about buying consolidated product suites: one great product and a bunch of mediocre ones. And we know what we don't like about buying best-of-breed: multiple vendors, multiple interfaces, and multiple products that don't work well together. The security industry has gone back and forth between the two, as a new generation of IT security professionals rediscovers the downsides of each solution.

The real problem is that neither solution really works, and we continually fool ourselves into believing whatever we don't have is better than what we have at the time. And the real solution is to buy results, not products.

Honestly, no one wants to buy IT security. People want to buy whatever they want—connectivity, a Web presence, email, networked applications, whatever—and they want it to be secure. That they're forced to spend money on IT security is an artifact of the youth of the computer industry. And sooner or later the need to buy security will disappear.

It will disappear because IT vendors are starting to realize they have to provide security as part of whatever they're selling. It will disappear because organizations are starting to buy services instead of products, and demanding security as part of those services. It will disappear because the security industry will disappear as a consumer category, and will instead market to the IT industry.

The critical driver here is outsourcing. Outsourcing is the ultimate consolidator, because the customer no longer cares about the details. If I buy my network services from a large IT infrastructure company, I don't care if it secures things by installing the hot new intrusion prevention systems, by configuring

the routers and servers as to obviate the need for network-based security, or if it uses magic security dust given to it by elven kings. I just want a contract that specifies a level and quality of service, and my vendor can figure it out.

IT is infrastructure. Infrastructure is always outsourced. And the details of how the infrastructure works are left to the companies that provide it.

This is the future of IT, and when that happens we're going to start to see a type of consolidation we haven't seen before. Instead of large security companies gobbling up small security companies, both large and small security companies will be gobbled up by non-security companies. It's already starting to happen. In 2006, IBM bought ISS. The same year BT bought my company, Counterpane, and last year it bought INS. These aren't large security companies buying small security companies; these are non-security companies buying large and small security companies.

If I were Symantec and McAfee, I would be preparing myself for a buyer.

This is good consolidation. Instead of having to choose between a single product suite that isn't very good or a best-of-breed set of products that don't work well together, we can ignore the issue completely. We can just find an infrastructure provider that will figure it out and make it work—who cares how?

Prediction: RSA Conference Will Shrink Like a Punctured Balloon

Originally published in Wired News, April 17, 2008

Last week was the RSA Conference, easily the largest information security conference in the world. More than 17,000 people descended on San Francisco's Moscone Center to hear some of the more than 250 talks, attend I-didn't-try-to-count parties, and try to evade over 350 exhibitors vying to sell them stuff.

Talk to the exhibitors, though, and the most common complaint is that the attendees aren't buying.

It's not the quality of the wares. The show floor is filled with new security products, new technologies, and new ideas. Many of these are products that will make the attendees' companies more secure in all sorts of different ways. The problem is that most of the people attending the RSA Conference can't understand what the products do or why they should buy them. So they don't.

I spoke with one person whose trip was paid for by a smallish security firm. He was one of the company's first customers, and the company was proud to parade him in front of the press. I asked him whether he walked through the show floor, looking at the company's competitors to see if there was any benefit to switching.

"I can't figure out what any of those companies do," he replied.

I believe him. The booths are filled with broad product claims, meaningless security platitudes and unintelligible marketing literature. You could walk into a booth, listen to a five-minute sales pitch by a marketing type, and still not know what the company does. Even seasoned security professionals are confused.

Commerce requires a meeting of the minds between buyer and seller, and it's just not happening. The sellers can't explain what they're selling to the buyers, and the buyers don't buy because they don't understand what the sellers are selling. There's a mismatch between the two; they're so far apart that they're barely speaking the same language.

This is a bad thing in the near term—some good companies will go bankrupt and some good security technologies won't get deployed—but it's a good thing in the long run. It demonstrates that the computer industry is maturing: IT is getting complicated and subtle, and users are starting to treat it like infrastructure.

For a while now I have predicted the death of the security industry. Not the death of information security as a vital requirement, of course, but the death of the end-user security industry that gathers at the RSA Conference. When something becomes infrastructure—power, water, cleaning service, tax preparation—customers care less about details and more about results. Technological innovations become something the infrastructure providers pay attention to, and they package it for their customers.

No one wants to buy security. They want to buy something truly useful—database management systems, Web 2.0 collaboration tools, a company-wide network—and they want it to be secure. They don't want to have to become IT security experts. They don't want to have to go to the RSA Conference. This is the future of IT security.

You can see it in the large IT outsourcing contracts that companies are signing—not security outsourcing contracts, but more general IT contracts that include security. You can see it in the current wave of industry consolidation: not large security companies buying small security companies, but non-security companies buying security companies. And you can see it in

the new popularity of software as a service: Customers want solutions; who cares about the details?

Imagine if the inventor of antilock brakes—or any automobile safety or security feature—had to sell them directly to the consumer. It would be an uphill battle convincing the average driver that he needed to buy them; maybe that technology would have succeeded and maybe it wouldn't. But that's not what happens. Antilock brakes, airbags and that annoying sensor that beeps when you're backing up too close to another object are sold to automobile companies, and those companies bundle them together into cars that are sold to consumers. This doesn't mean that automobile safety isn't important, and often these new features are touted by the car manufacturers.

The RSA Conference won't die, of course. Security is too important for that. There will still be new technologies, new products and new startups. But it will become inward-facing, slowly turning into an industry conference. It'll be security companies selling to the companies who sell to corporate and home users—and will no longer be a 17,000-person user conference.

How to Sell Security

Originally published in CIO, May 26, 2008

It's a truism in sales that it's easier to sell someone something he wants than a defense against something he wants to avoid. People are reluctant to buy insurance, or home security devices, or computer security anything. It's not they don't ever buy these things, but it's an uphill struggle.

The reason is psychological. And it's the same dynamic when it's a security vendor trying to sell its products or services, a CIO trying to convince senior management to invest in security or a security officer trying to implement a security policy with her company's employees.

It's also true that the better you understand your buyer, the better you can sell.

Why People Are Willing to Take Risks

First, a bit about Prospect Theory, the underlying theory behind the newly popular field of behavioral economics. Prospect Theory was developed by Daniel Kahneman and Amos Tversky in 1979 (Kahneman went on to win a Nobel Prize for this and other similar work) to explain how people make

trade-offs that involve risk. Before this work, economists had a model of “economic man,” a rational being who makes trade-offs based on some logical calculation. Kahneman and Tversky showed that real people are far more subtle and ornery.

Here’s an experiment that illustrates Prospect Theory. Take a roomful of subjects and divide them into two groups. Ask one group to choose between these two alternatives: a sure gain of \$500 and 50 percent chance of gaining \$1,000. Ask the other group to choose between these two alternatives: a sure loss of \$500 and a 50 percent chance of losing \$1,000.

These two trade-offs are very similar, and traditional economics predicts that whether you’re contemplating a gain or a loss doesn’t make a difference: People make trade-offs based on a straightforward calculation of the relative outcome. Some people prefer sure things and others prefer to take chances. Whether the outcome is a gain or a loss doesn’t affect the mathematics and therefore shouldn’t affect the results. This is traditional economics, and it’s called Utility Theory.

But Kahneman’s and Tversky’s experiments contradicted Utility Theory. When faced with a gain, about 85 percent of people chose the sure smaller gain over the risky larger gain. But when faced with a loss, about 70 percent chose the risky larger loss over the sure smaller loss.

This experiment, repeated again and again by many researchers, across ages, genders, cultures and even species, rocked economics, yielded the same result. Directly contradicting the traditional idea of “economic man,” Prospect Theory recognizes that people have subjective values for gains and losses. We have evolved a cognitive bias: a pair of heuristics. One, a sure gain is better than a chance at a greater gain, or “A bird in the hand is worth two in the bush.” And two, a sure loss is worse than a chance at a greater loss, or “Run away and live to fight another day.” Of course, these are not rigid rules. Only a fool would take a sure \$100 over a 50 percent chance at \$1,000,000. But all things being equal, we tend to be risk-averse when it comes to gains and risk-seeking when it comes to losses.

This cognitive bias is so powerful that it can lead to logically inconsistent results. Google the “Asian Disease Experiment” for an almost surreal example. Describing the same policy choice in different ways—either as “200 lives saved out of 600” or “400 lives lost out of 600”—yields wildly different risk reactions.

Evolutionarily, the bias makes sense. It’s a better survival strategy to accept small gains rather than risk them for larger ones, and to risk larger losses rather than accept smaller losses. Lions, for example, chase young or wounded

wildebeests because the investment needed to kill them is lower. Mature and healthy prey would probably be more nutritious, but there's a risk of missing lunch entirely if it gets away. And a small meal will tide the lion over until another day. Getting through today is more important than the possibility of having food tomorrow. Similarly, it is better to risk a larger loss than to accept a smaller loss. Because animals tend to live on the razor's edge between starvation and reproduction, any loss of food—whether small or large—can be equally bad. Because both can result in death, and the best option is to risk everything for the chance at no loss at all.

How to Sell Security

How does Prospect Theory explain the difficulty of selling the prevention of a security breach? It's a choice between a small sure loss—the cost of the security product—and a large risky loss: for example, the results of an attack on one's network. Of course there's a lot more to the sale. The buyer has to be convinced that the product works, and he has to understand the threats against him and the risk that something bad will happen. But all things being equal, buyers would rather take the chance that the attack won't happen than suffer the sure loss that comes from purchasing the security product.

Security sellers know this, even if they don't understand why, and are continually trying to frame their products in positive results. That's why you see slogans with the basic message, "We take care of security so you can focus on your business," or carefully crafted ROI models that demonstrate how profitable a security purchase can be. But these never seem to work. Security is fundamentally a negative sell.

One solution is to stoke fear. Fear is a primal emotion, far older than our ability to calculate trade-offs. And when people are truly scared, they're willing to do almost anything to make that feeling go away; lots of other psychological research supports that. Any burglar alarm salesman will tell you that people buy only after they've been robbed, or after one of their neighbors has been robbed. And the fears stoked by 9/11, and the politics surrounding 9/11, have fueled an entire industry devoted to counterterrorism. When emotion takes over like that, people are much less likely to think rationally.

Though effective, fear mongering is not very ethical. The better solution is not to sell security directly, but to include it as part of a more general product or service. Your car comes with safety and security features built in; they're not sold separately. Same with your house. And it should be the same with computers and networks. Vendors need to build security into the products

and services that customers actually want. CIOs should include security as an integral part of everything they budget for. Security shouldn't be a separate policy for employees to follow but part of overall IT policy.

Security is inherently about avoiding a negative, so you can never ignore the cognitive bias embedded so deeply in the human brain. But if you understand it, you have a better chance of overcoming it.

Why Do We Accept Signatures by Fax? —

Originally published in Wired News, May 29, 2008

Aren't fax signatures the weirdest thing? It's trivial to cut and paste—with real scissors and glue—anyone's signature onto a document so that it'll look real when faxed. There is so little security in fax signatures that it's mind-boggling that anyone accepts them.

Yet people do, all the time. I've signed book contracts, credit card authorizations, nondisclosure agreements and all sorts of financial documents—all by fax. I even have a scanned file of my signature on my computer, so I can virtually cut and paste it into documents and fax them directly from my computer without ever having to print them out. What in the world is going on here?

And, more importantly, why are fax signatures still being used after years of experience? Why aren't there many stories of signatures forged through the use of fax machines?

The answer comes from looking at fax signatures not as an isolated security measure, but in the context of the larger system. Fax signatures work because signed faxes exist within a broader communications context.

In a 2003 paper, *Economics, Psychology, and Sociology of Security*, professor Andrew Odlyzko looks at fax signatures and concludes:

Although fax signatures have become widespread, their usage is restricted. They are not used for final contracts of substantial value, such as home purchases. That means that the insecurity of fax communications is not easy to exploit for large gain. Additional protection against abuse of fax insecurity is provided by the context in which faxes are used. There are records of phone calls that carry the faxes, paper trails inside enterprises and so on. Furthermore, unexpected large financial transfers trigger scrutiny. As a result, successful frauds are not easy to carry out by purely technical means.

He's right. Thinking back, there really aren't ways in which a criminal could use a forged document sent by fax to defraud me. I suppose an unscrupulous consulting client could forge my signature on a non-disclosure agreement and then sue me, but that hardly seems worth the effort. And if my broker received a fax document from me authorizing a money transfer to a Nigerian bank account, he would certainly call me before completing it.

Credit card signatures aren't verified in person, either—and I can already buy things over the phone with a credit card—so there are no new risks there, and Visa knows how to monitor transactions for fraud. Lots of companies accept purchase orders via fax, even for large amounts of stuff, but there's a physical audit trail, and the goods are shipped to a physical address—probably one the seller has shipped to before. Signatures are kind of a business lubricant: mostly, they help move things along smoothly.

Except when they don't.

On October 30, 2004, Tristian Wilson was released from a Memphis jail on the authority of a forged fax message. It wasn't even a particularly good forgery. It wasn't on the standard letterhead of the West Memphis Police Department. The name of the policeman who signed the fax was misspelled. And the time stamp on the top of the fax clearly showed that it was sent from a local McDonald's.

The success of this hack has nothing to do with the fact that it was sent over by fax. It worked because the jail had lousy verification procedures. They didn't notice any discrepancies in the fax. They didn't notice the phone number from which the fax was sent. They didn't call and verify that it was official. The jail was accustomed to getting release orders via fax, and just acted on this one without thinking. Would it have been any different had the forged release form been sent by mail or courier?

Yes, fax signatures always exist in context, but sometimes they are the linchpin within that context. If you can mimic enough of the context, or if those on the receiving end become complacent, you can get away with mischief.

Arguably, this is part of the security process. Signatures themselves are poorly defined. Sometimes a document is valid even if not signed: A person with both hands in a cast can still buy a house. Sometimes a document is invalid even if signed: The signer might be drunk, or have a gun pointed at his head. Or he might be a minor. Sometimes a valid signature isn't enough; in the United States there is an entire infrastructure of "notary publics" who officially witness signed documents. When I started filing my tax returns electronically, I had to sign a document stating that I wouldn't be signing my income tax documents. And banks don't even bother verifying signatures on

checks less than \$30,000; it's cheaper to deal with fraud after the fact than prevent it.

Over the course of centuries, business and legal systems have slowly sorted out what types of additional controls are required around signatures, and in which circumstances.

Those same systems will be able to sort out fax signatures, too, but it'll be slow. And that's where there will be potential problems. Already fax is a declining technology. In a few years it'll be largely obsolete, replaced by PDFs sent over e-mail and other forms of electronic documentation. In the past, we've had time to figure out how to deal with new technologies. Now, by the time we institutionalize these measures, the technologies are likely to be obsolete.

What that means is people are likely to treat fax signatures—or whatever replaces them—exactly the same way as paper signatures. And sometimes that assumption will get them into trouble.

But it won't cause social havoc. Wilson's story is remarkable mostly because it's so exceptional. And even he was rearrested at his home less than a week later. Fax signatures may be new, but fake signatures have always been a possibility. Our legal and business systems need to deal with the underlying problem—false authentication—rather than focus on the technology of the moment. Systems need to defend themselves against the possibility of fake signatures, regardless of how they arrive.

The Pros and Cons of LifeLock

Originally published in Wired News, June 12, 2008

LifeLock, one of the companies that offers identity-theft protection in the United States, has been taking quite a beating recently. They're being sued by credit bureaus, competitors and lawyers in several states that are launching class action lawsuits. And the stories in the media. . . it's like a piranha feeding frenzy.

There are also a lot of errors and misconceptions. With its aggressive advertising campaign and a CEO who publishes his Social Security number and dares people to steal his identity—Todd Davis, 457-55-5462—LifeLock is a company that's easy to hate. But the company's story has some interesting security lessons, and it's worth understanding in some detail.

In December 2003, as part of the Fair and Accurate Credit Transactions Act, or FACTA, credit bureaus were forced to allow you to put a fraud alert on their credit reports, requiring lenders to verify your identity before issuing a credit

card in your name. This alert is temporary, and expires after 90 days. Several companies have sprung up—LifeLock, Debix, LoudSiren, TrustedID—that automatically renew these alerts and effectively make them permanent.

This service pisses off the credit bureaus and their financial customers. The reason lenders don't routinely verify your identity before issuing you credit is that it takes time, costs money and is one more hurdle between you and another credit card. (Buy, buy, buy—it's the American way.) So in the eyes of credit bureaus, LifeLock's customers are inferior goods; selling their data isn't as valuable. LifeLock also opts its customers out of pre-approved credit card offers, further making them less valuable in the eyes of credit bureaus.

And, so began a smear campaign on the part of the credit bureaus. You can read their points of view in this *New York Times* article, written by a reporter who didn't do much more than regurgitate their talking points. And the class action lawsuits have piled on, accusing LifeLock of deceptive business practices, fraudulent advertising and so on. The biggest smear is that LifeLock didn't even protect Todd Davis, and that his identity was allegedly stolen.

It wasn't. Someone in Texas used Davis's SSN to get a \$500 advance against his paycheck. It worked because the loan operation didn't check with any of the credit bureaus before approving the loan—perfectly reasonable for an amount this small. The payday-loan operation called Davis to collect, and LifeLock cleared up the problem. His credit report remains spotless.

The Experian credit bureau's lawsuit basically claims that fraud alerts are only for people who have been victims of identity theft. This seems spurious; the text of the law states that anyone “who asserts a good faith suspicion that the consumer has been or is about to become a victim of fraud or related crime” can request a fraud alert. It seems to me that includes anybody who has ever received one of those notices about their financial details being lost or stolen, which is everybody.

As to deceptive business practices and fraudulent advertising—those just seem like class action lawyers piling on. LifeLock's aggressive fear-based marketing doesn't seem any worse than a lot of other similar advertising campaigns. My guess is that the class action lawsuits won't go anywhere.

In reality, forcing lenders to verify identity before issuing credit is exactly the sort of thing we need to do to fight identity theft. Basically, there are two ways to deal with identity theft: Make personal information harder to steal, and make stolen personal information harder to use. We all know the former doesn't work, so that leaves the latter. If Congress wanted to solve the problem

for real, one of the things it would do is make fraud alerts permanent for everybody. But the credit industry's lobbyists would never allow that.

LifeLock does a bunch of other clever things. They monitor the national address database, and alert you if your address changes. They look for your credit and debit card numbers on hacker and criminal websites and such, and assist you in getting a new number if they see it. They have a million-dollar service guarantee—for complicated legal reasons, they can't call it insurance—to help you recover if your identity is ever stolen.

But even with all of this, I am not a LifeLock customer. At \$120 a year, it's just not worth it. You wouldn't know it from the press attention, but dealing with identity theft has become easier and more routine. Sure, it's a pervasive problem. The Federal Trade Commission reported that 8.3 million Americans were identity-theft victims in 2005. But that includes things like someone stealing your credit card and using it, something that rarely costs you any money and that LifeLock doesn't protect against. New account fraud is much less common, affecting 1.8 million Americans per year, or 0.8 percent of the adult population. The FTC hasn't published detailed numbers for 2006 or 2007, but the rate seems to be declining.

New card fraud is also not very damaging. The median amount of fraud the thief commits is \$1,350, but you're not liable for that. Some spectacularly horrible identity-theft stories notwithstanding, the financial industry is pretty good at quickly cleaning up the mess. The victim's median out-of-pocket cost for new account fraud is only \$40, plus ten hours of grief to clean up the problem. Even assuming your time is worth \$100 an hour, LifeLock isn't worth more than \$8 a year.

And it's hard to get any data on how effective LifeLock really is. They've been in business three years and have about a million customers, but most of them have joined up in the last year. They've paid out on their service guarantee 113 times, but a lot of those were for things that happened before their customers became customers. (It was easier to pay than argue, I assume.) But they don't know how often the fraud alerts actually catch an identity thief in the act. My guess is that it's less than the 0.8 percent fraud rate above.

LifeLock's business model is based more on the fear of identity theft than the actual risk.

It's pretty ironic of the credit bureaus to attack LifeLock on its marketing practices, since they know all about profiting from the fear of identity theft. FACTA also forced the credit bureaus to give Americans a free credit report

once a year upon request. Through deceptive marketing techniques, they've turned this requirement into a multimillion-dollar business.

Get LifeLock if you want, or one of its competitors if you prefer. But remember that you can do most of what these companies do yourself. You can put a fraud alert on your own account, but you have to remember to renew it every three months. You can also put a credit freeze on your account, which is more work for the average consumer but more effective if you're a privacy wonk—and the rules differ by state. And maybe someday Congress will do the right thing and put LifeLock out of business by forcing lenders to verify identity every time they issue credit in someone's name.

The Problem Is Information Insecurity

Originally published in Security Watch, August 10, 2008

Information insecurity is costing us billions. We pay for it in theft: information theft, financial theft. We pay for it in productivity loss, both when networks stop working and in the dozens of minor security inconveniences we all have to endure. We pay for it when we have to buy security products and services to reduce those other two losses. We pay for security, year after year.

The problem is that all the money we spend isn't fixing the problem. We're paying, but we still end up with insecurities.

The problem is insecure software. It's bad design, poorly implemented features, inadequate testing and security vulnerabilities from software bugs. The money we spend on security is to deal with the effects of insecure software.

And that's the problem. We're not paying to improve the security of the underlying software. We're paying to deal with the problem rather than to fix it.

The only way to fix this problem is for vendors to fix their software, and they won't do it until it's in their financial best interests to do so.

Today, the costs of insecure software aren't borne by the vendors that produce the software. In economics, this is known as an externality, the cost of a decision that's borne by people other than those making the decision.

There are no real consequences to the vendors for having bad security or low-quality software. Even worse, the marketplace often rewards low quality.

More precisely, it rewards additional features and timely release dates, even if they come at the expense of quality.

If we expect software vendors to reduce features, lengthen development cycles and invest in secure software development processes, it needs to be in their financial best interests to do so. If we expect corporations to spend significant resources on their own network security—especially the security of their customers—it also needs to be in their financial best interests.

Liability law is a way to make it in those organizations' best interests. Raising the risk of liability raises the costs of doing it wrong and therefore increases the amount of money a CEO is willing to spend to do it right. Security is risk management; liability fiddles with the risk equation.

Basically, we have to tweak the risk equation so the CEO cares about actually fixing the problem, and putting pressure on his balance sheet is the best way to do that.

Clearly, this isn't all or nothing. There are many parties involved in a typical software attack. There's the company that sold the software with the vulnerability in the first place. There's the person who wrote the attack tool. There's the attacker himself, who used the tool to break into a network.

There's the owner of the network, who was entrusted with defending that network. One hundred percent of the liability shouldn't fall on the shoulders of the software vendor, just as 100% shouldn't fall on the attacker or the network owner. But today, 100% of the cost falls directly on the network owner, and that just has to stop.

We will always pay for security. If software vendors have liability costs, they'll pass those on to us. It might not be cheaper than what we're paying today. But as long as we're going to pay, we might as well pay to fix the problem. Forcing the software vendor to pay to fix the problem and then pass those costs on to us means that the problem might actually get fixed.

Liability changes everything. Currently, there is no reason for a software company not to offer feature after feature after feature. Liability forces software companies to think twice before changing something. Liability forces companies to protect the data they're entrusted with. Liability means that those in the best position to fix the problem are actually responsible for the problem.

Information security isn't a technological problem. It's an economics problem. And the way to improve information technology is to fix the economics problem. Do that, and everything else will follow.

Security ROI: Fact or Fiction?

Originally published in CSO Magazine, September 2, 2008

Return on investment, or ROI, is a big deal in business. Any business venture needs to demonstrate a positive return on investment, and a good one at that, in order to be viable.

It's become a big deal in IT security, too. Many corporate customers are demanding ROI models to demonstrate that a particular security investment pays off. And in response, vendors are providing ROI models that demonstrate how their particular security solution provides the best return on investment.

It's a good idea in theory, but it's mostly bunk in practice.

Before I get into the details, there's one point I have to make. "ROI" as used in a security context is inaccurate. Security is not an investment that provides a return, like a new factory or a financial instrument. It's an expense that, hopefully, pays for itself in cost savings. Security is about loss prevention, not about earnings. The term just doesn't make sense in this context.

But as anyone who has lived through a company's vicious end-of-year budget-slashing exercises knows, when you're trying to make your numbers, cutting costs is the same as increasing revenues. So while security can't produce ROI, loss prevention most certainly affects a company's bottom line.

And a company should implement only security countermeasures that affect its bottom line positively. It shouldn't spend more on a security problem than the problem is worth. Conversely, it shouldn't ignore problems that are costing it money when there are cheaper mitigation alternatives. A smart company needs to approach security as it would any other business decision: costs versus benefits.

The classic methodology is called annualized loss expectancy (ALE), and it's straightforward. Calculate the cost of a security incident in both tangibles like time and money, and intangibles like reputation and competitive advantage. Multiply that by the chance the incident will occur in a year. That tells you how much you should spend to mitigate the risk. So, for example, if your store has a 10 percent chance of getting robbed and the cost of being robbed is \$10,000, then you should spend \$1,000 a year on security. Spend more than that, and you're wasting money. Spend less than that, and you're also wasting money.

Of course, that \$1,000 has to reduce the chance of being robbed to zero in order to be cost-effective. If a security measure cuts the chance of robbery by 40 percent—to 6 percent a year—then you should spend no more than

\$400 on it. If another security measure reduces it by 80 percent, it's worth \$800. And if two security measures both reduce the chance of being robbed by 50 percent and one costs \$300 and the other \$700, the first one is worth it and the second isn't.

The Data Imperative

The key to making this work is good data; the term of art is “actuarial tail.” If you're doing an ALE analysis of a security camera at a convenience store, you need to know the crime rate in the store's neighborhood and maybe have some idea of how much cameras improve the odds of convincing criminals to rob another store instead. You need to know how much a robbery costs: in merchandise, in time and annoyance, in lost sales due to spooked patrons, in employee morale. You need to know how much not having the cameras costs in terms of employee morale; maybe you're having trouble hiring salespeople to work the night shift. With all that data, you can figure out if the cost of the camera is cheaper than the loss of revenue if you close the store at night—assuming that the closed store won't get robbed as well. And then you can decide whether to install one.

Cybersecurity is considerably harder, because there just isn't enough good data. There aren't good crime rates for cyberspace, and we have a lot less data about how individual security countermeasures—or specific configurations of countermeasures—mitigate those risks. We don't even have data on incident costs.

One problem is that the threat moves too quickly. The characteristics of the things we're trying to prevent change so quickly that we can't accumulate data fast enough. By the time we get some data, there's a new threat model for which we don't have enough data. So we can't create ALE models.

But there's another problem, and it's that the math quickly falls apart when it comes to rare and expensive events. Imagine you calculate the cost—reputational costs, loss of customers, etc.—of having your company's name in the newspaper after an embarrassing cybersecurity event to be \$20 million. Also assume that the odds are 1 in 10,000 of that happening in any one year. ALE says you should spend no more than \$2,000 mitigating that risk.

So far, so good. But maybe your CFO thinks an incident would cost only \$10 million. You can't argue, since we're just estimating. But he just cut your security budget in half. A vendor trying to sell you a product finds a Web analysis claiming that the odds of this happening are actually 1 in 1,000.

Accept this new number, and suddenly a product costing 10 times as much is still a good investment.

It gets worse when you deal with even more rare and expensive events. Imagine you're in charge of terrorism mitigation at a chlorine plant. What's the cost to your company, in money and reputation, of a large and very deadly explosion? \$100 million? \$1 billion? \$10 billion? And the odds: 1 in a hundred thousand, 1 in a million, 1 in 10 million? Depending on how you answer those two questions—and any answer is really just a guess—you can justify spending anywhere from \$10 to \$100,000 annually to mitigate that risk.

Or take another example: airport security. Assume that all the new airport security measures increase the waiting time at airports by—and I'm making this up—30 minutes per passenger. There were 760 million passenger boardings in the United States in 2007. This means that the extra waiting time at airports has cost us a collective 43,000 years of extra waiting time. Assume a 70-year life expectancy, and the increased waiting time has “killed” 620 people per year—930 if you calculate the numbers based on 16 hours of awake time per day. So the question is: If we did away with increased airport security, would the result be more people dead from terrorism or fewer?

Caveat Emptor

This kind of thing is why most ROI models you get from security vendors are nonsense. Of course their model demonstrates that their product or service makes financial sense: They've jiggered the numbers so that they do.

This doesn't mean that ALE is useless, but it does mean you should 1) mistrust any analyses that come from people with an agenda and 2) use any results as a general guideline only. So when you get an ROI model from your vendor, take its framework and plug in your own numbers. Don't even show the vendor your improvements; it won't consider any changes that make its product or service less cost-effective to be an “improvement.” And use those results as a general guide, along with risk management and compliance analyses, when you're deciding what security products and services to buy.

Social Networking Risks

Originally published in Information Security, February 2009

This essay appeared as the first half of a point-counterpoint with Marcus Ranum.

Are employees blogging corporate secrets? It's not an unreasonable fear, actually. People have always talked about work to their friends. It's human nature for people to talk about what's going on in their lives, and work is a lot of most people's lives. Historically, organizations generally didn't care very much. The conversations were intimate and ephemeral, so the risk was small. Unless you worked for the military with actual national secrets, no one worried about it very much.

What has changed is the nature of how we interact with our friends. We talk about our lives on our blogs, on social networking sites such as Facebook and Twitter, and on message boards pertaining to the work we're doing. What was once intimate and ephemeral is now available to the whole world, indexed by Google, and archived for posterity. A good open-source intelligence gatherer can learn a lot about what a company is doing by monitoring its employees' online activities. It's no wonder some organizations are nervous.

So yes, organizations should be concerned about employees leaking corporate secrets on social networking sites. And, as much as I hate to admit it, disciplinary action against employees who reveal too much in public is probably in order. But actually policing employees is almost certainly more expensive and more trouble than it's worth. And when an organization catches an employee being a bit too chatty about work details, it should be as forgiving as possible.

That's because this sort of openness is the future of work, and the organizations that get used to it or—even better—embrace it, are going to do better in the long run than organizations that futilely try to fight it.

The Internet is the greatest generation gap since rock and roll, and what we're seeing here is one particular skirmish across that gap. The younger generation, used to spending a lot of its life in public, clashes with an older generation in charge of a corporate culture that presumes a greater degree of discretion and greater level of control.

There are two things that are always true about generation gaps. The first is that the elder generation is always right about the problems that will result from whatever new/different/bad thing the younger generation is doing. And the second is that the younger generation is always right that whatever they're doing will become the new normal. These things have to be true; the older generation understands the problems better, but they're the ones who fade away and die.

Living an increasingly public life on social networking sites is the new normal. More corporate—and government—transparency is becoming the new normal. CEOs who blog aren't yet the new normal, but will be eventually. And then what will corporate secrecy look like? Organizations will still have secrets, of course, but they will be more public and more open about what they're doing and what they're thinking of doing. It'll be different than it is now, but it most likely won't be any worse.

Today isn't that day yet, which is why it's still proper for organizations to worry about loose fingers uploading corporate secrets. But the sooner an organization can adapt to this new normal and figure out how to be successful within it, the better it will survive these transitions. In the near term, it will be more likely to attract the next-generation talent it needs to figure out how to thrive. In the long term. . .well, we don't know what it will mean yet.

Same with blocking those sites; yes, they're enormous time-wasters. But if an organization has a problem with employee productivity, they're not going to solve it by censoring Internet access. Focus on the actual problem, and don't waste time on the particulars of how the problem manifests itself.

Do You Know Where Your Data Are?

Originally published in the Wall Street Journal, April 28, 2009

Do you know what your data did last night? Almost none of the more than 27 million people who took the RealAge quiz realized that their personal health data was being used by drug companies to develop targeted e-mail marketing campaigns.

There's a basic consumer protection principle at work here, and it's the concept of "unfair and deceptive" trade practices. Basically, a company shouldn't be able to say one thing and do another: sell used goods as new, lie on ingredients lists, advertise prices that aren't generally available, claim features that don't exist, and so on.

Buried in RealAge's 2,400-word privacy policy is this disclosure: "If you elect to say yes to becoming a free RealAge Member, we will periodically send you free newsletters and e-mails that directly promote the use of our site(s) or the purchase of our products or services and may contain, in whole or in part, advertisements for third parties which relate to marketed products of selected RealAge partners."

They maintain that when you join the website, you consent to receiving pharmaceutical company spam. But since that isn't spelled out, it's not really informed consent. That's deceptive.

Cloud computing is another technology where users entrust their data to service providers. Salesforce.com, Gmail, and Google Docs are examples; your data isn't on your computer—it's out in the "cloud" somewhere—and you access it from your web browser. Cloud computing has significant benefits for customers and huge profit potential for providers. It's one of the fastest growing IT market segments—69% of Americans now use some sort of cloud computing services—but the business is rife with shady, if not outright deceptive, advertising.

Take Google, for example. Last month, the Electronic Privacy Information Center (I'm on its board of directors) filed a complaint with the Federal Trade Commission concerning Google's cloud computing services. On its website, Google repeatedly assures customers that their data is secure and private, while published vulnerabilities demonstrate that it is not. Google's not foolish, though; its Terms of Service explicitly disavow any warranty or any liability for harm that might result from Google's negligence, recklessness, malevolent intent, or even purposeful disregard of existing legal obligations to protect the privacy and security of user data. EPIC claims that's deceptive.

Facebook isn't much better. Its plainly written (and not legally binding) Statement of Principles contains an admirable set of goals, but its denser and more legalistic Statement of Rights and Responsibilities undermines a lot of it. One research group who studies these documents called it "democracy theater": Facebook wants the appearance of involving users in governance, without the messiness of actually having to do so. Deceptive.

These issues are not identical. RealAge is hiding what it does with your data. Google is trying to both assure you that your data is safe and duck any responsibility when it's not. Facebook wants to market a democracy but run a dictatorship. But they all involve trying to deceive the customer.

Cloud computing services like Google Docs, and social networking sites like RealAge and Facebook, bring with them significant privacy and security risks over and above traditional computing models. Unlike data on my own computer, which I can protect to whatever level I believe prudent, I have no control over any of these sites, nor any real knowledge of how these companies protect my privacy and security. I have to trust them.

This may be fine—the advantages might very well outweigh the risks—but users often can't weigh the trade-offs because these companies are going out of their way to hide the risks.

Of course, companies don't want people to make informed decisions about where to leave their personal data. RealAge wouldn't get 27 million members if its webpage clearly stated "you are signing up to receive e-mails containing advertising from pharmaceutical companies," and Google Docs wouldn't get five million users if its webpage said "We'll take some steps to protect your privacy, but you can't blame us if something goes wrong."

And of course, trust isn't black and white. If, for example, Amazon tried to use customer credit card info to buy itself office supplies, we'd all agree that that was wrong. If it used customer names to solicit new business from their friends, most of us would consider this wrong. When it uses buying history to try to sell customers new books, many of us appreciate the targeted marketing. Similarly, no one expects Google's security to be perfect. But if it didn't fix known vulnerabilities, most of us would consider that a problem.

This is why understanding is so important. For markets to work, consumers need to be able to make informed buying decisions. They need to understand both the costs and benefits of the products and services they buy. Allowing sellers to manipulate the market by outright lying, or even by hiding vital information, about their products breaks capitalism—and that's why the government has to step in to ensure markets work smoothly.

Last month, Mary K. Engle, Acting Deputy Director of the FTC's Bureau of Consumer Protection said: "a company's marketing materials must be consistent with the nature of the product being offered. It's not enough to disclose the information only in a fine print of a lengthy online user agreement." She was speaking about Digital Rights Management and, specifically, an incident where Sony used a music copy protection scheme without disclosing that it secretly installed software on customers' computers. DRM is different from cloud computing or even online surveys and quizzes, but the principle is the same.

Engle again: "if your advertising giveth and your EULA [license agreement] taketh away don't be surprised if the FTC comes calling." That's the right response from government.

Be Careful When You Come to Put Your Trust in the Clouds

Originally published in the Guardian, June 4, 2009

This year's overhyped IT concept is cloud computing. Also called software as a service (SaaS), cloud computing is when you run software over the Internet and access it via a browser. The salesforce.com customer management software is an example of this. So is Google Docs. If you believe the hype, cloud computing is the future.

But, hype aside, cloud computing is nothing new. It's the modern version of the timesharing model from the 1960s, which was eventually killed by the rise of the personal computer. It's what Hotmail and Gmail have been doing all these years, and it's social networking sites, remote backup companies, and remote email filtering companies such as MessageLabs. Any IT outsourcing—network infrastructure, security monitoring, remote hosting—is a form of cloud computing.

The old timesharing model arose because computers were expensive and hard to maintain. Modern computers and networks are drastically cheaper, but they're still hard to maintain. As networks have become faster, it is again easier to have someone else do the hard work. Computing has become more of a utility; users are more concerned with results than technical details, so the tech fades into the background.

But what about security? Isn't it more dangerous to have your email on Hotmail's servers, your spreadsheets on Google's, your personal conversations on Facebook's, and your company's sales prospects on salesforce.com's? Well, yes and no.

IT security is about trust. You have to trust your CPU manufacturer, your hardware, operating system and software vendors—and your ISP. Any one of these can undermine your security: crash your systems, corrupt data, allow an attacker to get access to systems. We've spent decades dealing with worms and rootkits that target software vulnerabilities. We've worried about infected chips. But in the end, we have no choice but to blindly trust the security of the IT providers we use.

SaaS moves the trust boundary out one step further—you now have to also trust your software service vendors—but it doesn't fundamentally change anything. It's just another vendor we need to trust.

There is one critical difference. When a computer is within your network, you can protect it with other security systems such as firewalls and IDSs. You can build a resilient system that works even if those vendors you have to trust may not be as trustworthy as you like. With any outsourcing model, whether it be cloud computing or something else, you can't. You have to trust your outsourcer completely. You not only have to trust the outsourcer's security, but its reliability, its availability, and its business continuity.

You don't want your critical data to be on some cloud computer that abruptly disappears because its owner goes bankrupt. You don't want the company you're using to be sold to your direct competitor. You don't want the company to cut corners, without warning, because times are tight. Or raise its prices and then refuse to let you have your data back. These things can happen with software vendors, but the results aren't as drastic.

There are two different types of cloud computing customers. The first only pays a nominal fee for these services—and uses them for free in exchange for ads: e.g., Gmail and Facebook. These customers have no leverage with their outsourcers. You can lose everything. Companies like Google and Amazon won't spend a lot of time caring. The second type of customer pays considerably for these services: to salesforce.com, MessageLabs, managed network companies, and so on. These customers have more leverage, providing they write their service contracts correctly. Still, nothing is guaranteed.

Trust is a concept as old as humanity, and the solutions are the same as they have always been. Be careful who you trust, be careful what you trust them with, and be careful how much you trust them. Outsourcing is the future of computing. Eventually we'll get this right, but you don't want to be a casualty along the way.

Is Perfect Access Control Possible?

Originally published in Information Security, September 2009

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

Access control is difficult in an organizational setting. On one hand, every employee needs enough access to do his job. On the other hand, every time you give an employee more access, there's more risk: he could abuse that access, or lose information he has access to, or be socially engineered into giving that access to a malfasant. So a smart, risk-conscious organization will give each employee the exact level of access he needs to do his job, and no more.

Over the years, there's been a lot of work put into role-based access control. But despite the large number of academic papers and high-profile security products, most organizations don't implement it—at all—with the predictable security problems as a result.

Regularly we read stories of employees abusing their database access-control privileges for personal reasons: medical records, tax records, passport records, police records. NSA eavesdroppers spy on their wives and girlfriends. Departing employees take corporate secrets.

A spectacular access control failure occurred in the UK in 2007. An employee of Her Majesty's Revenue & Customs had to send a couple of thousand sample records from a database on all children in the country to National Audit Office. But it was easier for him to copy the entire database of 25 million people onto a couple of disks and put it in the mail than it was to select out just the records needed. Unfortunately, the discs got lost in the mail, and the story was a huge embarrassment for the government.

Eric Johnson at Dartmouth's Tuck School of Business has been studying the problem, and his results won't startle anyone who has thought about it at all. RBAC is very hard to implement correctly. Organizations generally don't even know who has what role. The employee doesn't know, the boss doesn't know—and these days the employee might have more than one boss—and senior management certainly doesn't know. There's a reason RBAC came out of the military; in that world, command structures are simple and well-defined.

Even worse, employees' roles change all the time—Johnson chronicled one business group of 3,000 people that made 1,000 role changes in just three months—and it's often not obvious what information an employee needs until he actually needs it. And information simply isn't that granular. Just as it's much easier to give someone access to an entire file cabinet than to only the particular files he needs, it's much easier to give someone access to an entire database than only the particular records he needs.

This means that organizations either over-entitle or under-entitle employees. But since getting the job done is more important than anything else, organizations tend to over-entitle. Johnson estimates that 50 percent to 90 percent of employees are over-entitled in large organizations. In the uncommon instance where an employee needs access to something he normally doesn't have, there's generally some process for him to get it. And access is almost never revoked once it's been granted. In large formal organizations, Johnson was able to predict how long an employee had worked there based on how much access he had.

Clearly, organizations can do better. Johnson's current work involves building access-control systems with easy self-escalation, audit to make sure that power isn't abused, violation penalties (Intel, for example, issues "speeding tickets" to violators), and compliance rewards. His goal is to implement incentives and controls that manage access without making people too risk-averse.

In the end, a perfect access control system just isn't possible; organizations are simply too chaotic for it to work. And any good system will allow a certain number of access control violations, if they're made in good faith by people just trying to do their jobs. The "speeding ticket" analogy is better than it looks: we post limits of 55 miles per hour, but generally don't start ticketing people unless they're going over 70.

News Media Strategies for Survival for Journalists

Originally published in Twin Cities Daily Planet, November 14, 2009

Those of us living through the Internet-caused revolution in journalism can't see what's going to come out the other side: how readers will interact with journalism, what the sources of journalism will be, how journalists will make money. All we do know is that mass-market journalism is hurting, badly, and may not survive. And that we have no idea how to thrive in this new world of digital media.

I have five pieces of advice to those trying to survive and wanting to thrive: based both on experiences as a successful Internet pundit and blogger, and my observations of others, successful and unsuccessful. I'll talk about writing, but everything I say applies to audio and video as well.

One, be interesting. Yes, that's obvious. But the scale is different now. It used to be you could be interesting in aggregate; a few interesting articles or features could carry an entire publication. Now every single piece of writing has to be interesting; otherwise, it won't get read, passed around, or linked to. Have something to say. Pick a niche you can become known for.

Two, be entertaining. Interesting isn't enough; you have to entertain people as well. Internet readers live in a world where millions of things are constantly vying for their attention. Only the best individual pieces of content thrive in

this environment. Often, “best” means “most entertaining.” Opinions are dime a dozen on the Internet; you need to make sure yours are worth your readers’ time.

Three, be engaging. Readers want to be engaged. They want to be part of a community. They want to engage, with each other as well as with you, on their own terms. Engagement might involve comment or discussion areas, or ways people can follow your work. Anything that limits engagement inhibits community. What this means depends on context; sometimes you have to allow community to develop naturally, even if it’s in ways you don’t like. Sometimes you need to censor off-topic comments to prevent hateful or annoying commenters from driving others away. In general, though, you should allow anonymous comments. You should make your interface as easy as possible to use. You should reply to your readers. And you shouldn’t treat your readers solely as marketing opportunities. The more your writing fosters engagement, the more popular it will be.

Four, be available. Readers need to be able to interact with your writing on their own terms. This means you can’t make it difficult for them to find and link to your content. Make sure your content is accessible by any and every Internet device out there. Never take your old writing off the Internet. Never change your URLs. Never make it hard for them to find or link to a URL. Never put your writing behind a paywall. You’re part of an ecosystem now; fail to play by the rules and you quickly become isolated.

Five, be agile. The Internet changes all the time; what’s true today might not be true in two years. Don’t lock yourself in to a particular look, or a particular web technology. Simple interfaces are better than flashy complicated ones; I don’t care what your ad agency tells you. Agility applies to making money, too. We have no idea what financial models will thrive in the future, but it seems likely that it will be a portfolio of different things. You’ll be more likely to write for different publications. You’ll be more likely to figure out cross subsidies, so that some things pay for the others. I have a free blog and a free monthly newsletter, and charge for books, speaking engagements, and consulting. Your mix will be different. If you’re lucky, everything you do will augment everything else.

Revolutions are scary times. The old crumbles around us, and we have no idea what—if anything—will be built on its ruins. Remember, though, that human nature doesn’t change. People will always gravitate to the interesting, entertaining, engaging, and available, and the agile will be the first on the scene.

Security and Function Creep

*Originally published in IEEE Security & Privacy, January/
February 2010*

Security is rarely static. Technology changes the capabilities of both security systems and attackers. But there's something else that changes security's cost/benefit trade-off: how the underlying systems being secured are used. Far too often we build security for one purpose, only to find it being used for another purpose—one it wasn't suited for in the first place. And then the security system has to play catch-up.

Take driver's licenses, for example. Originally designed to demonstrate a credential—the ability to drive a car—they looked like other credentials: medical licenses or elevator certificates of inspection. They were wallet-sized, of course, but they didn't have much security associated with them. Then, slowly, driver's licenses took on a second application: they became age-verification tokens in bars and liquor stores. Of course the security wasn't up to the task—teenagers can be extraordinarily resourceful if they set their minds to it—and over the decades driver's licenses got photographs, tamper-resistant features (once, it was easy to modify the birth year), and technologies that made counterfeiting harder. There was little value in counterfeiting a driver's license, but a lot of value in counterfeiting an age-verification token.

Today, US driver's licenses are taking on yet another function: security against terrorists. The Real ID Act—the government's attempt to make driver's licenses even more secure—has nothing to do with driving or even with buying alcohol, and everything to do with trying to make that piece of plastic an effective way to verify that someone is not on the terrorist watch list. Whether this is a good idea, or actually improves security, is another matter entirely.

You can see this kind of function creep everywhere. Internet security systems designed for informational Web sites are suddenly expected to provide security for banking Web sites. Security systems that are good enough to protect cheap commodities from being stolen are suddenly ineffective once the price of those commodities rises high enough. Application security systems, designed for locally owned networks, are expected to work even when the application is moved to a cloud computing environment. And cloud computing security, designed for the needs of corporations, is expected to be suitable for government applications as well—maybe even military applications.

Sometimes it's obvious that security systems designed for one environment won't work in another. We don't arm our soldiers the same way we arm our policemen, and we can't take commercial vehicles and easily turn them into ones outfitted for the military. We understand that we might need to upgrade our home security system if we suddenly come into possession of a bag of diamonds. Yet many think the same security that protects our home computers will also protect voting machines, and the same operating systems that run our businesses are suitable for military uses.

But these are all conscious decisions, and we security professionals often know better. The real problems arise when the changes happen in the background, without any conscious thought. We build a network security system that's perfectly adequate for the threat and—like a driver's license becoming an age-verification token—the network accrues more and more functions. But because it has already been pronounced “secure,” we can't get any budget to re-evaluate and improve the security until after the bad guys have figured out the vulnerabilities and exploited them.

I don't like having to play catch-up in security, but we seem doomed to keep doing so.

Weighing the Risk of Hiring Hackers

Originally published in Information Security, June 2010

This essay previously appeared as the first half of a point-counterpoint with Marcus Ranum.

Any essay on hiring hackers quickly gets bogged down in definitions. What is a hacker, and how is he different from a cracker? I have my own definitions, but I'd rather define the issue more specifically: Would you hire someone convicted of a computer crime to fill a position of trust in your computer network? Or, more generally, would you hire someone convicted of a crime for a job related to that crime?

The answer, of course, is “it depends.” It depends on the specifics of the crime. It depends on the ethics involved. It depends on the recidivism rate of the type of criminal. It depends a whole lot on the individual.

Would you hire a convicted pedophile to work at a day care center? Would you hire Bernie Madoff to manage your investment fund? The answer is almost certainly no to those two—but you might hire a convicted bank robber to

consult on bank security. You might hire someone who was convicted of false advertising to write ad copy for your next marketing campaign. And you might hire someone who ran a chop shop to fix your car. It depends on the person and the crime.

It can get even murkier. Would you hire a CIA-trained assassin to be a bodyguard? Would you put a general who led a successful attack in charge of defense? What if they were both convicted of crimes in whatever country they were operating in? There are different legal and ethical issues, to be sure, but in both cases the people learned a certain set of skills regarding offense that could be transferable to defense.

Which brings us back to computers. Hacking is primarily a mindset: a way of thinking about security. Its primary focus is in attacking systems, but it's invaluable to the defense of those systems as well. Because computer systems are so complex, defending them often requires people who can think like attackers.

Admittedly, there's a difference between thinking like an attacker and acting like a criminal, and between researching vulnerabilities in fielded systems and exploiting those vulnerabilities for personal gain. But there is a huge variability in computer crime convictions, and—at least in the early days—many hacking convictions were unjust and unfair. And there's also a difference between someone's behavior as a teenager and his behavior later in life. Additionally, there might very well be a difference between someone's behavior before and after a hacking conviction. It all depends on the person.

An employer's goal should be to hire moral and ethical people with the skill set required to do the job. And while a hacking conviction is certainly a mark against a person, it isn't always grounds for complete non-consideration.

"We don't hire hackers" and "we don't hire felons" are coarse generalizations, in the same way that "we only hire people with this or that security certification" is. They work—you're less likely to hire the wrong person if you follow them—but they're both coarse and flawed. Just as all potential employees with certifications aren't automatically good hires, all potential employees with hacking convictions aren't automatically bad hires. Sure, it's easier to hire people based on things you can learn from checkboxes, but you won't get the best employees that way. It's far better to look at the individual, and put those check boxes into context. But we don't always have time to do that.

Last winter, a Minneapolis attorney who works to get felons a fair shake after they served their time told of a sign he saw: "Snow shovelers wanted. Felons need not apply." It's not good for society if felons who have served their time can't even get jobs shoveling snow.

Should Enterprises Give In to IT Consumerization at the Expense of Security?

Originally published in Information Security, September 2010

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

If you're a typical wired American, you've got a bunch of tech tools you like and a bunch more you covet. You have a cell phone that can easily text. You've got a laptop configured just the way you want it. Maybe you have a Kindle for reading, or an iPad. And when the next new thing comes along, some of you will line up on the first day it's available.

So why can't work keep up? Why are you forced to use an unfamiliar, and sometimes outdated, operating system? Why do you need a second laptop, maybe an older and clunkier one? Why do you need a second cell phone with a new interface, or a BlackBerry, when your phone already does e-mail? Or a second BlackBerry tied to corporate e-mail? Why can't you use the cool stuff you already have?

More and more companies are letting you. They're giving you an allowance and allowing you to buy whatever laptop you want, and to connect into the corporate network with whatever device you choose. They're allowing you to use whatever cell phone you have, whatever portable e-mail device you have, whatever you personally need to get your job done. And the security office is freaking.

You can't blame them, really. Security is hard enough when you have control of the hardware, operating system and software. Lose control of any of those things, and the difficulty goes through the roof. How do you ensure that the employee devices are secure, and have up-to-date security patches? How do you control what goes on them? How do you deal with the tech support issues when they fail? How do you even begin to manage this logistical nightmare? Better to dig your heels in and say "no."

But security is on the losing end of this argument, and the sooner it realizes that, the better.

The meta-trend here is consumerization: cool technologies show up for the consumer market before they're available to the business market. Every corporation is under pressure from its employees to allow them to use these

new technologies at work, and that pressure is only getting stronger. Younger employees simply aren't going to stand for using last year's stuff, and they're not going to carry around a second laptop. They're either going to figure out ways around the corporate security rules, or they're going to take another job with a more trendy company. Either way, senior management is going to tell security to get out of the way. It might even be the CEO, who wants to get to the company's databases from his brand new iPad, driving the change. Either way, it's going to be harder and harder to say no.

At the same time, cloud computing makes this easier. More and more, employee computing devices are nothing more than dumb terminals with a browser interface. When corporate e-mail is all webmail, corporate documents are all on GoogleDocs, and when all the specialized applications have a web interface, it's easier to allow employees to use any up-to-date browser. It's what companies are already doing with their partners, suppliers, and customers.

Also on the plus side, technology companies have woken up to this trend and—from Microsoft and Cisco on down to the startups—are trying to offer security solutions. Like everything else, it's a mixed bag: some of them will work and some of them won't, most of them will need careful configuration to work well, and few of them will get it right. The result is that we'll muddle through, as usual.

Security is always a tradeoff, and security decisions are often made for non-security reasons. In this case, the right decision is to sacrifice security for convenience and flexibility. Corporations want their employees to be able to work from anywhere, and they're going to have loosened control over the tools they allow in order to get it.

The Vulnerabilities Market and the Future of Security

Originally published in Forbes, May 30, 2012

Recently, there have been several articles about the new market in zero-day exploits: new and unpatched computer vulnerabilities. It's not just software companies, who sometimes pay bounties to researchers who alert them of security vulnerabilities so they can fix them. And it's not only criminal organizations that pay for vulnerabilities they can exploit. Now there are governments,

and companies who sell to governments, who buy vulnerabilities with the intent of keeping them secret so they can exploit them.

This market is larger than most people realize, and it's becoming even larger. Forbes recently published a price list for zero-day exploits, along with the story of a hacker who received \$250K from "a US government contractor." (At first I didn't believe the story or the price list, but I have been convinced that they both are true.) Forbes published a profile of a company called Vupen, whose business is selling zero-day exploits. Other companies doing this range from startups like Netragard and Endgame to large defense contractors like Northrop Grumman, General Dynamics, and Raytheon.

This is very different than in 2007, when researcher Charlie Miller wrote about his attempts to sell zero-day exploits; and a 2010 survey implied that there wasn't much money in selling zero days. The market has matured substantially in the past few years.

This new market perturbs the economics of finding security vulnerabilities. And it does so to the detriment of us all.

I've long argued that the process of finding vulnerabilities in software systems increases overall security. This is because the economics of vulnerability hunting favored disclosure. As long as the principal gain from finding a vulnerability was notoriety, publicly disclosing vulnerabilities was the only obvious path. In fact, it took years for our industry to move from a norm of full-disclosure—announcing the vulnerability publicly and damn the consequences—to something called "responsible disclosure": giving the software vendor a head start in fixing the vulnerability. Changing economics is what made the change stick: instead of just hacker notoriety, a successful vulnerability finder could land some lucrative consulting gigs, and being a responsible security researcher helped. But regardless of the motivations, a disclosed vulnerability is one that—at least in most cases—is patched. And a patched vulnerability makes us all more secure.

This is why the new market for vulnerabilities is so dangerous; it results in vulnerabilities remaining secret and unpatched. That it's even more lucrative than the public vulnerabilities market means that more hackers will choose this path. And unlike the previous reward of notoriety and consulting gigs, it gives software programmers within a company the incentive to deliberately create vulnerabilities in the products they're working on—and then secretly sell them to some government agency.

No commercial vendors perform the level of code review that would be necessary to detect, and prove mal-intent for, this kind of sabotage.

Even more importantly, the new market for security vulnerabilities results in a variety of government agencies around the world that have a strong interest in those vulnerabilities remaining unpatched. These range from law-enforcement agencies like the FBI and the German police who are trying to build targeted Internet surveillance tools, to intelligence agencies like the NSA who are trying to build mass Internet surveillance tools, to military organizations who are trying to build cyber-weapons.

All of these agencies have long had to wrestle with the choice of whether to use newly discovered vulnerabilities to protect or to attack. Inside the NSA, this was traditionally known as the “equities issue,” and the debate was between the COMSEC (communications security) side of the NSA and the SIGINT (signals intelligence) side. If they found a flaw in a popular cryptographic algorithm, they could either use that knowledge to fix the algorithm and make everyone’s communications more secure, or they could exploit the flaw to eavesdrop on others—while at the same time allowing even the people they wanted to protect to remain vulnerable. This debate raged through the decades inside the NSA. From what I’ve heard, by 2000, the COMSEC side had largely won, but things flipped completely around after 9/11.

The whole point of disclosing security vulnerabilities is to put pressure on vendors to release more secure software. It’s not just that they patch the vulnerabilities that are made public—the fear of bad press makes them implement more secure software development processes. It’s another economic process; the cost of designing software securely in the first place is less than the cost of the bad press after a vulnerability is announced plus the cost of writing and deploying the patch. I’d be the first to admit that this isn’t perfect—there’s a lot of very poorly written software still out there—but it’s the best incentive we have.

We’ve always expected the NSA, and those like them, to keep the vulnerabilities they discover secret. We have been counting on the public community to find and publicize vulnerabilities, forcing vendors to fix them. With the rise of these new pressures to keep zero-day exploits secret, and to sell them for exploitation, there will be even less incentive on software vendors to ensure the security of their products.

As the incentive for hackers to keep their vulnerabilities secret grows, the incentive for vendors to build secure software shrinks. As a recent EFF essay put it, this is “security for the 1%.” And it makes the rest of us less safe.

So You Want to Be a Security Expert

Originally published in Krebs on Security, July 12, 2012

This essay originally appeared as part of a series of advice columns on how to break into the field of security.

I regularly receive e-mail from people who want advice on how to learn more about computer security, either as a course of study in college or as an IT person considering it as a career choice.

First, know that there are many subspecialties in computer security. You can be an expert in keeping systems from being hacked, or in creating unhackable software. You can be an expert in finding security problems in software, or in networks. You can be an expert in viruses, or policies, or cryptography. There are many, many opportunities for many different skill sets. You don't have to be a coder to be a security expert.

In general, though, I have three pieces of advice to anyone who wants to learn computer security.

Study. Studying can take many forms. It can be classwork, either at universities or at training conferences like SANS and Offensive Security. (These are good self-starter resources.) It can be reading; there are a lot of excellent books out there—and blogs—that teach different aspects of computer security out there. Don't limit yourself to computer science, either. You can learn a lot by studying other areas of security, and soft sciences like economics, psychology, and sociology.

Do. Computer security is fundamentally a practitioner's art, and that requires practice. This means using what you've learned to configure security systems, design new security systems, and—yes—break existing security systems. This is why many courses have strong hands-on components; you won't learn much without it.

Show. It doesn't matter what you know or what you can do if you can't demonstrate it to someone who might want to hire you. This doesn't just mean sounding good in an interview. It means sounding good on mailing lists and in blog comments. You can show your expertise by making podcasts and writing your own blog. You can teach seminars at your local user group meetings. You can write papers for conferences, or books.

I am a fan of security certifications, which can often demonstrate all of these things to a potential employer quickly and easily.

I've really said nothing here that isn't also true for a gazillion other areas of study, but security also requires a particular mindset—one I consider essential for success in this field. I'm not sure it can be taught, but it certainly can be encouraged. "This kind of thinking is not natural for most people. It's not natural for engineers. Good engineering involves thinking about how things can be made to work; the security mindset involves thinking about how things can be made to fail. It involves thinking like an attacker, an adversary or a criminal. You don't have to exploit the vulnerabilities you find, but if you don't see the world that way, you'll never notice most security problems." This is especially true if you want to design security systems and not just implement them. Remember Schneier's Law: "Any person can invent a security system so clever that she or he can't think of how to break it." The only way your designs are going to be trusted is if you've made a name for yourself breaking other people's designs.

One final word about cryptography. Modern cryptography is particularly hard to learn. In addition to everything above, it requires graduate-level knowledge in mathematics. And, as in computer security in general, your prowess is demonstrated by what you can break. The field has progressed a lot since I wrote this guide and self-study cryptanalysis course a dozen years ago, but they're not bad places to start.

When It Comes to Security, We're Back to Feudalism

Originally published in Wired, November 26, 2012

Some of us have pledged our allegiance to Google: We have Gmail accounts, we use Google Calendar and Google Docs, and we have Android phones. Others have pledged allegiance to Apple: We have Macintosh laptops, iPhones, and iPads; and we let iCloud automatically synchronize and back up everything. Still others of us let Microsoft do it all. Or we buy our music and e-books from Amazon, which keeps records of what we own and allows downloading to a Kindle, computer, or phone. Some of us have pretty much abandoned e-mail altogether. . . for Facebook.

These vendors are becoming our feudal lords, and we are becoming their vassals. We might refuse to pledge allegiance to all of them—or to a particular one we don't like. Or we can spread our allegiance around. But either way, it's becoming increasingly difficult to not pledge allegiance to at least one of them.

Feudalism provides security. Classical medieval feudalism depended on overlapping, complex, hierarchical relationships. There were oaths and obligations: a series of rights and privileges. A critical aspect of this system was protection: vassals would pledge their allegiance to a lord, and in return, that lord would protect them from harm.

Of course, I'm romanticizing here; European history was never this simple, and the description is based on stories of that time, but that's the general model.

And it's this model that's starting to permeate computer security today.

I Pledge Allegiance to the United States of Convenience

Traditional computer security centered around users. Users had to purchase and install anti-virus software and firewalls, ensure their operating system and network were configured properly, update their software, and generally manage their own security.

This model is breaking, largely due to two developments:

1. New Internet-enabled devices where the vendor maintains more control over the hardware and software than we do—like the iPhone and Kindle; and
2. Services where the host maintains our data for us—like Flickr and Hotmail.

Now, we users must trust the security of these hardware manufacturers, software vendors, and cloud providers.

We choose to do it because of the convenience, redundancy, automation, and shareability. We like it when we can access our e-mail anywhere, from any computer. We like it when we can restore our contact lists after we've lost our phones. We want our calendar entries to automatically appear on all of our devices. These cloud storage sites do a better job of backing up our photos and files than we would manage by ourselves; Apple does a great job keeping malware out of its iPhone apps store.

In this new world of computing, we give up a certain amount of control, and in exchange we trust that our lords will both treat us well and protect us from harm. Not only will our software be continually updated with the newest and coolest functionality, but we trust it will happen without our being overtaxed by fees and required upgrades. We trust that our data and devices won't be exposed to hackers, criminals, and malware. We trust that governments won't be allowed to illegally spy on us.

Trust is our only option. In this system, we have no control over the security provided by our feudal lords. We don't know what sort of security methods they're using, or how they're configured. We mostly can't install our own security products on iPhones or Android phones; we certainly can't install them on Facebook, Gmail, or Twitter. Sometimes we have control over whether or not to accept the automatically flagged updates—iPhone, for example—but we rarely know what they're about or whether they'll break anything else. (On the Kindle, we don't even have that freedom.)

The Good, the Bad, and the Ugly

I'm not saying that feudal security is all bad. For the average user, giving up control is largely a good thing. These software vendors and cloud providers do a lot better job of security than the average computer user would. Automatic cloud backup saves a lot of data; automatic updates prevent a lot of malware. The network security at any of these providers is better than that of most home users.

Feudalism is good for the individual, for small startups, and for medium-sized businesses that can't afford to hire their own in-house or specialized expertise. Being a vassal has its advantages, after all.

For large organizations, however, it's more of a mixed bag. These organizations are used to trusting other companies with critical corporate functions: They've been outsourcing their payroll, tax preparation, and legal services for decades. But IT regulations often require audits. Our lords don't allow vassals to audit them, even if those vassals are themselves large and powerful.

Yet feudal security isn't without its risks.

Our lords can make mistakes with security, as recently happened with Apple, Facebook, and Photobucket. They can act arbitrarily and capriciously, as Amazon did when it cut off a Kindle user for living in the wrong country. They tether us like serfs; just try to take data from one digital lord to another.

Ultimately, they will always act in their own self-interest, as companies do when they mine our data in order to sell more advertising and make more

money. These companies own us, so they can sell us off—again, like serfs—to rival lords. . . or turn us into the authorities.

Historically, early feudal arrangements were ad hoc, and the more powerful party would often simply renege on his part of the bargain. Eventually, the arrangements were formalized and standardized: both parties had rights and privileges (things they could do) as well as protections (things they couldn't do to each other).

Today's Internet feudalism, however, is ad hoc and one-sided. We give companies our data and trust them with our security, but we receive very few assurances of protection in return, and those companies have very few restrictions on what they can do.

This needs to change. There should be limitations on what cloud vendors can do with our data; rights, like the requirement that they delete our data when we want them to; and liabilities when vendors mishandle our data.

Like everything else in security, it's a trade-off. We need to balance that trade-off. In Europe, it was the rise of the centralized state and the rule of law that undermined the ad hoc feudal system; it provided more security and stability for both lords and vassals. But these days, government has largely abdicated its role in cyberspace, and the result is a return to the feudal relationships of yore.

Perhaps instead of hoping that our Internet-era lords will be sufficiently clever and benevolent—or putting our faith in the Robin Hoods who block phone surveillance and circumvent DRM systems—it's time we step in in our role as governments (both national and international) to create the regulatory environments that protect us vassals (and the lords as well). Otherwise, we really are just serfs.

You Have No Control Over Security on the Feudal Internet

Originally published in Harvard Business Review, June 6, 2013

Facebook regularly abuses the privacy of its users. Google has stopped supporting its popular RSS feeder. Apple prohibits all iPhone apps that are political or sexual. Microsoft might be cooperating with some governments to spy on Skype calls, but we don't know which ones. Both Twitter and LinkedIn

have recently suffered security breaches that affected the data of hundreds of thousands of their users.

If you've started to think of yourself as a hapless peasant in a *Game of Thrones* power struggle, you're more right than you may realize. These are not traditional companies, and we are not traditional customers. These are feudal lords, and we are their vassals, peasants, and serfs.

Power has shifted in IT, in favor of both cloud-service providers and closed-platform vendors. This power shift affects many things, and it profoundly affects security.

Traditionally, computer security was the user's responsibility. Users purchased their own antivirus software and firewalls, and any breaches were blamed on their inattentiveness. It's kind of a crazy business model. Normally we expect the products and services we buy to be safe and secure, but in IT we tolerated lousy products and supported an enormous aftermarket for security.

Now that the IT industry has matured, we expect more security "out of the box." This has become possible largely because of two technology trends: cloud computing and vendor-controlled platforms. The first means that most of our data resides on other networks: Google Docs, Salesforce.com, Facebook, Gmail. The second means that our new Internet devices are both closed and controlled by the vendors, giving us limited configuration control: iPhones, ChromeBooks, Kindles, Blackberries. Meanwhile, our relationship with IT has changed. We used to use our computers to do things. We now use our vendor-controlled computing devices to go places. All of these places are owned by someone.

The new security model is that someone else takes care of it—without telling us any of the details. I have no control over the security of my Gmail or my photos on Flickr. I can't demand greater security for my presentations on Prezi or my task list on Trello, no matter how confidential they are. I can't audit any of these cloud services. I can't delete cookies on my iPad or ensure that files are securely erased. Updates on my Kindle happen automatically, without my knowledge or consent. I have so little visibility into the security of Facebook that I have no idea what operating system they're using.

There are a lot of good reasons why we're all flocking to these cloud services and vendor-controlled platforms. The benefits are enormous, from cost to convenience to reliability to security itself. But it is inherently a feudal relationship. We cede control of our data and computing platforms to these companies and trust that they will treat us well and protect us from harm. And if we pledge complete allegiance to them—if we let them control our email

and calendar and address book and photos and everything—we get even more benefits. We become their vassals; or, on a bad day, their serfs.

There are a lot of feudal lords out there. Google and Apple are the obvious ones, but Microsoft is trying to control both user data and the end-user platform as well. Facebook is another lord, controlling much of the socializing we do on the Internet. Other feudal lords are smaller and more specialized—Amazon, Yahoo, Verizon, and so on—but the model is the same.

To be sure, feudal security has its advantages. These companies are much better at security than the average user. Automatic backup has saved a lot of data after hardware failures, user mistakes, and malware infections. Automatic updates have increased security dramatically. This is also true for small organizations; they are more secure than they would be if they tried to do it themselves. For large corporations with dedicated IT security departments, the benefits are less clear. Sure, even large companies outsource critical functions like tax preparation and cleaning services, but large companies have specific requirements for security, data retention, audit, and so on—and that's just not possible with most of these feudal lords.

Feudal security also has its risks. Vendors can, and do, make security mistakes affecting hundreds of thousands of people. Vendors can lock people into relationships, making it hard for them to take their data and leave. Vendors can act arbitrarily, against our interests; Facebook regularly does this when it changes peoples' defaults, implements new features, or modifies its privacy policy. Many vendors give our data to the government without notice, consent, or a warrant; almost all sell it for profit. This isn't surprising, really; companies should be expected to act in their own self-interest and not in their users' best interest.

The feudal relationship is inherently based on power. In Medieval Europe, people would pledge their allegiance to a feudal lord in exchange for that lord's protection. This arrangement changed as the lords realized that they had all the power and could do whatever they wanted. Vassals were used and abused; peasants were tied to their land and became serfs.

It's the Internet lords' popularity and ubiquity that enable them to profit; laws and government relationships make it easier for them to hold onto power. These lords are vying with each other for profits and power. By spending time on their sites and giving them our personal information—whether through search queries, e-mails, status updates, likes, or simply our behavioral characteristics—we are providing the raw material for that struggle. In this way we are like serfs, toiling the land for our feudal lords. If you don't believe me, try

to take your data with you when you leave Facebook. And when war breaks out among the giants, we become collateral damage.

So how do we survive? Increasingly, we have little alternative but to trust *someone*, so we need to decide who we trust—and who we don't—and then act accordingly. This isn't easy; our feudal lords go out of their way not to be transparent about their actions, their security, or much of anything. Use whatever power you have—as individuals, none; as large corporations, more—to negotiate with your lords. And, finally, don't be extreme in any way: politically, socially, culturally. Yes, you can be shut down without recourse, but it's usually those on the edges that are affected. Not much solace, I agree, but it's something.

On the policy side, we have an action plan. In the short term, we need to keep circumvention—the ability to modify our hardware, software, and data files—legal and preserve net neutrality. Both of these things limit how much the lords can take advantage of us, and they increase the possibility that the market will force them to be more benevolent. The last thing we want is the government—that's us—spending resources to enforce one particular business model over another and stifling competition.

In the longer term, we all need to work to reduce the power imbalance. Medieval feudalism evolved into a more balanced relationship in which lords had responsibilities as well as rights. Today's Internet feudalism is both ad-hoc and one-sided. We have no choice but to trust the lords, but we receive very few assurances in return. The lords have a lot of rights, but few responsibilities or limits. We need to balance this relationship, and government intervention is the only way we're going to get it. In medieval Europe, the rise of the centralized state and the rule of law provided the stability that feudalism lacked. The Magna Carta first forced responsibilities on governments and put humans on the long road toward government by the people and for the people.

We need a similar process to rein in our Internet lords, and it's not something that market forces are likely to provide. The very definition of power is changing, and the issues are far bigger than the Internet and our relationships with our IT providers.

2

Crime, Terrorism, Spying, and War

America's Dilemma: Close Security Holes, or Exploit Them Ourselves

Originally published in Wired News, May 1, 2008

On April 27, 2007, Estonia was attacked in cyberspace. Following a diplomatic incident with Russia about the relocation of a Soviet World War II memorial, the networks of many Estonian organizations, including the Estonian parliament, banks, ministries, newspapers and broadcasters, were attacked and—in many cases—shut down. Estonia was quick to blame Russia, which was equally quick to deny any involvement.

It was hyped as the first cyberwar: Russia attacking Estonia in cyberspace. But nearly a year later, evidence that the Russian government was involved in the denial-of-service attacks still hasn't emerged. Though Russian hackers were indisputably the major instigators of the attack, the only individuals positively identified have been young ethnic Russians living inside Estonia, who were pissed off over the statue incident.

You know you've got a problem when you can't tell a hostile attack by another nation from bored kids with an axe to grind.

Separating cyberwar, cyberterrorism and cybercrime isn't easy; these days you need a scorecard to tell the difference. It's not just that it's hard to trace people in cyberspace, it's that military and civilian attacks—and defenses—look the same.

The traditional term for technology the military shares with civilians is “dual use.” Unlike hand grenades and tanks and missile targeting systems, dual-use technologies have both military and civilian applications. Dual-use

technologies used to be exceptions; even things you'd expect to be dual use, like radar systems and toilets, were designed differently for the military. But today, almost all information technology is dual use. We both use the same operating systems, the same networking protocols, the same applications, and even the same security software.

And attack technologies are the same. The recent spurt of targeted hacks against US military networks, commonly attributed to China, exploit the same vulnerabilities and use the same techniques as criminal attacks against corporate networks. Internet worms make the jump to classified military networks in less than 24 hours, even if those networks are physically separate. The Navy Cyber Defense Operations Command uses the same tools against the same threats as any large corporation.

Because attackers and defenders use the same IT technology, there is a fundamental tension between cyberattack and cyberdefense. The National Security Agency has referred to this as the “equities issue,” and it can be summarized as follows: When a military discovers a vulnerability in a dual-use technology, they can do one of two things. They can alert the manufacturer and fix the vulnerability, thereby protecting both the good guys and the bad guys. Or they can keep quiet about the vulnerability and not tell anyone, thereby leaving the good guys insecure but also leaving the bad guys insecure.

The equities issue has long been hotly debated inside the NSA. Basically, the NSA has two roles: eavesdrop on their stuff, and protect our stuff. When both sides use the same stuff, the agency has to decide whether to exploit vulnerabilities to eavesdrop on their stuff or close the same vulnerabilities to protect our stuff.

In the 1980s and before, the tendency of the NSA was to keep vulnerabilities to themselves. In the 1990s, the tide shifted, and the NSA was starting to open up and help us all improve our security defense. But after the attacks of 9/11, the NSA shifted back to the attack: vulnerabilities were to be hoarded in secret. Slowly, things in the US are shifting back again.

So now we're seeing the NSA help secure Windows Vista and releasing their own version of Linux. The DHS, meanwhile, is funding a project to secure popular open source software packages, and across the Atlantic the UK's GCHQ is finding bugs in PGPDisk and reporting them back to the company. (NSA is rumored to be doing the same thing with BitLocker.)

I'm in favor of this trend, because my security improves for free. Whenever the NSA finds a security problem and gets the vendor to fix it, our security gets better. It's a side-benefit of dual-use technologies.

But I want governments to do more. I want them to use their buying power to improve my security. I want them to offer countrywide contracts for software, both security and non-security, that have explicit security requirements. If these contracts are big enough, companies will work to modify their products to meet those requirements. And again, we all benefit from the security improvements.

The only example of this model I know about is a US government-wide procurement competition for full-disk encryption, but this can certainly be done with firewalls, intrusion detection systems, databases, networking hardware, even operating systems.

When it comes to IT technologies, the equities issue should be a no-brainer. The good uses of our common hardware, software, operating systems, network protocols, and everything else vastly outweigh the bad uses. It's time that the government used its immense knowledge and experience, as well as its buying power, to improve cybersecurity for all of us.

Are Photographers Really a Threat?

Originally published in the Guardian, June 4, 2008

What is it with photographers these days? Are they really all terrorists, or does everyone just think they are?

Since 9/11, there has been an increasing war on photography. Photographers have been harassed, questioned, detained, arrested or worse, and declared to be unwelcome. We've been repeatedly told to watch out for photographers, especially suspicious ones. Clearly any terrorist is going to first photograph his target, so vigilance is required.

Except that it's nonsense. The 9/11 terrorists didn't photograph anything. Nor did the London transport bombers, the Madrid subway bombers, or the liquid bombers arrested in 2006. Timothy McVeigh didn't photograph the Oklahoma City Federal Building. The Unabomber didn't photograph anything; neither did shoe-bomber Richard Reid. Photographs aren't being found amongst the papers of Palestinian suicide bombers. The IRA wasn't known for its photography. Even those manufactured terrorist plots that the US government likes to talk about—the Ft. Dix terrorists, the JFK airport bombers, the Miami 7, the Lackawanna 6—no photography.

Given that real terrorists, and even wannabe terrorists, don't seem to photograph anything, why is it such pervasive conventional wisdom that terrorists photograph their targets? Why are our fears so great that we have no choice but to be suspicious of any photographer?

Because it's a movie-plot threat.

A movie-plot threat is a specific threat, vivid in our minds like the plot of a movie. You remember them from the months after the 9/11 attacks: anthrax spread from crop dusters, a contaminated milk supply, terrorist scuba divers armed with almanacs. Our imaginations run wild with detailed and specific threats, from the news, and from actual movies and television shows. These movie plots resonate in our minds and in the minds of others we talk to. And many of us get scared.

Terrorists taking pictures is a quintessential detail in any good movie. Of course it makes sense that terrorists will take pictures of their targets. They have to do reconnaissance, don't they? We need 45 minutes of television action before the actual terrorist attack—90 minutes if it's a movie—and a photography scene is just perfect. It's our movie-plot terrorists that are photographers, even if the real-world ones are not.

The problem with movie-plot security is it only works if we guess the plot correctly. If we spend a zillion dollars defending Wimbledon and terrorists blow up a different sporting event, that's money wasted. If we post guards all over the Underground and terrorists bomb a crowded shopping area, that's also a waste. If we teach everyone to be alert for photographers, and terrorists don't take photographs, we've wasted money and effort, and taught people to fear something they shouldn't.

And even if terrorists did photograph their targets, the math doesn't make sense. Billions of photographs are taken by honest people every year, 50 billion by amateurs alone in the US. And the national monuments you imagine terrorists taking photographs of are the same ones tourists like to take pictures of. If you see someone taking one of those photographs, the odds are infinitesimal that he's a terrorist.

Of course, it's far easier to explain the problem than it is to fix it. Because we're a species of storytellers, we find movie-plot threats uniquely compelling. A single vivid scenario will do more to convince people that photographers might be terrorists than all the data I can muster to demonstrate that they're not.

Fear aside, there aren't many legal restrictions on what you can photograph from a public place that's already in public view. If you're harassed, it's almost certainly a law enforcement official, public or private, acting way beyond his authority. There's nothing in any post-9/11 law that restricts your right to photograph.

This is worth fighting. Search "photographer rights" on Google and download one of the several wallet documents that can help you if you get harassed; I found one for the UK, US, and Australia. Don't cede your right to photograph in public. Don't propagate the terrorist photographer story. Remind them that prohibiting photography was something we used to ridicule about the USSR. Eventually sanity will be restored, but it may take a while.

CCTV Doesn't Keep Us Safe, Yet the Cameras Are Everywhere

Originally published in the Guardian, June 26, 2008

Pervasive security cameras don't substantially reduce crime. There are exceptions, of course, and that's what gets the press. Most famously, CCTV cameras helped catch James Bulger's murderers in 1993. And earlier this year, they helped convict Steve Wright of murdering five women in the Ipswich area. But these are the well-publicized exceptions. Overall, CCTV cameras aren't very effective.

This fact has been demonstrated again and again: by a comprehensive study for the Home Office in 2005, by several studies in the US, and again with new data announced last month by New Scotland Yard. They actually solve very few crimes, and their deterrent effect is minimal.

Conventional wisdom predicts the opposite. But if that were true, then camera-happy London, with something like 500,000, would be the safest city on the planet. It isn't, of course, because of technological limitations of cameras, organizational limitations of police and the adaptive abilities of criminals.

To some, it's comforting to imagine vigilant police monitoring every camera, but the truth is very different. Most CCTV footage is never looked at until well after a crime is committed. When it is examined, it's very common for the viewers not to identify suspects. Lighting is bad and images are grainy, and

criminals tend not to stare helpfully at the lens. Cameras break far too often. The best camera systems can still be thwarted by sunglasses or hats. Even when they afford quick identification—think of the 2005 London transport bombers and the 9/11 terrorists—police are often able to identify suspects without the cameras. Cameras afford a false sense of security, encouraging laziness when we need police to be vigilant.

The solution isn't for police to watch the cameras. Unlike an officer walking the street, cameras only look in particular directions at particular locations. Criminals know this, and can easily adapt by moving their crimes to someplace not watched by a camera—and there will always be such places. Additionally, while a police officer on the street can respond to a crime in progress, the same officer in front of a CCTV screen can only dispatch another officer to arrive much later. By their very nature, cameras result in underused and misallocated police resources.

Cameras aren't completely ineffective, of course. In certain circumstances, they're effective in reducing crime in enclosed areas with minimal foot traffic. Combined with adequate lighting, they substantially reduce both personal attacks and auto-related crime in car parks. And from some perspectives, simply moving crime around is good enough. If a local Tesco installs cameras in its store, and a robber targets the store next door as a result, that's money well spent by Tesco. But it doesn't reduce the overall crime rate, so is a waste of money to the township.

But the question really isn't whether cameras reduce crime; the question is whether they're worth it. And given their cost (£500 m in the past 10 years), their limited effectiveness, the potential for abuse (spying on naked women in their own homes, sharing nude images, selling best-of videos, and even spying on national politicians) and their Orwellian effects on privacy and civil liberties, most of the time they're not. The funds spent on CCTV cameras would be far better spent on hiring experienced police officers.

We live in a unique time in our society: the cameras are everywhere, and we can still see them. Ten years ago, cameras were much rarer than they are today. And in 10 years, they'll be so small you won't even notice them. Already, companies like L-1 Security Solutions are developing police-state CCTV surveillance technologies like facial recognition for China, technology that will find their way into countries like the UK. The time to address appropriate limits on this technology is before the cameras fade from notice.

Chinese Cyberattacks: Myth or Menace?

Originally published in Information Security, July 2008

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

The popular media narrative is that there is a coordinated attempt by the Chinese government to hack into US computers—military, government, corporate—and steal secrets. The truth is a lot more complicated.

There certainly is a lot of hacking coming out of China. Any company that does security monitoring sees it all the time. Of course, they can't prove that it comes out of China. But the majority of servers used in the attacks are located in China, using DNS bouncers that can only be registered by people literate in Chinese. The hacker websites where different hackers and hacker groups brag about their exploits and sell hacker tools and how-to videos are written in Chinese. Technically, it's possible all the attackers are from, say, Canada and trying to disguise themselves, but it seems pretty unlikely.

These hacker groups seem not to be working for the Chinese government. They don't seem to be coordinated by the Chinese military. They're basically young, male, patriotic Chinese citizens, demonstrating they're as good as everyone else. Besides the American networks the media likes to talk about, their targets also include pro-Tibet, pro-Taiwan, Falun Gong and pro-Uyghur sites.

The hackers are in this for two reasons: fame and glory, and an attempt to make a living. The fame and glory comes from their nationalistic goals. Some of these hackers are heroes in China. They're upholding the country's honor against both anti-Chinese forces like the pro-Tibet movement and larger forces like the United States. And the money comes from several sources. The groups sell owned computers, malware services and data they steal on the black market. They sell hacker tools and videos to others wanting to pay. They even sell t-shirts, hats and other merchandise on their websites.

This is not to say the Chinese military ignores the hacker groups within their country. The People's Liberation Army has long had a doctrine of "informationization." It considers cyberwarfare a leapfrog technology, one that will

allow it to achieve military parity with the West without having to engage in an expensive missile-for-missile arms race like the one that bankrupted the Soviet Union. Certainly the Chinese government knows the leaders of the hacker movement and chooses to look the other way. It probably buys good stuff, and probably recruits for its organizations from this self-selecting pool of experienced hacking experts. It certainly learns from the hackers.

And some of the hackers are good. Scott Henderson has been tracking Chinese hacker groups for years and writes about them in his blog, www.thedarkvisitor.com, and his book of the same name. He's watched the hackers become more sophisticated in tools and techniques. They're stealthy. They do good network reconnaissance. My guess is what the Pentagon thinks is the problem is only a small percentage of the actual problem.

And they discover their own vulnerabilities. Earlier this year, F-Secure found an attack against a pro-Tibet network that used an unpatched zero-day vulnerability to install a backdoor. That same attack was used two weeks earlier against a large multinational defense contractor. They also hoard vulnerabilities. During the 1999 conflict over the two-states theory, in a heated exchange with a group of Taiwanese hackers, one Chinese group threatened to unleash multiple stockpiled worms at once. There was no reason to disbelieve this threat.

If anything, the fact that these groups aren't being run by the Chinese government makes the problem worse. Without central political coordination, they're likely to take more risks, do stupider things and generally ignore the political fallout of their actions. In this regard, they're more like a non-state actor. So while I'm perfectly happy that the US government is using the threat of Chinese hacking as an impetus to get its cybersecurity in order, and I hope it succeeds, I also hope the US government recognizes that these groups are not acting under the direction of the Chinese military and doesn't treat their actions as officially approved by the Chinese government.

How a Classic Man-in-the-Middle Attack Saved Colombian Hostages

Originally published in Wired News, July 10, 2008

Last week's dramatic rescue of 15 hostages held by the guerrilla organization FARC was the result of months of intricate deception on the part of

the Colombian government. At the center was a classic man-in-the-middle attack.

In a man-in-the-middle attack, the attacker inserts himself between two communicating parties. Both believe they're talking to each other, and the attacker can delete or modify the communications at will.

The *Wall Street Journal* reported how this gambit played out in Colombia: "The plan had a chance of working because, for months, in an operation one army officer likened to a "broken telephone," military intelligence had been able to convince Ms. Betancourt's captor, Gerardo Aguilar, a guerrilla known as "Cesar," that he was communicating with his top bosses in the guerrillas' seven-man secretariat. Army intelligence convinced top guerrilla leaders that they were talking to Cesar. In reality, both were talking to army intelligence."

This ploy worked because Cesar and his guerrilla bosses didn't know one another well. They didn't recognize one another's voices, and didn't have a friendship or shared history that could have tipped them off about the ruse. Man-in-the-middle is defeated by context, and the FARC guerrillas didn't have any.

And that's why man-in-the-middle, abbreviated MITM in the computer-security community, is such a problem online: Internet communication is often stripped of any context. There's no way to recognize someone's face. There's no way to recognize someone's voice. When you receive an e-mail purporting to come from a person or organization, you have no idea who actually sent it. When you visit a website, you have no idea if you're really visiting that website. We all like to pretend that we know who we're communicating with—and for the most part, of course, there isn't any attacker inserting himself into our communications—but in reality, we don't. And there are lots of hacker tools that exploit this unjustified trust, and implement MITM attacks.

Even with context, it's still possible for MITM to fool both sides—because electronic communications are often intermittent. Imagine that one of the FARC guerrillas became suspicious about who he was talking to. So he asks a question about their shared history as a test: "What did we have for dinner that time last year?" or something like that. On the telephone, the attacker wouldn't be able to answer quickly, so his ruse would be discovered. But e-mail conversation isn't synchronous. The attacker could simply pass that question through to the other end of the communications, and when he got the answer back, he would be able to reply.

This is the way MITM attacks work against web-based financial systems. A bank demands authentication from the user: a password, a one-time code

from a token or whatever. The attacker sitting in the middle receives the request from the bank and passes it to the user. The user responds to the attacker, who passes that response to the bank. Now the bank assumes it is talking to the legitimate user, and the attacker is free to send transactions directly to the bank. This kind of attack completely bypasses any two-factor authentication mechanisms, and is becoming a more popular identity-theft tactic.

There are cryptographic solutions to MITM attacks, and there are secure web protocols that implement them. Many of them require shared secrets, though, making them useful only in situations where people already know and trust one another.

The NSA-designed STU-III and STE secure telephones solve the MITM problem by embedding the identity of each phone together with its key. (The NSA creates all keys and is trusted by everyone, so this works.) When two phones talk to each other securely, they exchange keys and display the other phone's identity on a screen. Because the phone is in a secure location, the user now knows who he is talking to, and if the phone displays another organization—as it would if there were a MITM attack in progress—he should hang up.

Zfone, a secure VoIP system, protects against MITM attacks with a short authentication string. After two Zfone terminals exchange keys, both computers display a four-character string. The users are supposed to manually verify that both strings are the same—"my screen says 5C19; what does yours say?"—to ensure that the phones are communicating directly with each other and not with an MITM. The AT&T TSD-3600 worked similarly.

This sort of protection is embedded in SSL, although no one uses it. As it is normally used, SSL provides an encrypted communications link to whoever is at the other end: bank and phishing site alike. And the better phishing sites create valid SSL connections, so as to more effectively fool users. But if the user wanted to, he could manually check the SSL certificate to see if it was issued to "National Bank of Trustworthiness" or "Two Guys With a Computer in Nigeria."

No one does, though, because you have to both remember and be willing to do the work. (The browsers could make this easier if they wanted to, but they don't seem to want to.) In the real world, you can easily tell a branch of your bank from a money changer on a street corner. But on the Internet, a phishing site can be easily made to look like your bank's legitimate website. Any method of telling the two apart takes work. And that's the first step to fooling you with a MITM attack.

Man-in-the-middle isn't new, and it doesn't have to be technological. But the Internet makes the attacks easier and more powerful, and that's not going to change anytime soon.

How to Create the Perfect Fake Identity —

Originally published in Wired News, September 4, 2008

Let me start off by saying that I'm making this whole thing up.

Imagine you're in charge of infiltrating sleeper agents into the United States. The year is 1983, and the proliferation of identity databases is making it increasingly difficult to create fake credentials. Ten years ago, someone could have just shown up in the country and gotten a driver's license, Social Security card and bank account—possibly using the identity of someone roughly the same age who died as a young child—but it's getting harder. And you know that trend will only continue. So you decide to grow your own identities.

Call it "identity farming." You invent a handful of infants. You apply for Social Security numbers for them. Eventually, you open bank accounts for them, file tax returns for them, register them to vote, and apply for credit cards in their name. And now, 25 years later, you have a handful of identities ready and waiting for some real people to step into them.

There are some complications, of course. Maybe you need people to sign their name as parents—or, at least, mothers. Maybe you need doctors to fill out birth certificates. Maybe you need to fill out paperwork certifying that you're home-schooling these children. You'll certainly want to exercise their financial identity: depositing money into their bank accounts and withdrawing it from ATMs, using their credit cards and paying the bills, and so on. And you'll need to establish some sort of addresses for them, even if it is just a mail drop.

You won't be able to get driver's licenses or photo IDs in their name. That isn't critical, though; in the US, more than 20 million adult citizens don't have photo IDs. But other than that, I can't think of any reason why identity farming wouldn't work.

Here's the real question: Do you actually have to show up for any part of your life?

Again, I made this all up. I have no evidence that anyone is actually doing this. It's not something a criminal organization is likely to do; twenty-five years

is too distant a payoff horizon. The same logic holds true for terrorist organizations; it's not worth it. It might have been worth it to the KGB—although perhaps harder to justify after the Soviet Union broke up in 1991—and might be an attractive option for existing intelligence adversaries like China.

Immortals could also use this trick to self-perpetuate themselves, inventing their own children and gradually assuming their identity, then killing their parents off. They could even show up for their own driver's license photos, wearing a beard as the father and blue spiked hair as the son. I'm told this is a common idea in *Highlander* fan fiction.

The point isn't to create another movie plot threat, but to point out the central role that data has taken on in our lives. Previously, I've said that we all have a data shadow that follows us around, and that more and more institutions interact with our data shadows instead of with us. We only intersect with our data shadows once in a while—when we apply for a driver's license or passport, for example—and those interactions are authenticated by older, less-secure interactions. The rest of the world assumes that our photo IDs glue us to our data shadows, ignoring the rather flimsy connection between us and our plastic cards. (And, no, REAL-ID won't help.)

It seems to me that our data shadows are becoming increasingly distinct from us, almost with a life of their own. What's important now is our shadows; we're secondary. And as our society relies more and more on these shadows, we might even become unnecessary.

Our data shadows can live a perfectly normal life without us.

A Fetishistic Approach to Security Is a Perverse Way to Keep Us Safe

Originally published in the Guardian, September 4, 2008

We spend far more effort defending our countries against specific movie-plot threats, rather than the real, broad threats. In the US during the months after the 9/11 attacks, we feared terrorists with scuba gear, terrorists with crop dusters and terrorists contaminating our milk supply. Both the UK and the US fear terrorists with small bottles of liquid. Our imaginations run wild with vivid specific threats. Before long, we're envisioning an entire movie plot, without Bruce Willis saving the day. And we're scared.

It's not just terrorism; it's any rare risk in the news. The big fear in Canada right now, following a particularly gruesome incident, is random decapitations on intercity buses. In the US, fears of school shootings are much greater than the actual risks. In the UK, it's child predators. And people all over the world mistakenly fear flying more than driving. But the very definition of news is something that hardly ever happens. If an incident is in the news, we shouldn't worry about it. It's when something is so common that it's no longer news—car crashes, domestic violence—that we should worry. But that's not the way people think.

Psychologically, this makes sense. We are a species of storytellers. We have good imaginations and we respond more emotionally to stories than to data. We also judge the probability of something by how easy it is to imagine, so stories that are in the news feel more probable—and ominous—than stories that are not. As a result, we overreact to the rare risks we hear stories about, and fear specific plots more than general threats.

The problem with building security around specific targets and tactics is that it's only effective if we happen to guess the plot correctly. If we spend billions defending the Underground and terrorists bomb a school instead, we've wasted our money. If we focus on the World Cup and terrorists attack Wimbledon, we've wasted our money.

It's this fetish-like focus on tactics that results in the security follies at airports. We ban guns and knives, and terrorists use box-cutters. We take away box-cutters and corkscrews, so they put explosives in their shoes. We screen shoes, so they use liquids. We take away liquids, and they're going to do something else. Or they'll ignore airplanes entirely and attack a school, church, theatre, stadium, shopping mall, airport terminal outside the security area, or any of the other places where people pack together tightly.

These are stupid games, so let's stop playing. Some high-profile targets deserve special attention and some tactics are worse than others. Airplanes are particularly important targets because they are national symbols and because a small bomb can kill everyone aboard. Seats of government are also symbolic, and therefore attractive, targets. But targets and tactics are interchangeable.

The following three things are true about terrorism. One, the number of potential terrorist targets is infinite. Two, the odds of the terrorists going after any one target is zero. And three, the cost to the terrorist of switching targets is zero.

We need to defend against the broad threat of terrorism, not against specific movie plots. Security is most effective when it doesn't require us to guess. We need to focus resources on intelligence and investigation: identifying terrorists, cutting off their funding and stopping them regardless of what their plans are. We need to focus resources on emergency response: lessening the impact of a terrorist attack, regardless of what it is. And we need to face the geopolitical consequences of our foreign policy.

In 2006, UK police arrested the liquid bombers not through diligent airport security, but through intelligence and investigation. It didn't matter what the bombers' target was. It didn't matter what their tactic was. They would have been arrested regardless. That's smart security. Now we confiscate liquids at airports, just in case another group happens to attack the exact same target in exactly the same way. That's just illogical.

The Seven Habits of Highly Ineffective Terrorists

Originally published in Wired News, October 1, 2008

Most counterterrorism policies fail, not because of tactical problems, but because of a fundamental misunderstanding of what motivates terrorists in the first place. If we're ever going to defeat terrorism, we need to understand what drives people to become terrorists in the first place.

Conventional wisdom holds that terrorism is inherently political, and that people become terrorists for political reasons. This is the "strategic" model of terrorism, and it's basically an economic model. It posits that people resort to terrorism when they believe—rightly or wrongly—that terrorism is worth it; that is, when they believe the political gains of terrorism minus the political costs are greater than if they engaged in some other, more peaceful form of protest. It's assumed, for example, that people join Hamas to achieve a Palestinian state; that people join the PKK to attain a Kurdish national homeland; and that people join al-Qaida to, among other things, get the United States out of the Persian Gulf.

If you believe this model, the way to fight terrorism is to change that equation, and that's what most experts advocate. Governments tend to minimize the political gains of terrorism through a no-concessions policy; the international community tends to recommend reducing the political grievances of

terrorists via appeasement, in hopes of getting them to renounce violence. Both advocate policies to provide effective nonviolent alternatives, like free elections.

Historically, none of these solutions has worked with any regularity. Max Abrahms, a predoctoral fellow at Stanford University's Center for International Security and Cooperation, has studied dozens of terrorist groups from all over the world. He argues that the model is wrong. In a paper published this year in *International Security* that—sadly—doesn't have the title "Seven Habits of Highly Ineffective Terrorists," he discusses, well, seven habits of highly ineffective terrorists. These seven tendencies are seen in terrorist organizations all over the world, and they directly contradict the theory that terrorists are political maximizers.

Terrorists, he writes, (1) attack civilians, a policy that has a lousy track record of convincing those civilians to give the terrorists what they want; (2) treat terrorism as a first resort, not a last resort, failing to embrace nonviolent alternatives like elections; (3) don't compromise with their target country, even when those compromises are in their best interest politically; (4) have protean political platforms, which regularly, and sometimes radically, change; (5) often engage in anonymous attacks, which precludes the target countries making political concessions to them; (6) regularly attack other terrorist groups with the same political platform; and (7) resist disbanding, even when they consistently fail to achieve their political objectives or when their stated political objectives have been achieved.

Abrahms has an alternative model to explain all this: People turn to terrorism for social solidarity. He theorizes that people join terrorist organizations worldwide in order to be part of a community, much like the reason inner-city youths join gangs in the United States.

The evidence supports this. Individual terrorists often have no prior involvement with a group's political agenda, and often join multiple terrorist groups with incompatible platforms. Individuals who join terrorist groups are frequently not oppressed in any way, and often can't describe the political goals of their organizations. People who join terrorist groups most often have friends or relatives who are members of the group, and the great majority of terrorists are socially isolated: unmarried young men or widowed women who weren't working prior to joining. These things are true for members of terrorist groups as diverse as the IRA and al-Qaida.

For example, several of the 9/11 hijackers planned to fight in Chechnya, but they didn't have the right paperwork so they attacked America instead.

The mujahedeen had no idea whom they would attack after the Soviets withdrew from Afghanistan, so they sat around until they came up with a new enemy: America. Pakistani terrorists regularly defect to another terrorist group with a totally different political platform. Many new al-Qaida members say, unconvincingly, that they decided to become a jihadist after reading an extreme, anti-American blog, or after converting to Islam, sometimes just a few weeks before. These people know little about politics or Islam, and they frankly don't even seem to care much about learning more. The blogs they turn to don't have a lot of substance in these areas, even though more informative blogs do exist.

All of this explains the seven habits. It's not that they're ineffective; it's that they have a different goal. They might not be effective politically, but they are effective socially: They all help preserve the group's existence and cohesion.

This kind of analysis isn't just theoretical; it has practical implications for counterterrorism. Not only can we now better understand who is likely to become a terrorist, we can engage in strategies specifically designed to weaken the social bonds within terrorist organizations. Driving a wedge between group members—commuting prison sentences in exchange for actionable intelligence, planting more double agents within terrorist groups—will go a long way to weakening the social bonds within those groups.

We also need to pay more attention to the socially marginalized than to the politically downtrodden, like unassimilated communities in Western countries. We need to support vibrant, benign communities and organizations as alternative ways for potential terrorists to get the social cohesion they need. And finally, we need to minimize collateral damage in our counterterrorism operations, as well as clamping down on bigotry and hate crimes, which just creates more dislocation and social isolation, and the inevitable calls for revenge.

Why Society Should Pay the True Costs of Security

Originally published in the Guardian, October 2, 2008

It's not true that no one worries about terrorists attacking chemical plants. It's just that our politics seem to leave us unable to deal with the threat. Toxins such as ammonia, chlorine, propane and flammable mixtures are being produced or stored as a result of legitimate industrial processes. Chlorine gas

is particularly toxic; in addition to bombing a plant, someone could hijack a chlorine truck or blow up a railcar. Phosgene is even more dangerous. And many chemical plants are located in places where an act of sabotage—or an accident—could threaten thousands of people.

The problem of securing chemical plants is simple once you understand the underlying economics. Normally, we leave the security of something up to its owner. The basic idea is that the owner of each chemical plant 1) best understands the risks, and 2) is the one who loses out if security fails. Any outsider—i.e., regulatory agency—is just going to get it wrong. It's the basic free-market argument, and in most instances it makes a lot of sense.

And chemical plants have security. They have cameras, fences, guards. They have built-in fail-safe mechanisms. For example, many large chemical companies use hazardous substances like phosgene, methyl isocyanate and ethylene oxide in their plants, but don't ship them between locations. They minimize the amounts that are stored.

This is all good and right, and what free-market capitalism dictates. The problem is, that isn't enough. Any rational owner of a chemical plant will only secure the plant up to its value to him or her. That is, if the plant is worth \$100m (£55m), then it makes no sense to spend \$200m on securing it. If the odds of it being attacked are less than 1%, it doesn't even make sense to spend \$1m on securing it. The mathematics are more complicated than this, because you have to factor in such things as the reputational cost of having your name splashed all over the media after an incident, but that's the basic idea.

But to society, the cost of an attack can be much, much greater. If a terrorist blows up a particularly toxic plant in the middle of a densely populated area, deaths could be in the tens of thousands and damage could be in the hundreds of millions. Indirect economic damage could be in the billions. The owner of the chlorine plant would pay none of these potential costs.

Sure, the owner could be sued. But they're not at risk for more than the value of the company—and, in any case, they'd probably be smarter to take the chance. Expensive lawyers can work wonders, courts can be fickle and the government could step in and bail the company out. And a smart company can often protect itself by spinning off the risky asset in a subsidiary company, or selling it off completely. The overall result is that chemical plants are secured to a much smaller degree than the risk warrants.

If we—the community living near the chemical plant, or the nation as a whole—expect the owner of that plant to spend money for increased security to account for those externalities, we're going to have to pay for it.

We have three ways of doing that. One, we can do it ourselves, stationing government police or military or contractors around the chemical plants. Two, we can pay the owners to do it, subsidizing some sort of security standard. Or three, we could regulate security and force the companies to pay for it themselves.

There's no free lunch, of course. "We," as in society, still pay for it in increased prices for whatever the chemical plants are producing, but the cost is paid for by the consumers rather than by taxpayers. Asking nicely just isn't going to work.

Why Technology Won't Prevent Identity Theft

Originally published in the Wall Street Journal, January 9, 2009

Impersonation isn't new. In 1556, a Frenchman was executed for impersonating Martin Guerre and this week hackers impersonated Barack Obama on Twitter. It's not even unique to humans: mockingbirds, Viceroy butterflies, and the brown octopus all use impersonation as a survival strategy. For people, detecting impersonation is a hard problem for three reasons: we need to verify the identity of people we don't know, we interact with people through "narrow" communications channels like the telephone and Internet, and we want computerized systems to do the verification for us.

Traditional impersonation involves people fooling people. It's still done today: impersonating garbage men to collect tips, impersonating parking lot attendants to collect fees, or impersonating the French president to fool Sarah Palin. Impersonating people like policemen, security guards, and meter readers is a common criminal tactic.

These tricks work because we all regularly interact with people we don't know. No one could successfully impersonate your brother, your best friend, or your boss, because you know them intimately. But a policeman or a parking lot attendant? That's just someone with a badge or a uniform. But badges and ID cards only help if you know how to verify one. Do you know what a valid police ID looks like? Or how to tell a real telephone repairman's badge from a forged one?

Still, it's human nature to trust these credentials. We naturally trust uniforms, even though we know that anyone can wear one. When we visit a Web

site, we use the professionalism of the page to judge whether or not it's really legitimate—never mind that anyone can cut and paste graphics. Watch the next time someone other than law enforcement verifies your ID; most people barely look at it.

Impersonation is even easier over limited communications channels. On the telephone, how can you distinguish someone working at your credit card company from someone trying to steal your account details and login information? On e-mail, how can you distinguish someone from your company's tech support from a hacker trying to break into your network—or the mayor of Paris from an impersonator? Once in a while someone frees himself from jail by faxing a forged release order to his warden. This is social engineering: impersonating someone convincingly enough to fool the victim.

These days, a lot of identity verification happens with computers. Computers are fast at computation but not very good at judgment, and can be tricked. So people can fool speed cameras by taping a fake license plate over the real one, fingerprint readers with a piece of tape, or automatic face scanners with—and I'm not making this up—a photograph of a face held in front of their own. Even the most bored policeman wouldn't fall for any of those tricks.

This is why identity theft is such a big problem today. So much authentication happens online, with only a small amount of information: user ID, password, birth date, Social Security number, and so on. Anyone who gets that information can impersonate you to a computer, which doesn't know any better.

Despite all of these problems, most authentication systems work most of the time. Even something as ridiculous as faxed signatures work, and can be legally binding. But no authentication system is perfect, and impersonation is always possible.

This lack of perfection is okay, though. Security is a trade-off, and any well-designed authentication system balances security with ease of use, customer acceptance, cost, and so on. More authentication isn't always better. Banks make this trade-off when they don't bother authenticating signatures on checks under amounts like \$25,000; it's cheaper to deal with fraud after the fact. Web sites make this trade-off when they use simple passwords instead of something more secure, and merchants make this trade-off when they don't bother verifying your signature against your credit card. We make this trade-off when we accept police badges, Best Buy uniforms, and faxed signatures with only a cursory amount of verification.

Good authentication systems also balance false positives against false negatives. Impersonation is just one way these systems can fail; they can also fail

to authenticate the real person. An ATM is better off allowing occasional fraud than preventing legitimate account holders access to their money. On the other hand, a false positive in a nuclear launch system is much more dangerous; better to not launch the missiles.

Decentralized authentication systems work better than centralized ones. Open your wallet, and you'll see a variety of physical tokens used to identify you to different people and organizations: your bank, your credit card company, the library, your health club, and your employer, as well as a catch-all driver's license used to identify you in a variety of circumstances. That assortment is actually more secure than a single centralized identity card: each system must be broken individually, and breaking one doesn't give the attacker access to everything. This is one of the reasons that centralized systems like REAL-ID make us less secure.

Finally, any good authentication system uses defense in depth. Since no authentication system is perfect, there need to be other security measures in place if authentication fails. That's why all of a corporation's assets and information isn't available to anyone who can bluff his way into the corporate offices. That is why credit card companies have expert systems analyzing suspicious spending patterns. And it's why identity theft won't be solved by making personal information harder to steal.

We can reduce the risk of impersonation, but it will always be with us; technology cannot "solve" it in any absolute sense. Like any security, the trick is to balance the trade-offs. Too little security, and criminals withdraw money from all our bank accounts. Too much security and when Barack Obama calls to congratulate you on your reelection, you won't believe it's him.

Terrorists May Use Google Earth, but Fear Is No Reason to Ban It

Originally published in the Guardian, January 29, 2009

It regularly comes as a surprise to people that our own infrastructure can be used against us. And in the wake of terrorist attacks or plots, there are fear-induced calls to ban, disrupt or control that infrastructure. According to officials investigating the Mumbai attacks, the terrorists used images from Google Earth to help learn their way around. This isn't the first time Google

Earth has been charged with helping terrorists: in 2007, Google Earth images of British military bases were found in the homes of Iraqi insurgents. Incidents such as these have led many governments to demand that Google remove or blur images of sensitive locations: military bases, nuclear reactors, government buildings, and so on. An Indian court has been asked to ban Google Earth entirely.

This isn't the only way our information technology helps terrorists. Last year, a US army intelligence report worried that terrorists could plan their attacks using Twitter, and there are unconfirmed reports that the Mumbai terrorists read the Twitter feeds about their attacks to get real-time information they could use. British intelligence is worried that terrorists might use voice over IP services such as Skype to communicate. Terrorists may train on Second Life and World of Warcraft. We already know they use websites to spread their message and possibly even to recruit.

Of course, all of this is exacerbated by open-wireless access, which has been repeatedly labeled a terrorist tool and which has been the object of attempted bans.

Mobile phone networks help terrorists, too. The Mumbai terrorists used them to communicate with each other. This has led some cities, including London, to propose turning off mobile phone coverage in the event of a terrorist attack.

Let's all stop and take a deep breath. By its very nature, communications infrastructure is general. It can be used to plan both legal and illegal activities, and it's generally impossible to tell which is which. When I send and receive email, it looks exactly the same as a terrorist doing the same thing. To the mobile phone network, a call from one terrorist to another looks exactly the same as a mobile phone call from one victim to another. Any attempt to ban or limit infrastructure affects everybody. If India bans Google Earth, a future terrorist won't be able to use it to plan; nor will anybody else. Open Wi-Fi networks are useful for many reasons, the large majority of them positive, and closing them down affects all those reasons. Terrorist attacks are very rare, and it is almost always a bad trade-off to deny society the benefits of a communications technology just because the bad guys might use it too.

Communications infrastructure is especially valuable during a terrorist attack. Twitter was the best way for people to get real-time information about the attacks in Mumbai. If the Indian government shut Twitter down—or London blocked mobile phone coverage—during a terrorist attack, the lack of communications for everyone, not just the terrorists, would increase the

level of terror and could even increase the body count. Information lessens fear and makes people safer.

None of this is new. Criminals have used telephones and mobile phones since they were invented. Drug smugglers use airplanes and boats, radios and satellite phones. Bank robbers have long used cars and motorcycles as getaway vehicles, and horses before then. I haven't seen it talked about yet, but the Mumbai terrorists used boats as well. They also wore boots. They ate lunch at restaurants, drank bottled water, and breathed the air. Society survives all of this because the good uses of infrastructure far outweigh the bad uses, even though the good uses are—by and large—small and pedestrian and the bad uses are rare and spectacular. And while terrorism turns society's very infrastructure against itself, we only harm ourselves by dismantling that infrastructure in response—just as we would if we banned cars because bank robbers used them too.

Thwarting an Internal Hacker

*Originally published in the Wall Street Journal,
February 16, 2009*

Rajendrasinh Makwana was a UNIX contractor for Fannie Mae. On October 24, he was fired. Before he left, he slipped a logic bomb into the organization's network. The bomb would have “detonated” on January 31. It was programmed to disable access to the server on which it was running, block any network monitoring software, systematically and irretrievably erase everything—and then replicate itself on all 4,000 Fannie Mae servers. Court papers claim the damage would have been in the millions of dollars, a number that seems low. Fannie Mae would have been shut down for at least a week.

Luckily—and it does seem it was pure luck—another programmer discovered the script a week later, and disabled it.

Insiders are a perennial problem. They have access, and they're known by the system. They know how the system and its security works, and its weak points. They have opportunity. Bank heists, casino thefts, large-scale corporate fraud, train robberies: many of the most impressive criminal attacks involve insiders. And, like Makwana's attempt at revenge, these insiders can have pretty intense motives—motives that can only intensify as the economy continues to suffer and layoffs increase.

Insiders are especially pernicious attackers because they're trusted. They have access because they're *supposed* to have access. They have opportunity,

and an understanding of the system, because they use it—or they designed, built, or installed it. They're already inside the security system, making them much harder to defend against.

It's not possible to design a system without trusted people. They're everywhere. In offices, employees are trusted people given access to facilities and resources, and allowed to act—sometimes broadly, sometimes narrowly—in the company's name. In stores, employees are allowed access to the back room and the cash register; and customers are trusted to walk into the store and touch the merchandise. IRS employees are trusted with personal tax information; hospital employees are trusted with personal health information. Banks, airports, and prisons couldn't operate without trusted people.

Replacing trusted people with computers doesn't make the problem go away; it just moves it around and makes it even more complex. The computer, software, and network designers, implementers, coders, installers, maintainers, etc. are all trusted people. See any analysis of the security of electronic voting machines, or some of the frauds perpetrated against computerized gambling machines, for some graphic examples of the risks inherent in replacing people with computers.

Of course, this problem is much, much older than computers. And the solutions haven't changed much throughout history, either. There are five basic techniques to deal with trusted people:

1. Limit the number of trusted people. This one is obvious. The fewer people who have root access to the computer system, know the combination to the safe, or have the authority to sign checks, the more secure the system is.
2. Ensure that trusted people are also trustworthy. This is the idea behind background checks, lie detector tests, personality profiling, prohibiting convicted felons from getting certain jobs, limiting other jobs to citizens, the TSA's no-fly list, and so on, as well as behind bonding employees, which means there are deep pockets standing behind them if they turn out not to be trustworthy.
3. Limit the amount of trust each person has. This is compartmentalization; the idea here is to limit the amount of damage a person can do if he ends up not being trustworthy. This is the concept behind giving people keys that only unlock their office or passwords that only unlock their account, as well as "need to know" and other levels of security clearance.

4. Give people overlapping spheres of trust. This is what security professionals call defense in depth. It's why it takes two people with two separate keys to launch nuclear missiles, and two signatures on corporate checks over a certain value. It's the idea behind bank tellers requiring management overrides for high-value transactions, double-entry bookkeeping, and all those guards and cameras at casinos. It's why, when you go to a movie theater, one person sells you a ticket and another person standing a few yards away tears it in half: It makes it much harder for one employee to defraud the system. It's why key bank employees need to take their two-week vacations all at once—so their replacements have a chance to uncover any fraud.
5. Detect breaches of trust after the fact and prosecute the guilty. In the end, the four previous techniques can only do so well. Trusted people can subvert a system. Most of the time, we discover the security breach after the fact and then punish the perpetrator through the legal system: publicly, so as to provide a deterrence effect and increase the overall level of security in society. This is why audit is so vital.

These security techniques don't only protect against fraud or sabotage; they protect against the more common problem: mistakes. Trusted people aren't perfect; they can inadvertently cause damage. They can make a mistake, or they can be tricked into making a mistake through social engineering.

Good security systems use multiple measures, all working together. Fannie Mae certainly limits the number of people who have the ability to slip malicious scripts into their computer systems, and certainly limits the access that most of these people have. It probably has a hiring process that makes it less likely that malicious people come to work at Fannie Mae. It obviously doesn't have an audit process by which a change one person makes on the servers is checked by someone else; I'm sure that would be prohibitively expensive. Certainly the company's IT department should have terminated Makwana's network access as soon as he was fired, and not at the end of the day.

In the end, systems will always have trusted people who can subvert them. It's important to keep in mind that incidents like this don't happen very often; that most people are honest and honorable. Security is very much designed to protect against the dishonest minority. And often little things—like disabling access immediately upon termination—can go a long way.

An Enterprising Criminal Has Spotted a Gap in the Market

Originally published in the Guardian, April 2, 2009

Before his arrest, Tom Berge stole lead roof tiles from several buildings in south-east England, including the Honeywood Museum in Carshalton, the Croydon parish church, and the Sutton high school for girls. He then sold those tiles to scrap metal dealers.

As a security expert, I find this story interesting for two reasons. First, among attempts to ban, or at least censor, Google Earth, lest it help the terrorists, here is an actual crime that relied on the service: Berge needed Google Earth for reconnaissance.

But more interesting is the discrepancy between the value of the lead tiles to the original owner and to the thief. The Sutton school had to spend £10,000 to buy new lead tiles; the Croydon Church had to repair extensive water damage after the theft. But Berge only received £700 a tonne from London scrap metal dealers.

This isn't an isolated story; the same dynamic is in play with other commodities as well.

There is an epidemic of copper wiring thefts worldwide; copper is being stolen out of telephone and power stations—and off poles in the streets—and thieves have died because they didn't understand the dangers of high voltage. Homeowners are returning from holiday to find the copper pipes stolen from their houses. In 2001, scrap copper was worth 70c (50p) a pound in the US. In April 2008, it was worth \$4.

Gasoline siphoning became more common as pump prices rose. And used restaurant grease, formerly either given away or sold for pennies to farmers, is being stolen from restaurant carparks and turned into biofuels. Newspapers and other recyclables are stolen from pavements, and trees are stolen and resold as Christmas trees.

Iron fences have been stolen from buildings and houses, manhole covers have been stolen from the middle of streets, and aluminum guard rails have been stolen from roadways. Steel is being stolen for scrap, too. In 2004 in Ukraine, thieves stole an entire steel bridge.

These crimes are particularly expensive to society because the replacement cost is much higher than the thief's profit. A manhole is worth \$5–\$10 as scrap, but it costs \$500 to replace, including labor. A thief may take \$20 worth of copper from a construction site, but do \$10,000 in damage in the process. And the increased threat means more money being spent on security to protect those commodities in the first place.

Security can be viewed as a tax on the honest, and these thefts demonstrate that our taxes are going up. And unlike many taxes, we don't benefit from their collection. The cost to society of retrofitting manhole covers with locks, or replacing them with less resalable alternatives, is high; but there is no benefit other than reducing theft.

These crimes are a harbinger of the future: evolutionary pressure on our society, if you will. Criminals are often referred to as social parasites, but they are an early warning system of societal changes. Unfettered by laws or moral restrictions, they can be the first to respond to changes that the rest of society will be slower to pick up on. In fact, currently there's a reprieve. Scrap metal prices are all down from last year—copper is currently \$1.62 per pound, and lead is half what Berge got—and thefts are down too.

We've designed much of our infrastructure around the assumptions that commodities are cheap and theft is rare. We don't protect transmission lines, manhole covers, iron fences, or lead flashing on roofs. But if commodity prices really are headed for new higher stable points, society will eventually react and find alternatives for these items—or find ways to protect them. Criminals were the first to point this out, and will continue to exploit the system until it destabilizes.

We Shouldn't Poison Our Minds with Fear of Bioterrorism

Originally published in the Guardian, May 14, 2009

Terrorists attacking our food supply is a nightmare scenario that has been given new life during the recent swine flu outbreak. Although it seems easy to do, understanding why it hasn't happened is important. GR Dalziel, at the Nanyang Technological University in Singapore, has written a report chronicling every confirmed case of malicious food contamination in the world since

1950: 365 cases in all, plus 126 additional unconfirmed cases. What he found demonstrates the reality of terrorist food attacks.

It turns out 72% of the food poisonings occurred at the end of the food supply chain—at home—typically by a friend, relative, neighbor, or co-worker trying to kill or injure a specific person. A characteristic example is Heather Mook of York, who in 2007 tried to kill her husband by putting rat poison in his spaghetti.

Most of these cases resulted in fewer than five casualties—Mook only injured her husband in this incident—although 16% resulted in five or more. Of the 19 cases that claimed 10 or more lives, four involved serial killers operating over several years.

Another 23% of cases occurred at the retail or food service level. A 1998 incident in Japan, where someone put arsenic in a curry sold at a summer festival, killing four and hospitalizing 63, is a typical example. Only 11% of these incidents resulted in 100 or more casualties, while 44% resulted in none.

There are very few incidents of people contaminating the actual food supply. People deliberately contaminated a water supply seven times, resulting in three deaths. There is only one example of someone deliberately contaminating a crop before harvest—in Australia in 2006—and the crops were recalled before they could be sold. And in the three cases of someone deliberately contaminating food during packaging and distribution, including a 2005 case in the UK where glass and needles were baked into loaves of bread, no one died or was injured.

This isn't the stuff of bioterrorism. The closest example occurred in 1984 in the US, where members of a religious group known as the Rajneeshees contaminated several restaurant salad bars with *Salmonella enterica typhimurium*, sickening 751, hospitalizing 45, but killing no one. In fact, no one knew this was malicious until a year later, when one of the perpetrators admitted it.

Almost all of the food contaminations used conventional poisons such as cyanide, drain cleaner, mercury, or weed killer. There were nine incidents of biological agents, including salmonella, ricin, and fecal matter, and eight cases of radiological matter. The 2006 London poisoning of the former KGB agent Alexander Litvinenko with polonium-210 in his tea is an example of the latter.

And that assassination illustrates the real risk of malicious food poisonings. What is discussed in terrorist training manuals, and what the CIA is worried about, is the use of contaminated food in targeted assassinations. The quantities involved for mass poisonings are too great, the nature of the food

supply too vast and the details of any plot too complicated and unpredictable to be a real threat. That becomes crystal clear as you read the details of the different incidents: it's hard to kill one person, and very hard to kill dozens. Hundreds, thousands: it's just not going to happen any time soon. The fear of bioterror is much greater, and the panic from any bioterror scare will injure more people, than bioterrorism itself.

Far more dangerous are accidental contaminations due to negligent industry practices, such as the 2006 spinach E coli and, more recently, peanut salmonella contaminations in the US, the 2008 milk contaminations in China, and the BSE-infected beef from earlier this decade. And the systems we have in place to deal with these accidental contaminations also work to mitigate any intentional ones.

In 2004, the then US secretary of health and human services, Tommy Thompson, said on Fox News: "I cannot understand why terrorists have not attacked our food supply. Because it is so easy to do."

Guess what? It's not at all easy to do.

Raising the Cost of Paperwork Errors Will Improve Accuracy

Originally published in the Guardian, June 24, 2009

It's a sad, horrific story. Homeowner returns to find his house demolished. The demolition company was hired legitimately but there was a mistake and it demolished the wrong house. The demolition company relied on GPS coordinates, but requiring street addresses isn't a solution. A typo in the address is just as likely, and it would have demolished the house just as quickly. The problem is less how the demolishers knew which house to knock down, and more how they confirmed that knowledge. They trusted the paperwork, and the paperwork was wrong. Informality works when everybody knows everybody else. When merchants and customers know each other, government officials and citizens know each other, and people know their neighbors, people know what's going on. In that sort of milieu, if something goes wrong, people notice.

In our modern anonymous world, paperwork is how things get done. Traditionally, signatures, forms, and watermarks all made paperwork official.

Forgeries were possible but difficult. Today, there's still paperwork, but for the most part it only exists until the information makes its way into a computer database. Meanwhile, modern technology—computers, fax machines and desktop publishing software—has made it easy to forge paperwork. Every case of identity theft has, at its core, a paperwork failure. Fake work orders, purchase orders, and other documents are used to steal computers, equipment, and stock. Occasionally, fake faxes result in people being sprung from prison. Fake boarding passes can get you through airport security. This month hackers officially changed the name of a Swedish man.

A reporter even changed the ownership of the Empire State Building. Sure, it was a stunt, but this is a growing form of crime. Someone pretends to be you—preferably when you're away on holiday—and sells your home to someone else, forging your name on the paperwork. You return to find someone else living in your house, someone who thinks he legitimately bought it. In some senses, this isn't new. Paperwork mistakes and fraud have happened ever since there was paperwork. And the problem hasn't been fixed yet for several reasons.

One, our sloppy systems generally work fine, and it's how we get things done with minimum hassle. Most people's houses don't get demolished and most people's names don't get maliciously changed. As common as identity theft is, it doesn't happen to most of us. These stories are news because they are so rare. And in many cases, it's cheaper to pay for the occasional blunder than ensure it never happens.

Two, sometimes the incentives aren't in place for paperwork to be properly authenticated. The people who demolished that family home were just trying to get a job done. The same is true for government officials processing title and name changes. Banks get paid when money is transferred from one account to another, not when they find a paperwork problem. We're all irritated by forms stamped 17 times, and other mysterious bureaucratic processes, but these are actually designed to detect problems.

And three, there's a psychological mismatch: it is easy to fake paperwork, yet for the most part we act as if it has magical properties of authenticity. What's changed is scale. Fraud can be perpetrated against hundreds of thousands, automatically. Mistakes can affect that many people, too. What we need are laws that penalize people or companies—criminally or civilly—who make paperwork errors. This raises the cost of mistakes, making authenticating paperwork more attractive, which changes the incentives of those on the receiving end of the paperwork. And that will cause the market to devise

technologies to verify the provenance, accuracy, and integrity of information: telephone verification, addresses and GPS co-ordinates, cryptographic authentication, systems that double- and triple-check, and so on.

We can't reduce society's reliance on paperwork, and we can't eliminate errors based on it. But we can put economic incentives in place for people and companies to authenticate paperwork more.

So-Called Cyberattack Was Overblown —

Originally published in MPR News Q, July 13, 2009

To hear the media tell it, the United States suffered a major cyberattack last week. Stories were everywhere. “Cyber Blitz hits US, Korea” was the headline in Thursday's Wall Street Journal. North Korea was blamed.

Where were you when North Korea attacked America? Did you feel the fury of North Korea's armies? Were you fearful for your country? Or did your resolve strengthen, knowing that we would defend our homeland bravely and valiantly?

My guess is that you didn't even notice, that—if you didn't open a newspaper or read a news website—you had no idea anything was happening. Sure, a few government websites were knocked out, but that's not alarming or even uncommon. Other government websites were attacked but defended themselves, the sort of thing that happens all the time. If this is what an international cyberattack looks like, it hardly seems worth worrying about at all.

Politically motivated cyberattacks are nothing new. We've seen UK vs. Ireland. Israel vs. the Arab states. Russia vs. several former Soviet Republics. India vs. Pakistan, especially after the nuclear bomb tests in 1998. China vs. the United States, especially in 2001 when a US spy plane collided with a Chinese fighter jet. And so on and so on.

The big one happened in 2007, when the government of Estonia was attacked in cyberspace following a diplomatic incident with Russia about the relocation of a Soviet World War II memorial. The networks of many Estonian organizations, including the Estonian parliament, banks, ministries, newspapers and broadcasters, were attacked and—in many cases—shut down. Estonia was quick to blame Russia, which was equally quick to deny any involvement.

It was hyped as the first cyberwar, but after two years there is still no evidence that the Russian government was involved. Though Russian hackers

were indisputably the major instigators of the attack, the only individuals positively identified have been young ethnic Russians living inside Estonia, who were angry over the statue incident.

Poke at any of these international incidents, and what you find are kids playing politics. Last Wednesday, South Korea's National Intelligence Service admitted that it didn't actually know that North Korea was behind the attacks: "North Korea or North Korean sympathizers in the South" was what it said. Once again, it'll be kids playing politics.

This isn't to say that cyberattacks by governments aren't an issue, or that cyberwar is something to be ignored. The constant attacks by Chinese nationals against US networks may not be government-sponsored, but it's pretty clear that they're tacitly government-approved. Criminals, from lone hackers to organized crime syndicates, attack networks all the time. And war expands to fill every possible theater: land, sea, air, space, and now cyberspace. But cyberterrorism is nothing more than a media invention designed to scare people. And for there to be a cyberwar, there first needs to be a war.

Israel is currently considering attacking Iran in cyberspace, for example. If it tries, it'll discover that attacking computer networks is an inconvenience to the nuclear facilities it's targeting, but doesn't begin to substitute for bombing them.

In May, President Obama gave a major speech on cybersecurity. He was right when he said that cybersecurity is a national security issue, and that the government needs to step up and do more to prevent cyberattacks. But he couldn't resist hyping the threat with scare stories: "In one of the most serious cyber incidents to date against our military networks, several thousand computers were infected last year by malicious software—malware," he said. What he didn't add was that those infections occurred because the Air Force couldn't be bothered to keep its patches up to date.

This is the face of cyberwar: easily preventable attacks that, even when they succeed, only a few people notice. Even this current incident is turning out to be a sloppily modified five-year-old worm that no modern network should still be vulnerable to.

Securing our networks doesn't require some secret advanced NSA technology. It's the boring network security administration stuff we already know how to do: keep your patches up to date, install good anti-malware software, correctly configure your firewalls and intrusion-detection systems, monitor your networks. And while some government and corporate networks do a pretty good job at this, others fail again and again.

Enough of the hype and the bluster. The news isn't the attacks, but that some networks had security lousy enough to be vulnerable to them.

Why Framing Your Enemies Is Now Virtually Child's Play

Originally published in the Guardian, October 15, 2009

A few years ago, a company began to sell a liquid with identification codes suspended in it. The idea was that you would paint it on your stuff as proof of ownership. I commented that I would paint it on someone else's stuff, then call the police.

I was reminded of this recently when a group of Israeli scientists demonstrated that it's possible to fabricate DNA evidence. So now, instead of leaving your own DNA at a crime scene, you can leave fabricated DNA. And it isn't even necessary to fabricate. In Charlie Stross's novel *Halting State*, the bad guys foul a crime scene by blowing around the contents of a vacuum cleaner bag, containing the DNA of dozens, if not hundreds, of people.

This kind of thing has been going on forever. It's an arms race, and when technology changes, the balance between attacker and defender changes. But when automated systems do the detecting, the results are different. Face recognition software can be fooled by cosmetic surgery, or sometimes even just a photograph. And when fooling them becomes harder, the bad guys fool them on a different level. Computer-based detection gives the defender economies of scale, but the attacker can use those same economies of scale to defeat the detection system.

Google, for example, has anti-fraud systems that detect—and shut down—advertisers who try to inflate their revenue by repeatedly clicking on their own AdSense ads. So people built bots to repeatedly click on the AdSense ads of their competitors, trying to convince Google to kick them out of the system.

Similarly, when Google started penalizing a site's search engine rankings for having "bad neighbors"—backlinks from link farms, adult or gambling sites, or blog spam—people engaged in sabotage: they built link farms and left blog comment spam linking to their competitors' sites.

The same sort of thing is happening on Yahoo Answers. Initially, companies would leave answers pushing their products, but Yahoo started policing this.

So people have written bots to report abuse on all their competitors. There are Facebook bots doing the same sort of thing.

Last month, Google introduced Sidewiki, a browser feature that lets you read and post comments on virtually any webpage. People and industries are worried about the effects unrestrained commentary might have on their businesses, and how they might control the comments. I'm sure Google has sophisticated systems ready to detect commercial interests that try to take advantage of the system, but are they ready to deal with commercial interests that try to frame their competitors? And do we want to give one company the power to decide which comments should rise to the top and which get deleted?

Whenever you build a security system that relies on detection and identification, you invite the bad guys to subvert the system so it detects and identifies someone else. Sometimes this is hard—leaving someone else's fingerprints on a crime scene is hard, as is using a mask of someone else's face to fool a guard watching a security camera—and sometimes it's easy. But when automated systems are involved, it's often very easy. It's not just hardened criminals that try to frame each other, it's mainstream commercial interests.

With systems that police Internet comments and links, there's money involved in commercial messages—so you can be sure some will take advantage of it. This is the arms race. Build a detection system, and the bad guys try to frame someone else. Build a detection system to detect framing, and the bad guys try to frame someone else framing someone else. Build a detection system to detect framing of framing, and well, there's no end, really. Commercial speech is on the Internet to stay; we can only hope that they don't pollute the social systems we use so badly that they're no longer useful.

Beyond Security Theater

Originally published in New Internationalist, n. 427, November 2009, pp. 10–13

Terrorism is rare, far rarer than many people think. It's rare because very few people want to commit acts of terrorism, and executing a terrorist plot is much harder than television makes it appear. The best defenses against terrorism are largely invisible: investigation, intelligence, and emergency response. But even these are less effective at keeping us safe than our social and political

policies, both at home and abroad. However, our elected leaders don't think this way: they are far more likely to implement security theater against movie-plot threats.

A movie-plot threat is an overly specific attack scenario. Whether it's terrorists with crop dusters, terrorists contaminating the milk supply, or terrorists attacking the Olympics, specific stories affect our emotions more intensely than mere data does. Stories are what we fear. It's not just hypothetical stories: terrorists flying planes into buildings, terrorists with bombs in their shoes or in their water bottles, and terrorists with guns and bombs waging a coordinated attack against a city are even scarier movie-plot threats because they actually happened.

Security theater refers to security measures that make people feel more secure without doing anything to actually improve their security. An example: the photo ID checks that have sprung up in office buildings. No-one has ever explained why verifying that someone has a photo ID provides any actual security, but it looks like security to have a uniformed guard-for-hire looking at ID cards. Airport-security examples include the National Guard troops stationed at US airports in the months after 9/11—their guns had no bullets. The US color-coded system of threat levels, the pervasive harassment of photographers, and the metal detectors that are increasingly common in hotels and office buildings since the Mumbai terrorist attacks, are additional examples.

To be sure, reasonable arguments can be made that some terrorist targets are more attractive than others: airplanes because a small bomb can result in the death of everyone aboard, monuments because of their national significance, national events because of television coverage, and transportation because of the numbers of people who commute daily. But there are literally millions of potential targets in any large country (there are five million commercial buildings alone in the US), and hundreds of potential terrorist tactics; it's impossible to defend every place against everything, and it's impossible to predict which tactic and target terrorists will try next.

Feeling and Reality

Security is both a feeling and a reality. The propensity for security theater comes from the interplay between the public and its leaders. When people are scared, they need something done that will make them feel safe, even if it doesn't truly make them safer. Politicians naturally want to do something in response to crisis, even if that something doesn't make any sense.

Often, this “something” is directly related to the details of a recent event: we confiscate liquids, screen shoes, and ban box cutters on airplanes. But it’s not the target and tactics of the last attack that are important, but the next attack. These measures are only effective if we happen to guess what the next terrorists are planning. If we spend billions defending our rail systems, and the terrorists bomb a shopping mall instead, we’ve wasted our money. If we concentrate airport security on screening shoes and confiscating liquids, and the terrorists hide explosives in their brassieres and use solids, we’ve wasted our money. Terrorists don’t care what they blow up and it shouldn’t be our goal merely to force the terrorists to make a minor change in their tactics or targets.

Our penchant for movie plots blinds us to the broader threats. And security theater consumes resources that could better be spent elsewhere.

Any terrorist attack is a series of events: something like planning, recruiting, funding, practicing, executing, aftermath. Our most effective defenses are at the beginning and end of that process—intelligence, investigation, and emergency response—and least effective when they require us to guess the plot correctly. By intelligence and investigation, I don’t mean the broad data-mining or eavesdropping systems that have been proposed and in some cases implemented—those are also movie-plot stories without much basis in actual effectiveness—but instead the traditional “follow the evidence” type of investigation that has worked for decades.

Unfortunately for politicians, the security measures that work are largely invisible. Such measures include enhancing the intelligence-gathering abilities of the secret services, hiring cultural experts and Arabic translators, building bridges with Islamic communities both nationally and internationally, funding police capabilities—both investigative arms to prevent terrorist attacks, and emergency communications systems for after attacks occur—and arresting terrorist plotters without media fanfare. They do not include expansive new police or spying laws. Our police don’t need any new laws to deal with terrorism; rather, they need apolitical funding. These security measures don’t make good television, and they don’t help, come re-election time. But they work, addressing the reality of security instead of the feeling.

The arrest of the “liquid bombers” in London is an example: they were caught through old-fashioned intelligence and police work. Their choice of target (airplanes) and tactic (liquid explosives) didn’t matter; they would have been arrested regardless.

But even as we do all of this we cannot neglect the feeling of security, because it's how we collectively overcome the psychological damage that terrorism causes. It's not security theater we need, it's direct appeals to our feelings. The best way to help people feel secure is by acting secure around them. Instead of reacting to terrorism with fear, we—and our leaders—need to react with indomitability.

Refuse to Be Terrorized

By not overreacting, by not responding to movie-plot threats, and by not becoming defensive, we demonstrate the resilience of our society, in our laws, our culture, our freedoms. There is a difference between indomitability and arrogant “bring ‘em on” rhetoric. There's a difference between accepting the inherent risk that comes with a free and open society, and hyping the threats.

We should treat terrorists like common criminals and give them all the benefits of true and open justice—not merely because it demonstrates our indomitability, but because it makes us all safer. Once a society starts circumventing its own laws, the risks to its future stability are much greater than terrorism.

Supporting real security even though it's invisible, and demonstrating indomitability even though fear is more politically expedient, requires real courage. Demagoguery is easy. What we need is leaders willing both to do what's right and to speak the truth.

Despite fearful rhetoric to the contrary, terrorism is not a transcendent threat. A terrorist attack cannot possibly destroy a country's way of life; it's only our reaction to that attack that can do that kind of damage. The more we undermine our own laws, the more we convert our buildings into fortresses, the more we reduce the freedoms and liberties at the foundation of our societies, the more we're doing the terrorists' job for them.

We saw some of this in the Londoners' reaction to the 2005 transport bombings. Among the political and media hype and fearmongering, there was a thread of firm resolve. People didn't fall victim to fear. They rode the trains and buses the next day and continued their lives. Terrorism's goal isn't murder; terrorism attacks the mind, using victims as a prop. By refusing to be terrorized, we deny the terrorists their primary weapon: our own fear.

Today, we can project indomitability by rolling back all the fear-based post-9/11 security measures. Our leaders have lost credibility; getting it back requires a decrease in hyperbole. Ditch the invasive mass surveillance systems and new police state-like powers. Return airport security to pre-9/11 levels. Remove swagger from our foreign policies. Show the world that our legal system is up to the challenge of terrorism. Stop telling people to report all

suspicious activity; it does little but make us suspicious of each other, increasing both fear and helplessness.

Terrorism has always been rare, and for all we've heard about 9/11 changing the world, it's still rare. Even 9/11 failed to kill as many people as automobiles do in the US every single month. But there's a pervasive myth that terrorism is easy. It's easy to imagine terrorist plots, both large-scale "poison the food supply" and small-scale "10 guys with guns and cars." Movies and television bolster this myth, so many people are surprised that there have been so few attacks in Western cities since 9/11. Certainly intelligence and investigation successes have made it harder, but mostly it's because terrorist attacks are actually hard. It's hard to find willing recruits, to co-ordinate plans, and to execute those plans—and it's easy to make mistakes.

Counterterrorism is also hard, especially when we're psychologically prone to muck it up. Since 9/11, we've embarked on strategies of defending specific targets against specific tactics, overreacting to every terrorist video, stoking fear, demonizing ethnic groups, and treating the terrorists as if they were legitimate military opponents who could actually destroy a country or a way of life—all of this plays into the hands of terrorists. We'd do much better by leveraging the inherent strengths of our modern democracies and the natural advantages we have over the terrorists: our adaptability and survivability, our international network of laws and law enforcement, and the freedoms and liberties that make our society so enviable. The way we live is open enough to make terrorists rare; we are observant enough to prevent most of the terrorist plots that exist, and indomitable enough to survive the even fewer terrorist plots that actually succeed. We don't need to pretend otherwise.

Cold War Encryption Is Unrealistic in Today's Trenches

Originally published in the Japan Times, December 23, 2009

Sometimes mediocre encryption is better than strong encryption, and sometimes no encryption is better still.

The *Wall Street Journal* reported this week that Iraqi, and possibly also Afghan, militants are using commercial software to eavesdrop on US Predators, other unmanned aerial vehicles, or UAVs, and even piloted planes. The systems weren't "hacked"—the insurgents can't control them—but because the

downlink is unencrypted, they can watch the same video stream as the coalition troops on the ground.

The naive reaction is to ridicule the military. Encryption is so easy that HDTVs do it—just a software routine and you're done—and the Pentagon has known about this flaw since Bosnia in the 1990s. But encrypting the data is the easiest part; key management is the hard part. Each UAV needs to share a key with the ground station. These keys have to be produced, guarded, transported, used and then destroyed. And the equipment, both the Predators and the ground terminals, needs to be classified and controlled, and all the users need security clearance.

The command and control channel is, and always has been, encrypted—because that's both more important and easier to manage. UAVs are flown by air-men sitting at comfortable desks on US military bases, where key management is simpler. But the video feed is different. It needs to be available to all sorts of people, of varying nationalities and security clearances, on a variety of field terminals, in a variety of geographical areas, in all sorts of conditions—with everything constantly changing. Key management in this environment would be a nightmare.

Additionally, how valuable is this video downlink to the enemy? The primary fear seems to be that the militants watch the video, notice their compound being surveilled and flee before the missiles hit. Or notice a bunch of Marines walking through a recognizable area and attack them. This might make a great movie scene, but it's not very realistic. Without context, and just by peeking at random video streams, the risk caused by eavesdropping is low.

Contrast this with the additional risks if you encrypt: A soldier in the field doesn't have access to the real-time video because of a key management failure; a UAV can't be quickly deployed to a new area because the keys aren't in place; we can't share the video information with our allies because we can't give them the keys; most soldiers can't use this technology because they don't have the right clearances. Given this risk analysis, not encrypting the video is almost certainly the right decision.

There is another option, though. During the Cold War, the NSA's primary adversary was Soviet intelligence, and it developed its crypto solutions accordingly. Even though that level of security makes no sense in Bosnia, and

certainly not in Iraq and Afghanistan, it is what the NSA had to offer. If you encrypt, they said, you have to do it “right.”

The problem is, the world has changed. Today’s insurgent adversaries don’t have KGB-level intelligence gathering or cryptanalytic capabilities. At the same time, computer and network data gathering has become much cheaper and easier, so they have technical capabilities the Soviets could only dream of. Defending against these sorts of adversaries doesn’t require military-grade encryption only where it counts; it requires commercial-grade encryption everywhere possible.

This sort of solution would require the NSA to develop a whole new level of lightweight commercial-grade security systems for military applications—not just office-data “Sensitive but Unclassified” or “For Official Use Only” classifications. It would require the NSA to allow keys to be handed to uncleared UAV operators, and perhaps read over insecure phone lines and stored in people’s back pockets. It would require the sort of ad hoc key management systems you find in Internet protocols, or in DRM systems. It wouldn’t be anywhere near perfect, but it would be more commensurate with the actual threats.

And it would help defend against a completely different threat facing the Pentagon: The PR threat. Regardless of whether the people responsible made the right security decision when they rushed the Predator into production, or when they convinced themselves that local adversaries wouldn’t know how to exploit it, or when they forgot to update their Bosnia-era threat analysis to account for advances in technology, the story is now being played out in the press. The Pentagon is getting beaten up because it’s not protecting against the threat—because it’s easy to make a sound bite where the threat sounds really dire. And now it has to defend against the perceived threat to the troops, regardless of whether the defense actually protects the troops or not. Reminds me of the TSA, actually.

So the military is now committed to encrypting the video. . . eventually. The next generation Predators, called Reapers—Who names this stuff? Second-grade boys?—will have the same weakness. Maybe we’ll have encrypted video by 2010, or 2014, but I don’t think that’s even remotely possible unless the NSA relaxes its key management and classification requirements and embraces a lightweight, less secure encryption solution for these sorts of situations. The real failure here is the failure of the Cold War security model to deal with today’s threats.

Profiling Makes Us Less Safe

*Originally published in New York Times Room for Debate blog,
January 4, 2010*

There are two kinds of profiling. There's behavioral profiling based on how someone acts, and there's automatic profiling based on name, nationality, method of ticket purchase, and so on. The first one can be effective, but is very hard to do right. The second one makes us all less safe. The problem with automatic profiling is that it doesn't work.

Terrorists don't fit a profile and cannot be plucked out of crowds by computers. They're European, Asian, African, Hispanic, and Middle Eastern, male and female, young and old. Umar Farouk Abdul Mutallab was Nigerian. Richard Reid, the shoe bomber, was British with a Jamaican father. Germaine Lindsay, one of the 7/7 London bombers, was Afro-Caribbean. Dirty bomb suspect Jose Padilla was Hispanic-American. The 2002 Bali terrorists were Indonesian. Timothy McVeigh was a white American. So was the Unabomber. The Chechen terrorists who blew up two Russian planes in 2004 were female. Palestinian terrorists routinely recruit "clean" suicide bombers, and have used unsuspecting Westerners as bomb carriers.

Without an accurate profile, the system can be statistically demonstrated to be no more effective than random screening.

And, even worse, profiling creates two paths through security: one with less scrutiny and one with more. And once you do that, you invite the terrorists to take the path with less scrutiny. That is, a terrorist group can safely probe any profiling system and figure out how to beat the profile. And once they do, they're going to get through airport security with the minimum level of screening every time.

As counterintuitive as it may seem, we're all more secure when we randomly select people for secondary screening—even if it means occasionally screening wheelchair-bound grandmothers and innocent looking children. And, as an added bonus, it doesn't needlessly anger the ethnic groups we need on our side if we're going to be more secure against terrorism.

Fixing Intelligence Failures

Originally published in San Francisco Chronicle, January 15, 2010

President Obama in his speech last week rightly focused on fixing the intelligence failures that resulted in Umar Farouk Abdulmutallab being ignored, rather than on technologies targeted at the details of his underwear-bomb plot. But while Obama's instincts are right, reforming intelligence for this new century and its new threats is a more difficult task than he might like.

We don't need new technologies, new laws, new bureaucratic overlords, or—for heaven's sake—new agencies. What prevents information sharing among intelligence organizations is the culture of the generation that built those organizations.

The US intelligence system is a sprawling apparatus, spanning the FBI and the State Department, the CIA and the National Security Agency, and the Department of Homeland Security—itsself an amalgamation of two dozen different organizations—designed and optimized to fight the Cold War. The single, enormous adversary then was the Soviet Union: as bureaucratic as they come, with a huge budget, and capable of very sophisticated espionage operations. We needed to defend against technologically advanced electronic eavesdropping operations, their agents trying to bribe or seduce our agents, and a worldwide intelligence gathering capability that hung on our every word.

In that environment, secrecy was paramount. Information had to be protected by armed guards and double fences, shared only among those with appropriate security clearances and a legitimate “need to know,” and it was better not to transmit information at all than to transmit it insecurely.

Today's adversaries are different. There are still governments, like China, that are after our secrets. But the secrets they're after are more often corporate than military, and most of the other organizations of interest are like al Qaeda: decentralized, poorly funded and incapable of the intricate spy versus spy operations the Soviet Union could pull off.

Against these adversaries, sharing is far more important than secrecy. Our intelligence organizations need to trade techniques and expertise with

industry, and they need to share information among the different parts of themselves. Today's terrorist plots are loosely organized ad hoc affairs, and those dots that are so important for us to connect beforehand might be on different desks, in different buildings, owned by different organizations.

Critics have pointed to laws that prohibited inter-agency sharing but, as the 9/11 commission found, the law allows for far more sharing than goes on. It doesn't happen because of inter-agency rivalries, a reliance on outdated information systems, and a culture of secrecy. What we need is an intelligence community that shares ideas and hunches and facts on their versions of Facebook, Twitter and wikis. We need the bottom-up organization that has made the Internet the greatest collection of human knowledge and ideas ever assembled.

The problem is far more social than technological. Teaching your mom to "text" and your dad to Twitter doesn't make them part of the Internet generation, and giving all those cold warriors blogging lessons won't change their mentality—or the culture. The reason this continues to be a problem, the reason President George W. Bush couldn't change things even after the 9/11 commission came to much the same conclusions as President Obama's recent review did, is generational.

The Internet is the greatest generation gap since rock "n" roll, and it's just as true inside government as out. We might have to wait for the elders inside these agencies to retire and be replaced by people who grew up with the Internet.

Spy Cameras Won't Make Us Safer

Originally published in CNN, February 25, 2010

On January 19, a team of at least 15 people assassinated Hamas leader Mahmoud al-Mabhouh. The Dubai police released video footage of 11 of them. While it was obviously a very professional operation, the 27 minutes of video is fascinating in its banality. Team members walk through the airport, check in and out of hotels, get in and out of taxis. They make no effort to hide themselves from the cameras, sometimes seeming to stare directly into them. They obviously don't care that they're being recorded, and—in fact—the cameras didn't prevent the assassination, nor as far as we know have they helped as yet in identifying the killers.

Pervasive security cameras don't substantially reduce crime. This fact has been demonstrated repeatedly: in San Francisco public housing, in a New York apartment complex, in Philadelphia, in Washington, DC, in study after study in both the US and the UK. Nor are they instrumental in solving many crimes after the fact.

There are exceptions, of course, and proponents of cameras can always cherry-pick examples to bolster their argument. These success stories are what convince us; our brains are wired to respond more strongly to anecdotes than to data. But the data is clear: CCTV cameras have minimal value in the fight against crime.

While it's comforting to imagine vigilant police monitoring every camera, the truth is very different, for a variety of reasons: technological limitations of cameras, organizational limitations of police, and the adaptive abilities of criminals. No one looks at most CCTV footage until well after a crime is committed. And when the police do look at the recordings, it's very common for them to be unable to identify suspects. Criminals don't often stare helpfully at the lens, and—unlike the Dubai assassins—tend to wear sunglasses and hats. Cameras break far too often. Even when they afford quick identification—think of the footage of the 9/11 terrorists going through airport security, or the 7/7 London transport bombers just before the bombs exploded—police are often able to identify those suspects even without the cameras. Cameras afford a false sense of security, encouraging laziness when we need police to be vigilant.

The solution isn't for police to watch the cameras more diligently. Unlike an officer walking the street, cameras only look in particular directions at particular locations. Criminals know this, and can easily adapt by moving their crimes to places not watched by a camera—and there will always be such places. And while a police officer on the street can respond to a crime in progress, someone watching a CCTV screen can only dispatch an officer to arrive much later. By their very nature, cameras result in underused and misallocated police resources.

Cameras aren't completely ineffective, of course. Used properly, they're effective in reducing crime in enclosed areas with minimal foot traffic. Combined with adequate lighting, they substantially reduce both personal attacks and auto-related crime in multi-story parking garages. And sometimes it is cost-effective for a store to install cameras to catch shoplifters, or a casino to install

cameras to detect cheaters. But these are instances where there is a specific risk at a specific location.

But the important question isn't whether cameras solve past crime or deter future crime; it's whether they're a good use of resources. They're expensive, both in money and their Orwellian effects on privacy and civil liberties. Their inevitable misuse is another cost: police have already spied on naked women in their own homes, shared nude images, sold best-of videos, and spied on national politicians. While we might be willing to accept these downsides for a real increase in security, cameras don't provide that. Despite our predilection for preferring technological solutions over human ones, the funds now spent on CCTV cameras would be far better spent on hiring and training police officers.

We live in a unique time in our society: cameras are everywhere, but we can still see them. Ten years ago, cameras were much rarer than they are today. Ten years from now, they'll be so small you won't even notice them. Already, people can buy surveillance cameras in household objects to spy on their spouses and baby sitters—I particularly like the one hidden in a shower mirror—cameras in pens to spy on their colleagues, and remotely turn on laptop cameras to spy on anyone. Companies are developing police state-type CCTV surveillance technologies for China, technology that will find its way into countries like the US

If universal surveillance were the answer, lots of us would have moved to former East Germany. If surveillance cameras were the answer, camera-happy London, with something like 500,000 of them at a cost of \$700 million, would be the safest city on the planet. We didn't and it isn't, because surveillance and surveillance cameras don't make us safer. The money spent on cameras in London, and in cities across America, could be much better spent on actual policing.

Scanners, Sensors Are Wrong Way to Secure the Subway

Originally published in Daily News, April 7, 2010

People intent on preventing a Moscow-style terrorist attack against the New York subway system are proposing a range of expensive new underground security measures, some temporary and some permanent.

They should save their money—and instead invest every penny they’re considering pouring into new technologies into intelligence and old-fashioned policing.

Intensifying security at specific stations only works against terrorists who aren’t smart enough to move to another station. Cameras are useful only if all the stars align: The terrorists happen to walk into the frame, the video feeds are being watched in real time and the police can respond quickly enough to be effective. They’re much more useful *after* an attack, to figure out who pulled it off.

Installing biological and chemical detectors requires similarly implausible luck—plus a terrorist plot that includes the specific biological or chemical agent that is being detected.

What all these misguided reactions have in common is that they’re based on “movie-plot threats”: overly specific attack scenarios. They fill our imagination vividly, in full color with rich detail. Before long, we’re envisioning an entire story line, with or without Bruce Willis saving the day. And we’re scared.

It’s not that movie-plot threats are not worth worrying about. It’s that each one—Moscow’s subway attack, the bombing of the Oklahoma City federal building, etc.—is too specific. These threats are infinite, and the bad guys can easily switch among them.

New York has thousands of possible targets, and there are dozens of possible tactics. Implementing security against movie-plot threats is only effective if we correctly guess which specific threat to protect against. That’s unlikely.

A far better strategy is to spend our limited counterterrorism resources on investigation and intelligence—and on emergency response. These measures don’t hinge on any specific threat; they don’t require us to guess the tactic or target correctly. They’re effective in a variety of circumstances, even nonterrorist ones.

The result may not be flashy or outwardly reassuring—as are pricey new scanners in airports. But the strategy will save more lives.

The 2006 arrest of the liquid bombers—who wanted to detonate liquid explosives to be brought onboard airliners traveling from England to North America—serves as an excellent example. The plotters were arrested in their London apartments, and their attack was foiled before they ever got to the airport.

It didn’t matter if they were using liquids or solids or gases. It didn’t even matter if they were targeting airports or shopping malls or theaters. It was a straightforward, although hardly simple, matter of following leads.

Gimmicky security measures are tempting—but they’re distractions we can’t afford. The Christmas Day bomber chose his tactic because it would circumvent last year’s security measures, and the next attacker will choose his tactic—and target—according to similar criteria. Spend money on cameras and guards in the subways, and the terrorists will simply modify their plot to render those countermeasures ineffective.

Humans are a species of storytellers, and the Moscow story has obvious parallels in New York. When we read the word “subway,” we can’t help but think about the system we use every day. This is a natural response, but it doesn’t make for good public policy. We’d all be safer if we rose above the simple parallels and the need to calm our fears with expensive and seductive new technologies—and countered the threat the smart way.

Preventing Terrorist Attacks in Crowded Areas

*Originally published in New York Times Room for Debate blog,
May 3, 2010*

In the wake of Saturday’s failed Times Square car bombing, it’s natural to ask how we can prevent this sort of thing from happening again. The answer is stop focusing on the specifics of what actually happened, and instead think about the threat in general.

Think about the security measures commonly proposed. Cameras won’t help. They don’t prevent terrorist attacks, and their forensic value after the fact is minimal. In the Times Square case, surely there’s enough other evidence—the car’s identification number, the auto body shop the stolen license plates came from, the name of the fertilizer store—to identify the guy. We will almost certainly not need the camera footage. The images released so far, like the images in so many other terrorist attacks, may make for exciting television, but their value to law enforcement officers is limited.

Check points won’t help, either. You can’t check everybody and everything. There are too many people to check, and too many train stations, buses, theaters, department stores and other places where people congregate. Patrolling guards, bomb-sniffing dogs, chemical and biological weapons detectors: they all suffer from similar problems. In general, focusing on specific tactics or

defending specific targets doesn't make sense. They're inflexible; possibly effective if you guess the plot correctly, but completely ineffective if you don't. At best, the countermeasures just force the terrorists to make minor changes in their tactic and target.

It's much smarter to spend our limited counterterrorism resources on measures that don't focus on the specific. It's more efficient to spend money on investigating and stopping terrorist attacks before they happen, and responding effectively to any that occur. This approach works because it's flexible and adaptive; it's effective regardless of what the bad guys are planning for next time.

After the Christmas Day airplane bombing attempt, I was asked how we can better protect our airplanes from terrorist attacks. I pointed out that the event was a security success—the plane landed safely, nobody was hurt, a terrorist was in custody—and that the next attack would probably have nothing to do with explosive underwear. After the Moscow subway bombing, I wrote that overly specific security countermeasures like subway cameras and sensors were a waste of money.

Now we have a failed car bombing in Times Square. We can't protect against the next imagined movie-plot threat. Isn't it time to recognize that the bad guys are flexible and adaptive, and that we need the same quality in our countermeasures?

Where Are All the Terrorist Attacks?

Originally published in AOL News, May 4, 2010

As the details of the Times Square car bomb attempt emerge in the wake of Faisal Shahzad's arrest Monday night, one thing has already been made clear: Terrorism is fairly easy. All you need is a gun or a bomb, and a crowded target. Guns are easy to buy. Bombs are easy to make. Crowded targets—not only in New York, but all over the country—are easy to come by. If you're willing to die in the aftermath of your attack, you could launch a pretty effective terrorist attack with a few days of planning, maybe less.

But if it's so easy, why aren't there more terrorist attacks like the failed car bomb in New York's Times Square? Or the terrorist shootings in Mumbai? Or the Moscow subway bombings? After the enormous horror and tragedy of 9/11, why have the past eight years been so safe in the US?

There are actually several answers to this question. One, terrorist attacks are harder to pull off than popular imagination—and the movies—lead everyone to believe. Two, there are far fewer terrorists than the political rhetoric of the past eight years leads everyone to believe. And three, random minor terrorist attacks don't serve Islamic terrorists' interests right now.

Hard to Pull Off

Terrorism sounds easy, but the actual attack is the easiest part.

Putting together the people, the plot and the materials is hard. It's hard to sneak terrorists into the US. It's hard to grow your own inside the US. It's hard to operate; the general population, even the Muslim population, is against you.

Movies and television make terrorist plots look easier than they are. It's hard to hold conspiracies together. It's easy to make a mistake. Even 9/11, which was planned before the climate of fear that event engendered, just barely succeeded. Today, it's much harder to pull something like that off without slipping up and getting arrested.

Few Terrorists

But even more important than the difficulty of executing a terrorist attack, there aren't a lot of terrorists out there.

Al-Qaida isn't a well-organized global organization with movie-plot-villain capabilities; it's a loose collection of people using the same name. Despite the post-9/11 rhetoric, there isn't a terrorist cell in every major city. If you think about the major terrorist plots we've foiled in the US—the JFK bombers, the Fort Dix plotters—they were mostly amateur terrorist wannabes with no connection to any sort of al-Qaida central command, and mostly no ability to effectively carry out the attacks they planned.

The successful terrorist attacks—the Fort Hood shooter, the guy who flew his plane into the Austin IRS office, the anthrax mailer—were largely nut cases operating alone. Even the unsuccessful shoe bomber, and the equally unsuccessful Christmas Day underwear bomber, had minimal organized help—and that help originated outside the US.

Terrorism doesn't occur without terrorists, and they are far rarer than popular opinion would have it.

Small Attacks Aren't Enough

Lastly, and perhaps most subtly, there's not a lot of value in unspectacular terrorism anymore.

If you think about it, terrorism is essentially a PR stunt. The death of innocents and the destruction of property isn't the goal of terrorism; it's just the tactic used. And acts of terrorism are intended for two audiences: for the victims, who are supposed to be terrorized as a result, and for the allies and potential allies of the terrorists, who are supposed to give them more funding and generally support their efforts.

An act of terrorism that doesn't instill terror in the target population is a failure, even if people die. And an act of terrorism that doesn't impress the terrorists' allies is not very effective, either.

Fortunately for us and unfortunately for the terrorists, 9/11 upped the stakes. It's no longer enough to blow up something like the Oklahoma City Federal Building. Terrorists need to blow up airplanes or the Brooklyn Bridge or the Sears Tower or JFK airport—something big to impress the folks back home. Small no-name targets just don't cut it anymore.

Note that this is very different than terrorism by an occupied population: the IRA in Northern Ireland, Iraqis in Iraq, Palestinians in Israel. Setting aside the actual politics, all of these terrorists believe they are repelling foreign invaders. That's not the situation here in the US

So, to sum up: If you're just a loner wannabe who wants to go out with a bang, terrorism is easy. You're more likely to get caught if you take a long time to plan or involve a bunch of people, but you might succeed. If you're a representative of al-Qaida trying to make a statement in the US, it's much harder. You just don't have the people, and you're probably going to slip up and get caught.

Worst-Case Thinking Makes Us Nuts, Not Safe

Originally published in CNN, May 12, 2010

At a security conference recently, the moderator asked the panel of distinguished cybersecurity leaders what their nightmare scenario was. The answers were the predictable array of large-scale attacks: against our communications

infrastructure, against the power grid, against the financial system, in combination with a physical attack.

I didn't get to give my answer until the afternoon, which was: "My nightmare scenario is that people keep talking about their nightmare scenarios."

There's a certain blindness that comes from worst-case thinking. An extension of the precautionary principle, it involves imagining the worst possible outcome and then acting as if it were a certainty. It substitutes imagination for thinking, speculation for risk analysis and fear for reason. It fosters powerlessness and vulnerability and magnifies social paralysis. And it makes us more vulnerable to the effects of terrorism.

Worst-case thinking means generally bad decision making for several reasons. First, it's only half of the cost-benefit equation. Every decision has costs and benefits, risks and rewards. By speculating about what can possibly go wrong, and then acting as if that is likely to happen, worst-case thinking focuses only on the extreme but improbable risks and does a poor job at assessing outcomes.

Second, it's based on flawed logic. It begs the question by assuming that a proponent of an action must prove that the nightmare scenario is impossible.

Third, it can be used to support any position or its opposite. If we build a nuclear power plant, it could melt down. If we don't build it, we will run short of power and society will collapse into anarchy. If we allow flights near Iceland's volcanic ash, planes will crash and people will die. If we don't, organs won't arrive in time for transplant operations and people will die. If we don't invade Iraq, Saddam Hussein might use the nuclear weapons he might have. If we do, we might destabilize the Middle East, leading to widespread violence and death.

Of course, not all fears are equal. Those that we tend to exaggerate are more easily justified by worst-case thinking. So terrorism fears trump privacy fears, and almost everything else; technology is hard to understand and therefore scary; nuclear weapons are worse than conventional weapons; our children need to be protected at all costs; and annihilating the planet is bad. Basically, any fear that would make a good movie plot is amenable to worst-case thinking.

Fourth and finally, worst-case thinking validates ignorance. Instead of focusing on what we know, it focuses on what we don't know—and what we can imagine.

Remember Defense Secretary Donald Rumsfeld's quote? "Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things

we do not know. But there are also unknown unknowns—the ones we don't know we don't know." And this: "the absence of evidence is not evidence of absence." Ignorance isn't a cause for doubt; when you can fill that ignorance with imagination, it can be a call to action.

Even worse, it can lead to hasty and dangerous acts. You can't wait for a smoking gun, so you act as if the gun is about to go off. Rather than making us safer, worst-case thinking has the potential to cause dangerous escalation.

The new undercurrent in this is that our society no longer has the ability to calculate probabilities. Risk assessment is devalued. Probabilistic thinking is repudiated in favor of "possibilistic thinking": Since we can't know what's likely to go wrong, let's speculate about what can possibly go wrong.

Worst-case thinking leads to bad decisions, bad systems design, and bad security. And we all have direct experience with its effects: airline security and the TSA, which we make fun of when we're not appalled that they're harassing 93-year-old women or keeping first-graders off airplanes. You can't be too careful!

Actually, you can. You can refuse to fly because of the possibility of plane crashes. You can lock your children in the house because of the possibility of child predators. You can eschew all contact with people because of the possibility of hurt. Steven Hawking wants to avoid trying to communicate with aliens because they might be hostile; does he want to turn off all the planet's television broadcasts because they're radiating into space? It isn't hard to parody worst-case thinking, and at its extreme it's a psychological condition.

Frank Furedi, a sociology professor at the University of Kent, writes: "Worst-case thinking encourages society to adopt fear as one of the dominant principles around which the public, the government and institutions should organize their life. It institutionalizes insecurity and fosters a mood of confusion and powerlessness. Through popularizing the belief that worst cases are normal, it incites people to feel defenseless and vulnerable to a wide range of future threats."

Even worse, it plays directly into the hands of terrorists, creating a population that is easily terrorized—even by failed terrorist attacks like the Christmas Day underwear bomber and the Times Square SUV bomber.

When someone is proposing a change, the onus should be on them to justify it over the status quo. But worst case thinking is a way of looking at the world that exaggerates the rare and unusual and gives the rare much more credence than it deserves.

It isn't really a principle; it's a cheap trick to justify what you already believe. It lets lazy or biased people make what seem to be cogent arguments without understanding the whole issue. And when people don't need to refute counterarguments, there's no point in listening to them.

Threat of "Cyberwar" Has Been Hugely Hyped

Originally published in CNN, July 7, 2010

There's a power struggle going on in the US government right now.

It's about who is in charge of cyber security, and how much control the government will exert over civilian networks. And by beating the drums of war, the military is coming out on top.

"The United States is fighting a cyberwar today, and we are losing," said former NSA director—and current cyberwar contractor—Mike McConnell. "Cyber 9/11 has happened over the last ten years, but it happened slowly so we don't see it," said former National Cyber Security Division director Amit Yoran. Richard Clarke, whom Yoran replaced, wrote an entire book hyping the threat of cyberwar.

General Keith Alexander, the current commander of the US Cyber Command, hypes it every chance he gets. This isn't just rhetoric of a few over-eager government officials and headline writers; the entire national debate on cyberwar is plagued with exaggerations and hyperbole.

Googling those names and terms—as well as "cyber Pearl Harbor," "cyber Katrina," and even "cyber Armageddon"—gives some idea how pervasive these memes are. Prefix "cyber" to something scary, and you end up with something really scary.

Cyberspace has all sorts of threats, day in and day out. Cybercrime is by far the largest: fraud, through identity theft and other means, extortion, and so on. Cyber-espionage is another, both government- and corporate-sponsored. Traditional hacking, without a profit motive, is still a threat. So is cyber-activism: people, most often kids, playing politics by attacking government and corporate websites and networks.

These threats cover a wide variety of perpetrators, motivations, tactics, and goals. You can see this variety in what the media has mislabeled as "cyberwar." The attacks against Estonian websites in 2007 were simple hacking attacks by

ethnic Russians angry at anti-Russian policies; these were denial-of-service attacks, a normal risk in cyberspace and hardly unprecedented.

A real-world comparison might be if an army invaded a country, then all got in line in front of people at the DMV so they couldn't renew their licenses. If that's what war looks like in the 21st century, we have little to fear.

Similar attacks against Georgia, which accompanied an actual Russian invasion, were also probably the responsibility of citizen activists or organized crime. A series of power blackouts in Brazil was caused by criminal extortionists—or was it sooty insulators? China is engaging in espionage, not war, in cyberspace. And so on.

One problem is that there's no clear definition of "cyberwar." What does it look like? How does it start? When is it over? Even cybersecurity experts don't know the answers to these questions, and it's dangerous to broadly apply the term "war" unless we know a war is going on.

Yet recent news articles have claimed that China declared cyberwar on Google, that Germany attacked China, and that a group of young hackers declared cyberwar on Australia. (Yes, cyberwar is so easy that even kids can do it.) Clearly we're not talking about real war here, but a rhetorical war: like the war on terror.

We have a variety of institutions that can defend us when attacked: the police, the military, the Department of Homeland Security, various commercial products and services, and our own personal or corporate lawyers. The legal framework for any particular attack depends on two things: the attacker and the motive. Those are precisely the two things you don't know when you're being attacked on the Internet. We saw this on July 4 last year, when US and South Korean websites were attacked by unknown perpetrators from North Korea—or perhaps England. Or was it Florida?

We surely need to improve our cybersecurity. But words have meaning, and metaphors matter. There's a power struggle going on for control of our nation's cybersecurity strategy, and the NSA and DoD are winning. If we frame the debate in terms of war, if we accept the military's expansive cyberspace definition of "war," we feed our fears.

We reinforce the notion that we're helpless—what person or organization can defend itself in a war?—and others need to protect us. We invite the military to take over security, and to ignore the limits on power that often get jettisoned during wartime.

If, on the other hand, we use the more measured language of cybercrime, we change the debate. Crime fighting requires both resolve and resources, but it's done within the context of normal life. We willingly give our police

extraordinary powers of investigation and arrest, but we temper these powers with a judicial system and legal protections for citizens.

We need to be prepared for war, and a Cyber Command is just as vital as an Army or a Strategic Air Command. And because kid hackers and cyber-warriors use the same tactics, the defenses we build against crime and espionage will also protect us from more concerted attacks. But we're not fighting a cyberwar now, and the risks of a cyberwar are no greater than the risks of a ground invasion. We need peacetime cyber-security, administered within the myriad structure of public and private security institutions we already have.

Cyberwar and the Future of Cyber Conflict

Originally published in Financial Times, December 2, 2010

The world is gearing up for cyberwar. The US Cyber Command became operational in November. NATO has enshrined cyber security among its new strategic priorities. The head of Britain's armed forces said recently that boosting cyber capability is now a huge priority for the UK. And we know China is already engaged in broad cyber espionage attacks against the west. So how can we control a burgeoning cyber arms race?

We may already have seen early versions of cyberwars in Estonia and Georgia, possibly perpetrated by Russia. It's hard to know for certain, not only because such attacks are often impossible to trace, but because we have no clear definitions of what a cyberwar actually is.

Does the 2007 attacks against Estonia, traced to a young Russian man living in Tallinn and no one else, count? What about a virus from an unknown origin, possibly targeted at an Iranian nuclear complex? Or espionage from within China, but not specifically directed by its government? To such questions one must add even more basic issues, like when a cyberwar is understood to have begun, and how it ends. When even cyber security experts can't answer these questions, it's hard to expect much from policymakers.

We can set parameters. It is obviously not an act of war just to develop digital weapons targeting another country. Using cyberattacks to spy on another nation is a grey area, which gets greyer still when a country penetrates information networks, just to see if it can do so. Penetrating such networks and

leaving a backdoor open, or even leaving logic bombs behind to be used later, is a harder case—yet the US and China are doing this to each other right now.

And what about when one country deliberately damages the economy of another, as one of the WikiLeaks cables shows that a member of China's politburo did against Google in January 2010? Definitions and rules are hard not just because the tools of war have changed, but because cyberspace puts them into the hands of a broader group of people. Previously only the military had weapons. Now anyone with sufficient computer skills can take matters into their own hands.

There are more basic problems too. When a nation is attacked in a regular conflict, a variety of military and civil institutions respond. The legal framework for this depends on two things: the attacker and the motive. But when you're attacked on the Internet, those are precisely the two things you don't know. We don't know if Georgia was attacked by the Russian government, or just some hackers living in Russia. In spite of much speculation, we don't know the origin, or target, of Stuxnet. We don't even know if last July 4's attacks against US and South Korean computers originated in North Korea, China, England, or Florida.

When you don't know, it's easy to get it wrong; and to retaliate against the wrong target, or for the wrong reason. That means it is easy for things to get out of hand. So while it is legitimate for nations to build offensive and defensive cyberwar capabilities we also need to think now about what can be done to limit the risk of cyberwar.

A first step would be a hotline between the world's cyber commands, modeled after similar hotlines among nuclear commands. This would at least allow governments to talk to each other, rather than guess where an attack came from. More difficult, but more important, are new cyberwar treaties. These could stipulate a no first use policy, outlaw unaimed weapons, or mandate weapons that self-destruct at the end of hostilities. The Geneva Conventions need to be updated too.

Cyber weapons beg to be used, so limits on stockpiles, and restrictions on tactics, are a logical end point. International banking, for instance, could be declared off-limits. Whatever the specifics, such agreements are badly needed. Enforcement will be difficult, but that's not a reason not to try. It's not too late to reverse the cyber arms race currently under way. Otherwise, it is only a matter of time before something big happens: perhaps by the rash actions of a low level military officer, perhaps by a non-state actor, perhaps by accident. And if the target nation retaliates, we could actually find ourselves in a cyberwar.

Why Terror Alert Codes Never Made Sense

Originally published in CNN, January 28, 2011

The Department of Homeland Security is getting rid of the color-coded threat level system. It was introduced after 9/11, and was supposed to tell you how likely a terrorist attack might be. Except that it never did.

Attacks happened more often when the level was yellow (“significant risk”) than when it was orange (“high risk”). And the one time it was red (“severe risk”), nothing happened. It’s never been blue or green, the two least dangerous levels.

The system has been at yellow for the past four years, and before then the changes seemed more timed to political events than actual terrorist threats. Not that any of this matters. We all ignored the levels because they didn’t tell us anything useful.

The problem is that the color-coded threat levels were vague and long-term, and didn’t correspond to useful actions people can take. Compare the color-coded threat levels with the US military’s DEFCON system. At each DEFCON level, there are specific actions people have to take: The real details of those actions are secret, but at DEFCON 3—I’m making this up—you might call everyone back from leave, at another you fuel all the bombers, at another you arm the bombs, and so on.

Knowing the current DEFCON level is important for those in the military, because it dictates what actions you should be taking. What am I supposed to do when the terrorist threat level is yellow? Or orange? I have no idea.

And no one else did, either. Were there plane trips you delayed when the level was orange that you made when it was yellow? Did any company base business decisions on it? Do we think the president consulted the level every morning?

Consider hurricane warnings. Hurricanes are short-term events, and it’s obvious when the danger is imminent and when it’s over.

People can do useful things in response to a hurricane warning—board up their windows, stay in the basement, evacuate—and there is a discrete period when their lives are markedly different; there is utility in the higher alert mode, even if nothing comes of it.

Compare people's reactions to hurricane threats with their reactions to earthquake threats.

According to scientists, California could experience a huge earthquake sometime in the next 200 years. Even though the magnitude of the disaster will be enormous, people can't stay alert for two centuries.

The news seems to have generated the same levels of short-term fear and long-term apathy in Californians that the terrorist threat level system does. It's human nature; people simply can't be vigilant indefinitely, and if a stimulus is constant, people learn to ignore it.

Any alert system that instills a vague feeling of dread or panic, without giving people anything to do in response, is ineffective. And a terrorist threat level that warns of a "significant risk of terrorist attacks"—that's what yellow is—for years is ignored.

The DHS could have lowered the level to something more reasonable, but that would have been politically impossible. If there was a terrorist attack and the threat level had been blue or green, the DHS would have been blamed for not warning us. Keeping the level high might increase the general dread among some people and cause sniggering among others, but it helps protect the jobs of those charged with keeping us safe from terrorism.

So good riddance to the color-coded terrorist alert system, that relic of our immediate post-9/11 panic. It never did serve any security purpose. And it will be much easier for us to accept that we are, and have been for a while, at a "low risk of terrorist attacks"—threat level green—when no politician has to risk his job by formally admitting it.

Debate Club: An International Cyberwar Treaty Is the Only Way to Stem the Threat

Originally published in US News, June 8, 2012

We're in the early years of a cyberwar arms race. It's expensive, it's destabilizing, and it threatens the very fabric of the Internet we use every day. Cyberwar treaties, as imperfect as they might be, are the only way to contain the threat.

If you read the press and listen to government leaders, we're already in the middle of a cyberwar. By any normal definition of the word "war," this is ridiculous.

But the definition of cyberwar has been expanded to include government-sponsored espionage, potential terrorist attacks in cyberspace, large-scale criminal fraud, and even hacker kids attacking government networks and critical infrastructure. This definition is being pushed both by the military and by government contractors, who are gaining power and making money on cyberwar fear.

The danger is that military problems beg for military solutions. We're starting to see a power grab in cyberspace by the world's militaries: large-scale monitoring of networks, military control of Internet standards, even military takeover of cyberspace. Last year's debate over an "Internet kill switch" is an example of this; it's the sort of measure that might be deployed in wartime but makes no sense in peacetime. At the same time, countries are engaging in offensive actions in cyberspace, with tools like Stuxnet and Flame.

Arms races stem from ignorance and fear: ignorance of the other side's capabilities, and fear that their capabilities are greater than yours. Once cyberweapons exist, there will be an impetus to use them. Both Stuxnet and Flame damaged networks other than their intended targets. Any military-inserted backdoors in Internet systems make us more vulnerable to criminals and hackers. And it is only a matter of time before something big happens, perhaps by the rash actions of a low-level military officer, perhaps by a non-state actor, perhaps by accident. And if the target nation retaliates, we could find ourselves in a real cyberwar.

The cyberwar arms race is destabilizing.

International cooperation and treaties are the only way to reverse this. Banning cyberweapons entirely is a good goal, but almost certainly unachievable. More likely are treaties that stipulate a no-first-use policy, outlaw unaimed or broadly targeted weapons, and mandate weapons that self-destruct at the end of hostilities. Treaties that restrict tactics and limit stockpiles could be a next step. We could prohibit cyberattacks against civilian infrastructure; international banking, for example, could be declared off-limits.

Yes, enforcement will be difficult. Remember how easy it was to hide a chemical weapons facility? Hiding a cyberweapons facility will be even easier. But we've learned a lot from our Cold War experience in negotiating nuclear, chemical, and biological treaties. The very act of negotiating limits the arms race and paves the way to peace. And even if they're breached, the world is safer because the treaties exist.

There's a common belief within the US military that cyberweapons treaties are not in our best interest: that we currently have a military advantage in cyberspace that we should not squander. That's not true. We might have an

offensive advantage—although that’s debatable—but we certainly don’t have a defensive advantage. More importantly, as a heavily networked country, we are inherently vulnerable in cyberspace.

Cyberspace threats are real. Military threats might get the publicity, but the criminal threats are both more dangerous and more damaging. Militarizing cyberspace will do more harm than good. The value of a free and open Internet is enormous.

Stop cyberwar fear mongering. Ratchet down cyberspace saber rattling. Start negotiations on limiting the militarization of cyberspace and increasing international police cooperation. This won’t magically make us safe, but it will make us safer.

Overreaction and Overly Specific Reactions to Rare Risks

Originally published in CNN, July 31, 2012

Horrific events, such as the massacre in Aurora, can be catalysts for social and political change. Sometimes it seems that they’re the only catalyst; recall how drastically our policies toward terrorism changed after 9/11 despite how moribund they were before.

The problem is that fear can cloud our reasoning, causing us to overreact and to overly focus on the specifics. And the key is to steer our desire for change in that time of fear.

Our brains aren’t very good at probability and risk analysis. We tend to exaggerate spectacular, strange and rare events, and downplay ordinary, familiar and common ones. We think rare risks are more common than they are. We fear them more than probability indicates we should.

There is a lot of psychological research that tries to explain this, but one of the key findings is this: People tend to base risk analysis more on stories than on data. Stories engage us at a much more visceral level, especially stories that are vivid, exciting or personally involving.

If a friend tells you about getting mugged in a foreign country, that story is more likely to affect how safe you feel traveling to that country than reading a page of abstract crime statistics will.

Novelty plus dread plus a good story equals overreaction.

And who are the major storytellers these days? Television and the Internet. So when news programs and sites endlessly repeat the story from Aurora, with interviews with those in the theater, interviews with the families and commentary by anyone who has a point to make, we start to think this is something to fear, rather than a rare event that almost never happens and isn't worth worrying about. In other words, reading five stories about the same event feels somewhat like five separate events, and that skews our perceptions.

We see the effects of this all the time.

It's strangers by whom we fear being murdered, kidnapped, raped and assaulted, when it's far more likely that any perpetrator of such offenses is a relative or a friend. We worry about airplane crashes and rampaging shooters instead of automobile crashes and domestic violence—both of which are far more common and far, far more deadly.

Our greatest recent overreaction to a rare event was our response to the terrorist attacks of 9/11. I remember then-Attorney General John Ashcroft giving a speech in Minnesota—where I live—in 2003 in which he claimed that the fact there were no new terrorist attacks since 9/11 was proof that his policies were working. I remember thinking: “There were no terrorist attacks in the two years preceding 9/11, and you didn't have any policies. What does that prove?”

What it proves is that terrorist attacks are very rare, and perhaps our national response wasn't worth the enormous expense, loss of liberty, attacks on our Constitution and damage to our credibility on the world stage. Still, overreacting was the natural thing for us to do. Yes, it was security theater and not real security, but it made many of us feel safer.

The rarity of events such as the Aurora massacre doesn't mean we should ignore any lessons it might teach us. Because people overreact to rare events, they're useful catalysts for social introspection and policy change. The key here is to focus not on the details of the particular event but on the broader issues common to all similar events.

Installing metal detectors at movie theaters doesn't make sense—there's no reason to think the next crazy gunman will choose a movie theater as his venue, and how effectively would a metal detector deter a lone gunman anyway?—but understanding the reasons why the United States has so many gun deaths compared with other countries does. The particular motivations of alleged killer James Holmes aren't relevant—the next gunman will have different motivations—but the general state of mental health care in the United States is.

Even with this, the most important lesson of the Aurora massacre is how rare these events actually are. Our brains are primed to believe that movie

theaters are more dangerous than they used to be, but they're not. The riskiest part of the evening is still the car ride to and from the movie theater, and even that's very safe.

But wear a seat belt all the same.

Militarizing Cyberspace Will Do More Harm Than Good

Originally published in the Irish Times, November 29, 2012

We're in the early years of a cyberwar arms race. It's expensive, it's destabilizing and it threatens the very fabric of the Internet we use every day. Cyberwar treaties, as imperfect as they might be, are the only way to contain the threat.

If you read the press and listen to government leaders, we're already in the middle of a cyberwar. By any normal definition of the word "war," this is ridiculous. But the definition of cyberwar has been expanded to include government-sponsored espionage, potential terrorist attacks in cyberspace, large-scale criminal fraud and even hacker kids attacking government networks and critical infrastructure. This definition is being pushed by the military and government contractors, both of which are gaining power and making money from cyberwar fears.

The main problem is that there are no good definitions of war in cyberspace. Also, we are increasingly seeing war-like tactics used in broader cyber conflicts. Technology is spreading capability, and the same "weaponry" is being used by everyone, from hackers to criminals to national militaries. It used to be that you could figure out whether you were at war by the weaponry deployed—but that is no longer the case.

This is important. When you're being attacked, there are a variety of organizations you can call on to defend yourself: the police, the military, whoever does anti-terrorism security in your country, your corporate lawyers.

The legal regime in which that defense operates depends on two things: who's attacking you, and why. Unfortunately, when you're being attacked in cyberspace, the two things you don't know are who's attacking you, and why. That makes defense, and national cyber defense policy, difficult.

The easy reaction is to lump all of these unknown attacks into "cyberwar." The corresponding danger is that military problems beg for military solutions.

We're starting to see a power grab in cyberspace by the world's militaries: large-scale monitoring of networks, military control of Internet standards, even military takeover of cyberspace.

The debate in the US over an "Internet kill switch" is another example; it's the sort of measure that might be deployed in wartime but makes no sense in peacetime. At the same time, countries are increasingly engaging in offensive actions in cyberspace, with tools such as Ghostnet (China), Stuxnet (US and Israel) and Flame (origin unknown).

A lot of what is being called cyberwar is little more than hacktivism—what I think of as kids playing politics—or criminal activity. Yes, it causes damage. Yes, we need to more effectively police cyberspace. But "police" is the operative word here. These are not threats that require a military response.

Arms races stem from ignorance and fear: ignorance of the other side's capabilities and fear that its capabilities are greater than one's own. Once cyberweapons exist, there will be an impetus to use them. Both Stuxnet and Flame damaged networks other than their intended targets. Any military-inserted backdoors in Internet systems make us more vulnerable to criminals and hackers.

It is only a matter of time before something big happens, perhaps by the rash actions of a low-level military officer, perhaps by a non-state actor, perhaps by accident. If the target nation retaliates, we could find ourselves in a real cyberwar.

The cyberwar arms race is destabilizing. International co-operation and treaties are the only way to reverse this. Banning cyberweapons entirely is a good goal, but almost certainly unachievable.

More likely are treaties that stipulate a no-first-use policy, outlaw unaimed or broadly targeted weapons, and mandate weapons that self-destruct at the end of hostilities. Treaties that restrict tactics and limit stockpiles could be a next step. We could also prohibit cyberattacks against civilian infrastructure; international banking, for example, could be declared off-limits.

Yes, enforcement will be difficult. Remember how easy it was to hide a chemical weapons facility? Hiding a cyberweapons facility will be even easier. But we have learnt a lot from our cold war experience in negotiating nuclear, chemical and biological treaties. The very act of negotiating limits the arms race and paves the way to peace. Even if they're breached, the world is safer because the treaties exist.

There's a common belief within the US military that cyberweapons treaties are not in the nation's best interest: that Americans have a military advantage in cyberspace that it should not squander. That's not true.

The US might have an offensive advantage—although that’s debatable—but it certainly doesn’t have a defensive advantage. More importantly, any heavily networked country such as the US is inherently vulnerable in cyberspace.

Cyberspace threats are real but militarizing cyberspace will do more harm than good. The value of a free and open Internet is enormous.

Rhetoric of Cyber War Breeds Fear—and More Cyber War

Originally published in the Irish Times, March 14, 2013

Americans have a weird relationship with the word “war.” We hate using it to describe actual wars but we love using it in a rhetorical context. We had the war on poverty, the war on crime, the war on drugs and the war on terror.

One of the big “wars” we’re talking about now is cyber war and, in this case, the word is dangerous. It is both a rhetorical war as well as something with elements of actual combat. The word also confuses the political debate about how to deal with cyber security.

The danger is that words frame the debate. If we use the rhetoric of war, we invoke feelings of fear and helplessness. We understand that this is something nations do to each other and that it’s not “normal” time when we’re at war.

We accept a different set of security solutions, one that more easily ignores freedoms and liberties.

We are more willing to let the military take over our Internet infrastructure and spy on our citizens.

On the other hand, if we use the rhetoric of peacetime espionage we think more about the rule of law and allow a much more limited role for the military.

Attacks from China

If you have been paying attention to the press recently, you might think China just started a cyberwar, attacking everyone in cyberspace.

First the *New York Times* announced it was the victim of a sophisticated cyberattack from China, one intended to obtain the names of Chinese nationals co-operating with the press.

Then the security firm Mandiant released a report naming a particular Chinese military unit as the source of a large number of cyberattacks against targets around the world.

Meanwhile, US president Barack Obama has signed a new cyber-security directive, citing threats from China as one of the motivations behind this action.

First, we need to understand that there is no cyber war going on. We are not nations at war and claiming otherwise is destabilizing. This is all espionage, something that has been going on between nations ever since nations were invented—and the US is giving as good as it's getting.

Seymour Hersh has written in the *New Yorker* magazine about US military operations in China.

Meanwhile, the US Cyber Command recently announced that it is expanding from 900 people to almost 5,000, while the National Security Agency is building a massive new data center in Utah. I'm sure China is just as fearful of the US as the US is of China.

While there are certainly a lot of state-sponsored cyberattacks emanating from China, it is not really news.

We in the security industry have been writing about Chinese cyberattacks for years, in earnest since the mid-2000s. Certainly, Internet-enabled espionage has been going on ever since there was an Internet.

GhostNet

In 2010, Google announced it was the victim of a sophisticated series of cyberattacks from China.

As with the attacks against the *New York Times*, the hackers were looking for particular people—in this case, human rights activists. Like the others, this attack was directed at more than one company.

At least 20 other large companies were targeted as well: Internet and technological companies, media companies and traditional companies.

In 2009, security researchers discovered a very sophisticated surveillance network they called GhostNet. They found it during an audit of the Dalai Lama's computers.

When they unraveled the command and control network, they found it was operating against high-value political, economic and media targets in 103 countries.

While there is no direct evidence that the Chinese were behind this, the list of targets read like a Who's Who of targets on whom China wants to spy.

The Chinese may deny particular incidents but they make no secret about their general policies.

They have been writing about their doctrine of domination in cyberspace for even longer. From our perspective there is not much new in the revelations of the past month. Still, the increasingly hostile war rhetoric and saber-rattling is worrisome.

We are in the early years of a cyberwar arms race. Arms races are fuelled by two things: ignorance and fear.

We don't know about the enemy's capabilities and we're afraid they're greater than our own.

So we spend more on weaponry, then even more. The other side does the same and the result is both dangerous and destabilizing.

Profitable

On the other hand, it's very profitable—for some.

There is an enormous amount of money and power that results from escalating a cyberwar arms race: power for the military, power for law enforcement and power for the large government contractors that support these organizations.

These are the people pounding the drums of cyber war and making news headlines warning us of a cyber 9/11, a cyber Pearl Harbor or—my favorite—a cyber Armageddon.

As long as “cyber” remains a prefix that scares, it will continue to be used as a tool to influence policy.

The Boston Marathon Bombing: Keep Calm and Carry On

Originally published in the Atlantic, April 15, 2013

As the details about the bombings in Boston unfold, it'd be easy to be scared. It'd be easy to feel powerless and demand that our elected leaders do something—anything—to keep us safe.

It'd be easy, but it'd be wrong. We need to be angry and empathize with the victims without being scared. Our fears would play right into the perpetrators' hands—and magnify the power of their victory for whichever goals whatever group behind this, still to be uncovered, has. We don't have to be scared, and we're not powerless. We actually have all the power here, and there's one thing we can do to render terrorism ineffective: Refuse to be terrorized.

It's hard to do, because terrorism is designed precisely to scare people—far out of proportion to its actual danger. A huge amount of research on fear and the brain teaches us that we exaggerate threats that are rare, spectacular, immediate, random—in this case involving an innocent child—senseless, horrific and graphic. Terrorism pushes all of our fear buttons, really hard, and we overreact.

But our brains are fooling us. Even though this will be in the news for weeks, we should recognize this for what it is: a rare event. That's the very definition of news: something that is unusual—in this case, something that almost never happens.

Remember after 9/11 when people predicted we'd see these sorts of attacks every few months? That never happened, and it wasn't because the TSA confiscated knives and snow globes at airports. Give the FBI credit for rolling up terrorist networks and interdicting terrorist funding, but we also exaggerated the threat. We get our ideas about how easy it is to blow things up from television and the movies. It turns out that terrorism is much harder than most people think. It's hard to find willing terrorists, it's hard to put a plot together, it's hard to get materials, and it's hard to execute a workable plan. As a collective group, terrorists are dumb, and they make dumb mistakes; criminal masterminds are another myth from movies and comic books.

Even the 9/11 terrorists got lucky.

If it's hard for us to keep this in perspective, it will be even harder for our leaders. They'll be afraid that by speaking honestly about the impossibility of attaining absolute security or the inevitability of terrorism—or that some American ideals are worth maintaining even in the face of adversity—they will be branded as “soft on terror.” And they'll be afraid that Americans might vote them out of office. Perhaps they're right, but where are the leaders who aren't afraid? What has happened to “the only thing we have to fear is fear itself”?

Terrorism, even the terrorism of radical Islamists and right-wing extremists and lone actors all put together, is not an “existential threat” against our nation. Even the events of 9/11, as horrific as they were, didn't do existential

damage to our nation. Our society is more robust than it might seem from watching the news. We need to start acting that way.

There are things we can do to make us safer, mostly around investigation, intelligence, and emergency response, but we will never be 100-percent safe from terrorism; we need to accept that.

How well this attack succeeds depends much less on what happened in Boston than by our reactions in the coming weeks and months. Terrorism isn't primarily a crime against people or property. It's a crime against our minds, using the deaths of innocents and destruction of property as accomplices. When we react from fear, when we change our laws and policies to make our country less open, the terrorists succeed, even if their attacks fail. But when we refuse to be terrorized, when we're indomitable in the face of terror, the terrorists fail, even if their attacks succeed.

Don't glorify the terrorists and their actions by calling this part of a "war on terror." Wars involve two legitimate sides. There's only one legitimate side here; those on the other are criminals. They should be found, arrested, and punished. But we need to be vigilant not to weaken the very freedoms and liberties that make this country great, meanwhile, just because we're scared.

Empathize, but refuse to be terrorized. Instead, be indomitable—and support leaders who are as well. That's how to defeat terrorists.

Why FBI and CIA Didn't Connect the Dots

Originally published in CNN, May 2, 2013

The FBI and the CIA are being criticized for not keeping better track of Tamerlan Tsarnaev in the months before the Boston Marathon bombings. How could they have ignored such a dangerous person? How do we reform the intelligence community to ensure this kind of failure doesn't happen again?

It's an old song by now, one we heard after the 9/11 attacks in 2001 and after the Underwear Bomber's failed attack in 2009. The problem is that connecting the dots is a bad metaphor, and focusing on it makes us more likely to implement useless reforms.

Connecting the dots in a coloring book is easy and fun. They're right there on the page, and they're all numbered. All you have to do is move your pencil

from one dot to the next, and when you're done, you've drawn a sailboat. Or a tiger. It's so simple that 5-year-olds can do it.

But in real life, the dots can only be numbered after the fact. With the benefit of hindsight, it's easy to draw lines from a Russian request for information to a foreign visit to some other piece of information that might have been collected.

In hindsight, we know who the bad guys are. Before the fact, there are an enormous number of potential bad guys.

How many? We don't know. But we know that the no-fly list had 21,000 people on it last year. The Terrorist Identities Datamart Environment, also known as the watch list, has 700,000 names on it.

We have no idea how many potential "dots" the FBI, CIA, NSA and other agencies collect, but it's easily in the millions. It's easy to work backwards through the data and see all the obvious warning signs. But before a terrorist attack, when there are millions of dots—some important but the vast majority unimportant—uncovering plots is a lot harder.

Rather than thinking of intelligence as a simple connect-the-dots picture, think of it as a million unnumbered pictures superimposed on top of each other. Or a random-dot stereogram. Is it a sailboat, a puppy, two guys with pressure-cooker bombs or just an unintelligible mess of dots? You try to figure it out.

It's not a matter of not enough data, either.

Piling more data onto the mix makes it harder, not easier. The best way to think of it is a needle-in-a-haystack problem; the last thing you want to do is increase the amount of hay you have to search through.

The television show "Person of Interest" is fiction, not fact.

There's a name for this sort of logical fallacy: hindsight bias.

First explained by psychologists Daniel Kahneman and Amos Tversky, it's surprisingly common. Since what actually happened is so obvious once it happens, we overestimate how obvious it was before it happened.

We actually misremember what we once thought, believing that we knew all along that what happened would happen. It's a surprisingly strong tendency, one that has been observed in countless laboratory experiments and real-world examples of behavior. And it's what all the post-Boston-Marathon bombing dot-connectors are doing.

Before we start blaming agencies for failing to stop the Boston bombers, and before we push "intelligence reforms" that will shred civil liberties without making us any safer, we need to stop seeing the past as a bunch of obvious dots that need connecting.

Kahneman, a Nobel prize winner, wisely noted: “Actions that seemed prudent in foresight can look irresponsibly negligent in hindsight.” Kahneman calls it “the illusion of understanding,” explaining that the past is only so understandable because we have cast it as simple inevitable stories and leave out the rest.

Nassim Taleb, an expert on risk engineering, calls this tendency the “narrative fallacy.” We humans are natural storytellers, and the world of stories is much more tidy, predictable and coherent than the real world.

Millions of people behave strangely enough to warrant the FBI’s notice, and almost all of them are harmless. It is simply not possible to find every plot beforehand, especially when the perpetrators act alone and on impulse.

We have to accept that there always will be a risk of terrorism, and that when the occasional plot succeeds, it’s not necessarily because our law enforcement systems have failed.

The FBI’s New Wiretapping Plan Is Great News for Criminals

Originally published in Foreign Policy, May 29, 2013

The FBI wants a new law that will make it easier to wiretap the Internet. Although its claim is that the new law will only maintain the status quo, it’s really much worse than that. This law will result in less-secure Internet products and create a foreign industry in more-secure alternatives. It will impose costly burdens on affected companies. It will assist totalitarian governments in spying on their own citizens. And it won’t do much to hinder actual criminals and terrorists.

As the FBI sees it, the problem is that people are moving away from traditional communication systems like telephones onto computer systems like Skype. Eavesdropping on telephones used to be easy. The FBI would call the phone company, which would bring agents into a switching room and allow them to literally tap the wires with a pair of alligator clips and a tape recorder. In the 1990s, the government forced phone companies to provide an analog capability on digital switches; but today, more and more communications happens over the Internet.

What the FBI wants is the ability to eavesdrop on *everything*. Depending on the system, this ranges from easy to impossible. E-mail systems like Gmail

are easy. The mail resides in Google's servers, and the company has an office full of people who respond to requests for lawful access to individual accounts from governments all over the world. Encrypted voice systems like Silent Circle are impossible to eavesdrop on—the calls are encrypted from one computer to the other, and there's no central node to eavesdrop from. In those cases, the only way to make the system eavesdroppable is to add a backdoor to the user software. This is precisely the FBI's proposal. Companies that refuse to comply would be fined \$25,000 a day.

The FBI believes it can have it both ways: that it can open systems to its eavesdropping, but keep them secure from anyone else's eavesdropping. That's just not possible. It's impossible to build a communications system that allows the FBI surreptitious access but doesn't allow similar access by others. When it comes to security, we have two options: We can build our systems to be as secure as possible from eavesdropping, or we can deliberately weaken their security. We have to choose one or the other.

This is an old debate, and one we've been through many times. The NSA even has a name for it: the equities issue. In the 1980s, the equities debate was about export control of cryptography. The government deliberately weakened US cryptography products because it didn't want foreign groups to have access to secure systems. Two things resulted: fewer Internet products with cryptography, to the insecurity of everybody, and a vibrant foreign security industry based on the unofficial slogan "Don't buy the US stuff—it's lousy."

In 1993, the debate was about the Clipper Chip. This was another deliberately weakened security product, an encrypted telephone. The FBI convinced AT&T to add a backdoor that allowed for surreptitious wiretapping. The product was a complete failure. Again, why would anyone buy a deliberately weakened security system?

In 1994, the Communications Assistance for Law Enforcement Act mandated that US companies build eavesdropping capabilities into phone switches. These were sold internationally; some countries liked having the ability to spy on their citizens. Of course, so did criminals, and there were public scandals in Greece (2005) and Italy (2006) as a result.

In 2012, we learned that every phone switch sold to the Department of Defense had security vulnerabilities in its surveillance system. And just this May, we learned that Chinese hackers breached Google's system for providing surveillance data for the FBI.

The new FBI proposal will fail in all these ways and more. The bad guys will be able to get around the eavesdropping capability, either by building their own security systems—not very difficult—or buying the more-secure foreign products that will inevitably be made available. Most of the good guys, who don't understand the risks or the technology, will not know enough to bother and will be less secure. The eavesdropping functions will 1) result in more obscure—and less secure—product designs, and 2) be vulnerable to exploitation by criminals, spies, and everyone else. US companies will be forced to compete at a disadvantage; smart customers won't buy the substandard stuff when there are more-secure foreign alternatives. Even worse, there are lots of foreign governments who want to use these sorts of systems to spy on their own citizens. Do we really want to be exporting surveillance technology to the likes of China, Syria, and Saudi Arabia?

The FBI's short-sighted agenda also works against the parts of the government that are still working to secure the Internet for everyone. Initiatives within the NSA, the DOD, and DHS to do everything from securing computer operating systems to enabling anonymous web browsing will all be harmed by this.

What to do, then? The FBI claims that the Internet is “going dark,” and that it's simply trying to maintain the status quo of being able to eavesdrop. This characterization is disingenuous at best. We are entering a golden age of surveillance; there's more electronic communications available for eavesdropping than ever before, including whole new classes of information: location tracking, financial tracking, and vast databases of historical communications such as e-mails and text messages. The FBI's surveillance department has it better than ever. With regard to voice communications, yes, software phone calls will be harder to eavesdrop upon. (Although there are questions about Skype's security.) That's just part of the evolution of technology, and one that on balance is a positive thing.

Think of it this way: We don't hand the government copies of our house keys and safe combinations. If agents want access, they get a warrant and then pick the locks or bust open the doors, just as a criminal would do. A similar system would work on computers. The FBI, with its increasingly non-transparent procedures and systems, has failed to make the case that this isn't good enough.

Finally there's a general principle at work that's worth explicitly stating. All tools can be used by the good guys and the bad guys. Cars have enormous societal value, even though bank robbers can use them as getaway cars. Cash

is no different. Both good guys and bad guys send e-mails, use Skype, and eat at all-night restaurants. But because society consists overwhelmingly of good guys, the good uses of these dual-use technologies greatly outweigh the bad uses. Strong Internet security makes us all safer, even though it helps the bad guys as well. And it makes no sense to harm all of us in an attempt to harm a small subset of us.

US Offensive Cyberwar Policy

Originally published in CNN, June 18, 2013

Today, the United States is conducting offensive cyberwar actions around the world.

More than passively eavesdropping, we're penetrating and damaging foreign networks for both espionage and to ready them for attack. We're creating custom-designed Internet weapons, pre-targeted and ready to be "fired" against some piece of another country's electronic infrastructure on a moment's notice.

This is much worse than what we're accusing China of doing to us. We're pursuing policies that are both expensive and destabilizing and aren't making the Internet any safer. We're reacting from fear, and causing other countries to counter-react from fear. We're ignoring resilience in favor of offense.

Welcome to the cyberwar arms race, an arms race that will define the Internet in the 21st century.

Presidential Policy Directive 20, issued last October and released by Edward Snowden, outlines US cyberwar policy. Most of it isn't very interesting, but there are two paragraphs about "Offensive Cyber Effect Operations," or OCEO, that are intriguing:

OCEO can offer unique and unconventional capabilities to advance US national objectives around the world with little or no warning to the adversary or target and with potential effects ranging from subtle to severely damaging. The development and sustainment of OCEO capabilities, however, may require considerable time and effort if access and tools for a specific target do not already exist.

The United States Government shall identify potential targets of national importance where OCEO can offer a favorable balance of effectiveness and risk as compared with other instruments of national power, establish and

maintain OCEO capabilities integrated as appropriate with other US offensive capabilities, and execute those capabilities in a manner consistent with the provisions of this directive.

These two paragraphs, and another paragraph about OCEO, are the only parts of the document classified “top secret.” And that’s because what they’re saying is very dangerous.

Cyberattacks have the potential to be both immediate and devastating. They can disrupt communications systems, disable national infrastructure, or, as in the case of Stuxnet, destroy nuclear reactors; but only if they’ve been created and targeted beforehand. Before launching cyberattacks against another country, we have to go through several steps.

We have to study the details of the computer systems they’re running and determine the vulnerabilities of those systems. If we can’t find exploitable vulnerabilities, we need to create them: leaving “backdoors” in hacker speak. Then we have to build new cyberweapons designed specifically to attack those systems.

Sometimes we have to embed the hostile code in those networks—these are called “logic bombs”—to be unleashed in the future. And we have to keep penetrating those foreign networks, because computer systems always change and we need to ensure that the cyberweapons are still effective.

Like our nuclear arsenal during the Cold War, our cyberweapons arsenal must be pretargeted and ready to launch.

That’s what Obama directed the US Cyber Command to do. We can see glimpses of how effective we are in Snowden’s allegations that the NSA is currently penetrating foreign networks around the world: “We hack network backbones—like huge Internet routers, basically—that give us access to the communications of hundreds of thousands of computers without having to hack every single one.”

The NSA and the US Cyber Command are basically the same thing. They’re both at Fort Meade in Maryland, and they’re both led by Gen. Keith Alexander. The same people who hack network backbones are also building weapons to destroy those backbones. At a March Senate briefing, Alexander boasted of creating more than a dozen offensive cyber units.

Longtime NSA watcher James Bamford reached the same conclusion in his recent profile of Alexander and the US Cyber Command (written before the Snowden revelations). He discussed some of the many cyberweapons the US purchases:

According to Defense News' C4ISR Journal and Bloomberg Businessweek, Endgame also offers its intelligence clients—agencies like Cyber Command, the NSA, the CIA, and British intelligence—a unique map showing them exactly where their targets are located. Dubbed Bonesaw, the map displays the geolocation and digital address of basically every device connected to the Internet around the world, providing what's called network situational awareness. The client locates a region on the password-protected web-based map, then picks a country and city—say, Beijing, China. Next the client types in the name of the target organization, such as the Ministry of Public Security's No. 3 Research Institute, which is responsible for computer security—or simply enters its address, 6 Zhengyi Road. The map will then display what software is running on the computers inside the facility, what types of malware some may contain, and a menu of custom-designed exploits that can be used to secretly gain entry. It can also pinpoint those devices infected with malware, such as the Conficker worm, as well as networks turned into botnets and zombies—the equivalent of a back door left open. . .

The buying and using of such a subscription by nation-states could be seen as an act of war. "If you are engaged in reconnaissance on an adversary's systems, you are laying the electronic battlefield and preparing to use it" wrote Mike Jacobs, a former NSA director for information assurance, in a McAfee report on cyberwarfare. "In my opinion, these activities constitute acts of war, or at least a prelude to future acts of war." The question is, who else is on the secretive company's client list? Because there is as of yet no oversight or regulation of the cyberweapons trade, companies in the cyber-industrial complex are free to sell to whomever they wish. "It should be illegal," said the former senior intelligence official involved in cyberwarfare. "I knew about Endgame when I was in intelligence. The intelligence community didn't like it, but they're the largest consumer of that business."

That's the key question: How much of what the United States is currently doing is an act of war by international definitions? Already we're accusing China of penetrating our systems in order to map "military capabilities that

could be exploited during a crisis.” What PPD-20 and Snowden describe is much worse, and certainly China, and other countries, are doing the same.

All of this mapping of vulnerabilities and keeping them secret for offensive use makes the Internet less secure, and these pre-targeted, ready-to-unleash cyberweapons are destabilizing forces on international relationships. Rooting around other countries’ networks, analyzing vulnerabilities, creating backdoors, and leaving logic bombs could easily be construed as an act of war. And all it takes is one over-achieving national leader for this all to tumble into actual war.

It’s time to stop the madness. Yes, our military needs to invest in cyberwar capabilities, but we also need international rules of cyberwar, more transparency from our own government on what we are and are not doing, international cooperation between governments and viable cyberweapons treaties. Yes, these are difficult. Yes, it’s a long slow process. Yes, there won’t be international consensus, certainly not in the beginning. But even with all of those problems, it’s a better path to go down than the one we’re on now.

We can start by taking most of the money we’re investing in offensive cyberwar capabilities and spend them on national cyberspace resilience. MAD, mutually assured destruction, made sense because there were two superpowers opposing each other. On the Internet there are all sorts of different powers, from nation-states to much less organized groups. An arsenal of cyberweapons begs to be used, and, as we learned from Stuxnet, there’s always collateral damage to innocents when they are. We’re much safer with a strong defense than with a counterbalancing offense.

3

Human Aspects of Security

Secret Questions Blow a Hole in Security

Originally published in ComputerWeekly, April 4, 2008

It's a mystery to me why websites think "secret questions" are a good idea. We sign up for an online service, choose a hard-to-guess (and equally hard-to-remember) password, and are then presented with a "secret question" to answer.

Twenty years ago, there was just one secret question: What's your mother's maiden name? Today, there are several: What street did you grow up on? What's the name of your favorite teacher? What's your favorite color? Often, you get to choose.

The idea is to give customers a backup password. If you forget your password, then the secret question is a way to verify your identity. It's a great idea from a customer service perspective—users are less likely to forget their first pet's name than some random password—but terrible for security.

The answer to the secret question is much easier to guess than a good password, and the information is much more public. I'll bet my childhood address is in some database somewhere. And worse, everybody seems to use the same series of secret questions.

The result is that the normal security protocol (passwords) falls back to a much less secure protocol (secret questions). The security of the entire system suffers. I'm sure the designers of the system thought the fallback system would only be used rarely, when a user forgot their password. But any good security engineer realizes that bad guys can force the failure whenever they want, and

that the whole security of the system rests on the security of the weaker of the two subsystems.

What can be done? As a customer, my usual technique is to type a completely random answer for the security question. I madly slap at my keyboard for a few seconds, and then forget about it. This ensures that an attacker has little chance of bypassing the password protection by successfully guessing the answer to my secret question, but it is pretty unpleasant if I forget my password. The one time this happened to me, I had to call the company to get my password and question reset. Yes, it was a right pain.

Which is maybe what should happen in the first place. I like to think that if I forget my password, it is really hard to gain access to my account. I want it to be so hard that an attacker can't possibly do it. I know this is a customer service issue, but it's a security issue, too. And if the password is controlling access to something important—like my bank account—then the bypass mechanism should be harder, not easier.

Passwords have reached the end of their useful life. Today, they only work for low-security applications. The secret question is just one manifestation of that fact.

When You Lose a Piece of Kit, the Real Loss Is the Data It Contains

Originally published in the Guardian, December 4, 2008

These days, losing electronic devices is less about the hardware and more about the data. Hardly a week goes by without another newsworthy data loss. People leave thumb drives, memory sticks, mobile phones and even computers everywhere. And some of that data isn't easily replaceable. Sure, you can blame it on personal or organizational sloppiness, but part of the problem is that more and more information fits on smaller and smaller devices.

My primary computer is an ultraportable laptop. It contains every email I've sent and received over the past 12 years—I think of it as my backup brain—as well as an enormous amount of personal and work-related documents.

I have several USB thumb drives, including an 8GB drive that serves as my primary backup while travelling. It contains a complete copy of the past 12 months of my life. A larger USB portable drive serves as my primary storage

device for photographs; I carry that around regularly, too, as I like to edit my photos on flights.

My mobile phone is a Palm Treo smartphone. It holds not only my frequently called phone numbers, but my entire address book—including any personal notes I've made—my calendar for the past 10 years, hundreds of emails, all my text messages and a log of every phone call I've made and received. At least, it would if I didn't take specific pains to clean that information out once in a while.

Backup DVDs. iPods with calendars and address books. USB drives with portable desktops. I could go on. The upside to this is that so much of our information is at our fingertips. I travel so extensively that I need my office anywhere I am, so I want everything with me everywhere. The downside is that it's now amazingly easy to lose an enormous amount of information. And there are two problems with that. One, you've lost the information. And two, that perhaps someone else has found it.

The first problem is easily solvable with backup. Everything you own should be backed up regularly. Not just your computer, but your PDA and mobile phone and anything else with personal data. Backups should be tested regularly. There's nothing worse than losing something and having the backups fail when you try to restore to a replacement device.

The second problem is solvable several ways. The best way is encryption. On your computer, hard-disk encryption programs like PGPDisk or TrueCrypt allow you to encrypt files, folders or entire disk partitions. Several manufacturers market USB thumb drives with built-in encryption. Some PDA manufacturers are starting to add password protection—not as good as encryption, but at least it's something—to their devices, and there are a few aftermarket PDA encryption programs. I use these wherever possible, and I strongly recommend that everyone else do the same.

Where encryption isn't possible, pay attention and erase unneeded data. Delete old emails from your BlackBerry, texts from your cellphone and old data from your address books regularly. It can be difficult to know exactly what your PDA is storing, and how to erase it. Manufacturers could help with this by introducing better functionality and thereby making the devices easier to use.

Another thing manufacturers can do is to provide the option to delete the data remotely if the device is lost. This is still a new idea, but it's gaining traction in the corporate market. These systems frequently allow for remote backup of the data, solving both problems at once. One last piece of advice:

for work-related equipment, you should follow your company's backup and security policies.

The goal here is peace of mind. When people lose computers or phones or USB drives, the real loss isn't the physical object, but the data contained within it. And while we won't be able to make these devices harder to lose, especially as they continue to shrink in physical size and grow in data capacity, we can make their loss cost merely money, not information or privacy.

The Kindness of Strangers

Originally published in the Wall Street Journal, March 12, 2009

When I was growing up, children were commonly taught: "don't talk to strangers." Strangers might be bad, we were told, so it's prudent to steer clear of them.

And yet most people are honest, kind, and generous, especially when someone asks them for help. If a small child is in trouble, the smartest thing he can do is find a nice-looking stranger and talk to him.

These two pieces of advice may seem to contradict each other, but they don't. The difference is that in the second instance, the child is choosing which stranger to talk to. Given that the overwhelming majority of people will help, the child is likely to get help if he chooses a random stranger. But if a stranger comes up to a child and talks to him or her, it's not a random choice. It's more likely, although still unlikely, that the stranger is up to no good.

As a species, we tend to help each other, and a surprising amount of our security and safety comes from the kindness of strangers. During disasters: floods, earthquakes, hurricanes, bridge collapses. In times of personal tragedy. And even in normal times.

If you're sitting in a café working on your laptop and need to get up for a minute, ask the person sitting next to you to watch your stuff. He's very unlikely to steal anything. Or, if you're nervous about that, ask the three people sitting around you. Those three people don't know each other, and will not only watch your stuff, but they'll also watch each other to make sure no one steals anything.

Again, this works because you're selecting the people. If three people walk up to you in the café and offer to watch your computer while you go to the bathroom, don't take them up on that offer. Your odds of getting three honest people are much lower.

Some computer systems rely on the kindness of strangers, too. The Internet works because nodes benevolently forward packets to each other without any recompense from either the sender or receiver of those packets. Wikipedia works because strangers are willing to write for, and edit, an encyclopedia—with no recompense.

Collaborative spam filtering is another example. Basically, once someone notices a particular e-mail is spam, he marks it, and everyone else in the network is alerted that it's spam. Marking the e-mail is a completely altruistic task; the person doing it gets no benefit from the action. But he receives benefit from everyone else doing it for other e-mails.

Tor is a system for anonymous Web browsing. The details are complicated, but basically, a network of Tor servers passes Web traffic among each other in such a way as to anonymize where it came from. Think of it as a giant shell game. As a Web surfer, I put my Web query inside a shell and send it to a random Tor server. That server knows who I am but not what I am doing. It passes that shell to another Tor server, which passes it to a third. That third server—which knows what I am doing but not who I am—processes the Web query. When the Web page comes back to that third server, the process reverses itself and I get my Web page. Assuming enough Web surfers are sending enough shells through the system, even someone eavesdropping on the entire network can't figure out what I'm doing.

It's a very clever system, and it protects a lot of people, including journalists, human rights activists, whistleblowers, and ordinary people living in repressive regimes around the world. But it only works because of the kindness of strangers. No one gets any benefit from being a Tor server; it uses up bandwidth to forward other people's packets around. It's more efficient to be a Tor client and use the forwarding capabilities of others. But if there are no Tor servers, then there's no Tor. Tor works because people are willing to set themselves up as servers, at no benefit to them.

Alibi clubs work along similar lines. You can find them on the Internet, and they're loose collections of people willing to help each other out with alibis. Sign up, and you're in. You can ask someone to pretend to be your doctor and call your boss. Or someone to pretend to be your boss and call your spouse. Or maybe someone to pretend to be your spouse and call your boss. Whatever you want, just ask and some anonymous stranger will come to your rescue. And because your accomplice is an anonymous stranger, it's safer than asking a friend to participate in your ruse.

There are risks in these sorts of systems. Regularly, marketers and other people with agendas try to manipulate Wikipedia entries to suit their interests. Intelligence agencies can, and almost certainly have, set themselves up as Tor servers to better eavesdrop on traffic. And a do-gooder could join an alibi club just to expose other members. But for the most part, strangers are willing to help each other, and systems that harvest this kindness work very well on the Internet.

Blaming the User Is Easy—But It’s Better to Bypass Them Altogether

Originally published in the Guardian, March 12, 2009

Blaming the victim is common in IT: users are to blame because they don’t patch their systems, choose lousy passwords, fall for phishing attacks, and so on. But, while users are, and will continue to be, a major source of security problems, focusing on them is an unhelpful way to think.

People regularly don’t do things they are supposed to: changing the oil in their cars, going to the dentist, replacing the batteries in their smoke detectors. Why? Because people learn from experience. If something is immediately harmful, such as touching a hot stove or petting a live tiger, they quickly learn not to do it. But if someone skips an oil change, ignores a computer patch, or chooses a lousy password, it’s unlikely to matter. No feedback, no learning.

We’ve tried to solve this in several ways. We give people rules of thumb: oil change every 5,000 miles; secure password guidelines. Or we send notifications: smoke alarms beep at us, dentists send postcards, Google warns us if we are about to visit a website suspected of hosting malware. But, again, the effects of ignoring these aren’t generally felt immediately.

This makes security primarily a hindrance to the user. It’s a recurring obstacle: something that interferes with the seamless performance of the user’s task. And it’s human nature, wired into our reasoning skills, to remove recurring obstacles. So, if the consequences of bypassing security aren’t obvious, then people will naturally do it.

This is the problem with Microsoft’s User Account Control (UAC). Introduced in Vista, the idea is to improve security by limiting the privileges applications have when they’re running. But the security prompts pop up too

frequently, and there's rarely any ill-effect from ignoring them. So people do ignore them.

This doesn't mean user education is worthless. On the contrary, user education is an important part of any corporate security program. And at home, the more users understand security threats and hacker tactics, the more secure their systems are likely to be. But we should also recognize the limitations of education.

The solution is to better design security systems that assume uneducated users: to prevent them from changing security settings that would leave them exposed to undue risk, or—even better—to take security out of their hands entirely.

For example, we all know that backups are a good thing. But if you forget to do a backup this week, nothing terrible happens. In fact, nothing terrible happens for years on end when you forget. So, despite what you know, you start believing that backups aren't really that important. Apple got the solution right with its backup utility Time Machine. Install it, plug in an external hard drive, and you are automatically backed up against hardware failure and human error. It's easier to use it than not.

For its part, Microsoft has made great strides in securing its operating system, providing default security settings in Windows XP and even more in Windows Vista to ensure that, when a naive user plugs a computer in, it's not defenseless.

Unfortunately, blaming the user can be good business. Mobile phone companies save money if they can bill their customers when a calling card number is stolen and used fraudulently. British banks save money by blaming users when they are victims of chip-and-pin fraud. This is continuing, with some banks going so far as to accuse the victim of perpetrating the fraud, despite evidence of large-scale fraud by organized crime syndicates.

The legal system needs to fix the business problems, but system designers need to work on the technical problems. They must accept that security systems that require the user to do the right thing are doomed to fail. And then they must design resilient security nevertheless.

The Value of Self-Enforcing Protocols

Originally published in Threatpost, August 10, 2009

There are several ways two people can divide a piece of cake in half. One way is to find someone impartial to do it for them. This works, but it requires

another person. Another way is for one person to divide the piece, and the other person to complain (to the police, a judge, or his parents) if he doesn't think it's fair. This also works, but still requires another person—at least to resolve disputes. A third way is for one person to do the dividing, and for the other person to choose the half he wants.

That third way, known by kids, pot smokers, and everyone else who needs to divide something up quickly and fairly, is called cut-and-choose. People use it because it's a self-enforcing protocol: a protocol designed so that neither party can cheat.

Self-enforcing protocols are useful because they don't require trusted third parties. Modern systems for transferring money—checks, credit cards, PayPal—require trusted intermediaries like banks and credit card companies to facilitate the transfer. Even cash transfers require a trusted government to issue currency, and they take a cut in the form of seigniorage. Modern contract protocols require a legal system to resolve disputes. Modern commerce wasn't possible until those systems were in place and generally trusted, and complex business contracts still aren't possible in areas where there is no fair judicial system. Barter is a self-enforcing protocol: nobody needs to facilitate the transaction or resolve disputes. It just works.

Self-enforcing protocols are safer than other types because participants don't gain an advantage from cheating. Modern voting systems are rife with the potential for cheating, but an open show of hands in a room—one that everyone in the room can count for himself—is self-enforcing. On the other hand, there's no secret ballot, late voters are potentially subjected to coercion, and it doesn't scale well to large elections. But there are mathematical election protocols that have self-enforcing properties, and some cryptographers have suggested their use in elections.

Here's a self-enforcing protocol for determining property tax: the homeowner decides the value of the property and calculates the resultant tax, and the government can either accept the tax or buy the home for that price. Sounds unrealistic, but the Greek government implemented exactly that system for the taxation of antiquities. It was the easiest way to motivate people to accurately report the value of antiquities.

A VAT, or value-added tax, is a self-enforcing alternative to sales tax. Sales tax is collected on the entire value of the thing at the point of retail sale; both the customer and the store owner want to cheat the government. But VAT is collected at every step between raw materials and that final customer; it's the difference between the price of the materials sold and the materials bought.

Buyers want official receipts with as high a purchase price as possible, so each buyer along the chain keeps each seller honest. Yes, there's still an incentive to cheat on the final sale to the customer, but the amount of tax collected at that point is much lower.

Of course, self-enforcing protocols aren't perfect. For example, someone in a cut-and-choose can punch the other guy and run away with the entire piece of cake. But perfection isn't the goal here; the goal is to reduce cheating by taking away potential avenues of cheating. Self-enforcing protocols improve security not by implementing countermeasures that prevent cheating, but by leveraging economic incentives so that the parties don't want to cheat.

One more self-enforcing protocol. Imagine a pirate ship that encounters a storm. The pirates are all worried about their gold, so they put their personal bags of gold in the safe. During the storm, the safe cracks open, and all the gold mixes up and spills out on the floor. How do the pirates determine who owns what? They each announce to the group how much gold they had. If the total of all the announcements matches what's in the pile, it's divided as people announced. If it's different, then the captain keeps it all. I can think of all kinds of ways this can go wrong—the captain and one pirate can collude to throw off the total, for example—but it is self-enforcing against individual misreporting.

Reputation Is Everything in IT Security

Originally published in the Guardian, November 11, 2009

In the past, our relationship with our computers was technical. We cared what CPU they had and what software they ran. We understood our networks and how they worked. We were experts, or we depended on someone else for expertise. And security was part of that expertise.

This is changing. We access our email via the web, from any computer or from our phones. We use Facebook, Google Docs, even our corporate networks, regardless of hardware or network. We, especially the younger of us, no longer care about the technical details. Computing is infrastructure; it's a commodity. It's less about products and more about services; we simply expect it to work, like telephone service or electricity or a transportation network.

Infrastructures can be spread on a broad continuum, ranging from generic to highly specialized. Power and water are generic; who supplies them doesn't really matter. Mobile phone services, credit cards, ISPs, and airlines are

mostly generic. More specialized infrastructure services are restaurant meals, haircuts, and social networking sites. Highly specialized services include tax preparation for complex businesses, management consulting, legal services, and medical services.

Sales for these services are driven by two things: price and trust. The more generic the service is, the more price dominates. The more specialized it is, the more trust dominates. IT is something of a special case because so much of it is free. So, for both specialized IT services where price is less important and for generic IT services—think Facebook—where there is no price, trust will grow in importance. IT is becoming a reputation-based economy, and this has interesting ramifications for security.

Some years ago, the major credit card companies became concerned about the plethora of credit-card-number thefts from sellers' databases. They worried that these might undermine the public's trust in credit cards as a secure payment system for the Internet. They knew the sellers would only protect these databases up to the level of the threat to the seller, and not to the greater level of threat to the industry as a whole. So they banded together and produced a security standard called PCI. It's wholly industry-enforced—by an industry that realized its reputation was more valuable than the sellers' databases.

A reputation-based economy means that infrastructure providers care more about security than their customers do. I realized this 10 years ago with my own company. We provided network-monitoring services to large corporations, and our internal network security was much more extensive than our customers'. Our customers secured their networks—that's why they hired us, after all—but only up to the value of their networks. If we mishandled any of our customers' data, we would have lost the trust of all of our customers.

I heard the same story at an ENISA conference in London last June, when an IT consultant explained that he had begun encrypting his laptop years before his customers did. While his customers might decide that the risk of losing their data wasn't worth the hassle of dealing with encryption, he knew that if he lost data from one customer, he risked losing all of his customers.

As IT becomes more like infrastructure, more like a commodity, expect service providers to improve security to levels greater than their customers would have done themselves.

In IT, customers learn about company reputation from many sources: magazine articles, analyst reviews, recommendations from colleagues, awards, certifications, and so on. Of course, this only works if customers have accurate

information. In a reputation economy, companies have a motivation to hide their security problems.

You've all experienced a reputation economy: restaurants. Some restaurants have a good reputation, and are filled with regulars. When restaurants get a bad reputation, people stop coming and they close. Tourist restaurants—whose main attraction is their location, and whose customers frequently don't know anything about their reputation—can thrive even if they aren't any good. And sometimes a restaurant can keep its reputation—an award in a magazine, a special occasion restaurant that “everyone knows” is the place to go—long after its food and service have declined.

The reputation economy is far from perfect.

When to Change Passwords

Originally published in Dark Reading, November 10, 2010

How often should you change your password? I get asked that question a lot, usually by people annoyed at their employer's or bank's password expiration policy—people who finally memorized their current password and are realizing they'll have to write down their new one. How could that possibly be more secure, they want to know.

The answer depends on what the password is used for.

The downside of changing passwords is that it makes them harder to remember. And if you force people to change their passwords regularly, they're more likely to choose easy-to-remember—and easy-to-guess—passwords than they are if they can use the same passwords for many years. So any password-changing policy needs to be chosen with that consideration in mind.

The primary reason to give an authentication credential—not just a password, but any authentication credential—an expiration date is to limit the amount of time a lost, stolen, or forged credential can be used by someone else. If a membership card expires after a year, then if someone steals that card he can at most get a year's worth of benefit out of it. After that, it's useless.

This becomes less important when the credential contains a biometric—even a photograph—or is verified online. It's much less important for a credit card or passport to have an expiration date now that they're not so much bearer documents as just pointers to a database. If, for example, the credit

card database knows when a card is no longer valid, there's no reason to put an expiration date on the card. But the expiration date does mean that a forgery is only good for a limited length of time.

Passwords are no different. If a hacker gets your password either by guessing or stealing it, he can access your network as long as your password is valid. If you have to update your password every quarter, that significantly limits the utility of that password to the attacker.

At least, that's the traditional theory. It assumes a passive attacker, one who will eavesdrop over time without alerting you that he's there. In many cases today, though, that assumption no longer holds. An attacker who gets the password to your bank account by guessing or stealing it isn't going to eavesdrop. He's going to transfer money out of your account—and then you're going to notice. In this case, it doesn't make a lot of sense to change your password regularly—but it's vital to change it immediately after the fraud occurs.

Someone committing espionage in a private network is more likely to be stealthy. But he's also not likely to rely on the user credential he guessed and stole; he's going to install backdoor access or create his own account. Here again, forcing network users to regularly change their passwords is less important than forcing everyone to change their passwords immediately after the spy is detected and removed—you don't want him getting in again.

Social networking sites are somewhere in the middle. Most of the criminal attacks against Facebook users use the accounts for fraud. "Help! I'm in London and my wallet was stolen. Please wire money to this account. Thank you." Changing passwords periodically doesn't help against this attack, although—of course—change your password as soon as you regain control of your account. But if your kid sister has your password—or the tabloid press, if you're that kind of celebrity—they're going to listen in until you change it. And you might not find out about it for months.

So in general: you don't need to regularly change the password to your computer or online financial accounts (including the accounts at retail sites); definitely not for low-security accounts. You should change your corporate login password occasionally, and you need to take a good hard look at your friends, relatives, and paparazzi before deciding how often to change your Facebook password. But if you break up with someone you've shared a computer with, change them all.

Two final points. One, this advice is for login passwords. There's no reason to change any password that is a key to an encrypted file. Just keep the same password as long as you keep the file, unless you suspect it's been compromised.

And two, it's far more important to choose a good password for the sites that matter—don't worry about sites you don't care about that nonetheless demand that you register and choose a password—in the first place than it is to change it. So if you have to worry about something, worry about that. And write your passwords down, or use a program like Password Safe.

The Big Idea: Bruce Schneier

Originally published in Whatever, February 16, 2012

My big idea is a big question. Every cooperative system contains parasites. How do we ensure that society's parasites don't destroy society's systems?

It's all about trust, really. Not the intimate trust we have in our close friends and relatives, but the more impersonal trust we have in the various people and systems we interact with in society. I trust airline pilots, hotel clerks, ATMs, restaurant kitchens, and the company that built the computer I'm writing this short essay on. I trust that they have acted and will act in the ways I expect them to. This type of trust is more a matter of consistency or predictability than of intimacy.

Of course, all of these systems contain parasites. Most people are naturally trustworthy, but some are not. There are hotel clerks who will steal your credit card information. There are ATMs that have been hacked by criminals. Some restaurant kitchens serve tainted food. There was even an airline pilot who deliberately crashed his Boeing 767 into the Atlantic Ocean in 1999.

My central metaphor is the Prisoner's Dilemma, which nicely exposes the tension between group interest and self-interest. And the dilemma even gives us a terminology to use: cooperators act in the group interest, and defectors act in their own selfish interest, to the detriment of the group. Too many defectors, and everyone suffers—often catastrophically.

The Prisoner's Dilemma is not only useful in describing the problem, but also serves as a way to organize solutions. We humans have developed four basic mechanisms for ways to limit defectors: what I call societal pressure. We use morals, reputation, laws, and security systems. It's all coercion, really, although we don't call it that. I'll spare you the details; it would require a book to explain. And it did.

This book marks another chapter in my career's endless series of generalizations. From mathematical security—cryptography—to computer and network security; from there to security technology in general; then to the economics of

security and the psychology of security; and now to—I suppose—the sociology of security. The more I try to understand how security works, the more of the world I need to encompass within my model.

When I started out writing this book, I thought I'd be talking a lot about the global financial crisis of 2008. It's an excellent example of group interest vs. self-interest, and how a small minority of parasites almost destroyed the planet's financial system. I even had a great quote by former Federal Reserve Chairman Alan Greenspan, where he admitted a "flaw" in his worldview. The exchange, which took place when he was being questioned by Congressman Henry Waxman at a 2008 Congressional hearing, was once the opening paragraphs of my book. I called the defectors "the dishonest minority," which was my original title.

That unifying example eventually faded into the background, to be replaced by a lot of separate examples. I talk about overfishing, childhood immunizations, paying taxes, voting, stealing, airplane security, gay marriage, and a whole lot of other things. I dumped the phrase "dishonest minority" entirely, partly because I didn't need it and partly because a vocal few early readers were reading it not as "the small percentage of us that are dishonest" but as "the minority group that is dishonest"—not at all the meaning I was trying to convey.

I didn't even realize I was talking about trust until most of the way through. It was a couple of early readers who—coincidentally, on the same day—told me my book wasn't about security, it was about trust. More specifically, it was about how different societal pressures, security included, induce trust. This interplay between cooperators and defectors, trust and security, compliance and coercion, affects everything having to do with people.

In the book, I wander through a dizzying array of academic disciplines: experimental psychology, evolutionary psychology, sociology, economics, behavioral economics, evolutionary biology, neuroscience, game theory, systems dynamics, anthropology, archeology, history, political science, law, philosophy, theology, cognitive science, and computer security. It sometimes felt as if I were blundering through a university, kicking down doors and demanding answers. "You anthropologists: what can you tell me about early human transgressions and punishments?" "Okay neuroscientists, what's the brain chemistry of cooperation? And you evolutionary psychologists, how can you explain that?" "Hey philosophers, what have you got?" I downloaded thousands—literally—of academic papers. In pre-Internet days I would have had to move into an academic library.

What's really interesting to me is what this all means for the future. We've never been able to eliminate defections. No matter how much societal pressure we bring to bear, we can't bring the murder rate in society to zero. We'll never see the

end of bad corporate behavior, or embezzlement, or rude people who make cell phone calls in movie theaters. That's fine, but it starts getting interesting when technology makes each individual defection more dangerous. That is, fishermen will survive even if a few of them defect and overfish—until defectors can deploy driftnets and single-handedly collapse the fishing stock. The occasional terrorist with a machine gun isn't a problem for society in the overall scheme of things; but a terrorist with a nuclear weapon could be.

Also—and this is the final kicker—not all defectors are bad. If you think about the notions of cooperating and defecting, they're defined in terms of the societal norm. Cooperators are people who follow the formal or informal rules of society. Defectors are people who, for whatever reason, break the rules. That definition says nothing about the absolute morality of the society or its rules. When society is in the wrong, it's defectors who are in the vanguard for change. So it was defectors who helped escaped slaves in the antebellum American South. It's defectors who are agitating to overthrow repressive regimes in the Middle East. And it's defectors who are fueling the Occupy Wall Street movement. Without defectors, society stagnates.

We simultaneously need more societal pressure to deal with the effects of technology, and less societal pressure to ensure an open, free, and evolving society. This is our big challenge for the coming decade.

High-Tech Cheats in a World of Trust

Originally published in New Scientist, February 27, 2012

I can put my cash card into an ATM anywhere in the world and take out a fistful of local currency, while the corresponding amount is debited from my bank account at home. I don't even think twice: regardless of the country, I trust that the system will work.

The whole world runs on trust. We trust that people on the street won't rob us, that the bank we deposited money in last month returns it this month, that the justice system punishes the guilty and exonerates the innocent. We trust the food we buy won't poison us, and the people we let in to fix our boiler won't murder us.

My career has taken me from cryptography to information security, general security technology to the economics and psychology of security. Most recently, I have become interested in how we induce trustworthy behavior.

Society is, after all, an interdependent system that requires widespread cooperation to function. People need to act in ways that are expected of them, to be consistent and compliant. And not just individuals, but organizations and systems.

But in any cooperative system, there is an alternative, parasitic, strategy available—cheating. A parasite obtains the benefits of widespread cooperation while at the same time taking advantage of it. There are—and always will be—robbers, crooked banks and judges who take bribes. So how do we ensure that the parasites are kept to a small enough minority to not ruin everything for everyone?

Remember the variations on the Prisoner's Dilemma, the game theory scenarios framed by Merrill Flood and Melvin Dresher at the US RAND Corporation think tank in 1950 that show why two individuals might not cooperate, even if it looks to be in their best interest to do so. The paradox is that it is in our collective interest to be trustworthy and cooperate, while it is in our individual self-interest to be parasitic and defect, or cheat. If too many defect, society stops functioning, the crime rate soars, international banking collapses and judicial rulings become available for sale to the highest bidder. No one would trust anyone, because there wouldn't be enough trust to go around.

The way to solve this is to put our thumb on the scales. If we can increase the benefits of cooperation or the costs of defection, we can induce people to act in the group interest—because it is also in their self-interest. In my book *Liars and Outliers*, I call such mechanisms societal pressures. A bank's reputation in the community is a societal pressure. So is the lock on the ATM that keeps criminals out.

This problem isn't new, nor unique to us. Since all complex systems must deal with the problems resulting from parasites it is not surprising that we have a complex interplay of societal pressures. The most basic are moral systems regulating our own behavior, and reputational systems we use to regulate each other's behavior. Most of us try not to treat others unfairly because it makes us feel bad and we know they will treat us badly in return. Most don't steal because we feel guilty—and there are consequences when we are caught. We recognize it is in our long-term self-interest not to act in our immediate self-interest.

Morals and reputation worked well enough for primitive lifestyles, but these began to fail as society grew too large. Trust is personal and intimate among people who know each other, and morals and reputation are easily limited to

an in-group. Institutional systems—laws—formalized reputational systems, and security technologies allowed societal pressure to scale up as we expanded into ever-larger groups.

So my naive trust in ATMs turns out to be based on many complex things: the moral inclinations of most of the people involved in building and operating transfer systems; the fact that a financial institution with a reputation for cheating would probably lose its customers; the myriad banking laws and regulations that exist to punish fraudsters; and knowing that the very different security measures underpinning ATMs, bank transfers and banking will work properly even if some of those involved would prefer to cheat me.

This trust isn't absolute, of course. Not every societal pressure affects everyone equally. Some care more about their reputations, others are naturally law-abiding and still others are better at picking locks. But the goal isn't total compliance, just to limit the scope for defection. Criminals still target ATMs, and the occasional rogue bank employee steals money from accounts. But for the most part, societal pressures keep defector damage small enough to keep the system intact.

But sometimes the scope is too great and underlying systems come crashing down. Overfishing has destroyed breeding stocks in many places. Crime and corruption have devastated some countries. The international banking system almost collapsed in 2008. But in general, societal pressures work as a delicate balance between cooperation and defection: too little societal pressure and the scope of defection becomes too great; too much and security is too costly.

This balance isn't static—technological changes disrupt it all the time. The changes can be related to defecting, so ATM-based “card skimmers” make it easier for criminals to steal my codes and empty my bank account. Or they may be related to security, with computerized auditing technology making it more difficult for fraudulent transactions to go through the system unnoticed. Or they could be unrelated to either: cheap telecoms make it easier to interconnect bank networks globally. Like societal pressures, these things change the Prisoner's Dilemma calculations.

Life becomes dangerously insecure when new technologies, innovations and ideas increase the scope of defection. Defectors innovate. New attacks become possible. Existing attacks become easier, cheaper, more reliable or more devastating. More people may defect, simply because it's easier to. In response, society must also innovate, to reduce the scope of defection and restore the balance. This dynamic is as old as civilization.

Global banking, terrorists with nuclear weapons, genetic engineering, bio-weapons, pandemics: we now have such dangerous systems that a few defectors can wreak havoc so great that reactive rebalancing might not be enough. Worse still, by the time that society realizes that the scope of defection has increased and societal pressures need to be increased, irreversible damage may already have been done.

To add to the complexity, not all defectors are bad. Neither cooperation nor defection relate to any absolute standard of morality. It is defectors who are in the vanguard for change, such as those who helped escaped slaves in the US south before the civil war. It is defectors who agitate to overthrow repressive regimes in the Middle East—and defectors who fuel the Occupy movement. Without them, society stagnates.

How to achieve this balance is at the core of many of our policy debates about the Internet. Anonymity is essential to freedom and liberty and saves the lives of dissidents everywhere. Yet it also protects criminals. Copyright both protects and stifles innovation. And balance is central to debates about air security, terrorism in general, and protecting economies against financial fraud. The big challenge will be to understand how to simultaneously provide both more societal pressure to deal with the threats of technology, and less pressure to ensure an open, free and evolving society.

Detecting Cheaters

Originally published in IEEE Security & Privacy, March/April 2011

Our brains are specially designed to deal with cheating in social exchanges. The evolutionary psychology explanation is that we evolved brain heuristics for the social problems that our prehistoric ancestors had to deal with. Once humans became good at cheating, they then had to become good at detecting cheating—otherwise, the social group would fall apart.

Perhaps the most vivid demonstration of this can be seen with variations on what's known as the Wason selection task, named after the psychologist who first studied it. Back in the 1960s, it was a test of logical reasoning; today, it's used more as a demonstration of evolutionary psychology. But before we get to the experiment, let's get into the mathematical background.

Propositional calculus is a system for deducing conclusions from true premises. It uses variables for statements because the logic works regardless of what the statements are. College courses on the subject are taught by either the mathematics or the philosophy department, and they're not generally considered to be easy classes. Two particular rules of inference are relevant here: *modus ponens* and *modus tollens*. Both allow you to reason from a statement of the form, "if P, then Q." (If Socrates was a man, then Socrates was mortal. If you are to eat dessert, then you must first eat your vegetables. If it is raining, then Gwendolyn had Crunchy Wunchies for breakfast. That sort of thing.) *Modus ponens* goes like this:

If P, then Q. P. Therefore, Q.

In other words, if you assume the conditional rule is true, and if you assume the antecedent of that rule is true, then the consequent is true. So,

If Socrates was a man, then Socrates was mortal. Socrates was a man. Therefore, Socrates was mortal.

Modus tollens is more complicated:

If P, then Q. Not Q. Therefore, not P.

If Socrates was a man, then Socrates was mortal. Socrates was not mortal. Therefore, Socrates was not a man.

This makes sense: if Socrates was not mortal, then he was a demigod or a stone statue or something.

Both are valid forms of logical reasoning. If you know "if P, then Q" and "P," then you know "Q." If you know "if P, then Q" and "not Q," then you know "not P." (The other two similar forms don't work. If you know "if P, then Q" and "Q," you don't know anything about "P." And if you know "if P, then Q" and "not P," then you don't know anything about "Q.")

If I explained this in front of an audience full of normal people, not mathematicians or philosophers, most of them would be lost. Unsurprisingly, they would have trouble either explaining the rules or using them properly. Just ask any grad student who has had to teach a formal logic class; people have trouble with this.

Consider the Wason selection task. Subjects are presented with four cards next to each other on a table. Each card represents a person, with each side

listing some statement about that person. The subject is then given a general rule and asked which cards he would have to turn over to ensure that the four people satisfied that rule. For example, the general rule might be, “If a person travels to Boston, then he or she takes a plane.” The four cards might correspond to travelers and have a destination on one side and a mode of transport on the other. On the side facing the subject, they read: “went to Boston,” “went to New York,” “took a plane,” and “took a car.” Formal logic states that the rule is violated if someone goes to Boston without taking a plane. Translating into propositional calculus, there’s the general rule: if P, then Q. The four cards are “P,” “not P,” “Q,” and “not Q.” To verify that “if P, then Q” is a valid rule, you have to verify *modus ponens* by turning over the “P” card and making sure that the reverse says “Q.” To verify *modus tollens*, you turn over the “not Q” card and make sure that the reverse doesn’t say “P.”

Shifting back to the example, you need to turn over the “went to Boston” card to make sure that person took a plane, and you need to turn over the “took a car” card to make sure that person didn’t go to Boston. You don’t—as many people think—need to turn over the “took a plane” card to see if it says “went to Boston” because you don’t care. The person might have been flying to Boston, New York, San Francisco, or London. The rule only says that people go to Boston fly; it doesn’t break the rule if someone flies elsewhere.

If you’re confused, you aren’t alone. When Wason first did this study, fewer than 10 percent of his subjects got it right. Others replicated the study and got similar results. The best result I’ve seen is “fewer than 25 percent.” Training in formal logic doesn’t seem to help very much. Neither does ensuring that the example is drawn from events and topics with which the subjects are familiar. People are just bad at the Wason selection task. They also tend to only take college logic classes upon requirement.

This isn’t just another “math is hard” story. There’s a point to this. The one variation of this task that people are surprisingly good at getting right is when the rule has to do with cheating and privilege. For example, change the four cards to children in a family—“gets dessert,” “doesn’t get dessert,” “ate vegetables,” and “didn’t eat vegetables”—and change the rule to “If a child gets dessert, he or she ate his or her vegetables.” Many people—65 to 80 percent—get it right immediately. They turn over the “ate dessert” card, making sure the child ate his vegetables, and they turn over the “didn’t eat vegetables” card, making sure the child didn’t get dessert. Another way of saying this is that they turn over the “benefit received” card to make sure the cost was paid. And they turn over the “cost not paid” card to make sure no benefit was received. They look for cheaters.

The difference is startling. Subjects don't need formal logic training. They don't need math or philosophy. When asked to explain their reasoning, they say things like the answer "popped out at them."

Researchers, particularly evolutionary psychologists Leda Cosmides and John Tooby, have run this experiment with a variety of wordings and settings and on a variety of subjects: adults in the US, UK, Germany, Italy, France, and Hong Kong; Ecuadorian schoolchildren; and Shiriari tribesmen in Ecuador. The results are the same: people are bad at the Wason selection task, except when the wording involves cheating.

In the world of propositional calculus, there's absolutely no difference between a rule about traveling to Boston by plane and a rule about eating vegetables to get dessert. But in our brains, there's an enormous difference: the first is an arbitrary rule about the world, and the second is a rule of social exchange. It's of the form "If you take Benefit B, you must first satisfy Requirement R."

Our brains are optimized to detect cheaters in a social exchange. We're good at it. Even as children, we intuitively notice when someone gets a benefit he didn't pay the cost for. Those of us who grew up with a sibling have experienced how the one child not only knew that the other cheated, but felt compelled to announce it to the rest of the family. As adults, we might have learned that life isn't fair, but we still know who among our friends cheats in social exchanges. We know who doesn't pay his or her fair share of a group meal. At an airport, we might not notice the rule "If a plane is flying internationally, then it boards 15 minutes earlier than domestic flights." But we'll certainly notice who breaks the "If you board first, then you must be a first-class passenger" rule.

Lance Armstrong and the Prisoner's Dilemma of Doping in Professional Sports

Originally published in Wired, October 26, 2012

Doping in professional sports is back in the news, as the overwhelming evidence against Lance Armstrong led to his being stripped of his seven Tour de France titles and more. But instead of focusing on the issues of performance-enhancing drugs and whether professional athletes be allowed to take them, I'd like to talk about the security and economic aspects of the issue.

Because drug testing is a *security* issue. Various sports federations around the world do their best to detect illegal doping, and players do their best to evade the tests. It's a classic security arms race: Improvements in detection technologies lead to improvements in drug detection evasion, which in turn spur the development of better detection capabilities. Right now, it seems drugs are winning; in some places, these drug tests are described as “intelligence tests”—if you can't get around them, you don't deserve to play.

But unlike other security arms races, the detectors have the *ability to look into the past*. A laboratory tested Lance Armstrong's urine in 2005 and found traces of the banned substance erythropoietin (EPO). What's interesting is that the urine sample tested was from 1999. Back then, there weren't any good tests for EPO in urine. Today there are, and the lab took a frozen urine sample—who knew that labs save urine samples from athletes?—and tested it. Armstrong was later cleared (the lab procedures were sloppy), but I don't think people understood the real ramifications of the episode: Testing can go back in time.

The ability to test backward has two major effects:

1. While those who develop new performance-enhancing drugs know exactly what sorts of tests anti-doping laboratories are going to run, and can test their drugs' ability to evade detection beforehand—they don't know what sorts of tests will be developed in the future. Athletes can't assume that just because a drug is undetectable today it will remain so years later.
2. Athletes accused of doping based on years-old urine samples have no way of defending themselves. They can't resubmit to testing; it's too late. (Though if I were an athlete worried about this I would regularly deposit urine “in escrow” to gain some ability to contest accusations.)

The Doping Arms Race as Prisoner's Dilemma

The doping arms race will continue because of the incentives: It's a classic Prisoner's Dilemma. Consider for example competing athletes Alice and Bob, who are individually deciding whether to take drugs or not. Alice thinks:

If Bob doesn't take any drugs, then it will be in my best interest to take them. They will give me a performance edge against Bob. I have a better chance of winning.

Similarly, if Bob takes drugs, it's also in my interest to agree to take them. At least that way Bob won't have an advantage over me.

So even though I have no control over what Bob chooses to do, taking drugs gives me the better outcome, regardless of his action.

Unfortunately, Bob goes through exactly the same analysis. As a result, they both take performance-enhancing drugs and neither has the advantage over the other. If they could just trust each other, they could refrain from taking the drugs and maintain the same non-advantage status, without any legal or physical danger.

But competing athletes can't trust each other, and everyone feels he or she has to dope—continuing to search out newer and increasingly undetectable drugs so they can compete. And the arms race continues.

The Ever-Evolving Problem

It's been this way in bicycle racing for decades. In the 1970s, cyclists used corticosteroids and psychostimulants such as Ritalin, and newly developed norepinephrine-dopamine re-uptake inhibitors such as Pemoline. They were banned, and by the end of the decade assays were developed to detect those substances. In the 1980s, athletes turned to newly developed analogs of endogenous substances made possible through recombinant DNA technology, including human growth hormone, testosterone, anabolic steroids, and synthetic human EPO.

Because EPO is a glycoprotein hormone that controls red blood cell production, it acts to increase oxygenation—an effect valued as highly by endurance athletes as it was by people suffering from anemia. EPO use became rampant in cycling and other sports, and continues to be rampant in spite of bans since the early 1990s and the development in the late 1990s of carbon-isotope ratio tests. Such tests are capable of determining whether substances are made naturally by the body, or come from performance-enhancing drugs.

Next came analogs of analogs, such as darbepoetin alfa (Aranesp), a variation on EPO that became commercially available in 2001. It swiftly gained a following among bike racers and other endurance athletes, and a test to detect it soon followed in 2003. Yet another EPO replacement, Mircera, found its way to both the medical and sports markets in 2007, and assays to detect it were developed by 2008.

Norbolethone, first developed in 1966, was resurrected in the late 1990s and marketed as the first designer steroid by an entrepreneurial bodybuilder-turned-chemist intent on evading detection by the doping police. Its fingerprint was traceable by 2002. This scenario was replayed with tetrahydrogestrinone and madol, with assays developed within two years of their introduction into sports. The mid-to-late 2000s have seen an increase in blood doping through blood transfusions used to increase blood oxygen concentrations. This was soon followed by the development of flow cytometry tests to detect it.

The as-yet-unrealized prospect of gene doping has led some regulatory bodies to preemptively ban any non-therapeutic uses of genetic technology in sports. Presumably tests to detect athletes using them will follow.

Testing and Enforcing

Some sports are more vigilant about drug detection than others. European bicycle racing is particularly vigilant; so are the Olympics. This can lead to some perverse outcomes. In at least two instances, positive tests for norandrosterone, a steroid of which traces are found naturally in human urine, have been traced to adulterated supplements consumed by unsuspecting bicycle racers. Another athlete tested positive for benzodiazepine after consuming a Chinese herbal product. The most widely used urine test for EPO has been found to result in false positives in urine collected after strenuous physical exercise, though this conclusion has been hotly contested by the test's developer and others.

The most widely used tests—rapid-screen immunoassays—all too frequently yield false positives in individuals taking routine over-the-counter and prescription pain relievers, and allergy and acid reflux medications. Two days after winning the first British medal in Alpine skiing at the 2002 Winter Games in Salt Lake City, Alain Baxter was forced to return the bronze medal due to a positive test for methamphetamine... resulting from a Vicks Vapor Inhaler.

American professional sports are far more lenient, often trying to give the appearance of vigilance while still allowing athletes to use performance-enhancing drugs. They know that fans want to see beefy linebackers, powerful sluggers, and lightning-fast sprinters. So, with a wink and a nod, American enforcers only test for the easy stuff.

In the end, doping is all about economics. Athletes will continue to dope because the Prisoner's Dilemma forces them to do so. Sports authorities will either improve their detection capabilities or continue to pretend to do so,

because they depend on fans and associated revenues. And as technology continues to improve, professional athletes will become more like deliberately designed racing cars.

Trust and Society

Originally published in the Montréal Review, February 2013

This morning, I flew from Boston to New York. Before that, I woke up in a hotel, trusting everyone on the staff who has a master key. I took a Boston taxi to the airport, trusting not just the taxi driver, but everyone else on the road. At Boston's Logan Airport, I had to trust everyone who worked for the airline, everyone who worked at the airport, and the thousands of other passengers. I also had to trust everyone who came in contact with the food I bought and ate before boarding my plane. In New York, I similarly had to trust everyone at LaGuardia Airport, my New York taxi driver, and the staff at my new hotel—where I am right now, writing this.

If I had to count, I'd guess I easily had to trust a hundred thousand people—and that was all before 10:30 this morning.

Humans are a trusting species. There were 120 people on my plane, almost all of them strangers to each other, and at no point did anyone jump up and attack the person sitting next to them. It's absurd for me to even say it, but if we had been a planeload of chimpanzees, that would have been impossible. Trust is essential for society to function—our civilization would collapse completely without it—and the fact that we don't think about it is a measure of how well that trust works.

Liars and Outliers is a book about trust and society. It's a way of thinking about society, and it's a way of conceptualizing society's problems. It's not a book about why trust is important; lots of people write about that. It's a book about how we induce trust: about how security enables trust. There were a lot of complicated mechanisms in play this morning to ensure that no one mugged me on the street, my taxi driver didn't rob me on the way to the airport, and the plane was staffed with a competent pilot. *Liars and Outliers* is a book about those mechanisms.

Trust is a complicated concept, and the word is overloaded with many meanings. There's a personal and intimate type of trust. When I say that I trust a friend, or a spouse, I'm talking less about their specific actions and

more about them as a person. I have a general reliance that they will behave in a trustworthy manner. In other words, I know their intentions and I trust their actions will be informed by those intentions.

There's also a less intimate, less personal form of trust. When I got into the taxicab this morning, I didn't know the driver's intentions. For all I knew, he could have been a bank robber by night. But I did trust that, for the specific interaction of taxi driver and passenger, he would behave in a trustworthy manner. And he trusted me to do the same. This is a much more limited form of trust. Maybe it's more "confidence" than trust, and maybe the driver's actions are more like "compliance" than trustworthiness, but it's the type of trust that makes society function. I like to call this sort of trust "cooperation."

In today's society, we need to trust not only people, but institutions and systems. It's not so much that I trusted the particular pilot who flew my plane this morning, but the airline that produces well-trained and well-rested pilots according to some schedule. And it's not so much that I trusted the particular taxi driver, but instead the taxi licensing system and overall police system that produced him. Similarly, when I used an ATM this morning—another interesting exercise in trust—it's less that I trusted that particular machine, bank, and service company—but instead that I trusted the national banking system to debit the proper amount from my bank account back home.

Here's how I like to look at it. All complex ecosystems require cooperation. This is true for biological ecosystems, social systems, and sociotechnical systems. Also, in any cooperative system, there also exists an alternative parasitical strategy. Examples include tapeworms in your digestive tract, thieves in a market, spammers on e-mail, and people who refuse to pay their taxes. These parasites can only survive if they're not too successful. That is, if their number gets too large or too powerful, the underlying system collapses. If there are too many tapeworms in your digestive tract, you die and then they die. Too many thieves in a market, and no one visits the market anymore and the thieves starve. Too many people stop paying their taxes, and you get Greece.

This means there is a fundamental tension between cooperation and what I call defection. (Students of game theory will immediately recognize these terms.) It's a tension between us as individuals and us collectively as society. We might individually want each other's stuff, but we're all better off living in a theft-free society. We might individually not want to pay our taxes, but we're all better off if everyone does. And so on. More to the point, we're individually better off if everyone cooperates *except us*. If I can both 1) live in a theft-free

society, and 2) steal your stuff, I personally have the best of both worlds. But if everyone acts that way, society collapses.

Most of us recognize this: that it's not in our long-term best interest to act in our short-term self-interest. But not everyone does. That's why we need mechanisms to induce trust. That's why we need security. And that's what *Liars and Outliers* is about.

How Secure Is the Papal Election?

Originally published in CNN, February 21, 2013

As the College of Cardinals prepares to elect a new pope, security people like me wonder about the process. How does it work, and just how hard would it be to hack the vote?

The rules for papal elections are steeped in tradition. John Paul II last codified them in 1996, and Benedict XVI left the rules largely untouched. The “Universi Dominici Gregis on the Vacancy of the Apostolic See and the Election of the Roman Pontiff” is surprisingly detailed.

Every cardinal younger than 80 is eligible to vote. We expect 117 to be voting. The election takes place in the Sistine Chapel, directed by the church chamberlain. The ballot is entirely paper-based, and all ballot counting is done by hand. Votes are secret, but everything else is open.

First, there's the “pre-scrutiny” phase.

“At least two or three” paper ballots are given to each cardinal, presumably so that a cardinal has extras in case he makes a mistake. Then nine election officials are randomly selected from the cardinals: three “scrutineers,” who count the votes; three “revisers,” who verify the results of the scrutineers; and three “infirmarii,” who collect the votes from those too sick to be in the chapel. Different sets of officials are chosen randomly for each ballot.

Each cardinal, including the nine officials, writes his selection for pope on a rectangular ballot paper “as far as possible in handwriting that cannot be identified as his.” He then folds the paper lengthwise and holds it aloft for everyone to see.

When everyone has written his vote, the “scrutiny” phase of the election begins. The cardinals proceed to the altar one by one. On the altar is a large chalice with a paten—the shallow metal plate used to hold communion wafers

during Mass—resting on top of it. Each cardinal places his folded ballot on the paten. Then he picks up the paten and slides his ballot into the chalice.

If a cardinal cannot walk to the altar, one of the scrutineers—in full view of everyone—does this for him.

If any cardinals are too sick to be in the chapel, the scrutineers give the infirmarii a locked empty box with a slot, and the three infirmarii together collect those votes. If a cardinal is too sick to write, he asks one of the infirmarii to do it for him. The box is opened, and the ballots are placed onto the paten and into the chalice, one at a time.

When all the ballots are in the chalice, the first scrutineer shakes it several times to mix them. Then the third scrutineer transfers the ballots, one by one, from one chalice to another, counting them in the process. If the total number of ballots is not correct, the ballots are burned and everyone votes again.

To count the votes, each ballot is opened, and the vote is read by each scrutineer in turn, the third one aloud. Each scrutineer writes the vote on a tally sheet. This is all done in full view of the cardinals.

The total number of votes cast for each person is written on a separate sheet of paper. Ballots with more than one name (overvotes) are void, and I assume the same is true for ballots with no name written on them (undervotes). Illegible or ambiguous ballots are much more likely, and I presume they are discarded as well.

Then there's the "post-scrutiny" phase. The scrutineers tally the votes and determine whether there's a winner. We're not done yet, though.

The revisers verify the entire process: ballots, tallies, everything. And then the ballots are burned. That's where the smoke comes from: white if a pope has been elected, black if not—the black smoke is created by adding water or a special chemical to the ballots.

Being elected pope requires a two-thirds plus one vote majority. This is where Pope Benedict made a change. Traditionally a two-thirds majority had been required for election. Pope John Paul II changed the rules so that after roughly 12 days of fruitless votes, a simple majority was enough to elect a pope. Benedict reversed this rule.

How hard would this be to hack?

First, the system is entirely manual, making it immune to the sorts of technological attacks that make modern voting systems so risky.

Second, the small group of voters—all of whom know each other—makes it impossible for an outsider to affect the voting in any way. The chapel is cleared and locked before voting. No one is going to dress up as a cardinal and sneak

into the Sistine Chapel. In short, the voter verification process is about as good as you're ever going to find.

A cardinal can't stuff ballots when he votes. The complicated paten-and-chalice ritual ensures that each cardinal votes once—his ballot is visible—and also keeps his hand out of the chalice holding the other votes. Not that they haven't thought about this: The cardinals are in "choir dress" during the voting, which has translucent lace sleeves under a short red cape, making sleight-of-hand tricks much harder. Additionally, the total would be wrong.

The rules anticipate this in another way: "If during the opening of the ballots the scrutineers should discover two ballots folded in such a way that they appear to have been completed by one elector, if these ballots bear the same name, they are counted as one vote; if however they bear two different names, neither vote will be valid; however, in neither of the two cases is the voting session annulled." This surprises me, as if it seems more likely to happen by accident and result in two cardinals' votes not being counted.

Ballots from previous votes are burned, which makes it harder to use one to stuff the ballot box. But there's one wrinkle: "If however a second vote is to take place immediately, the ballots from the first vote will be burned only at the end, together with those from the second vote." I assume that's done so there's only one plume of smoke for the two elections, but it would be more secure to burn each set of ballots before the next round of voting.

The scrutineers are in the best position to modify votes, but it's difficult. The counting is conducted in public, and there are multiple people checking every step. It'd be possible for the first scrutineer, if he were good at sleight of hand, to swap one ballot paper for another before recording it. Or for the third scrutineer to swap ballots during the counting process. Making the ballots large would make these attacks harder. So would controlling the blank ballots better, and only distributing one to each cardinal per vote. Presumably cardinals change their mind more often during the voting process, so distributing extra blank ballots makes sense.

There's so much checking and rechecking that it's just not possible for a scrutineer to misrecord the votes. And since they're chosen randomly for each ballot, the probability of a cabal being selected is extremely low. More interesting would be to try to attack the system of selecting scrutineers, which isn't well-defined in the document. Influencing the selection of scrutineers and revisers seems a necessary first step toward influencing the election.

If there's a weak step, it's the counting of the ballots.

There's no real reason to do a precount, and it gives the scrutineer doing the transfer a chance to swap legitimate ballots with others he previously stuffed up his sleeve. Shaking the chalice to randomize the ballots is smart, but putting the ballots in a wire cage and spinning it around would be more secure—albeit less reverent.

I would also add some kind of white-glove treatment to prevent a scrutineer from hiding a pencil lead or pen tip under his fingernails. Although the requirement to write out the candidate's name in full provides some resistance against this sort of attack.

Probably the biggest risk is complacency. What might seem beautiful in its tradition and ritual during the first ballot could easily become cumbersome and annoying after the twentieth ballot, and there will be a temptation to cut corners to save time. If the Cardinals do that, the election process becomes more vulnerable.

A 1996 change in the process lets the cardinals go back and forth from the chapel to their dorm rooms, instead of being locked in the chapel the whole time, as was done previously. This makes the process slightly less secure but a lot more comfortable.

Of course, one of the *infirmarii* could do what he wanted when transcribing the vote of an infirm cardinal. There's no way to prevent that. If the infirm cardinal were concerned about that but not privacy, he could ask all three *infirmarii* to witness the ballot.

There are also enormous social—religious, actually—disincentives to hacking the vote. The election takes place in a chapel and at an altar. The cardinals swear an oath as they are casting their ballot—further discouragement. The chalice and paten are the implements used to celebrate the Eucharist, the holiest act of the Catholic Church. And the scrutineers are explicitly exhorted not to form any sort of cabal or make any plans to sway the election, under pain of excommunication.

The other major security risk in the process is eavesdropping from the outside world. The election is supposed to be a completely closed process, with nothing communicated to the world except a winner. In today's high-tech world, this is very difficult. The rules explicitly state that the chapel is to be checked for recording and transmission devices “with the help of trustworthy individuals of proven technical ability.” That was a lot easier in 2005 than it will be in 2013.

What are the lessons here?

First, open systems conducted within a known group make voting fraud much harder. Every step of the election process is observed by everyone, and everyone knows everyone, which makes it harder for someone to get away with anything.

Second, small and simple elections are easier to secure. This kind of process works to elect a pope or a club president, but quickly becomes unwieldy for a large-scale election. The only way manual systems could work for a larger group would be through a pyramid-like mechanism, with small groups reporting their manually obtained results up the chain to more central tabulating authorities.

And third: When an election process is left to develop over the course of a couple of thousand years, you end up with something surprisingly good.

The Court of Public Opinion

Originally published in Wired, February 26, 2013

Recently, Elon Musk and the *New York Times* took to Twitter and the Internet to argue the data—and their grievances—over a failed road test and car review. Meanwhile, an Applebee’s server is part of a Change.org petition to get her job back after posting a pastor’s no-tip receipt comment online. And when he wasn’t paid quickly enough, a local Fitness SF web developer rewrote the company’s webpage to air his complaint.

All of these “cases” are seeking their judgments in the court of public opinion. The court of public opinion has a full docket; even brick-and-mortar establishments aren’t immune.

More and more individuals—and companies—are augmenting, even bypassing entirely, traditional legal process hoping to get a more favorable hearing in public.

Every day we have to interact with thousands of strangers, from people we pass on the street to people who touch our food to people we enter short-term business relationships with. Even though most of us don’t have the ability to protect our interests with physical force, we can all be confident when dealing with these strangers because—at least in part—we trust that the legal system will intervene on our behalf in case of a problem. Sometimes that problem involves people who break the rules of society, and the criminal courts deal with them; when the problem is a disagreement between two parties, the civil courts will. Courts are an ancient system of justice, and modern society cannot function without them.

What matters in this system are the facts and the laws. Courts are intended to be impartial and fair in doling out their justice, and societies flourish based on the extent to which we approach this ideal. When courts are unfair—when

judges can be bribed, when the powerful are treated better, when more expensive lawyers produce more favorable outcomes—society is harmed. We become more fearful and less able to trust each other. We are less willing to enter into agreement with strangers, and we spend more effort protecting our own because we don't believe the system is there to back us up.

The court of public opinion is an alternative system of justice. It's very different from the traditional court system: This court is based on reputation, revenge, public shaming, and the whims of the crowd. Having a good story is more important than having the law on your side. Being a sympathetic underdog is more important than being fair. Facts matter, but there are no standards of accuracy. The speed of the Internet exacerbates this; a good story spreads faster than a bunch of facts.

This court delivers *reputational* justice. Arguments are measured in relation to reputation. If one party makes a claim against another that seems plausible, based on both of their reputations, then that claim is likely to be received favorably. If someone makes a claim that clashes with the reputations of the parties, then it's likely to be disbelieved. Reputation is, of course, a commodity, and loss of reputation is the penalty this court imposes. In that respect, it less often recompenses the injured party and more often exacts revenge or retribution. And while those losses may be brutal, the effects are usually short-lived.

The court of public opinion has significant limitations. It works better for revenge and justice than for dispute resolution. It can punish a company for unfairly firing one of its employees or lying in an automobile test drive, but it's less effective at unraveling a complicated patent litigation or navigating a bankruptcy proceeding.

In many ways, this is a return to a medieval notion of “fama,” or reputation. In other ways, it's like mob justice: sometimes benign and beneficial, sometimes terrible (think French Revolution). Trial by public opinion isn't new; remember Rodney King and O.J. Simpson?

Mass media has enabled this system for centuries. But the Internet, and social media in particular, has changed how it's being used.

Now it's being used more deliberately, more often, by more and more powerful entities as a redress mechanism. Perhaps because it's perceived to be more efficient or perhaps because one of the parties feels they can get a more favorable hearing in this new court, but it's being used instead of lawsuits. Instead of a sideshow to actual legal proceedings, it is turning into an alternate system of dispute resolution and justice.

Part of this trend is because the Internet makes taking a case in front of the court of public opinion so much easier. It used to be that the injured party had to convince a traditional media outlet to make his case public; now he can take his case directly to the people. And while it's still a surprise when some cases go viral while others languish in obscurity, it's simply more effective to present your case on Facebook or Twitter.

Another reason is that the traditional court system is increasingly viewed as unfair. Today, money *can* buy justice: not by directly bribing judges, but by hiring better lawyers and forcing the other side to spend more money than they are able to. We know that the courts treat the rich and the poor differently, that corporations can get away with crimes individuals cannot, and that the powerful can lobby to get the specific laws and regulations they want—irrespective of any notions of fairness.

Smart companies have already prepared for battles in the court of public opinion. They've hired policy experts. They've hired firms to monitor Facebook, Twitter, and other Internet venues where these battles originate. They have response strategies and communications plans in place. They've recognized that while this court is very different from the traditional legal system, money and power does count and that there are ways to tip the outcomes in their favor: For example, fake grassroots movements can be just as effective on the Internet as they can in the offline world.

It's time we recognize the court of public opinion for what it is—an alternative crowd-enabled system of justice. We need to start discussing its merits and flaws; we need to understand when it results in justice, and how it can be manipulated by the powerful. We also need to have a frank conversation about the failings of the traditional justice scheme, and why people are motivated to take their grievances to the public. Despite 24-hour PR firms and incident-response plans, this is a court where corporations and governments are at an inherent disadvantage. And because the weak will continue to run ahead of the powerful, those in power will prefer to use the more traditional mechanisms of government: police, courts, and laws.

Social-media researcher Danah Boyd had it right when she wrote here in *Wired*: “In a networked society, who among us gets to decide where the moral boundaries lie? This isn't an easy question and it's at the root of how we, as a society, conceptualize justice.” It's not an easy question, but it's the key question. The moral and ethical issues surrounding the court of public opinion are the real ones, and ones that society will have to tackle in the decades to come.

On Security Awareness Training

Originally published in Dark Reading, March 19, 2013

Should companies spend money on security awareness training for their employees? It's a contentious topic, with respected experts on both sides of the debate. I personally believe that training users in security is generally a waste of time, and that the money can be spent better elsewhere. Moreover, I believe that our industry's focus on training serves to obscure greater failings in security design.

In order to understand my argument, it's useful to look at training's successes and failures. One area where it doesn't work very well is health. We are forever trying to train people to have healthier lifestyles: eat better, exercise more, whatever. And people are forever ignoring the lessons. One basic reason is psychological: we just aren't very good at trading off immediate gratification for long-term benefit. A healthier you is an abstract eventually; sitting in front of the television all afternoon with a McDonald's Super Monster Meal sounds really good *right now*. Similarly, computer security is an abstract benefit that gets in the way of enjoying the Internet. Good practices might protect me from a theoretical attack at some time in the future, but they're a lot of bother right now and I have more fun things to think about. This is the same trick Facebook uses to get people to give away their privacy; no one reads through new privacy policies; it's much easier to just click "OK" and start chatting with your friends. In short: security is never salient.

Another reason health training works poorly is that it's hard to link behaviors with benefits. We can train anyone—even laboratory rats—with a simple reward mechanism: push the button, get a food pellet. But with health, the connection is more abstract. If you're unhealthy, what caused it? It might have been something you did or didn't do years ago, it might have been one of the dozen things you have been doing and not doing for months, or it might have been the genes you were born with. Computer security is a lot like this, too.

Training laypeople in pharmacology also isn't very effective. We expect people to make all sorts of medical decisions at the drugstore, and they're not very good at it. Turns out that it's hard to teach expertise. We can't expect every mother to have the knowledge of a doctor or pharmacist or RN, and we certainly can't expect her to become an expert when most of the advice she's exposed to comes from manufacturers' advertising. In computer security, too, a lot of advice comes from companies with products and services to sell.

One area of health that is a training success is HIV prevention. HIV may be very complicated, but the rules for preventing it are pretty simple. And aside from certain sub-Saharan countries, we have taught people a new model of their health, and have dramatically changed their behavior. This is important: most lay medical expertise stems from folk models of health. Similarly, people have folk models of computer security. Maybe they're right and maybe they're wrong, but they're how people organize their thinking. This points to a possible way that computer security training can succeed. We should stop trying to teach expertise, and pick a few simple metaphors of security and train people to make decisions using those metaphors.

On the other hand, we still have trouble teaching people to wash their hands—even though it's easy, fairly effective, and simple to explain. Notice the difference, though. The risks of catching HIV are huge, and the cause of the security failure is obvious. The risks of not washing your hands are low, and it's not easy to tie the resultant disease to a particular not-washing decision. Computer security is more like hand washing than HIV.

Another area where training works is driving. We trained, either through formal courses or one-on-one tutoring, and passed a government test, to be allowed to drive a car. One reason that works is because driving is a near-term, really cool, obtainable goal. Another reason is even though the technology of driving has changed dramatically over the past century, that complexity has been largely hidden behind a fairly static interface. You might have learned to drive thirty years ago, but that knowledge is still relevant today. On the other hand, password advice from ten years ago isn't relevant today. Can I bank from my browser? Are PDFs safe? Are untrusted networks okay? Is JavaScript good or bad? Are my photos more secure in the cloud or on my own hard drive? The 'interface' we use to interact with computers and the Internet changes all the time, along with best practices for computer security. This makes training a lot harder.

Food safety is my final example. We have a bunch of simple rules—cooking temperatures for meat, expiration dates on refrigerated goods, the three-second rule for food being dropped on the floor—that are mostly right, but often ignored. If we can't get people to follow these rules, what hope do we have for computer security training?

To those who think that training users in security is a good idea, I want to ask: "Have you ever met an actual user?" They're not experts, and we can't expect them to become experts. The threats change constantly, the likelihood of failure is low, and there is enough complexity that it's hard for people to

understand how to connect their behavior to eventual outcomes. So they turn to folk remedies that, while simple, don't really address the threats.

Even if we could invent an effective computer security training program, there's one last problem. HIV prevention training works because affecting what the average person does is valuable. Even if only half the population practices safe sex, those actions dramatically reduce the spread of HIV. But computer security is often only as strong as the weakest link. If four-fifths of company employees learn to choose better passwords, or not to click on dodgy links, one-fifth still get it wrong and the bad guys still get in. As long as we build systems that are vulnerable to the worst case, raising the average case won't make them more secure.

The whole concept of security awareness training demonstrates how the computer industry has failed. We should be designing systems that won't let users choose lousy passwords and don't care what links a user clicks on. We should be designing systems that conform to their folk beliefs of security, rather than forcing them to learn new ones. Microsoft has a great rule about system messages that require the user to make a decision. They should be NEAT: necessary, explained, actionable, and tested. That's how we should be designing security interfaces. And we should be spending money on security training for developers. These are people who can be taught expertise in a fast-changing environment, and this is a situation where raising the average behavior increases the security of the overall system.

If we security engineers do our job right, users will get their awareness training informally and organically, from their colleagues and friends. People will learn the correct folk models of security, and be able to make decisions using them. Then maybe an organization can spend an hour a year reminding their employees what good security means at that organization, both on the computer and off. That makes a whole lot more sense.

Our New Regimes of Trust

Originally published in the SciTech Lawyer, Winter/Spring 2013

Society runs on trust. Over the millennia, we've developed a variety of mechanisms to induce trustworthy behavior in society. These range from a sense of guilt when we cheat, to societal disapproval when we lie, to laws that arrest fraudsters, to door locks and burglar alarms that keep thieves out

of our homes. They're complicated and interrelated, but they tend to keep society humming along.

The information age is transforming our society. We're shifting from evolved social systems to deliberately created socio-technical systems. Instead of having conversations in offices, we use Facebook. Instead of meeting friends, we IM. We shop online. We let various companies and governments collect comprehensive dossiers on our movements, our friendships, and our interests. We let others censor what we see and read. I could go on for pages.

None of this is news to anyone. But what's important, and much harder to predict, are the social changes resulting from these technological changes. With the rapid proliferation of computers—both fixed and mobile—computing devices and in-the-cloud processing, new ways of socialization have emerged. Facebook friends are fundamentally different than in-person friends. IM conversations are fundamentally different than voice conversations. Twitter has no pre-Internet analog. More social changes are coming. These social changes affect trust, and trust affects everything.

This isn't just academic. There has always been a balance in society between the honest and the dishonest, and technology continually upsets that balance. Online banking results in new types of cyberfraud. Facebook posts become evidence in employment and legal disputes. Cell phone location tracking can be used to round up political dissidents. Random blogs and websites become trusted sources, abetting propaganda. Crime has changed: easier impersonation, action at a greater distance, automation, and so on. The more our nation's infrastructure relies on cyberspace, the more vulnerable we are to cyberattack.

Think of this as a *security gap*: the time lag between when the bad guys figure out how to exploit a new technology and when the good guys figure out how to restore society's balance.

Critically, the security gap is larger when there's more technology, and especially in times of rapid technological change. More importantly, it's larger in times of rapid social change due to the increased use of technology. This is our world today. We don't know *how* the proliferation of networked, mobile devices will affect the systems we have in place to enable trust, but we do know it *will* affect them.

Trust is as old as our species. It's something we do naturally, and informally. We don't trust doctors because we've vetted their credentials, but because they sound learned. We don't trust politicians because we've analyzed their positions, but because we generally agree with their political philosophy—or the buzzwords they use. We trust many things because our friends trust them.

It's the same with corporations, government organizations, strangers on the street: this thing that's critical to society's smooth functioning occurs largely through intuition and relationship. Unfortunately, these traditional and low-tech mechanisms are increasingly failing us. Understanding how trust is being, and will be, affected—probably not by predicting, but rather by recognizing effects as quickly as possible—and then deliberately creating mechanisms to induce trustworthiness and enable trust, is the only thing that will enable society to adapt.

If there's anything I've learned in all my years working at the intersection of security and technology, it's that technology is rarely more than a small piece of the solution. People are always the issue and we need to think as broadly as possible about solutions. So while laws are important, they don't work in isolation. Much of our security comes from the informal mechanisms we've evolved over the millennia: systems of morals and reputation.

There will exist new regimes of trust in the information age. They simply must evolve, or society will suffer unpredictably. We have already begun fleshing out such regimes, albeit in an ad hoc manner. It's time for us to deliberately think about how trust works in the information age, and use legal, social, and technological tools to enable this trust. We might get it right by accident, but it'll be a long and ugly iterative process getting there if we do.

4

Privacy and Surveillance

The Myth of the “Transparent Society” —

Originally published in Wired News, March 6, 2008

When I write and speak about privacy, I am regularly confronted with the mutual disclosure argument. Explained in books like David Brin’s *The Transparent Society*, the argument goes something like this: In a world of ubiquitous surveillance, you’ll know all about me, but I will also know all about you. The government will be watching us, but we’ll also be watching the government. This is different than before, but it’s not automatically worse. And because I know your secrets, you can’t use my secrets as a weapon against me.

This might not be everybody’s idea of utopia—and it certainly doesn’t address the inherent value of privacy—but this theory has a glossy appeal, and could easily be mistaken for a way out of the problem of technology’s continuing erosion of privacy. Except it doesn’t work, because it ignores the crucial dissimilarity of power.

You cannot evaluate the value of privacy and disclosure unless you account for the relative power levels of the discloser and the disclosee.

If I disclose information to you, your power with respect to me increases. One way to address this power imbalance is for you to similarly disclose information to me. We both have less privacy, but the balance of power is maintained. But this mechanism fails utterly if you and I have different power levels to begin with.

An example will make this clearer. You’re stopped by a police officer, who demands to see identification. Divulging your identity will give the officer enormous power over you: He or she can search police databases using the

information on your ID; he or she can create a police record attached to your name; he or she can put you on this or that secret terrorist watch list. Asking to see the officer's ID in return gives you no comparable power over him or her. The power imbalance is too great, and mutual disclosure does not make it OK.

You can think of your existing power as the exponent in an equation that determines the value, to you, of more information. The more power you have, the more additional power you derive from the new data.

Another example: When your doctor says "take off your clothes," it makes no sense for you to say, "You first, doc." The two of you are not engaging in an interaction of equals.

This is the principle that should guide decision-makers when they consider installing surveillance cameras or launching data-mining programs. It's not enough to open the efforts to public scrutiny. All aspects of government work best when the relative power between the governors and the governed remains as small as possible—when liberty is high and control is low. Forced openness in government reduces the relative power differential between the two, and is generally good. Forced openness in laypeople increases the relative power, and is generally bad.

Seventeen-year-old Erik Crespo was arrested in 2005 in connection with a shooting in a New York City elevator. There's no question that he committed the shooting; it was captured on surveillance-camera videotape. But he claimed that while being interrogated, Detective Christopher Perino tried to talk him out of getting a lawyer, and told him that he had to sign a confession before he could see a judge.

Perino denied, under oath, that he ever questioned Crespo. But Crespo had received an MP3 player as a Christmas gift, and surreptitiously recorded the questioning. The defense brought a transcript and CD into evidence. Shortly thereafter, the prosecution offered Crespo a better deal than originally proffered (seven years rather than 15). Crespo took the deal, and Perino was separately indicted on charges of perjury.

Without that recording, it was the detective's word against Crespo's. And who would believe a murder suspect over a New York City detective? That power imbalance was reduced only because Crespo was smart enough to press the "record" button on his MP3 player. Why aren't all interrogations recorded? Why don't defendants have the right to those recordings, just as they have the right to an attorney? Police routinely record traffic stops from their squad cars for their own protection; that video record shouldn't stop once the suspect is no longer a threat.

Cameras make sense when trained on police, and in offices where lawmakers meet with lobbyists, and wherever government officials wield power over the people. Open-government laws, giving the public access to government records and meetings of governmental bodies, also make sense. These all foster liberty.

Ubiquitous surveillance programs that affect everyone without probable cause or warrant, like the National Security Agency's warrantless eavesdropping programs or various proposals to monitor everything on the Internet, foster control. And no one is safer in a political system of control.

Our Data, Ourselves

Originally published in Wired News, May 15, 2008

In the information age, we all have a data shadow.

We leave data everywhere we go. It's not just our bank accounts and stock portfolios, or our itemized bills, listing every credit card purchase and telephone call we make. It's automatic road-toll collection systems, supermarket affinity cards, ATMs and so on.

It's also our lives. Our love letters and friendly chat. Our personal e-mails and SMS messages. Our business plans, strategies and offhand conversations. Our political leanings and positions. And this is just the data we interact with. We all have shadow selves living in the data banks of hundreds of corporations' information brokers—information about us that is both surprisingly personal and uncannily complete—except for the errors that you can neither see nor correct.

What happens to our data happens to ourselves.

This shadow self doesn't just sit there: It's constantly touched. It's examined and judged. When we apply for a bank loan, it's our data that determines whether or not we get it. When we try to board an airplane, it's our data that determines how thoroughly we get searched—or whether we get to board at all. If the government wants to investigate us, they're more likely to go through our data than they are to search our homes; for a lot of that data, they don't even need a warrant.

Who controls our data controls our lives.

It's true. Whoever controls our data can decide whether we can get a bank loan, on an airplane or into a country. Or what sort of discount we get from a

merchant, or even how we're treated by customer support. A potential employer can, illegally in the US, examine our medical data and decide whether or not to offer us a job. The police can mine our data and decide whether or not we're a terrorist risk. If a criminal can get hold of enough of our data, he can open credit cards in our names, siphon money out of our investment accounts, even sell our property. Identity theft is the ultimate proof that control of our data means control of our life.

We need to take back our data.

Our data is a part of us. It's intimate and personal, and we have basic rights to it. It should be protected from unwanted touch.

We need a comprehensive data privacy law. This law should protect all information about us, and not be limited merely to financial or health information. It should limit others' ability to buy and sell our information without our knowledge and consent. It should allow us to see information about us held by others, and correct any inaccuracies we find. It should prevent the government from going after our information without judicial oversight. It should enforce data deletion, and limit data collection, where necessary. And we need more than token penalties for deliberate violations.

This is a tall order, and it will take years for us to get there. It's easy to do nothing and let the market take over. But as we see with things like grocery store club cards and click-through privacy policies on websites, most people either don't realize the extent their privacy is being violated or don't have any real choice. And businesses, of course, are more than happy to collect, buy, and sell our most intimate information. But the long-term effects of this on society are toxic; we give up control of ourselves.

The Future of Ephemeral Conversation

Originally published in the Wall Street Journal, November 21, 2008

When he becomes president, Barack Obama will have to give up his BlackBerry. Aides are concerned that his unofficial conversations would become part of the presidential record, subject to subpoena and eventually made public as part of the country's historical record.

This reality of the information age might be particularly stark for the president, but it's no less true for all of us. Conversation used to be ephemeral.

Whether face-to-face or by phone, we could be reasonably sure that what we said disappeared as soon as we said it. Organized crime bosses worried about phone taps and room bugs, but that was the exception. Privacy was just assumed.

This has changed. We chat in e-mail, over SMS and IM, and on social networking websites like Facebook, MySpace, and LiveJournal. We blog and we Twitter. These conversations—with friends, lovers, colleagues, members of our cabinet—are not ephemeral; they leave their own electronic trails.

We know this intellectually, but we haven't truly internalized it. We type on, engrossed in conversation, forgetting we're being recorded and those recordings might come back to haunt us later.

Oliver North learned this, way back in 1987, when messages he thought he had deleted were saved by the White House PROFS system, and then subpoenaed in the Iran-Contra affair. Bill Gates learned this in 1998 when his conversational e-mails were provided to opposing counsel as part of the anti-trust litigation discovery process. Mark Foley learned this in 2006 when his instant messages were saved and made public by the underage men he talked to. Paris Hilton learned this in 2005 when her cell phone account was hacked, and Sarah Palin learned it earlier this year when her Yahoo e-mail account was hacked. Someone in George W. Bush's administration learned this, and millions of e-mails went mysteriously and conveniently missing.

Ephemeral conversation is dying.

Cardinal Richelieu famously said, "If one would give me six lines written by the hand of the most honest man, I would find something in them to have him hanged." When all our ephemeral conversations can be saved for later examination, different rules have to apply. Conversation is not the same thing as correspondence. Words uttered in haste over morning coffee, whether spoken in a coffee shop or thumbed on a Blackberry, are not official pronouncements. Discussions in a meeting, whether held in a boardroom or a chat room, are not the same as answers at a press conference. And privacy isn't just about having something to hide; it has enormous value to democracy, liberty, and our basic humanity.

We can't turn back technology; electronic communications are here to stay and even our voice conversations are threatened. But as technology makes our conversations less ephemeral, we need laws to step in and safeguard ephemeral conversation. We need a comprehensive data privacy law, protecting our data and communications regardless of where it is stored or how it is processed. We need laws forcing companies to keep it private and delete it as soon as it

is no longer needed. Laws requiring ISPs to store e-mails and other personal communications are exactly what we don't need.

Rules pertaining to government need to be different, because of the power differential. Subjecting the president's communications to eventual public review increases liberty because it reduces the government's power with respect to the people. Subjecting our communications to government review decreases liberty because it reduces our power with respect to the government. The president, as well as other members of government, need some ability to converse ephemerally—just as they're allowed to have unrecorded meetings and phone calls—but more of their actions need to be subject to public scrutiny.

But laws can only go so far. Law or no law, when something is made public it's too late. And many of us like having complete records of all our e-mail at our fingertips; it's like our offline brains.

In the end, this is cultural.

The Internet is the greatest generation gap since rock and roll. We're now witnessing one aspect of that generation gap: the younger generation chats digitally, and the older generation treats those chats as written correspondence. Until our CEOs blog, our Congressmen Twitter, and our world leaders send each other LOLcats—until we have a Presidential election where both candidates have a complete history on social networking sites from before they were teenagers— we aren't fully an information age society.

When everyone leaves a public digital trail of their personal thoughts since birth, no one will think twice about it being there. Obama might be on the younger side of the generation gap, but the rules he's operating under were written by the older side. It will take another generation before society's tolerance for digital ephemera changes.

How to Prevent Digital Snooping

Originally published in the Wall Street Journal, December 9, 2008

As the first digital president, Barack Obama is learning the hard way how difficult it can be to maintain privacy in the information age. Earlier this year, his passport file was snooped by contract workers in the State Department. In October, someone at Immigration and Customs Enforcement leaked

information about his aunt's immigration status. And in November, Verizon employees peeked at his cell phone records.

What these three incidents illustrate is not that computerized databases are vulnerable to hacking—we already knew that, and anyway the perpetrators all had legitimate access to the systems they used—but how important audit is as a security measure.

When we think about security, we commonly think about preventive measures: locks to keep burglars out of our homes, bank safes to keep thieves from our money, and airport screeners to keep guns and bombs off airplanes. We might also think of detection and response measures: alarms that go off when burglars pick our locks or dynamite open bank safes, sky marshals on airplanes who respond when a hijacker manages to sneak a gun through airport security. But audit, figuring out who did what after the fact, is often far more important than any of those other three.

Most security against crime comes from audit. Of course we use locks and alarms, but we don't wear bulletproof vests. The police provide for our safety by investigating crimes after the fact and prosecuting the guilty: that's audit.

Audit helps ensure that people don't abuse positions of trust. The cash register, for example, is basically an audit system. Cashiers have to handle the store's money. To ensure they don't skim from the till, the cash register keeps an audit trail of every transaction. The store owner can look at the register totals at the end of the day and make sure the amount of money in the register is the amount that should be there.

The same idea secures us from police abuse, too. The police have enormous power, including the ability to intrude into very intimate aspects of our life in order to solve crimes and keep the peace. This is generally a good thing, but to ensure that the police don't abuse this power, we put in place systems of audit like the warrant process.

The whole NSA warrantless eavesdropping scandal was about this. Some misleadingly painted it as allowing the government to eavesdrop on foreign terrorists, but the government always had that authority. What the government wanted was to not have to submit a warrant, even after the fact, to a secret FISA court. What they wanted was to not be subject to audit.

That would be an incredibly bad idea. Law enforcement systems that don't have good audit features designed in, or are exempt from this sort of audit-based oversight, are much more prone to abuse by those in power—because they can abuse the system without the risk of getting caught. Audit is essential as the NSA increases its domestic spying. And large police databases, like

the FBI Next Generation Identification System, need to have strong audit features built in.

For computerized database systems like that—systems entrusted with other people’s information—audit is a very important security mechanism. Hospitals need to keep databases of very personal health information, and doctors and nurses need to be able to access that information quickly and easily. A good audit record of who accessed what when is the best way to ensure that those trusted with our medical information don’t abuse that trust. It’s the same with IRS records, credit reports, police databases, telephone records—anything personal that someone might want to peek at during the course of his job.

Which brings us back to President Obama. In each of those three examples, someone in a position of trust inappropriately accessed personal information. The difference between how they played out is due to differences in audit. The State Department’s audit worked best; they had alarm systems in place that alerted superiors when Obama’s passport files were accessed and who accessed them. Verizon’s audit mechanisms worked less well; they discovered the inappropriate account access and have narrowed the culprits down to a few people. Audit at Immigration and Customs Enforcement was far less effective; they still don’t know who accessed the information.

Large databases filled with personal information, whether managed by governments or corporations, are an essential aspect of the information age. And they each need to be accessed, for legitimate purposes, by thousands or tens of thousands of people. The only way to ensure those people don’t abuse the power they’re entrusted with is through audit. Without it, we will simply never know who’s peeking at what.

Architecture of Privacy

*Originally published in IEEE Security & Privacy,
January/February 2009*

The Internet isn’t really for us. We’re here at the beginning, stumbling around, just figuring out what it’s good for and how to use it. The Internet is for those born into it, those who have woven it into their lives from the beginning. The Internet is the greatest generation gap since rock and roll, and only our children can hope to understand it.

Larry Lessig famously said that, on the Internet, code is law. Facebook's architecture limits what we can do there, just as gravity limits what we can do on Earth. The 140-character limit on SMSs is as effective as a legal ban on grammar, spelling, and long-winded sentences: KTHXBYE.

As architects of the Internet, we have a special responsibility to our children to build an Internet that future generations will be proud of, one that encompasses basic human rights and values. We do this when we build systems that offer universal access support, open interfaces, and net neutrality, bypass censorship, limit surveillance, fight repression, give people control over their digital presence and digital personas, and foster individual liberty and privacy—especially privacy.

This would all be easier if the choices we made were temporary. But if history is any guide, they're not. Architecture, both physical and virtual, stays around far longer than we intend it to. College campuses built in the 1970s to limit student protests are still standing, as are buildings designed to defend against medieval siege engines. ASCII and TCP/IP aren't going anywhere anytime soon; neither are domain names, email addresses, or HTML. It's been many years, and we still haven't managed to get either DNSSEC or IPV6 deployed. A "just for now" decision can easily remain for decades.

Business and political realities make privacy harder. Some business models depend on walled gardens or invasive digital rights management controls. Other business models depend on collecting and selling personal data. Some countries depend on censorship to enforce morality or keep ideas out, while others depend on surveillance to control their citizens.

The natural tendencies of the Internet make privacy harder. Technology is the friend of intrusive tools. Digital sensors become smaller and more plentiful. More data is collected and stored every year. Privacy isn't something that occurs naturally online, it must be deliberately architected.

Companies that retain personal information put their customers at risk. Security breaches, court orders, and disgruntled employees are just a few of the ways to lose control of data. Good architectures that minimize data collection reduce these risks, just like guardrails on highways prevent more serious accidents when drivers lose control of their vehicles.

We need to be more deliberate. A lot of information-age architecture is about data: what is collected, who controls it, and how it is used. Data is the lifeblood of the information age, but much of it is very personal. We need to design systems that limit unnecessary data collection, give individuals control

over their data, and limit the ability of those in power to use that data for mass surveillance.

Data is the pollution of the information age. It's a byproduct of every computer-mediated interaction; all processes produce it. It stays around forever, unless it's disposed of. It can be recycled, but it has to be done carefully. And, like physical pollution during the early decades of the industrial age, most people completely ignore the problem.

Just as we look back at the beginning of the previous century and shake our heads at how the titans of the industrial age could ignore the pollution they caused, future generations will look back at us—in the early decades of the information age—and judge our architecture, and what we did to foster freedom, liberty, and democracy. Did we build information technologies that protected people's freedoms even during times when society tried to subvert them? Or did we build technologies that could easily be modified to watch and control? History will record our choices.

Privacy in the Age of Persistence

Originally published in BBC News, February 26, 2009

Welcome to the future, where everything about you is saved. A future where your actions are recorded, your movements are tracked, and your conversations are no longer ephemeral. A future brought to you not by some 1984-like dystopia, but by the natural tendencies of computers to produce data.

Data is the pollution of the information age. It's a natural by-product of every computer-mediated interaction. It stays around forever, unless it's disposed of. It is valuable when reused, but it must be done carefully. Otherwise, its after-effects are toxic.

And just as 100 years ago people ignored pollution in our rush to build the Industrial Age, today we're ignoring data in our rush to build the Information Age.

Increasingly, you leave a trail of digital footprints throughout your day. Once you walked into a bookstore and bought a book with cash. Now you visit Amazon, and all of your browsing and purchases are recorded. You used to buy a train ticket with coins; now your electronic fare card is tied to your bank account. Your store affinity cards give you discounts; merchants use the data on them to reveal detailed purchasing patterns.

Data about you is collected when you make a phone call, send an e-mail message, use a credit card, or visit a website. A national ID card will only exacerbate this.

More computerized systems are watching you. Cameras are ubiquitous in some cities, and eventually face recognition technology will be able to identify individuals. Automatic license plate scanners track vehicles in parking lots and cities. Color printers, digital cameras, and some photocopy machines have embedded identification codes. Aerial surveillance is used by cities to find building permit violators and by marketers to learn about home and garden size.

As RFID chips become more common, they'll be tracked, too. Already you can be followed by your cell phone, even if you never make a call. This is wholesale surveillance; not "follow that car," but "follow every car."

Computers are mediating conversation as well. Face-to-face conversations are ephemeral. Years ago, telephone companies might have known who you called and how long you talked, but not what you said. Today you chat in e-mail, by text message, and on social networking sites. You blog and you Twitter. These conversations—with family, friends, and colleagues—can be recorded and stored.

It used to be too expensive to save this data, but computer memory is now cheaper. Computer processing power is cheaper, too; more data is cross-indexed and correlated, and then used for secondary purposes. What was once ephemeral is now permanent.

Who collects and uses this data depends on local laws. In the US, corporations collect, then buy and sell, much of this information for marketing purposes. In Europe, governments collect more of it than corporations. On both continents, law enforcement wants access to as much of it as possible for both investigation and data mining.

Regardless of country, more organizations are collecting, storing, and sharing more of it.

More is coming. Keyboard logging programs and devices can already record everything you type; recording everything you say on your cell phone is only a few years away.

A "life recorder" you can clip to your lapel that'll record everything you see and hear isn't far behind. It'll be sold as a security device, so that no-one can attack you without being recorded. When that happens, will not wearing a life recorder be used as evidence that someone is up to no good, just as prosecutors

today use the fact that someone left his cell phone at home as evidence that he didn't want to be tracked?

You're living in a unique time in history: the technology is here, but it's not yet seamless. Identification checks are common, but you still have to show your ID. Soon it'll happen automatically, either by remotely querying a chip in your wallets or by recognizing your face on camera.

And all those cameras, now visible, will shrink to the point where you won't even see them. Ephemeral conversation will all but disappear, and you'll think it normal. Already your children live much more of their lives in public than you do. Your future has no privacy, not because of some police-state governmental tendencies or corporate malfeasance, but because computers naturally produce data.

Cardinal Richelieu famously said: "If one would give me six lines written by the hand of the most honest man, I would find something in them to have him hanged." When all your words and actions can be saved for later examination, different rules have to apply.

Society works precisely because conversation is ephemeral; because people forget, and because people don't have to justify every word they utter.

Conversation is not the same thing as correspondence. Words uttered in haste over morning coffee, whether spoken in a coffee shop or thumbed on a BlackBerry, are not official correspondence. A data pattern indicating "terrorist tendencies" is no substitute for a real investigation. Being constantly scrutinized undermines our social norms; furthermore, it's creepy. Privacy isn't just about having something to hide; it's a basic right that has enormous value to democracy, liberty, and our humanity.

We're not going to stop the march of technology, just as we cannot un-invent the automobile or the coal furnace. We spent the industrial age relying on fossil fuels that polluted our air and transformed our climate. Now we are working to address the consequences. (While still using said fossil fuels, of course.) This time around, maybe we can be a little more proactive.

Just as we look back at the beginning of the previous century and shake our heads at how people could ignore the pollution they caused, future generations will look back at us—living in the early decades of the information age—and judge our solutions to the proliferation of data.

We must, all of us together, start discussing this major societal change and what it means. And we must work out a way to create a future that our grandchildren will be proud of.

Should We Have an Expectation of Online Privacy?

Originally published in Information Security, May 2009

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

If your data is online, it is not private. Oh, maybe it seems private. Certainly, only you have access to your e-mail. Well, you and your ISP. And the sender's ISP. And any backbone provider who happens to route that mail from the sender to you. And, if you read your personal mail from work, your company. And, if they have taps at the correct points, the NSA and any other sufficiently well-funded government intelligence organization—domestic and international.

You could encrypt your mail, of course, but few of us do that. Most of us now use webmail. The general problem is that, for the most part, your online data is not under your control. Cloud computing and software as a service exacerbate this problem even more.

Your webmail is less under your control than it would be if you downloaded your mail to your computer. If you use Salesforce.com, you're relying on that company to keep your data private. If you use Google Docs, you're relying on Google. This is why the Electronic Privacy Information Center recently filed a complaint with the Federal Trade Commission: many of us are relying on Google's security, but we don't know what it is.

This is new. Twenty years ago, if someone wanted to look through your correspondence, he had to break into your house. Now, he can just break into your ISP. Ten years ago, your voicemail was on an answering machine in your office; now it's on a computer owned by a telephone company. Your financial accounts are on remote websites protected only by passwords; your credit history is collected, stored, and sold by companies you don't even know exist.

And more data is being generated. Lists of books you buy, as well as the books you look at, are stored in the computers of online booksellers. Your affinity card tells your supermarket what foods you like. What were cash transactions are now credit card transactions. What used to be an anonymous coin tossed into a toll booth is now an EZ Pass record of which highway you were on, and when. What used to be a face-to-face chat is now an e-mail, IM, or SMS conversation—or maybe a conversation inside Facebook.

Remember when Facebook recently changed its terms of service to take further control over your data? They can do that whenever they want, you know.

We have no choice but to trust these companies with our security and privacy, even though they have little incentive to protect them. Neither ChoicePoint, Lexis Nexis, Bank of America, nor T-Mobile bears the costs of privacy violations or any resultant identity theft.

This loss of control over our data has other effects, too. Our protections against police abuse have been severely watered down. The courts have ruled that the police can search your data without a warrant, as long as others hold that data. If the police want to read the e-mail on your computer, they need a warrant; but they don't need one to read it from the backup tapes at your ISP.

This isn't a technological problem; it's a legal problem. The courts need to recognize that in the information age, virtual privacy and physical privacy don't have the same boundaries. We should be able to control our own data, regardless of where it is stored. We should be able to make decisions about the security and privacy of that data, and have legal recourse should companies fail to honor those decisions. And just as the Supreme Court eventually ruled that tapping a telephone was a Fourth Amendment search, requiring a warrant—even though it occurred at the phone company switching office and not in the target's home or office—the Supreme Court must recognize that reading personal e-mail at an ISP is no different.

Offhand but On Record

Originally published in the Japan Times, August 19, 2009

Facebook recently made changes to its service agreement in order to make members' data more accessible to other computer users. Amuse, Inc. announced last week that hackers stole credit-card information from about 150,000 clients. Hackers broke into the social network Twitter's system and stole documents.

Your online data is not private. It may seem private, but it's not. Take e-mail, for example. You might be the only person who knows your e-mail password, but you're not the only person who can read your e-mail. Your e-mail provider can read it too—along with anyone he gives access to. That can include any backbone provider who happened to route that mail from the sender to you. In addition, if you read your e-mail from work, various people at your company have access to it, too. And, if they have taps at the correct points, so can the

police, the US National Security Agency, and any other well-funded national intelligence organization—along with any hackers or criminals sufficiently skilled to break into one of these sites.

Think about your Mixi or Facebook site. You're the only one with your password, but lots of other people can read your updates and look at your pictures. Your friends can see a lot of information about you—that's the whole point of these sites—and you don't really know who they share their information with. A lot of your stuff is public by default, and you probably keep it that way. You might respond to quizzes, and who knows where that data goes or who can see it. Workers at Mixi and Facebook can see everything, of course. They also grant access to portions of your data to third parties who want to sell their products to you.

You could set every privacy setting on your Mixi or Facebook site to maximum, but few of us do that—most of us don't even know how. You could encrypt your e-mail, but almost no one does that—and, anyway, that doesn't work with Webmail very easily. Maintaining your privacy is hard, even if you're an expert.

Cloud computing exacerbates this problem. If your company uses software-as-a-service providers such as Salesforce.com, contact management, or MessageLabs e-mail filtering, those companies have access to your data. If you use Google Docs, Google has access to your data. But even if you leave your data in your computer at home, you have to worry about your family or roommates, burglars, police with warrants, and Internet hackers and other criminals as well.

It's not just your online data that is at risk. It's your cell phone data—both the phone numbers you call and who call you, and the SMS messages you send and receive. It's your buying history, sitting in some credit card company's database. It's your medical records. It's the itemized list of everything you buy when you use a card that identifies you.

These risks are new. Twenty years ago, if someone wanted to look through your correspondence, they had to break into your house. Now, they can just break into your ISP. Ten years ago, your voicemail was on an answering machine in your office; now it's on a computer owned by a telephone company. Your financial accounts are on remote Web sites protected only by passwords; your credit history is collected, stored and sold by companies whose names you probably don't even know. Your digital data is no longer under your control.

And more data is being generated. Lists of everything you buy, and everything you look at but choose not to buy, are stored by online merchants both

in Japan and abroad. A record of everything you browse can be stored by your ISP if they choose to. What were cash transactions are now credit card transactions. What used to be a face-to-face chat is now an e-mail, instant message, or SMS conversation—or maybe a conversation within Mixi or Facebook.

Think of the number of people and companies that can know your location. Your cell phone knows where you are. Your air-travel history is stored in various airline databases, and unless you buy your tickets anonymously, your rail travel history is stored in JR's and other databases. Even your credit card company can reconstruct your whereabouts from your purchases.

All these systems are ostensibly private and secure, but many people have legitimate access and even more—such as hackers and criminals—can get illegitimate access. Japan's Personal Information Protection Act provides only some protections and may not apply if the computers that store your information are located in some other country.

Anonymity doesn't help much. Mixi might not know your real name and address, but there are many ways to link your identity to your account. Maybe your e-mail address identifies you or your ISP knows who you are. Your cell phone identifies you and your computer might, too. Use a credit card from your account and that identifies you. True anonymity is very difficult; we regularly identify ourselves online even if we think we do not.

The lesson in all of this is that little we do is ephemeral anymore. We leave electronic audit trails everywhere we go, with everything we do. This won't change: We can't turn back technology. But as technology makes our conversations less ephemeral, we need laws to step in and safeguard our privacy. We need comprehensive data privacy laws, protecting our data and communications regardless of where it is stored or how it is processed. We need laws forcing companies to keep it private and delete it as soon as it is no longer needed, and laws giving us the right to delete our data from third-party sites. And we need international cooperation to ensure that companies cannot flaunt data privacy laws simply by moving themselves offshore.

Laws can only go so far, though. Law or no law, when something is made public, it's too late. And many of us like having complete records of all our e-mail at our fingertips; it's like our offline memory.

In the end, this is a cultural issue.

The Internet is creating the greatest generation gap since rock 'n' roll. We're now witnessing one aspect of that generation gap: The younger generation chats digitally, and the older generation treats those chats as written correspondence. Until our CEOs blog, our Diet members all Twitter, and our world

leaders send each other LOLcats—until we have a national election where all the candidates have a complete history on social networking sites from before they were teenagers—we aren't fully an information age society.

When everyone leaves a public digital trail of their personal thoughts since birth, no one will think twice about it being there. Some of us might be on the younger side of the generation gap, but the rules we're operating under were written by the older side. It will take another generation before our privacy laws catch up with the death of the ephemeral conversation. Until then, we're just going to have to live with this loss of privacy.

Google's and Facebook's Privacy Illusion —

Originally published in Forbes, April 6, 2010

In January, Facebook CEO, Mark Zuckerberg, declared the age of privacy to be over. A month earlier, Google Chief Eric Schmidt expressed a similar sentiment. Add Scott McNealy's and Larry Ellison's comments from a few years earlier, and you've got a whole lot of tech CEOs proclaiming the death of privacy—especially when it comes to young people.

It's just not true. People, including the younger generation, still care about privacy. Yes, they're far more public on the Internet than their parents: writing personal details on Facebook, posting embarrassing photos on Flickr and having intimate conversations on Twitter. But they take steps to protect their privacy and vociferously complain when they feel it violated. They're not technically sophisticated about privacy and make mistakes all the time, but that's mostly the fault of companies and Web sites that try to manipulate them for financial gain.

To the older generation, privacy is about secrecy. And, as the Supreme Court said, once something is no longer secret, it's no longer private. But that's not how privacy works, and it's not how the younger generation thinks about it. Privacy is about control. When your health records are sold to a pharmaceutical company without your permission; when a social-networking site changes your privacy settings to make what used to be visible only to your friends visible to everyone; when the NSA eavesdrops on everyone's e-mail conversations—your loss of control over that information is the issue. We may not mind sharing our personal lives and thoughts, but we want to control how, where and with whom. A privacy failure is a control failure.

People's relationship with privacy is socially complicated. Saliency matters: People are more likely to protect their privacy if they're thinking about it, and less likely to if they're thinking about something else. Social-networking sites know this, constantly reminding people about how much fun it is to share photos and comments and conversations while downplaying the privacy risks. Some sites go even further, deliberately hiding information about how little control—and privacy—users have over their data. We all give up our privacy when we're not thinking about it.

Group behavior matters; we're more likely to expose personal information when our peers are doing it. We object more to losing privacy than we value its return once it's gone. Even if we don't have control over our data, an illusion of control reassures us. And we are poor judges of risk. All sorts of academic research backs up these findings.

Here's the problem: The very companies whose CEOs eulogize privacy make their money by controlling vast amounts of their users' information. Whether through targeted advertising, cross-selling or simply convincing their users to spend more time on their site and sign up their friends, more information shared in more ways, more publicly means more profits. This means these companies are motivated to continually ratchet down the privacy of their services, while at the same time pronouncing privacy erosions as inevitable and giving users the illusion of control.

You can see these forces in play with Google's launch of Buzz. Buzz is a Twitter-like chatting service, and when Google launched it in February, the defaults were set so people would follow the people they corresponded with frequently in Gmail, with the list publicly available. Yes, users could change these options, but—and Google knew this—changing options is hard and most people accept the defaults, especially when they're trying out something new. People were upset that their previously private e-mail contacts list was suddenly public. A Federal Trade Commission commissioner even threatened penalties. And though Google changed its defaults, resentment remained.

Facebook tried a similar control grab when it changed people's default privacy settings last December to make them more public. While users could, in theory, keep their previous settings, it took an effort. Many people just wanted to chat with their friends and clicked through the new defaults without realizing it.

Facebook has a history of this sort of thing. In 2006, it introduced News Feeds, which changed the way people viewed information about their friends.

There was no true privacy change in that users could not see more information than before; the change was in control—or arguably, just in the illusion of control. Still, there was a large uproar. And Facebook is doing it again; last month, the company announced new privacy changes that will make it easier for it to collect location data on users and sell that data to third parties.

With all this privacy erosion, those CEOs may actually be right—but only because they're working to kill privacy. On the Internet, our privacy options are limited to the options those companies give us and how easy they are to find. We have Gmail and Facebook accounts because that's where we socialize these days, and it's hard—especially for the younger generation—to opt out. As long as privacy isn't salient, and as long as these companies are allowed to forcibly change social norms by limiting options, people will increasingly get used to less and less privacy. There's no malice on anyone's part here; it's just market forces in action. If we believe privacy is a social good, something necessary for democracy, liberty and human dignity, then we can't rely on market forces to maintain it. Broad legislation protecting personal privacy by giving people control over their personal data is the only solution.

The Internet: Anonymous Forever

Originally published in Forbes, May 12, 2010

Universal identification is portrayed by some as the holy grail of Internet security. Anonymity is bad, the argument goes; and if we abolish it, we can ensure only the proper people have access to their own information. We'll know who is sending us spam and who is trying to hack into corporate networks. And when there are massive denial-of-service attacks, such as those against Estonia or Georgia or South Korea, we'll know who was responsible and take action accordingly.

The problem is that it won't work. Any design of the Internet must allow for anonymity. Universal identification is impossible. Even attribution—knowing who is responsible for particular Internet packets—is impossible. Attempting to build such a system is futile, and will only give criminals and hackers new ways to hide.

Imagine a magic world in which every Internet packet could be traced to its origin. Even in this world, our Internet security problems wouldn't be solved.

There's a huge gap between proving that a packet came from a particular computer and that a packet was directed by a particular person. This is the exact problem we have with botnets, or pedophiles storing child porn on innocents' computers. In these cases, we know the origins of the DDoS packets and the spam; they're from legitimate machines that have been hacked. Attribution isn't as valuable as you might think.

Implementing an Internet without anonymity is very difficult, and causes its own problems. In order to have perfect attribution, we'd need agencies—real-world organizations—to provide Internet identity credentials based on other identification systems: passports, national identity cards, driver's licenses, whatever. Sloppier identification systems, based on things such as credit cards, are simply too easy to subvert. We have nothing that comes close to this global identification infrastructure. Moreover, centralizing information like this actually hurts security because it makes identity theft that much more profitable a crime.

And realistically, any theoretical ideal Internet would need to allow people access even without their magic credentials. People would still use the Internet at public kiosks and at friends' houses. People would lose their magic Internet tokens just like they lose their driver's licenses and passports today. The legitimate bypass mechanisms would allow even more ways for criminals and hackers to subvert the system.

On top of all this, the magic attribution technology doesn't exist. Bits are bits; they don't come with identity information attached to them. Every software system we've ever invented has been successfully hacked, repeatedly. We simply don't have anywhere near the expertise to build an airtight attribution system.

Not that it really matters. Even if everyone could trace all packets perfectly, to the person or origin and not just the computer, anonymity would still be possible. It would just take one person to set up an anonymity server. If I wanted to send a packet anonymously to someone else, I'd just route it through that server. For even greater anonymity, I could route it through multiple servers. This is called onion routing and, with appropriate cryptography and enough users, it adds anonymity back to any communications system that prohibits it.

Attempts to banish anonymity from the Internet won't affect those savvy enough to bypass it, would cost billions, and would have only a negligible effect on security. What such attempts would do is affect the average user's access to free speech, including those who use the Internet's anonymity to survive: dissidents in Iran, China, and elsewhere.

Mandating universal identity and attribution is the wrong goal. Accept that there will always be anonymous speech on the Internet. Accept that you'll

never truly know where a packet came from. Work on the problems you can solve: software that's secure in the face of whatever packet it receives, identification systems that are secure enough in the face of the risks. We can do far better at these things than we're doing, and they'll do more to improve security than trying to fix insoluble problems.

The whole attribution problem is very similar to the copy-protection/digital-rights-management problem. Just as it's impossible to make specific bits not copyable, it's impossible to know where specific bits came from. Bits are bits. They don't naturally come with restrictions on their use attached to them, and they don't naturally come with author information attached to them. Any attempts to circumvent this limitation will fail, and will increasingly need to be backed up by the sort of real-world police-state measures that the entertainment industry is demanding in order to make copy-protection work. That's how China does it: police, informants, and fear.

Just as the music industry needs to learn that the world of bits requires a different business model, law enforcement and others need to understand that the old ideas of identification don't work on the Internet. For good or for bad, whether you like it or not, there's always going to be anonymity on the Internet.

A Taxonomy of Social Networking Data —

*Originally published in IEEE Security & Privacy,
July/August 2010*

Lately I've been reading about user security and privacy—control, really—on social networking sites. The issues are hard and the solutions harder, but I'm seeing a lot of confusion in even forming the questions. Social networking sites deal with several different types of user data, and it's essential to separate them.

Below is my taxonomy of social networking data, which I first presented at the Internet Governance Forum meeting last November, and again—revised—at an OECD workshop on the role of Internet intermediaries in June.

- Service data is the data you give to a social networking site in order to use it. Such data might include your legal name, your age, and your credit-card number.
- Disclosed data is what you post on your own pages: blog entries, photographs, messages, comments, and so on.

- Entrusted data is what you post on other people's pages. It's basically the same stuff as disclosed data, but the difference is that you don't have control over the data once you post it—another user does.
- Incidental data is what other people post about you: a paragraph about you that someone else writes, a picture of you that someone else takes and posts. Again, it's basically the same stuff as disclosed data, but the difference is that you don't have control over it, and you didn't create it in the first place.
- Behavioral data is data the site collects about your habits by recording what you do and who you do it with. It might include games you play, topics you write about, news articles you access (and what that says about your political leanings), and so on.
- Derived data is data about you that is derived from all the other data. For example, if 80 percent of your friends self-identify as gay, you're likely gay yourself.

There are other ways to look at user data. Some of it you give to the social networking site in confidence, expecting the site to safeguard the data. Some of it you publish openly and others use it to find you. And some of it you share only within an enumerated circle of other users. At the receiving end, social networking sites can monetize all of it: generally by selling targeted advertising.

Different social networking sites give users different rights for each data type. Some are always private, some can be made private, and some are always public. Some can be edited or deleted—I know one site that allows entrusted data to be edited or deleted within a 24-hour period—and some cannot. Some can be viewed and some cannot.

It's also clear that users should have different rights with respect to each data type. We should be allowed to export, change, and delete disclosed data, even if the social networking sites don't want us to. It's less clear what rights we have for entrusted data—and far less clear for incidental data. If you post pictures from a party with me in them, can I demand you remove those pictures—or at least blur out my face? (Go look up the conviction of three Google executives in Italian court over a YouTube video.) And what about behavioral data? It's frequently a critical part of a social networking site's business model. We often don't mind if a site uses it to target advertisements, but are less sanguine when it sells data to third parties.

As we continue our conversations about what sorts of fundamental rights people have with respect to their data, and more countries contemplate regulation on social networking sites and user data, it will be important to keep this taxonomy in mind. The sorts of things that would be suitable for one type of data might be completely unworkable and inappropriate for another.

The Difficulty of Surveillance Crowdsourcing

Originally published in Threatpost, November 8, 2010

Internet Eyes is a UK startup designed to crowdsource digital surveillance. People pay a small fee to become a “Viewer.” Once they do, they can log onto the site and view live anonymous feeds from surveillance cameras at retail stores. If they notice someone shoplifting, they can alert the store owner. Viewers get rated on their ability to differentiate real shoplifting from false alarms, can win 1,000 pounds if they detect the most shoplifting in some time interval, and otherwise get paid a wage that most likely won’t cover their initial fee.

Although the system has some nod towards privacy, groups like Privacy International oppose the system for fostering a culture of citizen spies. More fundamentally, though, I don’t think the system will work. Internet Eyes is primarily relying on voyeurism to compensate its Viewers. But most of what goes on in a retail store is incredibly boring. Some of it is actually voyeuristic, and very little of it is criminal. The incentives just aren’t there for Viewers to do more than peek, and there’s no obvious way to discouraging them from siding with the shoplifter and just watch the scenario unfold.

This isn’t the first time groups have tried to crowdsource surveillance camera monitoring. Texas’s Virtual Border Patrol tried the same thing: deputizing the general public to monitor the Texas-Mexico border. It ran out of money last year, and was widely criticized as a joke.

This system suffered the same problems as Internet Eyes—not enough incentive to do a good job, boredom because crime is the rare exception—as well as the fact that false alarms were very expensive to deal with.

Both of these systems remind me of the one time this idea was conceptualized correctly. Invented in 2003 by my friend and colleague Jay Walker, US HomeGuard also tried to crowdsource surveillance camera monitoring. But this system focused on one very specific security concern: people in no-man's areas. These are areas between fences at nuclear power plants or oil refineries, border zones, areas around dams and reservoirs, and so on: areas where there should never be anyone.

The idea is that people would register to become “spotters.” They would get paid a decent wage (that and patriotism was the incentive), receive a stream of still photos, and be asked a very simple question: “Is there a person or a vehicle in this picture?” If a monitor clicked “yes,” the photo—and the camera—would be referred to whatever professional response the camera owner had set up.

HomeGuard would monitor the monitors in two ways. One, by sending stored, known, photos to people regularly to verify that they were paying attention. And two, by sending live photos to multiple spotters and correlating the results, to many more monitors if a spotter claimed to have spotted a person or vehicle.

Just knowing that there's a person or a vehicle in a no-man's area is only the first step in a useful response, and HomeGuard envisioned a bunch of enhancements to the rest of that system. Flagged photos could be sent to the digital phones of patrolling guards, cameras could be controlled remotely by those guards, and speakers in the cameras could issue warnings. Remote citizen spotters were only useful for that first step, looking for a person or a vehicle in a photo that shouldn't contain any. Only real guards at the site itself could tell an intruder from the occasional maintenance person.

Of course the system isn't perfect. A would-be infiltrator could sneak past the spotters by holding a bush in front of him, or disguising himself as a vending machine. But it does fill in a gap in what fully automated systems can do, at least until image processing and artificial intelligence get significantly better.

HomeGuard never got off the ground. There was never any good data about whether spotters were more effective than motion sensors as a first level of defense. But more importantly, Walker says that the politics surrounding homeland security money post-9/11 was just too great to penetrate, and that as an outsider he couldn't get his ideas heard. Today, probably, the patriotic fervor that gripped so many people post-9/11 has dampened, and he'd probably have to pay his spotters more than he envisioned seven years ago. Still, I thought it was a clever idea then and I still think it's a clever idea—and it's an example of how to do surveillance crowdsourcing correctly.

Making the system more general runs into all sorts of problems. An amateur can spot a person or vehicle pretty easily, but is much harder pressed to notice a shoplifter. The privacy implications of showing random people pictures of no-man's lands is minimal, while a busy store is another matter—stores have enough individuality to be identifiable, as do people. Public photo tagging will even allow the process to be automated. And, of course, the normalization of a spy-on-your-neighbor surveillance society where it's perfectly reasonable to watch each other on cameras just in case one of us does something wrong.

The Internet Is a Surveillance State

Originally published in CNN, March 16, 2013

I'm going to start with three data points.

One: Some of the Chinese military hackers who were implicated in a broad set of attacks against the US government and corporations were identified because they accessed Facebook from the same network infrastructure they used to carry out their attacks.

Two: Hector Monsegur, one of the leaders of the LulzSec hacker movement, was identified and arrested last year by the FBI. Although he practiced good computer security and used an anonymous relay service to protect his identity, he slipped up.

And three: Paula Broadwell, who had an affair with CIA director David Petraeus, similarly took extensive precautions to hide her identity. She never logged in to her anonymous e-mail service from her home network. Instead, she used hotel and other public networks when she e-mailed him. The FBI correlated hotel registration data from several different hotels—and hers was the common name.

The Internet is a surveillance state. Whether we admit it to ourselves or not, and whether we like it or not, we're being tracked all the time. Google tracks us, both on its pages and on other pages it has access to. Facebook does the same; it even tracks non-Facebook users. Apple tracks us on our iPhones and iPads. One reporter used a tool called Collusion to track who was tracking him; 105 companies tracked his Internet use during one 36-hour period.

Increasingly, what we do on the Internet is being combined with other data about us. Unmasking Broadwell's identity involved correlating her Internet

activity with her hotel stays. Everything we do now involves computers, and computers produce data as a natural by-product. Everything is now being saved and correlated, and many big-data companies make money by building up intimate profiles of our lives from a variety of sources.

Facebook, for example, correlates your online behavior with your purchasing habits offline. And there's more. There's location data from your cell phone, there's a record of your movements from closed-circuit TVs.

This is ubiquitous surveillance: All of us being watched, all the time, and that data being stored forever. This is what a surveillance state looks like, and it's efficient beyond the wildest dreams of George Orwell.

Sure, we can take measures to prevent this. We can limit what we search on Google from our iPhones, and instead use computer web browsers that allow us to delete cookies. We can use an alias on Facebook. We can turn our cell phones off and spend cash. But increasingly, none of it matters.

There are simply too many ways to be tracked. The Internet, e-mail, cell phones, web browsers, social networking sites, search engines: these have become necessities, and it's fanciful to expect people to simply refuse to use them just because they don't like the spying, especially since the full extent of such spying is deliberately hidden from us and there are few alternatives being marketed by companies that don't spy.

This isn't something the free market can fix. We consumers have no choice in the matter. All the major companies that provide us with Internet services are interested in tracking us. Visit a website and it will almost certainly know who you are; there are lots of ways to be tracked without cookies. Cell phone companies routinely undo the web's privacy protection. One experiment at Carnegie Mellon took real-time videos of students on campus and was able to identify one-third of them by comparing their photos with publicly available tagged Facebook photos.

Maintaining privacy on the Internet is nearly impossible. If you forget even once to enable your protections, or click on the wrong link, or type the wrong thing, you've permanently attached your name to whatever anonymous service you're using. Monsegur slipped up once, and the FBI got him. If the director of the CIA can't maintain his privacy on the Internet, we've got no hope.

In today's world, governments and corporations are working together to keep things that way. Governments are happy to use the data corporations collect—occasionally demanding that they collect more and save it longer—to spy on us. And corporations are happy to buy data from governments. Together

the powerful spy on the powerless, and they're not going to give up their positions of power, despite what the people want.

Fixing this requires strong government will, but they're just as punch-drunk on data as the corporations. Slap-on-the-wrist fines notwithstanding, no one is agitating for better privacy laws.

So, we're done. Welcome to a world where Google knows exactly what sort of porn you all like, and more about your interests than your spouse does. Welcome to a world where your cell phone company knows exactly where you are all the time. Welcome to the end of private conversations, because increasingly your conversations are conducted by e-mail, text, or social networking sites.

And welcome to a world where all of this, and everything else that you do or is done on a computer, is saved, correlated, studied, passed around from company to company without your knowledge or consent; and where the government accesses it at will without a warrant.

Welcome to an Internet without privacy, and we've ended up here with hardly a fight.

Surveillance and the Internet of Things —

Originally published in the Guardian, May 16, 2013

The Internet has turned into a massive surveillance tool. We're constantly monitored on the Internet by hundreds of companies—both familiar and unfamiliar. Everything we do there is recorded, collected, and collated—sometimes by corporations wanting to sell us stuff and sometimes by governments wanting to keep an eye on us.

Ephemeral conversation is over. Wholesale surveillance is the norm. Maintaining privacy from these powerful entities is basically impossible, and any illusion of privacy we maintain is based either on ignorance or on our unwillingness to accept what's really going on.

It's about to get worse, though. Companies such as Google may know more about your personal interests than your spouse, but so far it's been limited by the fact that these companies only see computer data. And even though your computer habits are increasingly being linked to your offline behavior, it's still only behavior that involves computers.

The Internet of Things refers to a world where much more than our computers and cell phones is Internet-enabled. Soon there will be Internet-connected modules on our cars and home appliances. Internet-enabled medical devices will collect real-time health data about us. There'll be Internet-connected tags on our clothing. In its extreme, *everything* can be connected to the Internet. It's really just a matter of time, as these self-powered wireless-enabled computers become smaller and cheaper.

Lots has been written about the "Internet of Things" and how it will change society for the better. It's true that it will make a lot of wonderful things possible, but the "Internet of Things" will also allow for an even greater amount of surveillance than there is today. The Internet of Things gives the governments and corporations that follow our every move something they don't yet have: eyes and ears.

Soon everything we do, both online and offline, will be recorded and stored forever. The only question remaining is who will have access to all of this information, and under what rules.

We're seeing an initial glimmer of this from how location sensors on your mobile phone are being used to track you. Of course your cell provider needs to know where you are; it can't route your phone calls to your phone otherwise. But most of us broadcast our location information to many other companies whose apps we've installed on our phone. Google Maps certainly, but also a surprising number of app vendors who collect that information. It can be used to determine where you live, where you work, and who you spend time with.

Another early adopter was Nike, whose Nike+ shoes communicate with your iPod or iPhone and track your exercising. More generally, medical devices are starting to be Internet-enabled, collecting and reporting a variety of health data. Wiring appliances to the Internet is one of the pillars of the smart electric grid. Yes, there are huge potential savings associated with the smart grid, but it will also allow power companies—and anyone they decide to sell the data to—to monitor how people move about their house and how they spend their time.

Drones are another "thing" moving onto the Internet. As their price continues to drop and their capabilities increase, they will become a very powerful surveillance tool. Their cameras are powerful enough to see faces clearly, and there are enough tagged photographs on the Internet to identify many of us. We're not yet up to a real-time Google Earth equivalent, but it's not more than a few years away. And drones are just a specific application of CCTV cameras, which have been monitoring us for years, and will increasingly be networked.

Google's Internet-enabled glasses—Google Glass—are another major step down this path of surveillance. Their ability to record both audio and video will bring ubiquitous surveillance to the next level. Once they're common, you might never know when you're being recorded in both audio and video. You might as well assume that everything you do and say will be recorded and saved forever.

In the near term, at least, the sheer volume of data will limit the sorts of conclusions that can be drawn. The invasiveness of these technologies depends on asking the right questions. For example, if a private investigator is watching you in the physical world, she or he might observe odd behavior and investigate further based on that. Such serendipitous observations are harder to achieve when you're filtering databases based on pre-programmed queries. In other words, it's easier to ask questions about what you purchased and where you were than to ask what you did with your purchases and why you went where you did. These analytical limitations also mean that companies like Google and Facebook will benefit more from the Internet of Things than individuals—not only because they have access to more data, but also because they have more sophisticated query technology. And as technology continues to improve, the ability to automatically analyze this massive data stream will improve.

In the longer term, the Internet of Things means ubiquitous surveillance. If an object “knows” you have purchased it, and communicates via either Wi-Fi or the mobile network, then whoever or whatever it is communicating with will know where you are. Your car will know who is in it, who is driving, and what traffic laws that driver is following or ignoring. No need to show ID; your identity will already be known. Store clerks could know your name, address, and income level as soon as you walk through the door. Billboards will tailor ads to you, and record how you respond to them. Fast food restaurants will know what you usually order, and exactly how to entice you to order more. Lots of companies will know whom you spend your days—and nights—with. Facebook will know about any new relationship status before you bother to change it on your profile. And all of this information will all be saved, correlated, and studied. Even now, it feels a lot like science fiction.

Will you know any of this? Will your friends? It depends. Lots of these devices have, and will have, privacy settings. But these settings are remarkable not in how much privacy they afford, but in how much they deny. Access will likely be similar to your browsing habits, your files stored on Dropbox, your searches on Google, and your text messages from your phone. All of your data is saved by those companies—and many others—correlated, and then bought and sold without your knowledge or consent. You'd think that your privacy

settings would keep random strangers from learning everything about you, but it only keeps random strangers who *don't pay for the privilege*—or don't work for the government and have the ability to demand the data. Power is what matters here: you'll be able to keep the powerless from invading your privacy, but you'll have no ability to prevent the powerful from doing it again and again.

Government Secrets and the Need for Whistleblowers

Originally published in the Guardian, June 6, 2013

Yesterday, we learned that the NSA received all calling records from Verizon customers for a three-month period starting in April. That's everything except the voice content: who called who, where they were, how long the call lasted—for millions of people, both Americans and foreigners. This “metadata” allows the government to track the movements of everyone during that period, and build a detailed picture of who talks to whom. It's exactly the same data the Justice Department collected about AP journalists.

The *Guardian* delivered this revelation after receiving a copy of a secret memo about this—presumably from a whistleblower. We don't know if the other phone companies handed data to the NSA too. We don't know if this was a one-off demand or a continuously renewed demand; the order started a few days after the Boston bombers were captured by police.

We don't know a lot about how the government spies on us, but we know some things. We know the FBI has issued tens of thousands of ultra-secret National Security Letters to collect all sorts of data on people—we believe on millions of people—and has been abusing them to spy on cloud-computer users. We know it can collect a wide array of personal data from the Internet without a warrant. We also know that the FBI has been intercepting cell-phone data, all but voice content, for the past 20 years without a warrant, and can use the microphone on some powered-off cell phones as a room bug—presumably only with a warrant.

We know that the NSA has many domestic-surveillance and data-mining programs with codenames like Trailblazer, Stellar Wind, and Ragtime—deliberately using different codenames for similar programs to stymie oversight and conceal what's really going on. We know that the NSA is building an

enormous computer facility in Utah to store all this data, as well as faster computer networks to process it all. We know the US Cyber Command employs 4,000 people.

We know that the DHS is also collecting a massive amount of data on people, and that local police departments are running “fusion centers” to collect and analyze this data, and covering up its failures. This is all part of the militarization of the police.

Remember in 2003, when Congress defunded the decidedly creepy Total Information Awareness program? It didn't die; it just changed names and split into many smaller programs. We know that corporations are doing an enormous amount of spying on behalf of the government: all parts.

We know all of this not because the government is honest and forthcoming, but mostly through three backchannels—inadvertent hints or outright admissions by government officials in hearings and court cases, information gleaned from government documents received under FOIA, and government whistleblowers.

There's much more we don't know, and often what we know is obsolete. We know quite a bit about the NSA's ECHELON program from a 2000 European investigation, and about the DHS's plans for Total Information Awareness from 2002, but much less about how these programs have evolved. We can make inferences about the NSA's Utah facility based on the theoretical amount of data from various sources, the cost of computation, and the power requirements from the facility, but those are rough guesses at best. For a lot of this, we're completely in the dark.

And that's wrong.

The US government is on a secrecy binge. It overclassifies more information than ever. And we learn, again and again, that our government regularly classifies things not because they need to be secret, but because their release would be embarrassing.

Knowing how the government spies on us is important. Not only because so much of it is illegal—or, to be as charitable as possible, based on novel interpretations of the law—but because we have a right to know. Democracy requires an informed citizenry in order to function properly, and transparency and accountability are essential parts of that. That means knowing what our government is doing to us, in our name. That means knowing that the government is operating within the constraints of the law. Otherwise, we're living in a police state.

We need whistleblowers.

Leaking information without getting caught is difficult. It's almost impossible to maintain privacy in the Internet Age. The WikiLeaks platform seems to have been secure—Bradley Manning was caught not because of a technological flaw, but because someone he trusted betrayed him—but the US government seems to have successfully destroyed it as a platform. None of the spin-offs have risen to become viable yet. The *New Yorker* recently unveiled its Strongbox platform for leaking material, which is still new but looks good. This link contains the best advice on how to leak information to the press via phone, email, or the post office. The National Whistleblowers Center has a page on national-security whistleblowers and their rights.

Leaking information is also very dangerous. The Obama administration has embarked on a war on whistleblowers, pursuing them—both legally and through intimidation—further than any previous administration has done. Mark Klein, Thomas Drake, and William Binney have all been persecuted for exposing technical details of our surveillance state. Bradley Manning has been treated cruelly and inhumanly—and possibly tortured—for his more-indiscriminate leaking of State Department secrets.

The Obama administration's actions against the Associated Press, its persecution of Julian Assange, and its unprecedented prosecution of Manning on charges of “aiding the enemy” demonstrate how far it's willing to go to intimidate whistleblowers—as well as the journalists who talk to them.

But whistleblowing is vital, even more broadly than in government spying. It's necessary for good government, and to protect us from abuse of power.

We need details on the full extent of the FBI's spying capabilities. We don't know what information it routinely collects on American citizens, what extra information it collects on those on various watch lists, and what legal justifications it invokes for its actions. We don't know its plans for future data collection. We don't know what scandals and illegal actions—either past or present—are currently being covered up.

We also need information about what data the NSA gathers, either domestically or internationally. We don't know how much it collects surreptitiously, and how much it relies on arrangements with various companies. We don't know how much it uses password cracking to get at encrypted data, and how much it exploits existing system vulnerabilities. We don't know whether it deliberately inserts backdoors into systems it wants to monitor, either with or without the permission of the communications-system vendors.

And we need details about the sorts of analysis the organizations perform. We don't know what they quickly cull at the point of collection, and what they

store for later analysis—and how long they store it. We don't know what sort of database profiling they do, how extensive their CCTV and surveillance-drone analysis is, how much they perform behavioral analysis, or how extensively they trace friends of people on their watch lists.

We don't know how big the US surveillance apparatus is today, either in terms of money and people or in terms of how many people are monitored or how much data is collected. Modern technology makes it possible to monitor vastly more people—yesterday's NSA revelations demonstrate that they could easily surveil *everyone*—than could ever be done manually.

Whistleblowing is the moral response to immoral activity by those in power. What's important here are government programs and methods, not data about individuals. I understand I am asking for people to engage in illegal and dangerous behavior. Do it carefully and do it safely, but—and I am talking directly to you, person working on one of these secret and probably illegal programs—do it.

If you see something, say something. There are many people in the US that will appreciate and admire you.

For the rest of us, we can help by protesting this war on whistleblowers. We need to force our politicians not to punish them—to investigate the abuses and not the messengers—and to ensure that those unjustly persecuted can obtain redress.

Our government is putting its own self-interest ahead of the interests of the country. That needs to change.

Before Prosecuting, Investigate the Government

*Originally published in New York Times Room for Debate blog,
June 11, 2013*

Edward Snowden broke the law by releasing classified information. This isn't under debate; it's something everyone with a security clearance knows. It's written in plain English on the documents you have to sign when you get a security clearance, and it's part of the culture. The law is there for a good reason, and secrecy has an important role in military defense.

But before the Justice Department prosecutes Snowden, there are some other investigations that ought to happen.

We need to determine whether these National Security Agency programs are themselves legal. The administration has successfully barred anyone from bringing a lawsuit challenging these laws, on the grounds of national secrecy. Now that we know those arguments are without merit, it's time for those court challenges.

It's clear that some of the NSA programs exposed by Snowden violate the Constitution and others violate existing laws. Other people have an opposite view. The courts need to decide.

We need to determine whether classifying these programs is legal. Keeping things secret from the people is a very dangerous practice in a democracy, and the government is permitted to do so only under very specific circumstances. Reading the documents leaked so far, I don't see anything that needs to be kept secret. The argument that exposing these documents helps the terrorists doesn't even pass the laugh test; there's nothing here that changes anything any potential terrorist would do or not do. But in any case, now that the documents are public, the courts need to rule on the legality of their secrecy.

And we need to determine how we treat whistleblowers in this country. We have whistleblower protection laws that apply in some cases, particularly when exposing fraud, and other illegal behavior. NSA officials have repeatedly lied about the existence, and details, of these programs to Congress.

Only after all of these legal issues have been resolved should any prosecution of Snowden move forward. Because only then will we know the full extent of what he did, and how much of it is justified.

I believe that history will hail Snowden as a hero—his whistleblowing exposed a surveillance state and a secrecy machine run amok. I'm less optimistic of how the present day will treat him, and hope that the debate right now is less about the man and more about the government he exposed.

5

Psychology of Security

The Security Mindset

Originally published in Wired News, March 20, 2008

UnCLE Milton Industries has been selling ant farms to children since 1956. Some years ago, I remember opening one up with a friend. There were no actual ants included in the box. Instead, there was a card that you filled in with your address, and the company would mail you some ants. My friend expressed surprise that you could get ants sent to you in the mail.

I replied: “What’s really interesting is that these people will send a tube of live ants to anyone you tell them to.”

Security requires a particular mindset. Security professionals—at least the good ones—see the world differently. They can’t walk into a store without noticing how they might shoplift. They can’t use a computer without wondering about the security vulnerabilities. They can’t vote without trying to figure out how to vote twice. They just can’t help it.

SmartWater is a liquid with a unique identifier linked to a particular owner. “The idea is for me to paint this stuff on my valuables as proof of ownership,” I wrote when I first learned about the idea. “I think a better idea would be for me to paint it on *your* valuables, and then call the police.”

Really, we can’t help it.

This kind of thinking is not natural for most people. It’s not natural for engineers. Good engineering involves thinking about how things can be made to work; the security mindset involves thinking about how things can be made to fail. It involves thinking like an attacker, an adversary or a criminal. You don’t have to exploit the vulnerabilities you find, but if you don’t see the world that way, you’ll never notice most security problems.

I've often speculated about how much of this is innate, and how much is teachable. In general, I think it's a particular way of looking at the world, and that it's far easier to teach someone domain expertise—cryptography or software security or safecracking or document forgery—than it is to teach someone a security mindset.

Which is why CSE 484, an undergraduate computer-security course taught this quarter at the University of Washington, is so interesting to watch. Professor Tadayoshi Kohno is trying to teach a security mindset.

You can see the results in the blog the students are keeping. They're encouraged to post security reviews about random things: smart pill boxes, Quiet Care Elder Care monitors, Apple's Time Capsule, GM's OnStar, traffic lights, safe deposit boxes, and dorm-room security.

The most recent one is about an automobile dealership. The poster described how she was able to retrieve her car after service just by giving the attendant her last name. Now any normal car owner would be happy about how easy it was to get her car back, but someone with a security mindset immediately thinks: "Can I really get a car just by knowing the last name of someone whose car is being serviced?"

The rest of the blog post speculates on how someone could steal a car by exploiting this security vulnerability, and whether it makes sense for the dealership to have this lax security. You can quibble with the analysis—I'm curious about the liability that the dealership has, and whether their insurance would cover any losses—but that's all domain expertise. The important point is to notice, and then question, the security in the first place.

The lack of a security mindset explains a lot of bad security out there: voting machines, electronic payment cards, medical devices, ID cards, Internet protocols. The designers are so busy making these systems work that they don't stop to notice how they might fail or be made to fail, and then how those failures might be exploited. Teaching designers a security mindset will go a long way toward making future technological systems more secure.

That part's obvious, but I think the security mindset is beneficial in many more ways. If people can learn how to think outside their narrow focus and see a bigger picture, whether in technology or politics or their everyday lives, they'll be more sophisticated consumers, more skeptical citizens, less gullible people.

If more people had a security mindset, services that compromise privacy wouldn't have such a sizable market share—and Facebook would be totally different. Laptops wouldn't be lost with millions of unencrypted Social Security numbers on them, and we'd all learn a lot fewer security lessons the hard way. The power grid would be more secure. Identity theft would go way down. Medical records would be more private. If people had the security mindset, they wouldn't have tried to look at since they would have realized that they would be caught.

There's nothing magical about this particular university class; anyone can exercise his security mindset simply by trying to look at the world from an attacker's perspective. If I wanted to evade this particular security device, how would I do it? Could I follow the letter of this law but get around the spirit? If the person who wrote this advertisement, essay, article or television documentary were unscrupulous, what could he have done? And then, how can I protect myself from these attacks?

The security mindset is a valuable skill that everyone can benefit from, regardless of career path.

The Difference between Feeling and Reality in Security

Originally published in Wired News, April 3, 2008

Security is both a feeling and a reality, and they're different. You can feel secure even though you're not, and you can be secure even though you don't feel it. There are two different concepts mapped onto the same word—the English language isn't working very well for us here—and it can be hard to know which one we're talking about when we use the word.

There is considerable value in separating out the two concepts: in explaining how the two are different, and understanding when we're referring to one and when the other. There is value as well in recognizing when the two converge, understanding why they diverge, and knowing how they can be made to converge again.

Some fundamentals first. Viewed from the perspective of economics, security is a trade-off. There's no such thing as absolute security, and any security you get has some cost: in money, in convenience, in capabilities, in insecurities somewhere else, whatever. Every time someone makes a decision about security—computer security, community security, national security—he makes a trade-off.

People make these trade-offs as individuals. We all get to decide, individually, if the expense and inconvenience of having a home burglar alarm is worth the security. We all get to decide if wearing a bulletproof vest is worth the cost and tacky appearance. We all get to decide if we're getting our money's worth from the billions of dollars we're spending combating terrorism, and if invading Iraq was the best use of our counterterrorism resources. We might not have the power to *implement* our opinion, but we get to decide if we think it's worth it.

Now we may or may not have the expertise to make those trade-offs intelligently, but we make them anyway. All of us. People have a natural intuition about security trade-offs, and we make them, large and small, dozens of times throughout the day. We can't help it: It's part of being alive.

Imagine a rabbit, sitting in a field eating grass. And he sees a fox. He's going to make a security trade-off: Should he stay or should he flee? Over time, the rabbits that are good at making that trade-off will tend to reproduce, while the rabbits that are bad at it will tend to get eaten or starve.

So, as a successful species on the planet, you'd expect that human beings would be really good at making security trade-offs. Yet, at the same time, we can be hopelessly bad at it. We spend more money on terrorism than the data warrants. We fear flying and choose to drive instead. Why?

The short answer is that people make most trade-offs based on the *feeling* of security and not the reality.

I've written a lot about how people get security trade-offs wrong, and the cognitive biases that cause us to make mistakes. Humans have developed these biases because they make evolutionary sense. And most of the time, they work.

Most of the time—and this is important—our feeling of security matches the reality of security. Certainly, this is true of prehistory. Modern times are harder. Blame technology, blame the media, blame whatever. Our brains are much better optimized for the security trade-offs endemic to living in small family groups in the East African highlands in 100,000 B.C. than to those endemic to living in 2008 New York.

If we make security trade-offs based on the feeling of security rather than the reality, we choose security that makes us *feel* more secure over security

that actually makes us more secure. And that's what governments, companies, family members and everyone else provide. Of course, there are two ways to make people feel more secure. The first is to make people actually more secure and hope they notice. The second is to make people feel more secure without making them actually more secure, and hope they don't notice.

The key here is whether we notice. The feeling and reality of security tend to converge when we take notice, and diverge when we don't. People notice when 1) there are enough positive and negative examples to draw a conclusion, and 2) there isn't too much emotion clouding the issue.

Both elements are important. If someone tries to convince us to spend money on a new type of home burglar alarm, we as society will know pretty quickly if he's got a clever security device or if he's a charlatan; we can monitor crime rates. But if that same person advocates a new national antiterrorism system, and there weren't any terrorist attacks before it was implemented, and there weren't any after it was implemented, how do we know if his system was effective?

People are more likely to realistically assess these incidents if they don't contradict preconceived notions about how the world works. For example: It's obvious that a wall keeps people out, so arguing against building a wall across America's southern border to keep illegal immigrants out is harder to do.

The other thing that matters is agenda. There are lots of people, politicians, companies and so on who deliberately try to manipulate your feeling of security for their own gain. They try to cause fear. They invent threats. They take minor threats and make them major. And when they talk about rare risks with only a few incidents to base an assessment on—terrorism is the big example here—they are more likely to succeed.

Unfortunately, there's no obvious antidote. Information is important. We can't understand security unless we understand it. But that's not enough: Few of us really understand cancer, yet we regularly make security decisions based on its risk. What we do is accept that there are experts who understand the risks of cancer, and trust them to make the security trade-offs for us.

There are some complex feedback loops going on here, between emotion and reason, between reality and our knowledge of it, between feeling and familiarity, and between the understanding of how we reason and feel about security and our analyses and feelings. We're never going to stop making security trade-offs based on the feeling of security, and we're never going to completely prevent those with specific agendas from trying to take care of us. But the more we know, the better trade-offs we'll make.

How the Human Brain Buys Security

*Originally published in IEEE Security & Privacy,
July/August 2008*

People tend to be risk-averse when it comes to gains, and risk-seeking when it comes to losses. If you give people a choice between a \$500 sure gain and a coin-flip chance of a \$1,000 gain, about 75 percent will pick the sure gain. But give people a choice between a \$500 sure loss and a coin-flip chance of a \$1,000 loss, about 75 percent will pick the coin flip.

People don't have a standard mathematical model of risk in their heads. Their trade-offs are more subtle, and result from how our brains have developed. A computer might not see the difference between the two choices—it's simply a measure of how risk-averse you are—but humans do.

This fact might not seem like a big deal, but it overturned standard economic theory when it was first proposed in 1979. It's called "prospect theory," and was developed by Daniel Kahneman and Amos Tversky to explain how people make trade-offs that involve risk.

Evolutionarily, it makes sense. It's a better survival strategy to accept sure small gains rather than risk them for larger ones, and risk larger losses rather than accept smaller ones. Lions, for example, chase young or wounded wildebeest because the investment needed to kill them is lower. Mature and healthy prey would probably be more nutritious, but there's a risk of missing lunch entirely if it gets away. Because animals tend to live on the razor's edge between starvation and reproduction, any loss of food can result in death, and the best option is to risk everything for the chance of no loss at all.

This cognitive bias, demonstrated again and again by many researchers—across ages, genders, cultures, and even species—is so powerful that it can lead to logically inconsistent results. Google "Asian disease experiment" for an almost surreal example. Describing the same policy in two different ways, either as "200 lives saved out of 600" or "400 lives lost out of 600," yields wildly different risk reactions.

Prospect theory explains one of the biggest problems our industry has with selling security: no one actually wants to buy it. Salespeople have long known there are basically two motivations to get people to buy: greed and fear. Either buyers want something—and thus spend to get it—or don't want something, and spend to help prevent it. It's much easier to sell greed than fear.

Security is a fear sell. It's a choice between a small sure loss—the cost of the security product—and a large risky loss—the potential results of an attack on a network. Of course, there's a lot more to the sale. Buyers must be convinced that the product works, and they must understand the threats and the risk that something bad will happen. But all things being equal, buyers would rather take the chance than buy the security.

Sellers know this, even if they don't understand why, and are continually trying to frame security products in positive terms: slogans like “We take care of security so you can focus on your business,” or carefully crafted ROI models that demonstrate how profitable a security purchase can be.

Another option is to push the fear button really hard. Our brains might prefer risky large losses to sure smaller losses, but when we're really scared we'll do almost anything to make that feeling go away. In our industry, we call it FUD—fear, uncertainty, and doubt. We've seen fear alter the political landscape in several countries following the 9/11 terrorist attacks.

The better solution is not to sell security directly, but to include it as part of a more general product or service. Your car comes with safety and security features built in; they're not sold separately. And it should be the same with computers and networks. Vendors need to build security into the products and services that customers actually want. Security is inherently about avoiding a negative, so you can never ignore the cognitive bias embedded so deeply in the human brain. But if you understand it, you have a better chance of overcoming it.

Does Risk Management Make Sense?

Originally published in Information Security, October 2008

This essay appeared as the first half of a point-counterpoint with Marcus Ranum.

We engage in risk management all the time, but it only makes sense if we do it right.

“Risk management” is just a fancy term for the cost-benefit tradeoff associated with any security decision. It's what we do when we react to fear, or try to make ourselves feel secure. It's the fight-or-flight reflex that evolved in primitive fish and remains in all vertebrates. It's instinctual, intuitive and fundamental to life, and one of the brain's primary functions.

Some have hypothesized that humans have a “risk thermostat” that tries to maintain some optimal risk level. It explains why we drive our motorcycles

faster when we wear a helmet, or are more likely to take up smoking during wartime. It's our natural risk management in action.

The problem is our brains are intuitively suited to the sorts of risk management decisions endemic to living in small family groups in the East African highlands in 100,000 BC, and not to living in the New York City of 2008. We make systematic risk management mistakes—miscalculating the probability of rare events, reacting more to stories than data, responding to the feeling of security rather than reality, and making decisions based on irrelevant context. And that risk cockpit of ours? It's not nearly as finely tuned as we might like it to be.

Like a rabbit that responds to an oncoming car with its default predator avoidance behavior—dart left, dart right, dart left, and at the last moment jump—instead of just getting out of the way, our Stone Age intuition doesn't serve us well in a modern technological society. So when we in the security industry use the term “risk management,” we don't want you to do it by trusting your gut. We want you to do risk management consciously and intelligently, to analyze the tradeoff and make the best decision.

This means balancing the costs and benefits of any security decision—buying and installing a new technology, implementing a new procedure or forgoing a common precaution. It means allocating a security budget to mitigate different risks by different amounts. It means buying insurance to transfer some risks to others. It's what businesses do, all the time, about everything. IT security has its own risk management decisions, based on the threats and the technologies.

There's never just one risk, of course, and bad risk management decisions often carry an underlying tradeoff. Terrorism policy in the US. is based more on politics than actual security risk, but the politicians who make these decisions are concerned about the risks of not being re-elected.

Many corporate security decisions are made to mitigate the risk of lawsuits rather than address the risk of any actual security breach. And individuals make risk management decisions that consider not only the risks to the corporation, but the risks to their departments' budgets, and to their careers.

You can't completely remove emotion from risk management decisions, but the best way to keep risk management focused on the data is to formalize the methodology. That's what companies that manage risk for a living—insurance companies, financial trading firms and arbitrageurs—try to do. They try to replace intuition with models, and hunches with mathematics.

The problem in the security world is we often lack the data to do risk management well. Technological risks are complicated and subtle. We don't know how

well our network security will keep the bad guys out, and we don't know the cost to the company if we don't keep them out. And the risks change all the time, making the calculations even harder. But this doesn't mean we shouldn't try.

You can't avoid risk management; it's fundamental to business just as to life. The question is whether you're going to try to use data or whether you're going to just react based on emotions, hunches and anecdotes.

How the Great Conficker Panic Hacked into Human Credulity

Originally published in the Guardian, April 23, 2009

Conficker's April Fool's joke—the huge, menacing build-up and then nothing—is a good case study on how we think about risks, one whose lessons are applicable far outside computer security. Generally, our brains aren't very good at probability and risk analysis. We tend to use cognitive shortcuts instead of thoughtful analysis. This worked fine for the simple risks we encountered for most of our species' existence, but it's less effective against the complex risks society forces us to face today.

We tend to judge the probability of something happening on how easily we can bring examples to mind. It's why people tend to buy earthquake insurance after an earthquake, when the risk is lowest. It's why those of us who have been the victims of a crime tend to fear crime more than those who haven't. And it's why we fear a repeat of 9/11 more than other types of terrorism.

We fear being murdered, kidnapped, raped, and assaulted by strangers, when friends and relatives are far more likely to do those things to us. We worry about plane crashes instead of car crashes, which are far more common. We tend to exaggerate spectacular, strange, and rare events, and downplay more ordinary, familiar, and common ones.

We also respond more to stories than to data. If I show you statistics on crime in New York, you'll probably shrug and continue your vacation planning. But if a close friend gets mugged there, you're more likely to cancel your trip.

And specific stories are more convincing than general ones. That is why we buy more insurance against plane accidents than against travel accidents, or accidents in general. Or why, when surveyed, we are willing to pay more for air travel insurance covering “terrorist acts” than “all possible causes.” That

is why, in experiments, people judge specific scenarios more likely than more general ones, even if the general ones include the specific.

Conficker's 1 April deadline was precisely the sort of event humans tend to overreact to. It's a specific threat, which convinces us that it's credible. It's a specific date, which focuses our fear. Our natural tendency to exaggerate makes it more spectacular, which further increases our fear. Its repetition by the media makes it even easier to bring to mind. As the story becomes more vivid, it becomes more convincing.

The *New York Times* called it an "unthinkable disaster", the television news show *60 Minutes* said it could "disrupt the entire Internet," and we at the *Guardian* warned that it might be a "deadly threat." Naysayers were few and drowned out.

The first of April passed without incident, but Conficker is no less dangerous today. About 2.2m computers worldwide are still infected with Conficker.A and B, and about 1.3m more are infected with the nastier Conficker.C. It's true that on 1 April Conficker.C tried a new trick to update itself, but its authors could have updated the worm using another mechanism any day. In fact, they updated it on 8 April, and can do so again.

And Conficker is just one of many, many dangerous worms being run by criminal organizations. It came with a date and got a lot of press—that 1 April date was more hype than reality—but it's not particularly special. In short, there are many criminal organizations on the Internet using worms and other forms of malware to infect computers. They then use those computers to send spam, commit fraud, and infect more computers. The risks are real and serious. Luckily, keeping your anti-virus software up-to-date and not clicking on strange attachments can keep you pretty secure. Conficker spreads through a Windows vulnerability that was patched in October. You do have automatic update turned on, right?

But people being people, it takes a specific story for us to protect ourselves.

How Science Fiction Writers Can Help, or Hurt, Homeland Security

Originally published in Wired News, June 18, 2009

A couple of years ago, the Department of Homeland Security hired a bunch of science fiction writers to come in for a day and think of ways terrorists could

attack America. If our inability to prevent 9/11 marked a failure of imagination, as some said at the time, then who better than science fiction writers to inject a little imagination into counterterrorism planning?

I discounted the exercise at the time, calling it “embarrassing.” I never thought that 9/11 was a failure of imagination. I thought, and still think, that 9/11 was primarily a confluence of three things: the dual failure of centralized coordination and local control within the FBI, and some lucky breaks on the part of the attackers. More imagination leads to more movie-plot threats—which contributes to overall fear and overestimation of the risks. And that doesn’t help keep us safe at all.

Recently, I read a paper by Magne Jørgensen that provides some insight into why this is so. Titled *More Risk Analysis Can Lead to Increased Over-Optimism and Over-Confidence*, the paper isn’t about terrorism at all. It’s about software projects.

Most software development project plans are overly optimistic, and most planners are overconfident about their overoptimistic plans. Jørgensen studied how risk analysis affected this. He conducted four separate experiments on software engineers, and concluded (though there are lots of caveats in the paper, and more research needs to be done) that performing more risk analysis can make engineers more overoptimistic instead of more realistic.

Potential explanations all come from behavioral economics: cognitive biases that affect how we think and make decisions. (I’ve written about some of these biases and how they affect security decisions, and there’s a great book on the topic as well.)

First, there’s a control bias. We tend to underestimate risks in situations where we are in control, and overestimate risks in situations when we are not in control. Driving versus flying is a common example. This bias becomes stronger with familiarity, involvement and a desire to experience control, all of which increase with increased risk analysis. So the more risk analysis, the greater the control bias, and the greater the underestimation of risk.

The second explanation is the availability heuristic. Basically, we judge the importance or likelihood of something happening by the ease of bringing instances of that thing to mind. So we tend to overestimate the probability of a rare risk that is seen in a news headline, because it is so easy to imagine. Likewise, we underestimate the probability of things occurring that don’t happen to be in the news. A corollary of this phenomenon is that, if we’re asked to think about a series of things, we overestimate the probability of the last thing thought about because it’s more easily remembered.

According to Jørgensen's reasoning, people tend to do software risk analysis by thinking of the severe risks first, and then the more manageable risks. So the more risk analysis that's done, the less severe the last risk imagined, and thus the greater the underestimation of the total risk.

The third explanation is similar: the peak end rule. When thinking about a total experience, people tend to place too much weight on the last part of the experience. In one experiment, people had to hold their hands under cold water for one minute. Then, they had to hold their hands under cold water for one minute again, then keep their hands in the water for an additional 30 seconds while the temperature was gradually raised. When asked about it afterwards, most people preferred the second option to the first, even though the second had more total discomfort. (An intrusive medical device was redesigned along these lines, resulting in a longer period of discomfort but a relatively comfortable final few seconds. People liked it a lot better.) This means, like the second explanation, that the least severe last risk imagined gets greater weight than it deserves.

Fascinating stuff. But the biases produce the reverse effect when it comes to movie-plot threats. The more you think about far-fetched terrorism possibilities, the more outlandish and scary they become, and the less control you think you have. This causes us to overestimate the risks.

Think about this in the context of terrorism. If you're asked to come up with threats, you'll think of the significant ones first. If you're pushed to find more, if you hire science-fiction writers to dream them up, you'll quickly get into the low-probability movie plot threats. But since they're the last ones generated, they're more available. (They're also more vivid—science fiction writers are good at that—which also leads us to overestimate their probability.) They also suggest we're even less in control of the situation than we believed. Spending too much time imagining disaster scenarios leads people to overestimate the risks of disaster.

I'm sure there's also an anchoring effect in operation. This is another cognitive bias, where people's numerical estimates of things are affected by numbers they've most recently thought about, even random ones. People who are given a list of three risks will think the total number of risks are lower than people who are given a list of 12 risks. So if the science fiction writers come up with 137 risks, people will believe that the number of risks is higher than they otherwise would—even if they recognize the 137 number is absurd.

Jørgensen does not believe risk analysis is useless in software projects, and I don't believe scenario brainstorming is useless in counterterrorism. Both can lead to new insights and, as a result, a more intelligent analysis of both specific risks and general risk. But an over-reliance on either can be detrimental.

Last month, at the 2009 Homeland Security Science & Technology Stakeholders Conference in Washington D.C., science fiction writers helped the attendees think differently about security. This seems like a far better use of their talents than imagining some of the zillions of ways terrorists can attack America.

Privacy Salience and Social Networking Sites

Originally published in the Guardian, July 15, 2009

Reassuring people about privacy makes them more, not less, concerned. It's called "privacy salience," and Leslie John, Alessandro Acquisti, and George Loewenstein—all at Carnegie Mellon University—demonstrated this in a series of clever experiments. In one, subjects completed an online survey consisting of a series of questions about their academic behavior—"Have you ever cheated on an exam?" for example. Half of the subjects were first required to sign a consent warning—designed to make privacy concerns more salient—while the other half did not. Also, subjects were randomly assigned to receive either a privacy confidentiality assurance, or no such assurance. When the privacy concern was made salient (through the consent warning), people reacted negatively to the subsequent confidentiality assurance and were less likely to reveal personal information.

In another experiment, subjects completed an online survey where they were asked a series of personal questions, such as "Have you ever tried cocaine?" Half of the subjects completed a frivolous-looking survey—"How BAD are U??"—with a picture of a cute devil. The other half completed the same survey with the title "Carnegie Mellon University Survey of Ethical Standards," complete with a university seal and official privacy assurances. The results showed that people who were reminded about privacy were less likely to reveal personal information than those who were not.

Privacy salience does a lot to explain social networking sites and their attitudes towards privacy. From a business perspective, social networking sites don't want their members to exercise their privacy rights very much. They want members to be comfortable disclosing a lot of data about themselves.

Joseph Bonneau and Soeren Preibusch of Cambridge University have been studying privacy on 45 popular social networking sites around the world. (You may not have realized that there *are* 45 popular social networking sites around the world.) They found that privacy settings were often confusing and hard to access; Facebook, with its 61 privacy settings, is the worst. To understand some of the settings, they had to create accounts with different settings so they could compare the results. Privacy tends to increase with the age and popularity of a site. General-use sites tend to have more privacy features than niche sites.

But their most interesting finding was that sites consistently hide any mentions of privacy. Their splash pages talk about connecting with friends, meeting new people, sharing pictures: the benefits of disclosing personal data.

These sites do talk about privacy, but only on hard-to-find privacy policy pages. There, the sites give strong reassurances about their privacy controls and the safety of data members choose to disclose on the site. There, the sites display third-party privacy seals and other icons designed to assuage any fears members have.

It's the Carnegie Mellon experimental result in the real world. Users care about privacy, but don't really think about it day to day. The social networking sites don't want to remind users about privacy, even if they talk about it positively, because any reminder will result in users remembering their privacy fears and becoming more cautious about sharing personal data. But the sites also need to reassure those "privacy fundamentalists" for whom privacy is always salient, so they have very strong pro-privacy rhetoric for those who take the time to search them out. The two different marketing messages are for two different audiences.

Social networking sites are improving their privacy controls as a result of public pressure. At the same time, there is a counterbalancing business pressure to decrease privacy; watch what's going on right now on Facebook, for example. Naively, we should expect companies to make their privacy policies clear to allow customers to make an informed choice. But the marketing need to reduce privacy salience will frustrate market solutions to improve privacy; sites would much rather obfuscate the issue than compete on it as a feature.

Security, Group Size, and the Human Brain

*Originally published in IEEE Security & Privacy,
July/August 2009*

If the size of your company grows past 150 people, it's time to get name badges. It's not that larger groups are somehow less secure, it's just that 150 is the cognitive limit to the number of people a human brain can maintain a coherent social relationship with.

Primatologist Robin Dunbar derived this number by comparing neocortex—the “thinking” part of the mammalian brain—volume with the size of primate social groups. By analyzing data from 38 primate genera and extrapolating to the human neocortex size, he predicted a human “mean group size” of roughly 150.

This number appears regularly in human society; it's the estimated size of a Neolithic farming village, the size at which Hittite settlements split, and the basic unit in professional armies from Roman times to the present day. Larger group sizes aren't as stable because their members don't know each other well enough. Instead of thinking of the members as people, we think of them as groups of people. For such groups to function well, they need externally imposed structure, such as name badges.

Of course, badges aren't the only way to determine in-group/out-group status. Other markers include insignia, uniforms, and secret handshakes. They have different security properties and some make more sense than others at different levels of technology, but once a group reaches 150 people, it has to do something.

More generally, there are several layers of natural human group size that increase with a ratio of approximately three: 5, 15, 50, 150, 500, and 1500—although, really, the numbers aren't as precise as all that, and groups that are less focused on survival tend to be smaller. The layers relate to both the intensity and intimacy of relationship and the frequency of contact.

The smallest, three to five, is a “clique”: the number of people from whom you would seek help in times of severe emotional distress. The twelve to 20

group is the “sympathy group”: people with which you have special ties. After that, 30 to 50 is the typical size of hunter-gatherer overnight camps, generally drawn from the same pool of 150 people. No matter what size company you work for, there are only about 150 people you consider to be “co-workers.” (In small companies, Alice and Bob handle accounting. In larger companies, it’s the accounting department—and maybe you know someone there personally.) The 500-person group is the “megaband,” and the 1,500-person group is the “tribe.” Fifteen hundred is roughly the number of faces we can put names to, and the typical size of a hunter-gatherer society.

These numbers are reflected in military organization throughout history: squads of 10 to 15 organized into platoons of three to four squads, organized into companies of three to four platoons, organized into battalions of three to four companies, organized into regiments of three to four battalions, organized into divisions of two to three regiments, and organized into corps of two to three divisions.

Coherence can become a real problem once organizations get above about 150 in size. So as group sizes grow across these boundaries, they have more externally imposed infrastructure—and more formalized security systems. In intimate groups, pretty much all security is ad hoc. Companies smaller than 150 don’t bother with name badges; companies greater than 500 hire a guard to sit in the lobby and check badges. The military have had centuries of experience with this under rather trying circumstances, but even there the real commitment and bonding invariably occurs at the company level. Above that you need to have rank imposed by discipline.

The whole brain-size comparison might be bunk, and a lot of evolutionary psychologists disagree with it. But certainly security systems become more formalized as groups grow larger and their members less known to each other. When do more formal dispute resolution systems arise: town elders, magistrates, judges? At what size boundary are formal authentication schemes required? Small companies can get by without the internal forms, memos, and procedures that large companies require; when does what tend to appear? How does punishment formalize as group size increases? And how do all these things affect group coherence? People act differently on social networking sites like Facebook when their list of “friends” grows larger and less intimate. Local merchants sometimes let known regulars run up tabs. I lend books to friends with much less formality than a public library. What examples have you seen?

People Understand Risks—But Do Security Staff Understand People?

Originally published in the Guardian, August 5, 2009

People have a natural intuition about risk, and in many ways it's very good. It fails at times due to a variety of cognitive biases, but for normal risks that people regularly encounter, it works surprisingly well: often better than we give it credit for.

This struck me as I listened to yet another conference presenter complaining about security awareness training. He was talking about the difficulty of getting employees at his company to actually follow his security policies: encrypting data on memory sticks, not sharing passwords, not logging in from untrusted wireless networks. "We have to make people understand the risks," he said.

It seems to me that his co-workers understand the risks better than he does. They know what the real risks are at work, and that they all revolve around not getting the job done. Those risks are real and tangible, and employees feel them all the time. The risks of not following security procedures are much less real. Maybe the employee will get caught, but probably not. And even if he does get caught, the penalties aren't serious.

Given this accurate risk analysis, any rational employee will regularly circumvent security to get his or her job done. That's what the company rewards, and that's what the company actually wants.

"Fire someone who breaks security procedure, quickly and publicly," I suggested to the presenter. "That'll increase security awareness faster than any of your posters or lectures or newsletters." If the risks are real, people will get it.

You see the same sort of risk intuition on motorways. People are less careful about posted speed limits than they are about the actual speeds police issue tickets for. It's also true on the streets: people respond to real crime rates, not public officials proclaiming that a neighborhood is safe.

The warning stickers on ladders might make you think the things are considerably riskier than they are, but people have a good intuition about ladders and ignore most of the warnings. (This isn't to say that some people don't do stupid things around ladders, but for the most part they're safe. The warnings are more about the risk of lawsuits to ladder manufacturers than risks to people who climb ladders.)

As a species, we are naturally tuned in to the risks inherent in our environment. Throughout our evolution, our survival depended on making reasonably accurate risk management decisions intuitively, and we're so good at it, we don't even realize we're doing it.

Parents know this. Children have surprisingly perceptive risk intuition. They know when parents are serious about a threat and when their threats are empty. And they respond to the real risks of parental punishment, not the inflated risks based on parental rhetoric. Again, awareness training lectures don't work; there have to be real consequences.

It gets even weirder. The University College London professor John Adams popularized the metaphor of a mental risk thermostat. We tend to seek some natural level of risk, and if something becomes less risky, we tend to make it more risky. Motorcycle riders who wear helmets drive faster than riders who don't.

Our risk thermostats aren't perfect (that newly helmeted motorcycle rider will still decrease his overall risk) and will tend to remain within the same domain (he might drive faster, but he won't increase his risk by taking up smoking), but in general, people demonstrate an innate and finely tuned ability to understand and respond to risks.

Of course, our risk intuition fails spectacularly and often, with regards to rare risks, unknown risks, voluntary risks, and so on. But when it comes to the common risks we face every day—the kinds of risks our evolutionary survival depended on—we're pretty good.

So whenever you see someone in a situation who you think doesn't understand the risks, stop first and make sure you understand the risks. You might be surprised.

Nature's Fears Extend to Online Behavior —

Originally published in the Japan Times, November 3, 2009

It's hard work being prey. Watch the birds at a feeder. They're constantly on alert, and will fly away from food—from easy nutrition—at the slightest movement or sound. Given that I've never, ever seen a bird plucked from a feeder by a predator, it seems like a whole lot of wasted effort against a small threat.

Assessing and reacting to risk is one of the most important things a living creature has to deal with. The amygdala, an ancient part of the brain that first

evolved in primitive fishes, has that job. It's what's responsible for the fight-or-flight reflex. Adrenaline in the bloodstream, increased heart rate, increased muscle tension, sweaty palms; that's the amygdala in action. You notice it when you fear a dark alley, have vague fears of terrorism, or worry about predators stalking your children on the Internet. And it works fast, faster than consciousnesses: show someone a snake and their amygdala will react before their conscious brain registers that they're looking at a snake.

Fear motivates all sorts of animal behaviors. Schooling, flocking, and herding are all security measures. Not only is it less likely that any member of the group will be eaten, but each member of the group has to spend less time watching out for predators. Animals as diverse as bumblebees and monkeys both avoid food in areas where predators are common. Different prey species have developed various alarm calls, some surprisingly specific. And some prey species have even evolved to react to the alarms given off by other species.

Evolutionary biologist Randolph Nesse has studied animal defenses, particularly those that seem to be overreactions. These defenses are mostly all-or-nothing; a creature can't do them halfway. Birds flying off, sea cucumbers expelling their stomachs, and vomiting are examples. Using signal-detection theory, Nesse showed that all-or-nothing defenses are expected to have many false alarms. "The smoke detector principle shows that the overresponsiveness of many defenses is an illusion. The defenses appear overresponsive because they are 'inexpensive' compared to the harms they protect against and because errors of too little defense are often more costly than errors of too much defense." So, according to the theory, if flight costs 100 calories, both in flying and lost eating time, and there's a 1 in 100 chance of being eaten if you don't fly away, it's smarter for survival to use up 10,000 calories repeatedly flying at the slightest movement even though there's a 99 percent false-alarm rate. Whatever the numbers happen to be for a particular species, it has evolved to get the trade-off right.

This makes sense, until the conditions that the species evolved under change quicker than evolution can react to. Even though there are far fewer predators in the city, birds at my feeder react as if they were in the primeval forest. Even birds safe in a zoo's aviary don't realize that the situation has changed.

Humans are both no different and very different. We, too, feel fear and react with our amygdala, but we also have a conscious brain that can override those reactions. And we too live in a world very different from the one we evolved in. Our reflexive defenses might be optimized for the risks endemic to living in

small family groups in the East African highlands in 100,000 B.C.—not Tokyo in 2009. But we can go beyond fear, and actually think sensibly about security.

Far too often, we don't. We tend to be poor judges of risk. We overreact to rare risks, we ignore long-term risks, we magnify risks that are also morally offensive. We get risks wrong—threats, probabilities and costs—all the time. When we're afraid, really afraid, we'll do almost anything to make that fear go away. Politicians and marketers, both, have learned to push that fear button to get us to do what they want.

One night last month, I was woken from my hotel-room sleep by a loud, piercing alarm. There was no way I could ignore it, but I weighed the risks and did what any reasonable person would do under the circumstances: I stayed in bed and waited for the alarm to be turned off. No point getting dressed, walking down 10 flights of stairs, and going outside into the cold for what invariably would be a false alarm—serious hotel fires are very rare. Unlike the bird in an aviary, I knew better.

You can disagree with my risk calculus, and I'm sure many hotel guests walked downstairs and outside to the designated assembly point. But it's important to recognize that the ability to have this sort of discussion is uniquely human. And we need to have the discussion repeatedly, whether the topic is the monitoring of our children's Web-surfing habits, outsourcing our corporate IT infrastructure, or even the potential military invasion of another country. These things aren't part of our evolutionary history; we have no natural sense of how to respond to them. Our fears are often calibrated wrong, and reason is the only way we can override them.

6

Security and Technology

The Ethics of Vulnerability Research —————

Originally published in Information Security, May 2008

This was originally published as the first half of a point/counterpoint with Marcus Ranum.

The standard way to take control of someone else's computer is by exploiting a vulnerability in a software program on it. This was true in the 1960s when buffer overflows were first exploited to attack computers. It was true in 1988 when the Morris worm exploited a Unix vulnerability to attack computers on the Internet, and it's still how most modern malware works.

Vulnerabilities are software mistakes—mistakes in specification and design, but mostly mistakes in programming. Any large software package will have thousands of mistakes. These vulnerabilities lie dormant in our software systems, waiting to be discovered. Once discovered, they can be used to attack systems. This is the point of security patching: eliminating known vulnerabilities. But many systems don't get patched, so the Internet is filled with known, exploitable vulnerabilities.

New vulnerabilities are hot commodities. A hacker who discovers one can sell it on the black market, blackmail the vendor with disclosure, or simply publish it without regard to the consequences. Even if he does none of these, the mere fact the vulnerability is known by someone increases the risk to every user of that software. Given that, is it ethical to research new vulnerabilities?

Unequivocally, yes. Despite the risks, vulnerability research is enormously valuable. Security is a mindset, and looking for vulnerabilities nurtures that

mindset. Deny practitioners this vital learning tool, and security suffers accordingly.

Security engineers see the world differently than other engineers. Instead of focusing on how systems work, they focus on how systems fail, how they can be made to fail, and how to prevent—or protect against—those failures. Most software vulnerabilities don't ever appear in normal operations, only when an attacker deliberately exploits them. So security engineers need to think like attackers.

People without the mindset sometimes think they can design security products, but they can't. And you see the results all over society—in snake-oil cryptography, software, Internet protocols, voting machines, and fare card and other payment systems. Many of these systems had someone in charge of “security” on their teams, but it wasn't someone who thought like an attacker.

This mindset is difficult to teach, and may be something you're born with or not. But in order to train people possessing the mindset, they need to search for and find security vulnerabilities—again and again and again. And this is true regardless of the domain. Good cryptographers discover vulnerabilities in others' algorithms and protocols. Good software security experts find vulnerabilities in others' code. Good airport security designers figure out new ways to subvert airport security. And so on.

This is so important that when someone shows me a security design by someone I don't know, my first question is, “What has the designer broken?” Anyone can design a security system that he cannot break. So when someone announces, “Here's my security system, and I can't break it,” your first reaction should be, “Who are you?” If he's someone who has broken dozens of similar systems, his system is worth looking at. If he's never broken anything, the chance is zero that it will be any good.

Vulnerability research is vital because it trains our next generation of computer security experts. Yes, newly discovered vulnerabilities in software and airports put us at risk, but they also give us more realistic information about how good the security actually is. And yes, there are more and less responsible—and more and less legal—ways to handle a new vulnerability. But the bad guys are constantly searching for new vulnerabilities, and if we have any hope of securing our systems, we need the good guys to be at least as competent. To me, the question isn't whether it's ethical to do vulnerability research. If someone has the skill to analyze and provide better insights into the problem, the question is whether it is ethical for him not to do vulnerability research.

I've Seen the Future, and It Has a Kill Switch

Originally published in Wired News, June 26, 2008

It used to be that just the entertainment industries wanted to control your computers—and televisions and iPods and everything else—to ensure that you didn't violate any copyright rules. But now everyone else wants to get their hooks into your gear.

OnStar will soon include the ability for the police to shut off your engine remotely. Buses are getting the same capability, in case terrorists want to reenact the movie *Speed*. The Pentagon wants a kill switch installed on airplanes, and is worried about potential enemies installing kill switches on their own equipment.

Microsoft is doing some of the most creative thinking along these lines, with something it's calling "Digital Manners Policies." According to its patent application, DMP-enabled devices would accept broadcast "orders" limiting their capabilities. Cellphones could be remotely set to vibrate mode in restaurants and concert halls, and be turned off on airplanes and in hospitals. Cameras could be prohibited from taking pictures in locker rooms and museums, and recording equipment could be disabled in theaters. Professors finally could prevent students from texting one another during class.

The possibilities are endless, and very dangerous. Making this work involves building a nearly flawless hierarchical system of authority. That's a difficult security problem even in its simplest form. Distributing that system among a variety of different devices—computers, phones, PDAs, cameras, recorders—with different firmware and manufacturers, is even more difficult. Not to mention delegating different levels of authority to various agencies, enterprises, industries and individuals, and then enforcing the necessary safeguards.

Once we go down this path—giving one device authority over other devices—the security problems start piling up. Who has the authority to limit functionality of my devices, and how do they get that authority? What prevents them from abusing that power? Do I get the ability to override their limitations? In what circumstances, and how? Can they override my override?

How do we prevent this from being abused? Can a burglar, for example, enforce a "no photography" rule and prevent security cameras from working? Can the police enforce the same rule to avoid another Rodney King incident?

Do the police get “superuser” devices that cannot be limited, and do they get “supercontroller” devices that can limit anything? How do we ensure that only they get them, and what do we do when the devices inevitably fall into the wrong hands?

It’s comparatively easy to make this work in closed specialized systems—OnStar, airplane avionics, military hardware—but much more difficult in open-ended systems. If you think Microsoft’s vision could possibly be securely designed, all you have to do is look at the dismal effectiveness of the various copy-protection and digital-rights-management systems we’ve seen over the years. That’s a similar capabilities-enforcement mechanism, albeit simpler than these more general systems.

And that’s the key to understanding this system. Don’t be fooled by the scare stories of wireless devices on airplanes and in hospitals, or visions of a world where no one is yammering loudly on their cellphones in posh restaurants. This is really about media companies wanting to exert their control further over your electronics. They not only want to prevent you from surreptitiously recording movies and concerts, they want your new television to enforce good “manners” on your computer, and not allow it to record any programs. They want your iPod to politely refuse to copy music to a computer other than your own. They want to enforce *their* legislated definition of manners: to control what you do and when you do it, and to charge you repeatedly for the privilege whenever possible.

“Digital Manners Policies” is a marketing term. Let’s call this what it really is: Selective Device Jamming. It’s not polite, it’s dangerous. It won’t make anyone more secure—or more polite.

Software Makers Should Take Responsibility

Originally published in the Guardian, July 17, 2008

A recent study of Internet browsers worldwide discovered that over half—52%—of Internet Explorer users weren’t using the current version of the software. For other browsers the numbers were better, but not much: 17% of Firefox users, 35% of Safari users, and 44% of Opera users were using an old version.

This is particularly important because browsers are an increasingly common vector for Internet attacks, and old versions of browsers don't have all their security patches up to date. They're open to attack through vulnerabilities the vendors have already fixed.

Security professionals are quick to blame users who don't use the latest update and install every patch. "Keeping up is critical for security," they say, and "if someone doesn't update their system, it's their own fault that they get hacked." This sounds a lot like blaming the victim: "He should have known not to walk down that deserted street; it's his own fault he was mugged." Of course the victim could have—and quite possibly should have—taken further precautions, but the real blame lies elsewhere.

It's not as if patching is easy. Even in a corporate setting, systems administrators have trouble keeping up with the never-ending flow of software patches. There could easily be dozens per week across all operating systems and applications, and far too often they break things. Microsoft's Automatic Update feature has automated the process, but that's the exception. Patching is triage, and administrators are constantly prioritizing it along with everything else they're doing.

It's the system that's broken. There's no other industry where shoddy products are sold to a public that expects regular problems, and where consumers are the ones who have to learn how to fix them. If an automobile manufacturer has a problem with a car and issues a recall notice, it's a rare occurrence and a big deal—and you can take your car in and get it fixed for free. Computers are the only mass-market consumer item that pushes this burden onto the consumer, requiring him to have a high level of technical sophistication just to survive.

It doesn't have to be this way. It is possible to write quality software. It is possible to sell software products that work properly, and don't need to be constantly patched. The problem is that it's expensive and time consuming. Software vendors won't do it, of course, because the marketplace won't reward it.

The key to fixing this is software liabilities. Computers are also the only mass-market consumer item where the vendors accept no liability for faults. The reason automobiles are so well designed is that manufacturers face liabilities if they screw up. A lack of software liability is effectively a vast government subsidy of the computer industry. It allows them to produce more products faster, with less concern about safety, security, and quality.

Last summer, the House of Lords Science and Technology Committee issued a report on "Personal Internet Security." I was invited to give testimony for that

report, and one of my recommendations was that software vendors be held liable when they are at fault. Their final report included that recommendation. The government rejected the recommendations in that report last autumn, and last week the committee issued a report on their follow-up inquiry, which still recommends software liabilities.

Good for them.

I'm not implying that liabilities are easy, or that all the liability for security vulnerabilities should fall on the vendor. But the courts are good at partial liability. Any automobile liability suit has many potential responsible parties: the car, the driver, the road, the weather, possibly another driver and another car, and so on. Similarly, a computer failure has several parties who may be partially responsible: the software vendor, the computer vendor, the network vendor, the user, possibly another hacker, and so on. But we're never going to get there until we start. Software liability is the market force that will motivate companies to improve their software quality—and everyone's security.

Lesson from the DNS Bug: Patching Isn't Enough

Originally published in Wired News, July 23, 2008

Despite the best efforts of the security community, the details of a critical Internet vulnerability discovered by Dan Kaminsky about six months ago have leaked. Hackers are racing to produce exploit code, and network operators who haven't already patched the hole are scrambling to catch up. The whole mess is a good illustration of the problems with researching and disclosing flaws like this.

The details of the vulnerability aren't important, but basically it's a form of DNS cache poisoning. The DNS system is what translates domain names people understand, like `www.schneier.com`, to IP addresses computers understand: `204.11.246.1`. There is a whole family of vulnerabilities where the DNS system on your computer is fooled into thinking that the IP address for `www.badsite.com` is really the IP address for `www.goodsite.com`—there's no way for you to tell the difference—and that allows the criminals at `www.badsite.com` to trick you into doing all sorts of things, like giving up your bank account details.

Kaminsky discovered a particularly nasty variant of this cache-poisoning attack.

Here's the way the timeline was supposed to work: Kaminsky discovered the vulnerability about six months ago, and quietly worked with vendors to patch it. (There's a fairly straightforward fix, although the implementation nuances are complicated.) Of course, this meant describing the vulnerability to them; why would companies like Microsoft and Cisco believe him otherwise? On July 8, he held a press conference to announce the vulnerability—but not the details—and reveal that a patch was available from a long list of vendors. We would all have a month to patch, and Kaminsky would release details of the vulnerability at the BlackHat conference early next month.

Of course, the details leaked. How isn't important; it could have leaked a zillion different ways. Too many people knew about it for it to remain secret. Others who knew the general idea were too smart not to speculate on the details. I'm kind of amazed the details remained secret for this long; undoubtedly it had leaked into the underground community before the public leak two days ago. So now everyone who back-burnered the problem is rushing to patch, while the hacker community is racing to produce working exploits.

What's the moral here? It's easy to condemn Kaminsky: If he had shut up about the problem, we wouldn't be in this mess. But that's just wrong. Kaminsky found the vulnerability by accident. There's no reason to believe he was the first one to find it, and it's ridiculous to believe he would be the last. Don't shoot the messenger. The problem is with the DNS protocol; it's insecure.

The real lesson is that the patch treadmill doesn't work, and it hasn't for years. This cycle of finding security holes and rushing to patch them before the bad guys exploit those vulnerabilities is expensive, inefficient and incomplete. We need to design security into our systems right from the beginning. We need assurance. We need security engineers involved in system design. This process won't prevent every vulnerability, but it's much more secure—and cheaper—than the patch treadmill we're all on now.

What a security engineer brings to the problem is a particular mindset. He thinks about systems from a security perspective. It's not that he discovers all possible attacks before the bad guys do; it's more that he anticipates potential types of attacks, and defends against them even if he doesn't know their details. I see this all the time in good cryptographic designs. It's over-engineering based on intuition, but if the security engineer has good intuition, it generally works.

Kaminsky's vulnerability is a perfect example of this. Years ago, cryptographer Daniel J. Bernstein looked at DNS security and decided that Source Port Randomization was a smart design choice. That's exactly the work-around being rolled out now following Kaminsky's discovery. Bernstein didn't discover Kaminsky's attack; instead, he saw a general class of attacks and realized that this enhancement could protect against them. Consequently, the DNS program he wrote in 2000, `djbdns`, doesn't need to be patched; it's already immune to Kaminsky's attack.

That's what a good design looks like. It's not just secure against known attacks; it's also secure against unknown attacks. We need more of this, not just on the Internet but in voting machines, ID cards, transportation payment cards... everywhere. Stop assuming that systems are secure unless demonstrated insecure; start assuming that systems are insecure unless designed securely.

Why Being Open about Security Makes Us All Safer in the Long Run

Originally published in the Guardian, August 7, 2008

London's Oyster card has been cracked, and the final details will become public in October. NXP Semiconductors, the Philips spin-off that makes the system, lost a court battle to prevent the researchers from publishing. People might be able to use this information to ride for free, but the sky won't be falling. And the publication of this serious vulnerability actually makes us all safer in the long run.

Here's the story. Every Oyster card has a radio-frequency identification chip that communicates with readers mounted on the ticket barrier. That chip, the "Mifare Classic" chip, is used in hundreds of other transport systems as well—Boston, Los Angeles, Brisbane, Amsterdam, Taipei, Shanghai, Rio de Janeiro—and as an access pass in thousands of companies, schools, hospitals, and government buildings around Britain and the rest of the world.

The security of Mifare Classic is terrible. This is not an exaggeration; it's kindergarten cryptography. Anyone with any security experience would be embarrassed to put his name to the design. NXP attempted to deal with this embarrassment by keeping the design secret.

The group that broke Mifare Classic is from Radboud University Nijmegen in the Netherlands. They demonstrated the attack by riding the Underground for free, and by breaking into a building. Their two papers (one is already online) will be published at two conferences this autumn.

The second paper is the one that NXP sued over. They called disclosure of the attack “irresponsible,” warned that it will cause “immense damages,” and claimed that it “will jeopardize the security of assets protected with systems incorporating the Mifare IC.” The Dutch court would have none of it: “Damage to NXP is not the result of the publication of the article but of the production and sale of a chip that appears to have shortcomings.”

Exactly right. More generally, the notion that secrecy supports security is inherently flawed. Whenever you see an organization claiming that design secrecy is necessary for security—in ID cards, in voting machines, in airport security—it invariably means that its security is lousy and it has no choice but to hide it. Any competent cryptographer would have designed Mifare’s security with an open and public design.

Secrecy is fragile. Mifare’s security was based on the belief that no one would discover how it worked; that’s why NXP had to muzzle the Dutch researchers. But that’s just wrong. Reverse-engineering isn’t hard. Other researchers had already exposed Mifare’s lousy security. A Chinese company even sells a compatible chip. Is there any doubt that the bad guys already know about this, or will soon enough?

Publication of this attack might be expensive for NXP and its customers, but it’s good for security overall. Companies will only design security as good as their customers know to ask for. NXP’s security was so bad because customers didn’t know how to evaluate security: either they don’t know what questions to ask, or didn’t know enough to distrust the marketing answers they were given. This court ruling encourages companies to build security properly rather than relying on shoddy design and secrecy, and discourages them from promising security based on their ability to threaten researchers.

It’s unclear how this break will affect Transport for London. Cloning takes only a few seconds, and the thief only has to brush up against someone carrying a legitimate Oyster card. But it requires an RFID reader and a small piece of software which, while feasible for a techie, are too complicated for the average fare dodger. The police are likely to quickly arrest anyone who tries to sell cloned cards on any scale. TfL promises to turn off any cloned cards within 24 hours, but that will hurt the innocent victim who had his card cloned more than the thief.

The vulnerability is far more serious to the companies that use Mifare Classic as an access pass. It would be very interesting to know how NXP presented the system's security to them.

And while these attacks only pertain to the Mifare Classic chip, it makes me suspicious of the entire product line. NXP sells a more secure chip and has another on the way, but given the number of basic cryptography mistakes NXP made with Mifare Classic, one has to wonder whether the "more secure" versions will be sufficiently so.

Boston Court's Meddling with "Full Disclosure" Is Unwelcome

Originally published in Wired News, August 21, 2008

In eerily similar cases in the Netherlands and the United States, courts have recently grappled with the computer-security norm of "full disclosure," asking whether researchers should be permitted to disclose details of a fare-card vulnerability that allows people to ride the subway for free.

The "Oyster card" used on the London Tube was at issue in the Dutch case, and a similar fare card used on the Boston "T" was the center of the US case. The Dutch court got it right, and the American court, in Boston, got it wrong from the start—despite facing an open-and-shut case of First Amendment prior restraint.

The US court has since seen the error of its ways—but the damage is done. The MIT security researchers who were prepared to discuss their Boston findings at the DefCon security conference were prevented from giving their talk.

The ethics of full disclosure are intimately familiar to those of us in the computer-security field. Before full disclosure became the norm, researchers would quietly disclose vulnerabilities to the vendors—who would routinely ignore them. Sometimes vendors would even threaten researchers with legal action if they disclosed the vulnerabilities.

Later on, researchers started disclosing the existence of a vulnerability but not the details. Vendors responded by denying the security holes' existence, or calling them just theoretical. It wasn't until full disclosure became the norm that vendors began consistently fixing vulnerabilities quickly. Now that vendors routinely patch vulnerabilities, researchers generally give them advance notice to allow them to patch their systems before the vulnerability is

published. But even with this “responsible disclosure” protocol, it’s the threat of disclosure that motivates them to patch their systems. Full disclosure is the mechanism by which computer security improves.

Outside of computer security, secrecy is much more the norm. Some security communities, like locksmiths, behave much like medieval guilds, divulging the secrets of their profession only to those within it. These communities hate open research, and have responded with surprising vitriol to researchers who have found serious vulnerabilities in bicycle locks, combination safes, master-key systems, and many other security devices.

Researchers have received a similar reaction from other communities more used to secrecy than openness. Researchers—sometimes young students—who discovered and published flaws in copyright-protection schemes, voting-machine security, and now wireless access cards have all suffered recriminations and sometimes lawsuits for not keeping the vulnerabilities secret. When Christopher Soghoian created a website allowing people to print fake airline boarding passes, he got several unpleasant visits from the FBI.

This preference for secrecy comes from confusing a vulnerability with information *about* that vulnerability. Using secrecy as a security measure is fundamentally fragile. It assumes that the bad guys don’t do their own security research. It assumes that no one else will find the same vulnerability. It assumes that information won’t leak out even if the research results are suppressed. These assumptions are all incorrect.

The problem isn’t the researchers; it’s the products themselves. Companies will only design security as good as what their customers know to ask for. Full disclosure helps customers evaluate the security of the products they buy, and educates them in how to ask for better security. The Dutch court got it exactly right when it wrote: “Damage to NXP is not the result of the publication of the article but of the production and sale of a chip that appears to have shortcomings.”

In a world of forced secrecy, vendors make inflated claims about their products, vulnerabilities don’t get fixed, and customers are no wiser. Security research is stifled, and security technology doesn’t improve. The only beneficiaries are the bad guys.

If you’ll forgive the analogy, the ethics of full disclosure parallel the ethics of not paying kidnapping ransoms. We all know why we don’t pay kidnappers: It encourages more kidnappings. Yet in every kidnapping case, there’s someone—a spouse, a parent, an employer—with a good reason why, in this one case, we should make an exception.

The reason we want researchers to publish vulnerabilities is because that's how security improves. But in every case there's someone—the Massachusetts Bay Transit Authority, the locksmiths, an election machine manufacturer—who argues that, in this one case, we should make an exception.

We shouldn't. The benefits of responsibly publishing attacks greatly outweigh the potential harm. Disclosure encourages companies to build security properly rather than relying on shoddy design and secrecy, and discourages them from promising security based on their ability to threaten researchers. It's how we learn about security, and how we improve future security.

Quantum Cryptography: As Awesome as It Is Pointless

Originally published in Wired News, October 16, 2008

Quantum cryptography is back in the news, and the basic idea is still unbelievably cool, in theory, and nearly useless in real life.

The idea behind quantum crypto is that two people communicating using a quantum channel can be absolutely sure no one is eavesdropping. Heisenberg's uncertainty principle requires anyone measuring a quantum system to disturb it, and that disturbance alerts legitimate users as to the eavesdropper's presence. No disturbance, no eavesdropper—period.

This month we've seen reports on a new working quantum-key distribution network in Vienna, and a new quantum-key distribution technique out of Britain. Great stuff, but headlines like the BBC's "Unbreakable" encryption unveiled" are a bit much.

The basic science behind quantum crypto was developed, and prototypes built, in the early 1980s by Charles Bennett and Giles Brassard, and there have been steady advances in engineering since then. I describe basically how it all works in *Applied Cryptography, 2nd Edition* (pages 554–7). At least one company already sells quantum-key distribution products.

Note that this is totally separate from quantum computing, which also has implications for cryptography. Several groups are working on designing and

building a quantum computer, which is fundamentally different from a classical computer. If one were built—and we’re talking science fiction here—then it could factor numbers and solve discrete-logarithm problems very quickly. In other words, it could break all of our commonly used public-key algorithms. For symmetric cryptography it’s not that dire: A quantum computer would effectively halve the key length, so that a 256-bit key would be only as secure as a 128-bit key today. Pretty serious stuff, but years away from being practical. I think the best quantum computer today can factor the number 15.

While I like the science of quantum cryptography—my undergraduate degree was in physics—I don’t see any commercial value in it. I don’t believe it solves any security problem that needs solving. I don’t believe that it’s worth paying for, and I can’t imagine anyone but a few technophiles buying and deploying it. Systems that use it don’t magically become unbreakable, because the quantum part doesn’t address the weak points of the system.

Security is a chain; it’s as strong as the weakest link. Mathematical cryptography, as bad as it sometimes is, is the strongest link in most security chains. Our symmetric and public-key algorithms are pretty good, even though they’re not based on much rigorous mathematical theory. The real problems are elsewhere: computer security, network security, user interface and so on.

Cryptography is the one area of security that we can get right. We already have good encryption algorithms, good authentication algorithms and good key-agreement protocols. Maybe quantum cryptography can make that link stronger, but why would anyone bother? There are far more serious security problems to worry about, and it makes much more sense to spend effort securing those.

As I’ve often said, it’s like defending yourself against an approaching attacker by putting a huge stake in the ground. It’s useless to argue about whether the stake should be 50 feet tall or 100 feet tall, because either way, the attacker is going to go around it. Even quantum cryptography doesn’t “solve” all of cryptography: The keys are exchanged with photons, but a conventional mathematical algorithm takes over for the actual encryption.

I’m always in favor of security research, and I have enjoyed following the developments in quantum cryptography. But as a product, it has no future. It’s not that quantum cryptography might be insecure; it’s that cryptography is already sufficiently secure.

Passwords Are Not Broken, but How We Choose Them Sure Is

Originally published in the Guardian, November 13, 2008

I've been reading a lot about how passwords are no longer good security. The reality is more complicated. Passwords are still secure enough for many applications, but you have to choose a good one. And that's hard. The best way to explain how to choose a good password is to describe how they're broken. The most serious attack is called offline password guessing. There are commercial programs that do this, sold primarily to police departments. There are also hacker tools that do the same thing.

As computers have become faster, the guessers have got better, sometimes being able to test hundreds of thousands of passwords per second. These guessers might run for months on many machines simultaneously.

They guess intelligently. They don't run through every eight-letter combination from "aaaaaaaa" to "zzzzzzzz" in order. That's 200bn possible passwords, most of them very unlikely. They try the most common password first: "password1." (Don't laugh; the most common password used to be "password.")

A typical password consists of a root plus an appendage. The root isn't necessarily a dictionary word, but it's something pronounceable. An appendage is either a suffix (90% of the time) or a prefix (10% of the time). One guesser I studied starts with a dictionary of about 1,000 common passwords, things like "letmein," "temp," "123456," and so on. Then it tests them each with about 100 common suffix appendages: "1," "4u," "69," "abc," "!" and so on. It recovers about 24% of all passwords with just these 100,000 combinations.

Then the guesser tries different dictionaries: English words, names, foreign words, phonetic patterns and so on for roots; two digits, dates, single symbols and so on for appendages. It runs the dictionaries with various capitalizations and common substitutions: "\$" for "s," "@" for "a," "1" for "l" and so on. With a couple of weeks to a month's worth of time, this guessing strategy breaks about two-thirds of all passwords. But that assumes no biographical data. Any smart guesser collects whatever personal information it can on the subject before beginning. Postal codes are common appendages, so they're tested.

It also tests names and addresses from the address book, meaningful dates, and any other personal information. If it can, the guesser indexes the target hard drive and creates a dictionary out of every printable string, including

deleted files. If you ever kept an email with your password, or saved it in an obscure file somewhere, or if your program ever stored it in memory, this process will grab it. And it will recover your password faster.

So if you want your password to be hard to guess, you should choose something that this process will miss. My advice is to take a sentence and turn it into a password. Something like “This little piggy went to market” might become “tlpWENT2m.” That nine-character password won’t be in anyone’s dictionary. Of course, don’t use this one, because I’ve written about it. Choose your own sentence—something personal.

Strong passwords can still fail because people are sloppy. They write them on Post-it notes stuck to their monitors, share them with friends, or choose the same passwords for multiple applications. (I don’t care about low-security passwords here, only about ones that matter: your bank accounts, your credit cards, etc.) Websites are sloppy, too, allowing people to set up easy-to-guess “secret questions” as a backup password or email them to customers.

If you can’t remember your passwords, write them down and put the paper in your wallet. But just write the sentence—or better yet—a hint that will help you remember your sentence. Or use a free program like Password Safe, which I designed to help people securely store all their passwords. Don’t feel this is a failure; most of us have far too many passwords to be able to remember them all.

Passwords can still provide good authentication if used properly. The rise of alternate forms of authentication is more because people don’t use passwords securely, and less because they don’t work anymore.

America’s Next Top Hash Function Begins

Originally published in Wired News, November 19, 2008

You might not have realized it, but the next great battle of cryptography began this month. It’s not a political battle over export laws or key escrow or NSA eavesdropping, but an academic battle over who gets to be the creator of the next hash standard.

Hash functions are the most commonly used cryptographic primitive, and the most poorly understood. You can think of them as fingerprint functions: They take an arbitrary long data stream and return a fixed length, and

effectively unique, string. The security comes from the fact that while it's easy to generate the fingerprint from a file, it's infeasible to go the other way and generate a file given a fingerprint.

Originally created to make digital signatures more efficient, hashes are now used to secure the very fundamentals of our information infrastructure: in password logins, secure web connections, encryption key management, virus and malware scanning, and almost every cryptographic protocol in current use. Without cryptographic hash functions, the Internet would simply not work. At the same time, there isn't a good theory of hash functions. Unlike encryption algorithms, there are no secret keys involved; this makes it harder to mathematically define exactly what hash functions are.

The National Institute of Standards and Technology, NIST, is holding a competition to replace the SHA family of hash functions. "SHA" stands for "Secure Hash Algorithm." It was developed by the NSA in 1993 to replace the commercial MD4 and MD5 algorithms, and has been updated several times since then. All the SHA algorithms are very similar, and have been increasingly under attack, so NIST wants to replace them.

The competition is important because, unlike other technological standards, committee design—balancing the interests of diverse constituents—isn't conducive to good security. Security is best when it's designed by expert teams and then subjected to public review. And cryptography is best when it's chosen by competition.

In 1997, NIST held a competition for a block cipher to replace DES. Fifteen candidates and three-and-a-half years later, Rijndael became the new Advanced Encryption Standard—AES. NIST is doing the same thing for what it's calling SHA-3 (not, for some unexplained reason, the Advanced Hash Standard or AHS).

The deadline was October 31, and NIST received 64 submissions. This isn't surprising—I predicted 80—as most of the 15 AES submitters were professors, whose students at the time have become professors themselves, with their own students. (If NIST does a stream cipher competition in another ten years, they should expect about 256 submissions.) These submissions came from academia, from industry, and from hobbyists. *CIO* magazine recently interviewed one of the submitters, who is 15. Twenty-eight submissions have been made public by the submitters, and six of those have been broken.

NIST is going through all the submissions right now, making sure they are complete and proper. Their goal is to publish all accepted submissions by the end of November, in advance of the First Hash Function Candidate

Conference, to be held in Belgium right after the Fast Software Encryption workshop in February.

The group expects to quickly make a first cut of algorithms—hopefully to about a dozen—and give the community a year of cryptanalysis before making a second cut in 2010. After another year of cryptanalysis, NIST will choose a winner in 2011. Expect a final standard by 2012.

My advice for software developers is to let the process run its course. While it's tempting to use the new cool algorithms in your designs, it's far too soon to trust any of them. This process is likely to result in all sorts of new research results in hash function security, and some real cryptanalytic surprises. Give the community a few years to figure out which ones are good and which aren't.

I've previously called this sort of thing a cryptographic demolition derby: The last one left standing wins. But that's only partially true. Certainly all the groups will spend the next few years trying to cryptanalyze each other, but in the end there will be a bunch of unbroken algorithms. NIST will select one based on performance and features.

NIST has stated that the goal of this process is not to choose the best standard but to choose a good standard. I think that's smart; in this process, the best is the enemy of the good. While there's no rush to choose a new standard—the SHA-2 algorithms will remain secure for the foreseeable future—we don't want to analyze the candidates forever.

Personally, I was part of a group of eight cryptographers that submitted Skein to the competition. A decade ago, writing Twofish and participating in the AES process was the most fun I had ever had in cryptography. These next few years promise to be even more fun.

Tigers Use Scent, Birds Use Calls— Biometrics Are Just Animal Instinct

Originally published in the Guardian, January 8, 2009

Biometrics may seem new, but they're the oldest form of identification. Tigers recognize each other's scent; penguins recognize calls. Humans recognize each other by sight from across the room, voices on the phone, signatures on contracts and photographs on drivers' licenses. Fingerprints have been used to identify people at crime scenes for more than 100 years.

What is new about biometrics is that computers are now doing the recognizing: thumbprints, retinal scans, voiceprints, and typing patterns. There's a lot of technology involved here, in trying to both limit the number of false positives (someone else being mistakenly recognized as you) and false negatives (you being mistakenly not recognized). Generally, a system can choose to have less of one or the other; less of both is very hard.

Biometrics can vastly improve security, especially when paired with another form of authentication such as passwords. But it's important to understand their limitations as well as their strengths. On the strength side, biometrics are hard to forge. It's hard to affix a fake fingerprint to your finger or make your retina look like someone else's. Some people can mimic voices, and make-up artists can change people's faces, but these are specialized skills.

On the other hand, biometrics are easy to steal. You leave your fingerprints everywhere you touch, your retinal scan everywhere you look. Regularly, hackers have copied the prints of officials from objects they've touched and posted them on the Internet. We haven't yet had an example of a large biometric database being hacked into, but the possibility is there. Biometrics are unique identifiers, but they're not secrets.

And a stolen biometric can fool some systems. It can be as easy as cutting out a signature, pasting it on to a contract and then faxing the page to someone. The person on the other end doesn't know that the signature isn't valid because he didn't see it fixed on to the page. Remote logins by fingerprint fail in the same way. If there's no way to verify the print came from an actual reader, not from a stored computer file, the system is much less secure.

A more secure system is to use a fingerprint to unlock your mobile phone or computer. Because there is a trusted path from the fingerprint reader to the stored fingerprint the system uses to compare, an attacker can't inject a previously stored print as easily as he can cut and paste a signature. A photo on an ID card works the same way: the verifier can compare the face in front of him with the face on the card.

Fingerprints on ID cards are more problematic, because the attacker can try to fool the fingerprint reader. Researchers have made false fingers out of rubber or glycerin. Manufacturers have responded by building readers that also detect pores or a pulse.

The lesson is that biometrics work best if the system can verify that the biometric came from the person at the time of verification. The biometric identification system at the gates of the CIA headquarters works because there's a guard with a large gun making sure no one is trying to fool the system.

Of course, not all systems need that level of security. At Counterpane, the security company I founded, we installed hand geometry readers at the access doors to the operations center. Hand geometry is a hard biometric to copy, and the system was closed and didn't allow electronic forgeries. It worked very well.

One more problem with biometrics: they don't fail well. Passwords can be changed, but if someone copies your thumbprint, you're out of luck: you can't update your thumb. Passwords can be backed up, but if you alter your thumbprint in an accident, you're stuck. The failures don't have to be this spectacular: a voice print reader might not recognize someone with a sore throat, or a fingerprint reader might fail outside in freezing weather. Biometric systems need to be analyzed in light of these possibilities.

Biometrics are easy, convenient, and when used properly, very secure; they're just not a panacea. Understanding how they work and fail is critical to understanding when they improve security and when they don't.

The Secret Question Is: Why Do IT Systems Use Insecure Passwords? ---

Originally published in the Guardian, February 19, 2009

Since January, the Conficker.B worm has been spreading like wildfire across the Internet, infecting the French navy, hospitals in Sheffield, the court system in Houston, Texas, and millions of computers worldwide. One of the ways it spreads is by cracking administrator passwords on networks. Which leads to the important question: why are IT administrators still using easy-to-guess passwords?

Computer authentication systems have two basic requirements. They need to keep the bad guys from accessing your account, and they need to allow you to access your account. Both are important, and every system is a balancing act between the two. Too little security, and the bad guys will get in too easily. But if the authentication system is too complicated, restrictive, or hard to use, you won't be able, or won't bother, to use it.

Passwords are the most common authentication system. They're easy to implement and use, which is why they're so popular. But, as computers have become faster, password-guessing has become easier. Most people don't choose complicated enough passwords to remain secure against modern

password-guessing attacks. Conficker.B is even less clever—it just tries a list of about 200 common passwords.

To combat password-guessing, many systems force users to choose harder-to-guess passwords—requiring minimum lengths, non-alpha-numeric characters, etc.—and change their passwords more frequently. The first makes guessing harder, and the second makes a guessed password less valuable. This, of course, makes the system more annoying, so users respond by writing their passwords down and taping them to their monitors, or simply forgetting them more often. Smarter users use a secure password database such as Password Safe.

Users forgetting their passwords can be expensive—customer service reps have to field phone calls and reset passwords—so some systems include a backup authentication system: a secret question. If you forget your password, you can authenticate yourself with some personal information that only you know, such as your mother's maiden name, your favorite schoolteacher, the street you grew up on, the name of your first pet and so on. This may make the system more usable, but it also makes it much less secure: answers can be easily guessed, and are often known by people close to you.

A common enhancement is a one-time password generator, such as a SecurID token. This is a small device with a screen that displays a password that changes every time the button is pressed. This is called two-factor authentication, and is much more secure, because this token—"something you have"—is combined with a password—"something you know." But it's less usable, because the tokens have to be purchased and distributed to all users, and far too often it's "something you lost or forgot." And it costs money. Tokens are more frequently used in corporate environments, but banks and some online gaming worlds have taken to using them, although sometimes only as an option, because people don't like them.

In most cases, how an authentication system works when a legitimate user tries to log on is much more important than how it works when an impostor tries to log on. No security system is perfect, and there is some level of fraud associated with any of these authentication methods. But the instances of fraud are rare compared to the number of times someone tries to log on legitimately. If a given authentication system lets the bad guys in one in a 100 times, a bank could decide to live with the problem, or try to solve it in some other way. But if the same authentication system prevented legitimate customers from logging on even one in 1,000 times, the number of complaints would be enormous and the system wouldn't survive one week.

Balancing security and usability is hard, and many organizations get it wrong. But it's also evolving; organizations need to tighten their security and continue to push more involved authentication methods; and more savvy Internet users will then be willing to accept them. And IT administrators need to be leading that evolutionary change.

The Pros and Cons of Password Masking

Originally published in Schneier on Security, July 3, 2009

Usability guru Jakob Nielsen opened up a can of worms when he made the case for unmasking passwords in his blog. I chimed in that I agreed. Almost 165 comments on my blog (and several articles, essays, and many other blog posts) later, the consensus is that we were wrong.

I was certainly too glib. Like any security countermeasure, password masking has value. But like any countermeasure, password masking is not a panacea. And the costs of password masking need to be balanced with the benefits.

The cost is accuracy. When users don't get visual feedback from what they're typing, they're more prone to make mistakes. This is especially true with character strings that have non-standard characters and capitalization. This has several ancillary costs:

- Users get pissed off.
- Users are more likely to choose easy-to-type passwords, reducing both mistakes and security. Removing password masking will make people more comfortable with complicated passwords: they'll become easier to memorize and easier to use.

The benefits of password masking are more obvious:

- Security from shoulder surfing. If people can't look over your shoulder and see what you're typing, they're much less likely to be able to steal your password. Yes, they can look at your fingers instead, but that's *much* harder than looking at the screen. Surveillance cameras are also an issue: it's easier to watch someone's fingers on recorded video, but reading a cleartext password off a screen is trivial.

In some situations, there is a trust dynamic involved. Do you type your password while your boss is standing over your shoulder watching? How about your spouse or partner? Your parent or child? Your teacher or students? At ATMs, there's a social convention of standing away from someone using the machine, but that convention doesn't apply to computers. You might not trust the person standing next to you enough to let him see your password, but don't feel comfortable telling him to look away. Password masking solves that social awkwardness.

- Security from screen scraping malware. This is less of an issue; keyboard loggers are more common and unaffected by password masking. And if you have that kind of malware on your computer, you've got all sorts of problems.
- A security "signal." Password masking alerts users, and I'm thinking users who aren't particularly security savvy, that passwords are a secret.

I believe that shoulder surfing isn't nearly the problem it's made out to be. One, lots of people use their computers in private, with no one looking over their shoulders. Two, personal handheld devices are used very close to the body, making shoulder surfing all that much harder. Three, it's hard to quickly and accurately memorize a random non-alphanumeric string that flashes on the screen for a second or so.

This is not to say that shoulder surfing isn't a threat. It is. And, as many readers pointed out, password masking is one of the reasons it isn't more of a threat. And the threat is greater for those who are not fluent computer users: slow typists and people who are likely to choose bad passwords. But I believe that the risks are overstated.

Password masking is definitely important on public terminals with short PINs. (I'm thinking of ATMs.) The value of the PIN is large, shoulder surfing is more common, and a four-digit PIN is easy to remember in any case.

And lastly, this problem largely disappears on the Internet on your personal computer. Most browsers include the ability to save and then automatically populate password fields, making the usability problem go away at the expense of another security problem (the security of the password becomes the security of the computer). There's a Firefox plugin that gets rid of password masking. And programs like my own Password Safe allow passwords to be cut and pasted into applications, also eliminating the usability problem.

One approach is to make it a configurable option. High-risk banking applications could turn password masking on by default; other applications could turn it off by default. Browsers in public locations could turn it on by default. I like this, but it complicates the user interface.

A reader mentioned BlackBerry's solution, which is to display each character briefly before masking it; that seems like an excellent compromise.

I, for one, would like the option. I cannot type complicated WEP keys into Windows—twice! what's the deal with that?—without making mistakes. I cannot type my rarely used and very complicated PGP keys without making a mistake unless I turn off password masking. That's what I was reacting to when I said "I agree."

So was I wrong? Maybe. Okay, probably. Password masking definitely improves security; many readers pointed out that they regularly use their computer in crowded environments, and rely on password masking to protect their passwords. On the other hand, password masking reduces accuracy and makes it less likely that users will choose secure and hard-to-remember passwords, I will concede that the password masking trade-off is more beneficial than I thought in my snap reaction, but also that the answer is not nearly as obvious as we have historically assumed.

Technology Shouldn't Give Big Brother a Head Start

Originally published in MPR News Q, July 31, 2009

China is the world's most successful Internet censor. While the Great Firewall of China isn't perfect, it effectively limits information flowing in and out of the country. But now the Chinese government is taking things one step further.

Under a requirement taking effect soon, every computer sold in China will have to contain the Green Dam Youth Escort software package. Ostensibly a pornography filter, it is government spyware that will watch every citizen on the Internet.

Green Dam has many uses. It can police a list of forbidden Web sites. It can monitor a user's reading habits. It can even enlist the computer in some massive botnet attack, as part of a hypothetical future cyberwar.

China's actions may be extreme, but they're not unique. Democratic governments around the world—Sweden, Canada and the United Kingdom, for example—are rushing to pass laws giving their police new powers of Internet surveillance, in many cases requiring communications system providers to redesign products and services they sell.

Many are passing data retention laws, forcing companies to keep information on their customers. Just recently, the German government proposed giving itself the power to censor the Internet.

The United States is no exception. The 1994 CALEA law required phone companies to facilitate FBI eavesdropping, and since 2001, the NSA has built substantial eavesdropping systems in the United States. The government has repeatedly proposed Internet data retention laws, allowing surveillance into past activities as well as present.

Systems like this invite criminal appropriation and government abuse. New police powers, enacted to fight terrorism, are already used in situations of normal crime. Internet surveillance and control will be no different.

Official misuses are bad enough, but the unofficial uses worry me more. Any surveillance and control system must itself be secured. An infrastructure conducive to surveillance and control invites surveillance and control, both by the people you expect and by the people you don't.

China's government designed Green Dam for its own use, but it's already been subverted. Why does anyone think that criminals won't be able to use it to steal bank account and credit card information, use it to launch other attacks, or turn it into a massive spam-sending botnet?

Why does anyone think that only authorized law enforcement will mine collected Internet data or eavesdrop on phone and IM conversations?

These risks are not theoretical. After 9/11, the National Security Agency built a surveillance infrastructure to eavesdrop on telephone calls and e-mails within the United States.

Although procedural rules stated that only non-Americans and international phone calls were to be listened to, actual practice didn't always match those rules. NSA analysts collected more data than they were authorized to, and used the system to spy on wives, girlfriends and famous people like former President Bill Clinton.

But that's not the most serious misuse of a telecommunications surveillance infrastructure. In Greece, between June 2004 and March 2005, someone wire-tapped more than 100 cell phones belonging to members of the Greek government—the prime minister and the ministers of defense, foreign affairs and justice.

Ericsson built this wiretapping capability into Vodafone's products, and enabled it only for governments that requested it. Greece wasn't one of those governments, but someone still unknown—a rival political party? organized crime?—figured out how to surreptitiously turn the feature on.

Researchers have already found security flaws in Green Dam that would allow hackers to take over the computers. Of course there are additional flaws, and criminals are looking for them.

Surveillance infrastructure can be exported, which also aids totalitarianism around the world. Western companies like Siemens, Nokia, and Secure Computing built Iran's surveillance infrastructure. US companies helped build China's electronic police state. Twitter's anonymity saved the lives of Iranian dissidents—anonymity that many governments want to eliminate.

Every year brings more Internet censorship and control—not just in countries like China and Iran, but in the United States, the United Kingdom, Canada and other free countries.

The control movement is egged on by both law enforcement, trying to catch terrorists, child pornographers and other criminals, and by media companies, trying to stop file sharers.

It's bad civic hygiene to build technologies that could someday be used to facilitate a police state. No matter what the eavesdroppers and censors say, these systems put us all at greater risk. Communications systems that have no inherent eavesdropping capabilities are more secure than systems with those capabilities built in.

Lockpicking and the Internet

Originally published in Dark Reading, August 10, 2009

Physical locks aren't very good. They keep the honest out, but any burglar worth his salt can pick the common door lock pretty quickly.

It used to be that most people didn't know this. Sure, we all watched television criminals and private detectives pick locks with an ease only found on television and thought it realistic, but somehow we still held onto the belief that our own locks kept us safe from intruders.

The Internet changed that.

First was the *MIT Guide to Lockpicking*, written by the late Bob ("Ted the Tool") Baldwin. Then came Matt Blaze's 2003 paper on breaking master key

systems. After that, came a flood of lock picking information on the Net: opening a bicycle lock with a Bic pen, key bumping, and more. Many of these techniques were already known in both the criminal and locksmith communities. The locksmiths tried to suppress the knowledge, believing their guildlike secrecy was better than openness. But they've lost: never has there been more public information about lock picking—or safecracking, for that matter.

Lock companies have responded with more complicated locks, and more complicated disinformation campaigns.

There seems to be a limit to how secure you can make a wholly mechanical lock, as well as a limit to how large and unwieldy a key the public will accept. As a result, there is increasing interest in other lock technologies.

As a security technologist, I worry that if we don't fully understand these technologies and the new sorts of vulnerabilities they bring, we may be trading a flawed technology for an even worse one. Electronic locks are vulnerable to attack, often in new and surprising ways.

Start with keypads, more and more common on house doors. These have the benefit that you don't have to carry a physical key around, but there's the problem that you can't give someone the key for a day and then take it away when that day is over. As such, the security decays over time—the longer the keypad is in use, the more people know how to get in. More complicated electronic keypads have a variety of options for dealing with this, but electronic keypads work only when the power is on, and battery-powered locks have their own failure modes. Plus, far too many people never bother to change the default entry code.

Keypads have other security failures, as well. I regularly see keypads where four of the 10 buttons are more worn than the other six. They're worn from use, of course, and instead of 10,000 possible entry codes, I now have to try only 24.

Fingerprint readers are another technology, but there are many known security problems with those. And there are operational problems, too: They're hard to use in the cold or with sweaty hands; and leaving a key with a neighbor to let the plumber in starts having a spy-versus-spy feel.

Some companies are going even further. Earlier this year, Schlage launched a series of locks that can be opened either by a key, a four-digit code, or the Internet. That's right: The lock is online. You can send the lock SMS messages or talk to it via a website, and the lock can send you messages when someone opens it—or even when someone tries to open it and fails.

Sounds nifty, but putting a lock on the Internet opens up a whole new set of problems, none of which we fully understand. Even worse: Security is only as strong as the weakest link. Schlage's system combines the inherent "pickability" of a physical lock, the new vulnerabilities of electronic keypads, and the hacking risk of online. For most applications, that's simply too much risk.

The Battle Is On against Facebook and Co. to Regain Control of Our Files

Originally published in the Guardian, September 9, 2009

File deletion is all about control. This used to not be an issue. Your data was on your computer, and you decided when and how to delete a file. You could use the delete function if you didn't care about whether the file could be recovered or not, and a file erase program—I use BCWipe for Windows—if you wanted to ensure no one could ever recover the file.

As we move more of our data onto cloud computing platforms such as Gmail and Facebook, and closed proprietary platforms such as the Kindle and the iPhone deleting data is much harder.

You have to trust that these companies will delete your data when you ask them to, but they're generally not interested in doing so. Sites like these are more likely to make your data inaccessible than they are to physically delete it. Facebook is a known culprit: actually deleting your data from its servers requires a complicated procedure that may or may not work. And even if you do manage to delete your data, copies are certain to remain in the companies' backup systems. Gmail explicitly says this in its privacy notice.

Online backups, SMS messages, photos on photo sharing sites, smartphone applications that store your data in the network: you have no idea what really happens when you delete pieces of data or your entire account, because you're not in control of the computers that are storing the data.

This notion of control also explains how Amazon was able to delete a book that people had previously purchased on their Kindle e-book readers. The legalities are debatable, but Amazon had the technical ability to delete the file because it controls all Kindles. It has designed the Kindle so that it determines

when to update the software, whether people are allowed to buy Kindle books, and when to turn off people's Kindles entirely.

Vanish is a research project by Roxana Geambasu and colleagues at the University of Washington. They designed a prototype system that automatically deletes data after a set time interval. So you can send an email, create a Google Doc, post an update to Facebook, or upload a photo to Flickr, all designed to disappear after a set period of time. And after it disappears, no one—not anyone who downloaded the data, not the site that hosted the data, not anyone who intercepted the data in transit, not even you—will be able to read it. If the police arrive at Facebook or Google or Flickr with a warrant, they won't be able to read it.

The details are complicated, but Vanish breaks the data's decryption key into a bunch of pieces and scatters them around the web using a peer-to-peer network. Then it uses the natural turnover in these networks—machines constantly join and leave—to make the data disappear. Unlike previous programs that supported file deletion, this one doesn't require you to trust any company, organization, or website. It just happens.

Of course, Vanish doesn't prevent the recipient of an email or the reader of a Facebook page from copying the data and pasting it into another file, just as Kindle's deletion feature doesn't prevent people from copying a book's files and saving them on their computers. Vanish is just a prototype at this point, and it only works if all the people who read your Facebook entries or view your Flickr pictures have it installed on their computers as well; but it's a good demonstration of how control affects file deletion. And while it's a step in the right direction, it's also new and therefore deserves further security analysis before being adopted on a wide scale.

We've lost the control of data on some of the computers we own, and we've lost control of our data in the cloud. We're not going to stop using Facebook and Twitter just because they're not going to delete our data when we ask them to, and we're not going to stop using Kindles and iPhones because they may delete our data when we don't want them to. But we need to take back control of data in the cloud, and projects like Vanish show us how we can.

Now we need something that will protect our data when a large corporation decides to delete it.

The Difficulty of Un-Authentication

Originally published in Threatpost, September 28, 2009

In computer security, a lot of effort is spent on the authentication problem. Whether it's passwords, secure tokens, secret questions, image mnemonics, or something else, engineers are continually coming up with more complicated—and hopefully more secure—ways for you to prove you are who you say you are over the Internet.

This is important stuff, as anyone with an online bank account or remote corporate network knows. But a lot less thought and work have gone into the other end of the problem: how do you tell the system on the other end of the line that you're no longer there? How do you unauthenticate yourself?

My home computer requires me to log out or turn my computer off when I want to unauthenticate. This works for me because I know enough to do it, but lots of people just leave their computer on and running when they walk away. As a result, many office computers are left logged in when people go to lunch, or when they go home for the night. This, obviously, is a security vulnerability.

The most common way to combat this is by having the system time out. I could have my computer log me out automatically after a certain period of inactivity—five minutes, for example. Getting it right requires some fine tuning, though. Log the person out too quickly, and he gets annoyed; wait too long before logging him out, and the system could be vulnerable during that time. My corporate e-mail server logs me out after 10 minutes or so, and I regularly get annoyed at my corporate e-mail system.

Some systems have experimented with a token: a USB authentication token that has to be plugged in for the computer to operate, or an RFID token that logs people out automatically when the token moves more than a certain distance from the computer. Of course, people will be prone to just leave the token plugged in to their computer all the time; but if you attach it to their car keys or the badge they have to wear at all times when walking around the office, the risk is minimized.

That's expensive, though. A research project used a Bluetooth device, like a cell phone, and measured its proximity to a computer. The system could be programmed to lock the computer if the Bluetooth device moved out of range.

Some systems log people out after every transaction. This wouldn't work for computers, but it can work for ATMs. The machine spits my card out before it gives me my cash, or just requires a card swipe, and makes sure I take it out of the machine. If I want to perform another transaction, I have to reinsert my card and enter my PIN a second time.

There's a physical analogue that everyone can explain: door locks. Does your door lock behind you when you close the door, or does it remain unlocked until you lock it? The first instance is a system that automatically logs you out, and the second requires you to log out manually. Both types of locks are sold and used, and which one you choose depends on both how you use the door and who you expect to try to break in.

Designing systems for usability is hard, especially when security is involved. Almost by definition, making something secure makes it less usable. Choosing an unauthentication method depends a lot on how the system is used as well as the threat model. You have to balance increasing security with pissing the users off, and getting that balance right takes time and testing, and is much more an art than a science.

Is Antivirus Dead?

Originally published in Information Security, November 2009

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

Security is never black and white. If someone asks, “for best security, should I do A or B?” the answer almost invariably is both. But security is always a trade-off. Often it's impossible to do both A and B—there's no time to do both, it's too expensive to do both, or whatever—and you have to choose. In that case, you look at A and B and you make your best choice. But it's almost always more secure to do both.

Yes, antivirus programs have been getting less effective as new viruses are more frequent and existing viruses mutate faster. Yes, antivirus companies are forever playing catch-up, trying to create signatures for new viruses. Yes, signature-based antivirus software won't protect you when a virus is new, before the signature is added to the detection program. Antivirus is by no means a panacea.

On the other hand, an antivirus program with up-to-date signatures will protect you from a lot of threats. It'll protect you against viruses, against spyware, against Trojans—against all sorts of malware. It'll run in the background, automatically, and you won't notice any performance degradation at all. And—here's the best part—it can be free. AVG won't cost you a penny. To me, this is an easy trade-off, certainly for the average computer user who clicks on attachments he probably shouldn't click on, downloads things he probably shouldn't download, and doesn't understand the finer workings of Windows Personal Firewall.

Certainly security would be improved if people used whitelisting programs such as Bit9 Parity and Savant Protection—and I personally recommend Malwarebytes' Anti-Malware—but a lot of users are going to have trouble with this. The average user will probably just swat away the “you're trying to run a program not on your whitelist” warning message or—even worse—wonder why his computer is broken when he tries to run a new piece of software. The average corporate IT department doesn't have a good idea of what software is running on all the computers within the corporation, and doesn't want the administrative overhead of managing all the change requests. And whitelists aren't a panacea, either: they don't defend against malware that attaches itself to data files (think Word macro viruses), for example.

One of the newest trends in IT is consumerization, and if you don't already know about it, you soon will. It's the idea that new technologies, the cool stuff people want, will become available for the consumer market before they become available for the business market. What it means to business is that people—employees, customers, partners—will access business networks from wherever they happen to be, with whatever hardware and software they have. Maybe it'll be the computer you gave them when you hired them. Maybe it'll be their home computer, the one their kids use. Maybe it'll be their cell phone or PDA, or a computer in a hotel's business center. Your business will have no way to know what they're using, and—more importantly—you'll have no control.

In this kind of environment, computers are going to connect to each other without a whole lot of trust between them. Untrusted computers are going to connect to untrusted networks. Trusted computers are going to connect to untrusted networks. The whole idea of “safe computing” is going to take on a whole new meaning—every man for himself. A corporate network is going to need a simple, dumb, signature-based antivirus product at the gateway of its network. And a user is going to need a similar program to protect his computer.

Bottom line: antivirus software is neither necessary nor sufficient for security, but it's still a good idea. It's not a panacea that magically makes you safe,

nor is it is obsolete in the face of current threats. As countermeasures go, it's cheap, it's easy, and it's effective. I haven't dumped my antivirus program, and I have no intention of doing so anytime soon.

I don't even want an Xbox.

Virus and Protocol Scares Happen Every Day—but Don't Let Them Worry You

Originally published in the Guardian, December 9, 2009

Last month, researchers found a security flaw in the SSL protocol, which is used to protect sensitive web data. The protocol is used for online commerce, webmail, and social networking sites. Basically, hackers could hijack an SSL session and execute commands without the knowledge of either the client or the server. The list of affected products is enormous.

If this sounds serious to you, you're right. It is serious. Given that, what should you do now? Should you not use SSL until it's fixed, and only pay for Internet purchases over the phone? Should you download some kind of protection? Should you take some other remedial action? What?

If you read the IT press regularly, you'll see this sort of question again and again. The answer for this particular vulnerability, as for pretty much any other vulnerability you read about, is the same: do nothing. That's right, nothing. Don't panic. Don't change your behavior. Ignore the problem, and let the vendors figure it out.

There are several reasons for this. One, it's hard to figure out which vulnerabilities are serious and which are not. Vulnerabilities such as this happen multiple times a month. They affect different software, different operating systems, and different web protocols. The press either mentions them or not, somewhat randomly; just because it's in the news doesn't mean it's serious.

Two, it's hard to figure out if there's anything you can do. Many vulnerabilities affect operating systems or Internet protocols. The only sure fix would be to avoid using your computer. Some vulnerabilities have surprising consequences. The SSL vulnerability mentioned above could be used to hack Twitter. Did you expect that? I sure didn't.

Three, the odds of a particular vulnerability affecting you are small. There are a lot of fish in the Internet, and you're just one of billions.

Four, often you can't do anything. These vulnerabilities affect clients and servers, individuals and corporations. A lot of your data isn't under your direct control—it's on your web-based email servers, in some corporate database, or in a cloud computing application. If a vulnerability affects the computers running Facebook, for example, your data is at risk, whether you log in to Facebook or not.

It's much smarter to have a reasonable set of default security practices and continue doing them. This includes:

1. Install an antivirus program if you run Windows, and configure it to update daily. It doesn't matter which one you use; they're all about the same. For Windows, I like the free version of AVG Internet Security. Apple Mac and Linux users can ignore this, as virus writers target the operating system with the largest market share.
2. Configure your OS and network router properly. Microsoft's operating systems come with a lot of security enabled by default; this is good. But have someone who knows what they're doing check the configuration of your router, too.
3. Turn on automatic software updates. This is the mechanism by which your software patches itself in the background, without you having to do anything. Make sure it's turned on for your computer, OS, security software, and any applications that have the option. Yes, you have to do it for everything, as they often have separate mechanisms.
4. Show common sense regarding the Internet. This might be the hardest thing, and the most important. Know when an email is real, and when you shouldn't click on the link. Know when a website is suspicious. Know when something is amiss.
5. Perform regular backups. This is vital. If you're infected with something, you may have to reinstall your operating system and applications. Good backups ensure you don't lose your data—documents, photographs, music—if that becomes necessary.

That's basically it. I could give a longer list of safe computing practices, but this short one is likely to keep you safe. After that, trust the vendors. They spent all last month scrambling to fix the SSL vulnerability, and they'll spend

all this month scrambling to fix whatever new vulnerabilities are discovered. Let that be their problem.

The Failure of Cryptography to Secure Modern Networks

Originally published in Dark Reading, June 30, 2010

For a while now, I've pointed out that cryptography is singularly ill-suited to solve the major network security problems of today: denial-of-service attacks, website defacement, theft of credit card numbers, identity theft, viruses and worms, DNS attacks, network penetration, and so on.

Cryptography was invented to protect communications: data in motion. This is how cryptography was used throughout most of history, and this is how the militaries of the world developed the science. Alice was the sender, Bob the receiver, and Eve the eavesdropper. Even when cryptography was used to protect stored data—data at rest—it was viewed as a form of communication. In *Applied Cryptography*, I described encrypting stored data in this way: “a stored message is a way for someone to communicate with himself through time.” Data storage was just a subset of data communication.

In modern networks, the difference is much more profound. Communications are immediate and instantaneous. Encryption keys can be ephemeral, and systems like the STU-III telephone can be designed such that encryption keys are created at the beginning of a call and destroyed as soon as the call is completed. Data storage, on the other hand, occurs over time. Any encryption keys must exist as long as the encrypted data exists. And storing those keys becomes as important as storing the unencrypted data was. In a way, encryption doesn't reduce the number of secrets that must be stored securely; it just makes them much smaller.

Historically, the reason key management worked for stored data was that the key could be stored in a secure location: the human brain. People would remember keys and, barring physical and emotional attacks on the people themselves, would not divulge them. In a sense, the keys were stored in a “computer” that was not attached to any network. And there they were safe.

This whole model falls apart on the Internet. Much of the data stored on the Internet is only peripherally intended for use by people; it's primarily intended

for use by other computers. And therein lies the problem. Keys can no longer be stored in people's brains. They need to be stored on the same computer, or at least the network, that the data resides on. And that is much riskier.

Let's take a concrete example: credit card databases associated with websites. Those databases are not encrypted because it doesn't make any sense. The whole point of storing credit card numbers on a website is so it's accessible—so each time I buy something, I don't have to type it in again. The website needs to dynamically query the database and retrieve the numbers, millions of times a day. If the database were encrypted, the website would need the key. But if the key were on the same network as the data, what would be the point of encrypting it? Access to the website equals access to the database in either case. Security is achieved by good access control on the website and database, not by encrypting the data.

The same reasoning holds true elsewhere on the Internet as well. Much of the Internet's infrastructure happens automatically, without human intervention. This means that any encryption keys need to reside in software on the network, making them vulnerable to attack. In many cases, the databases are queried so often that they are simply left in plaintext, because doing otherwise would cause significant performance degradation. Real security in these contexts comes from traditional computer security techniques, not from cryptography.

Cryptography has inherent mathematical properties that greatly favor the defender. Adding a single bit to the length of a key adds only a slight amount of work for the defender, but doubles the amount of work the attacker has to do. Doubling the key length doubles the amount of work the defender has to do (if that—I'm being approximate here), but increases the attacker's workload exponentially. For many years, we have exploited that mathematical imbalance.

Computer security is much more balanced. There'll be a new attack, and a new defense, and a new attack, and a new defense. It's an arms race between attacker and defender. And it's a very fast arms race. New vulnerabilities are discovered all the time. The balance can tip from defender to attacker overnight, and back again the night after. Computer security defenses are inherently very fragile.

Unfortunately, this is the model we're stuck with. No matter how good the cryptography is, there is some other way to break into the system. Recall how the FBI read the PGP-encrypted email of a suspected Mafia boss several years ago. They didn't try to break PGP; they simply installed a keyboard sniffer on the target's computer. Notice that SSL- and TLS-encrypted web

communications are increasingly irrelevant in protecting credit card numbers; criminals prefer to steal them by the hundreds of thousands from back-end databases.

On the Internet, communications security is much less important than the security of the endpoints. And increasingly, we can't rely on cryptography to solve our security problems.

The Story behind the Stuxnet Virus

Originally published in Forbes, October 7, 2010

Computer security experts are often surprised at which stories get picked up by the mainstream media. Sometimes it makes no sense. Why this particular data breach, vulnerability, or worm and not others? Sometimes it's obvious. In the case of Stuxnet, there's a great story.

As the story goes, the Stuxnet worm was designed and released by a government—the US and Israel are the most common suspects—specifically to attack the Bushehr nuclear power plant in Iran. How could anyone not report that? It combines computer attacks, nuclear power, spy agencies and a country that's a pariah to much of the world. The only problem with the story is that it's almost entirely speculation.

Here's what we do know: Stuxnet is an Internet worm that infects Windows computers. It primarily spreads via USB sticks, which allows it to get into computers and networks not normally connected to the Internet. Once inside a network, it uses a variety of mechanisms to propagate to other machines within that network and gain privilege once it has infected those machines. These mechanisms include both known and patched vulnerabilities, and four “zero-day exploits”: vulnerabilities that were unknown and unpatched when the worm was released. (All the infection vulnerabilities have since been patched.)

Stuxnet doesn't actually do anything on those infected Windows computers, because they're not the real target. What Stuxnet looks for is a particular model of Programmable Logic Controller (PLC) made by Siemens (the press often refers to these as SCADA systems, which is technically incorrect). These are small embedded industrial control systems that run all sorts of automated processes: on factory floors, in chemical plants, in oil refineries, at pipelines—and, yes, in nuclear power plants. These PLCs are often

controlled by computers, and Stuxnet looks for Siemens SIMATIC WinCC/Step 7 controller software.

If it doesn't find one, it does nothing. If it does, it infects it using yet another unknown and unpatched vulnerability, this one in the controller software. Then it reads and changes particular bits of data in the controlled PLCs. It's impossible to predict the effects of this without knowing what the PLC is doing and how it is programmed, and that programming can be unique based on the application. But the changes are very specific, leading many to believe that Stuxnet is targeting a specific PLC, or a specific group of PLCs, performing a specific function in a specific location—and that Stuxnet's authors knew exactly what they were targeting.

It's already infected more than 50,000 Windows computers, and Siemens has reported 14 infected control systems, many in Germany. (These numbers were certainly out of date as soon as I typed them.) We don't know of any physical damage Stuxnet has caused, although there are rumors that it was responsible for the failure of India's INSAT-4B satellite in July. We believe that it did infect the Bushehr plant.

All the anti-virus programs detect and remove Stuxnet from Windows systems.

Stuxnet was first discovered in late June, although there's speculation that it was released a year earlier. As worms go, it's very complex and got more complex over time. In addition to the multiple vulnerabilities that it exploits, it installs its own driver into Windows. These have to be signed, of course, but Stuxnet used a stolen legitimate certificate. Interestingly, the stolen certificate was revoked on July 16, and a Stuxnet variant with a different stolen certificate was discovered on July 17.

Over time the attackers swapped out modules that didn't work and replaced them with new ones—perhaps as Stuxnet made its way to its intended target. Those certificates first appeared in January. USB propagation, in March.

Stuxnet has two ways to update itself. It checks back to two control servers, one in Malaysia and the other in Denmark, but also uses a peer-to-peer update system: When two Stuxnet infections encounter each other, they compare versions and make sure they both have the most recent one. It also has a kill date of June 24, 2012. On that date, the worm will stop spreading and delete itself.

We don't know who wrote Stuxnet. We don't know why. We don't know what the target is, or if Stuxnet reached it. But you can see why there is so much speculation that it was created by a government.

Stuxnet doesn't act like a criminal worm. It doesn't spread indiscriminately. It doesn't steal credit card information or account login credentials. It doesn't herd infected computers into a botnet. It uses multiple zero-day vulnerabilities. A criminal group would be smarter to create different worm variants and use one in each. Stuxnet performs sabotage. It doesn't threaten sabotage, like a criminal organization intent on extortion might.

Stuxnet was expensive to create. Estimates are that it took 8 to 10 people six months to write. There's also the lab setup—surely any organization that goes to all this trouble would test the thing before releasing it—and the intelligence gathering to know exactly how to target it. Additionally, zero-day exploits are valuable. They're hard to find, and they can only be used once. Whoever wrote Stuxnet was willing to spend a lot of money to ensure that whatever job it was intended to do would be done.

None of this points to the Bushehr nuclear power plant in Iran, though. Best I can tell, this rumor was started by Ralph Langner, a security researcher from Germany. He labeled his theory “highly speculative,” and based it primarily on the facts that Iran had an usually high number of infections (the rumor that it had the most infections of any country seems not to be true), that the Bushehr nuclear plant is a juicy target, and that some of the other countries with high infection rates—India, Indonesia, and Pakistan—are countries where the same Russian contractor involved in Bushehr is also involved. This rumor moved into the computer press and then into the mainstream press, where it became the accepted story, without any of the original caveats.

Once a theory takes hold, though, it's easy to find more evidence. The word “myrtus” appears in the worm: an artifact that the compiler left, possibly by accident. That's the myrtle plant. Of course, that doesn't mean that druids wrote Stuxnet. According to the story, it refers to Queen Esther, also known as Hadassah; she saved the Persian Jews from genocide in the 4th century B.C. “Hadassah” means “myrtle” in Hebrew.

Stuxnet also sets a registry value of “19790509” to alert new copies of Stuxnet that the computer has already been infected. It's rather obviously a date, but instead of looking at the gazillion things—large and small—that happened on that the date, the story insists it refers to the date Persian Jew Habib Elghanain was executed in Tehran for spying for Israel.

Sure, these markers could point to Israel as the author. On the other hand, Stuxnet's authors were uncommonly thorough about not leaving clues in their code; the markers could have been deliberately planted by someone

who wanted to frame Israel. Or they could have been deliberately planted by Israel, who wanted us to think they were planted by someone who wanted to frame Israel. Once you start walking down this road, it's impossible to know when to stop.

Another number found in Stuxnet is 0xDEADF007. Perhaps that means “Dead Fool” or “Dead Foot,” a term that refers to an airplane engine failure. Perhaps this means Stuxnet is trying to cause the targeted system to fail. Or perhaps not. Still, a targeted worm designed to cause a specific sabotage seems to be the most likely explanation.

If that's the case, why is Stuxnet so sloppily targeted? Why doesn't Stuxnet erase itself when it realizes it's not in the targeted network? When it infects a network via USB stick, it's supposed to only spread to three additional computers and to erase itself after 21 days—but it doesn't do that. A mistake in programming, or a feature in the code not enabled? Maybe we're not supposed to reverse engineer the target. By allowing Stuxnet to spread globally, its authors committed collateral damage worldwide. From a foreign policy perspective, that seems dumb. But maybe Stuxnet's authors didn't care.

My guess is that Stuxnet's authors, and its target, will forever remain a mystery.

The Dangers of a Software Monoculture —

Originally published in Information Security, November 2010

This essay appeared as the first half of a point/counterpoint with Marcus Ranum.

In 2003, a group of security experts—myself included—published a paper saying that 1) software monocultures are dangerous and 2) Microsoft, being the largest creator of monocultures out there, is the most dangerous. Marcus Ranum responded with an essay that basically said we were full of it. Now, eight years later, Marcus and I thought it would be interesting to revisit the debate.

The basic problem with a monoculture is that it's all vulnerable to the same attack. The Irish Potato Famine of 1845–9 is perhaps the most famous monoculture-related disaster. The Irish planted only one variety of potato, and the genetically identical potatoes succumbed to a rot caused by *Phytophthora infestans*. Compare that with the diversity of potatoes traditionally grown in

South America, each one adapted to the particular soil and climate of its home, and you can see the security value in heterogeneity.

Similar risks exist in networked computer systems. If everyone is using the same operating system or the same applications software or the same networking protocol, and a security vulnerability is discovered in that OS or software or protocol, a single exploit can affect everyone. This is the problem of large-scale Internet worms: many have affected millions of computers on the Internet.

If our networking environment weren't homogeneous, a single worm couldn't do so much damage. We'd be more like South America's potato crop than Ireland's. Conclusion: monoculture is bad; embrace diversity or die along with everyone else.

This analysis makes sense as far as it goes, but suffers from three basic flaws. The first is the assumption that our IT monoculture is as simple as the potato's. When the particularly virulent Storm worm hit, it only affected from 1–10 million of its billion-plus possible victims. Why? Because some computers were running updated antivirus software, or were within locked-down networks, or whatever. Two computers might be running the same OS or applications software, but they'll be inside different networks with different firewalls and IDSs and router policies, they'll have different antivirus programs and different patch levels and different configurations, and they'll be in different parts of the Internet connected to different servers running different services. As Marcus pointed out back in 2003, they'll be a little bit different themselves. That's one of the reasons large-scale Internet worms don't infect everyone—as well as the network's ability to quickly develop and deploy patches, new antivirus signatures, new IPS signatures, and so on.

The second flaw in the monoculture analysis is that it downplays the cost of diversity. Sure, it would be great if a corporate IT department ran half Windows and half Linux, or half Apache and half Microsoft IIS, but doing so would require more expertise and cost more money. It wouldn't cost twice the expertise and money—there is some overlap—but there are significant economies of scale that result from everyone using the same software and configuration. A single operating system locked down by experts is far more secure than two operating systems configured by sysadmins who aren't so expert. Sometimes, as Mark Twain said: "Put all your eggs in one basket, and then guard that basket!"

The third flaw is that you can only get a limited amount of diversity by using two operating systems, or routers from three vendors. South American potato diversity comes from hundreds of different varieties. Genetic diversity comes

from millions of different genomes. In monoculture terms, two is little better than one. Even worse, since a network's security is primarily the minimum of the security of its components, a diverse network is less secure because it is vulnerable to attacks against any of its heterogeneous components.

Some monoculture is necessary in computer networks. As long as we have to talk to each other, we're all going to have to use TCP/IP, HTML, PDF, and all sorts of other standards and protocols that guarantee interoperability. Yes, there will be different implementations of the same protocol—and this is a good thing—but that won't protect you completely. You can't be too different from everyone else on the Internet, because if you were, you couldn't be on the Internet.

Species basically have two options for propagating their genes: the lobster strategy and the avian strategy. Lobsters lay 5,000 to 40,000 eggs at a time, and essentially ignore them. Only a minuscule percentage of the hatchlings live to be four weeks old, but that's sufficient to ensure gene propagation; from every 50,000 eggs, an average of two lobsters is expected to survive to legal size. Conversely, birds produce only a few eggs at a time, then spend a lot of effort ensuring that most of the hatchlings survive. In ecology, this is known as *r/K* selection theory. In either case, each of those offspring varies slightly genetically, so if a new threat arises, some of them will be more likely to survive. But even so, extinctions happen regularly on our planet; neither strategy is foolproof.

Our IT infrastructure is a lot more like a bird than a lobster. Yes, monoculture is dangerous and diversity is important. But investing time and effort in ensuring our current infrastructure's survival is even more important.

How Changing Technology Affects Security

*Originally published in IEEE Security & Privacy,
March/April 2012*

Security is a tradeoff, a balancing act between attacker and defender. Unfortunately, that balance is never static. Changes in technology affect both sides. Society uses new technologies to decrease what I call the *scope of defection*—what attackers can get away with—and attackers use new

technologies to increase it. What's interesting is the difference between how the two groups incorporate new technologies.

Changes in security systems can be slow. Society has to implement any new security technology as a group, which implies agreement and coordination and—in some instances—a lengthy bureaucratic procurement process. Meanwhile, an attacker can just use the new technology. For example, at the end of the horse-and-buggy era, it was easier for a bank robber to use his new motorcar as a getaway vehicle than it was for a town's police department to decide it needed a police car, get the budget to buy one, choose which one to buy, buy it, and then develop training and policies for it. And if only one police department did this, the bank robber could just move to another town. Defectors are more agile and adaptable, making them much better at being early adopters of new technology.

We saw it in law enforcement's initial inability to deal with Internet crime. Criminals were simply more flexible. Traditional criminal organizations like the Mafia didn't immediately move onto the Internet; instead, new Internet-savvy criminals sprung up. They set up websites like CardersMarket and DarkMarket, and established new organized crime groups within a decade or so of the Internet's commercialization. Meanwhile, law enforcement simply didn't have the organizational fluidity to adapt as quickly. Cities couldn't fire their old-school detectives and replace them with people who understood the Internet. The detectives' natural inertia and tendency to sweep problems under the rug slowed things even more. They spent the better part of a decade playing catch-up.

There's one more problem: defenders are in what military strategist Carl von Clausewitz calls "the position of the interior." They have to defend against every possible attack, while the defector only has to find one flaw that allows one way through the defenses. As systems get more complicated due to technology, more attacks become possible. This means defectors have a first-mover advantage; they get to try the new attack first. Consequently, society is constantly responding: shoe scanners in response to the shoe bomber, harder-to-counterfeit money in response to better counterfeiting technologies, better antivirus software to combat new computer viruses, and so on. The attacker's clear advantage increases the scope of defection even further.

Of course, there are exceptions. There are technologies that immediately benefit the defender and are of no use at all to the attacker—for example, fingerprint technology allowed police to identify suspects after they left the crime scene and didn't provide any corresponding benefit to criminals. The same thing happened with immobilizing technology for cars, alarm systems

for houses, and computer authentication technologies. Some technologies benefit both but still give more advantage to the defenders. The radio allowed street policemen to communicate remotely, which increased our level of safety more than the corresponding downside of criminals communicating remotely endangers us.

Still, we tend to be reactive in security, and only implement new measures in response to an increased scope of defection. We're slow about doing it and even slower about getting it right.

The Importance of Security Engineering —

Originally published in Schneier on Security, August 28, 2012

A shorter version of this essay appeared in the September/October 2012 issue of IEEE Security & Privacy.

In May, neuroscientist and popular author Sam Harris and I debated the issue of profiling Muslims at airport security. We each wrote essays, then went back and forth on the issue. I don't recommend reading the entire discussion; we spent 14,000 words talking past each other. But what's interesting is how our debate illustrates the differences between a security engineer and an intelligent layman. Harris was uninterested in the detailed analysis required to understand a security system and unwilling to accept that security engineering is a specialized discipline with a body of knowledge and relevant expertise. He trusted his intuition.

Many people have researched how intuition fails us in security: Paul Slovic and Bill Burns on risk perception, Daniel Kahneman on cognitive biases in general, Rick Walsh on folk computer-security models. I've written about the psychology of security, and Daniel Gartner has written more. Basically, our intuitions are based on things like antiquated fight-or-flight models, and these increasingly fail in our technological world.

This problem isn't unique to computer security, or even security in general. But this misperception about security matters now more than it ever has. We're no longer asking people to make security choices only for themselves and their businesses; we need them to make security choices as a matter of public policy. And getting it wrong has increasingly bad consequences.

Computers and the Internet have collided with public policy. The entertainment industry wants to enforce copyright. Internet companies want to continue freely spying on users. Law-enforcement wants its own laws imposed

on the Internet: laws that make surveillance easier, prohibit anonymity, mandate the removal of objectionable images and texts, and require ISPs to retain data about their customers' Internet activities. Militaries want laws regarding cyber weapons, laws enabling wholesale surveillance, and laws mandating an Internet kill switch. "Security" is now a catch-all excuse for all sorts of authoritarianism, as well as for boondoggles and corporate profiteering.

Cory Doctorow recently spoke about the coming war on general-purpose computing. I talked about it in terms of the entertainment industry and Jonathan Zittrain discussed it more generally, but Doctorow sees it as a much broader issue. Preventing people from copying digital files is only the first skirmish; just wait until the DEA wants to prevent chemical printers from making certain drugs, or the FBI wants to prevent 3D printers from making guns.

I'm not here to debate the merits of any of these policies, but instead to point out that people will debate them. Elected officials will be expected to understand security implications, both good and bad, and will make laws based on that understanding. And if they aren't able to understand security engineering, or even accept that there is such a thing, the result will be ineffective and harmful policies.

So what do we do? We need to establish security engineering as a valid profession in the minds of the public and policy makers. This is less about certifications and (heaven forbid) licensing, and more about perception—and cultivating a security mindset. Amateurs produce amateur security, which costs more in dollars, time, liberty, and dignity while giving us less—or even no—security. We need everyone to know that.

We also need to engage with real-world security problems, and apply our expertise to the variety of technical and socio-technical systems that affect broader society. Everything involves computers, and almost everything involves the Internet. More and more, computer security *is* security.

Finally, and perhaps most importantly, we need to learn how to talk about security engineering to a non-technical audience. We need to convince policy makers to follow a logical approach instead of an emotional one—an approach that includes threat modeling, failure analysis, searching for unintended consequences, and everything else in an engineer's approach to design. Powerful lobbying forces are attempting to force security policies on society, largely for non-security reasons, and sometimes in secret. We need to stand up for security.

Technologies of Surveillance

Originally published in Schneier on Security, March 5, 2013

An edited version of this essay appeared in the New York Daily News.

It's a new day for the New York Police Department, with technology increasingly informing the way cops do their jobs. With innovation comes new possibilities but also new concerns.

For one, the NYPD is testing a new type of security apparatus that uses terahertz radiation to detect guns under clothing from a distance. As Police Commissioner Ray Kelly explained to the *Daily News* back in January, If something is obstructing the flow of that radiation—a weapon, for example—the device will highlight that object.

Ignore, for a moment, the glaring constitutional concerns, which make the stop-and-frisk debate pale in comparison: virtual strip-searching, evasion of probable cause, potential racial profiling. Organizations like the American Civil Liberties Union are all over those, even though their opposition probably won't make a difference. We're scared of both terrorism and crime, even as the risks decrease; and when we're scared, we're willing to give up all sorts of freedoms to assuage our fears. Often, the courts go along.

A more pressing question is the effectiveness of technologies that are supposed to make us safer. These include the NYPD's Domain Awareness System, developed by Microsoft, which aims to integrate massive quantities of data to alert cops when a crime may be taking place. Other innovations are surely in the pipeline, all promising to make the city safer. But are we being sold a bill of goods?

For example, press reports make the gun-detection machine look good. We see images from the camera that pretty clearly show a gun outlined under someone's clothing. From that, we can imagine how this technology can spot gun-toting criminals as they enter government buildings or terrorize neighborhoods. Given the right inputs, we naturally construct these stories in our heads. The technology seems like a good idea, we conclude.

The reality is that we reach these conclusions much in the same way we decide that, say, drinking Mountain Dew makes you look cool. These are, after all, the products of for-profit companies, pushed by vendors looking to

make sales. As such, they're marketed no less aggressively than soda pop and deodorant. Those images of criminals with concealed weapons were carefully created both to demonstrate maximum effectiveness and push our fear buttons. These companies deliberately craft stories of their effectiveness, both through advertising and placement on television and movies, where police are often showed using high-powered tools to catch high-value targets with minimum complication.

The truth is that many of these technologies are nowhere near as reliable as claimed. They end up costing us gazillions of dollars and open the door for significant abuse. Of course, the vendors hope that by the time we realize this, they're too embedded in our security culture to be removed.

The current poster child for this sort of morass is the airport full-body scanner. Rushed into airports after the underwear bomber Umar Farouk Abdulmutallab nearly blew up a Northwest Airlines flight in 2009, they made us feel better, even though they don't work very well and, ironically, wouldn't have caught Abdulmutallab with his underwear bomb. Both the Transportation Security Administration and vendors repeatedly lied about their effectiveness, whether they stored images, and how safe they were. In January, finally, backscatter X-ray scanners were removed from airports because the company who made them couldn't sufficiently blur the images so they didn't show travelers naked. Now, only millimeter-wave full-body scanners remain.

Another example is closed-circuit television (CCTV) cameras. These have been marketed as a technological solution to both crime and understaffed police and security organizations. London, for example, is rife with them, and New York has plenty of its own. To many, it seems apparent that they make us safer, despite cries of Big Brother. The problem is that in study after study, researchers have concluded that they don't.

Counterterrorist data mining and fusion centers: nowhere near as useful as those selling the technologies claimed. It's the same with DNA testing and fingerprint technologies: both are far less accurate than most people believe. Even torture has been oversold as a security system—this time by a government instead of a company—despite decades of evidence that it doesn't work and makes us all less safe.

It's not that these technologies are totally useless. It's that they're expensive, and none of them is a panacea. Maybe there's a use for a terahertz radar, and maybe the benefits of the technology are worth the costs. But we should not forget that there's a profit motive at work, too.

When Technology Overtakes Security

Originally published in Wired, March 14, 2013

A core, not side, effect of technology is its ability to magnify power and multiply force—for both attackers and defenders. One side creates ceramic handguns, laser-guided missiles, and new-identity theft techniques, while the other side creates anti-missile defense systems, fingerprint databases, and automatic facial recognition systems.

The problem is that it's not balanced: Attackers generally benefit from new security technologies before defenders do. They have a first-mover advantage. They're more nimble and adaptable than defensive institutions like police forces. They're not limited by bureaucracy, laws, or ethics. They can evolve faster. And entropy is on their side—it's easier to destroy something than it is to prevent, defend against, or recover from that destruction.

For the most part, though, society still wins. The bad guys simply can't do enough damage to destroy the underlying social system. The question for us is: can society still maintain security as technology becomes more advanced?

I don't think it can.

Because the damage attackers can cause becomes greater as technology becomes more powerful. Guns become more harmful, explosions become bigger, malware becomes more pernicious... and so on. A single attacker, or small group of attackers, can cause more destruction than ever before.

This is exactly why the whole post-9/11 weapons-of-mass-destruction debate was so overwrought: Terrorists are scary, terrorists flying airplanes into buildings are even scarier, and the thought of a terrorist with a nuclear bomb is absolutely terrifying.

As the destructive power of individual actors and fringe groups increases, so do the calls for—and society's acceptance of—increased security.

Rethinking Security

Traditional security largely works “after the fact.” We tend not to ban or restrict the objects that can do harm; instead, we punish the people who do harm with objects. There are exceptions, of course, but they're exactly that: exceptions. This system works as long as society can tolerate the destructive effects of those objects (for example, allowing people to own baseball bats and arresting

them after they use them in a riot is only viable if society can tolerate the potential for riots).

When that isn't enough, we resort to "before-the-fact" security measures. These come in two basic varieties: **general surveillance** of people in an effort to stop them before they do damage, and **specific interdictions** in an effort to stop people from using those technologies to do damage.

But these measures work better at keeping dangerous technologies out of the hands of amateurs than at keeping them out of the hands of professionals.

And in the global interconnected world we live in, they're not anywhere close to foolproof. Still, a climate of fear causes governments to try. Lots of technologies are already restricted: entire classes of drugs, entire classes of munitions, explosive materials, biological agents. There are age restrictions on vehicles and training restrictions on complex systems like aircraft. We're already almost entirely living in a surveillance state, though we don't realize it or won't admit it to ourselves. This will only get worse as technology advances... today's Ph.D. theses are tomorrow's high-school science-fair projects.

Increasingly, broad prohibitions on technologies, constant ubiquitous surveillance, and *Minority Report*-like preemptive security will become the norm. We can debate the effectiveness of various security measures in different circumstances. But the problem isn't that these security measures won't work—even as they shred our freedoms and liberties—it's that no security is perfect.

Because sooner or later, the technology will exist for a hobbyist to explode a nuclear weapon, print a lethal virus from a bio-printer, or turn our electronic infrastructure into a vehicle for large-scale murder. We'll have the technology eventually to annihilate ourselves in great numbers, and sometime after, that technology will become cheap enough to be easy.

As it gets easier for one member of a group to destroy the entire group, and the group size gets larger, the odds of *someone* in the group doing it approaches certainty. Our global interconnectedness means that our group size encompasses everyone on the planet, and since government hasn't kept up, we have to worry about the weakest-controlled member of the weakest-controlled country. Is this a fundamental limitation of technological advancement, one that could end civilization? First our fears grip us so strongly that, thinking about the short term, we willingly embrace a police state in a desperate attempt to keep us safe; then, someone goes off and destroys us anyway?

If security won't work in the end, what is the solution?

Resilience—building systems able to survive unexpected and devastating attacks—is the best answer we have right now. We need to recognize that large-scale attacks will happen, that society can survive more than we give it credit for, and that we can design systems to survive these sorts of attacks. Calling terrorism an existential threat is ridiculous in a country where more people die each month in car crashes than died in the 9/11 terrorist attacks.

If the US can survive the destruction of an entire city—witness New Orleans after Hurricane Katrina or even New York after Sandy—we need to start acting like it, and planning for it. Still, it's hard to see how resilience buys us anything but additional time. Technology will continue to advance, and right now we don't know how to adapt any defenses—including resilience—fast enough.

We need a more flexible and rationally reactive approach to these problems and new regimes of trust for our information-interconnected world. We're going to have to figure this out if we want to survive, and I'm not sure how many decades we have left.

7

Travel and Security

Crossing Borders with Laptops and PDAs —

Originally published in the Guardian, May 15, 2008

Last month a US court ruled that border agents can search your laptop, or any other electronic device, when you're entering the country. They can take your computer and download its entire contents, or keep it for several days. Customs and Border Patrol has not published any rules regarding this practice, and I and others have written a letter to Congress urging it to investigate and regulate this practice.

But the US is not alone. British customs agents search laptops for pornography. And there are reports on the Internet of this sort of thing happening at other borders, too. You might not like it, but it's a fact. So how do you protect yourself?

Encrypting your entire hard drive, something you should certainly do for security in case your computer is lost or stolen, won't work here. The border agent is likely to start this whole process with a "please type in your password." Of course you can refuse, but the agent can search you further, detain you longer, refuse you entry into the country and otherwise ruin your day.

You're going to have to hide your data. Set a portion of your hard drive to be encrypted with a different key—even if you also encrypt your entire hard drive—and keep your sensitive data there. Lots of programs allow you to do this. I use PGPDisk (from pgp.com). TrueCrypt (truecrypt.org) is also good, and free.

While customs agents might poke around on your laptop, they're unlikely to find the encrypted partition. (You can make the icon invisible, for some

added protection.) And if they download the contents of your hard drive to examine later, you won't care.

Be sure to choose a strong encryption password. Details are too complicated for a quick tip, but basically anything easy to remember is easy to guess. (My advice is at tinyurl.com/4f8z4n.) Unfortunately, this isn't a perfect solution. Your computer might have left a copy of the password on the disk somewhere, and (as I also describe at the above link) smart forensic software will find it.

So your best defense is to clean up your laptop. A customs agent can't read what you don't have. You don't need five years' worth of email and client data. You don't need your old love letters and those photos (you know the ones I'm talking about). Delete everything you don't absolutely need. And use a secure file erasure program to do it. While you're at it, delete your browser's cookies, cache and browsing history. It's nobody's business what websites you've visited. And turn your computer off—don't just put it to sleep—before you go through customs; that deletes other things. Think of all this as the last thing to do before you stow your electronic devices for landing. Some companies now give their employees forensically clean laptops for travel, and have them download any sensitive data over a virtual private network once they've entered the country. They send any work back the same way, and delete everything again before crossing the border to go home. This is a good idea if you can do it.

If you can't, consider putting your sensitive data on a USB drive or even a camera memory card: even 16GB cards are reasonably priced these days. Encrypt it, of course, because it's easy to lose something that small. Slip it in your pocket, and it's likely to remain unnoticed even if the customs agent pokes through your laptop. If someone does discover it, you can try saying: "I don't know what's on there. My boss told me to give it to the head of the New York office." If you've chosen a strong encryption password, you won't care if he confiscates it.

Lastly, don't forget your phone and PDA. Customs agents can search those too: emails, your phone book, your calendar. Unfortunately, there's nothing you can do here except delete things.

I know this all sounds like work, and that it's easier to just ignore everything here and hope you don't get searched. Today, the odds are in your favor. But new forensic tools are making automatic searches easier and easier, and the recent US court ruling is likely to embolden other countries. It's better to be safe than sorry.

The TSA's Useless Photo ID Rules

Originally published in the Los Angeles Times, August 28, 2008

The TSA is tightening its photo ID rules at airport security. Previously, people with expired IDs or who claimed to have lost their IDs were subjected to secondary screening. Then the Transportation Security Administration realized that meant someone on the government's no-fly list—the list that is supposed to keep our planes safe from terrorists—could just fly with no ID.

Now, people without ID must also answer personal questions from their credit history to ascertain their identity. The TSA will keep records of who those ID-less people are, too, in case they're trying to probe the system.

This may seem like an improvement, except that the photo ID requirement is a joke. Anyone on the no-fly list can easily fly whenever he wants. Even worse, the whole concept of matching passenger names against a list of bad guys has negligible security value.

How to fly, even if you are on the no-fly list: Buy a ticket in some innocent person's name. At home, before your flight, check in online and print out your boarding pass. Then, save that web page as a PDF and use Adobe Acrobat to change the name on the boarding pass to your own. Print it again. At the airport, use the fake boarding pass and your valid ID to get through security. At the gate, use the real boarding pass in the fake name to board your flight.

The problem is that it is unverified passenger names that get checked against the no-fly list. At security checkpoints, the TSA just matches IDs to whatever is printed on the boarding passes. The airline checks boarding passes against tickets when people board the plane. But because no one checks ticketed names against IDs, the security breaks down.

This vulnerability isn't new. It isn't even subtle. I wrote about it in 2003, and again in 2006. I asked Kip Hawley, who runs the TSA, about it in 2007. Today, any terrorist smart enough to Google "print your own boarding pass" can bypass the no-fly list.

This gaping security hole would bother me more if the very idea of a no-fly list weren't so ineffective. The system is based on the faulty notion that the feds have this master list of terrorists, and all we have to do is keep the people on the list off the planes.

That's just not true. The no-fly list—a list of people so dangerous they are not allowed to fly yet so innocent we can't arrest them—and the less dangerous

“watch list” contain a combined 1 million names representing the identities and aliases of an estimated 400,000 people. There aren’t that many terrorists out there; if there were, we would be feeling their effects.

Almost all of the people stopped by the no-fly list are false positives. It catches innocents such as Ted Kennedy, whose name is similar to someone’s on the list, and Yusuf Islam (formerly Cat Stevens), who was on the list but no one knew why.

The no-fly list is a Kafkaesque nightmare for the thousands of innocent Americans who are harassed and detained every time they fly. Put on the list by unidentified government officials, they can’t get off. They can’t challenge the TSA about their status or prove their innocence. (The US 9th Circuit Court of Appeals decided this month that no-fly passengers can sue the FBI, but that strategy hasn’t been tried yet.)

But even if these lists were complete and accurate, they wouldn’t work. Timothy McVeigh, the Unabomber, the D.C. snipers, the London subway bombers and most of the 9/11 terrorists weren’t on any list before they committed their terrorist acts. And if a terrorist wants to know if he’s on a list, the TSA has approved a convenient, \$100 service that allows him to figure it out: the Clear program, which issues IDs to “trusted travelers” to speed them through security lines. Just apply for a Clear card; if you get one, you’re not on the list.

In the end, the photo ID requirement is based on the myth that we can somehow correlate identity with intent. We can’t. And instead of wasting money trying, we would be far safer as a nation if we invested in intelligence, investigation and emergency response—security measures that aren’t based on a guess about a terrorist target or tactic.

That’s the TSA: Not doing the right things. Not even doing right the things it does.

The Two Classes of Airport Contraband —

Originally published in Wired News, September 18, 2008

Airport security found a jar of pasta sauce in my luggage last month. It was a 6-ounce jar, above the limit; the official confiscated it, because allowing it on the airplane with me would have been too dangerous. And to demonstrate

how dangerous he really thought that jar was, he blithely tossed it in a nearby bin of similar liquid bottles and sent me on my way.

There are two classes of contraband at airport security checkpoints: the class that will get you in trouble if you try to bring it on an airplane, and the class that will cheerily be taken away from you if you try to bring it on an airplane. This difference is important: Making security screeners confiscate anything from that second class is a waste of time. All it does is harm innocents; it doesn't stop terrorists at all.

Let me explain. If you're caught at airport security with a bomb or a gun, the screeners aren't just going to take it away from you. They're going to call the police, and you're going to be stuck for a few hours answering a lot of awkward questions. You may be arrested, and you'll almost certainly miss your flight. At best, you're going to have a very unpleasant day.

This is why articles about how screeners don't catch every—or even a majority—of guns and bombs that go through the checkpoints don't bother me. The screeners don't have to be perfect; they just have to be good enough. No terrorist is going to base his plot on getting a gun through airport security if there's a decent chance of getting caught, because the consequences of getting caught are too great.

Contrast that with a terrorist plot that requires a 12-ounce bottle of liquid. There's no evidence that the London liquid bombers actually had a workable plot, but assume for the moment they did. If some copycat terrorists try to bring their liquid bomb through airport security and the screeners catch them—like they caught me with my bottle of pasta sauce—the terrorists can simply try again. They can try again and again. They can keep trying until they succeed. Because there are no consequences to trying and failing, the screeners have to be 100 percent effective. Even if they slip up one in a hundred times, the plot can succeed.

The same is true for knitting needles, pocketknives, scissors, corkscrews, cigarette lighters and whatever else the airport screeners are confiscating this week. If there's no consequence to getting caught with it, then confiscating it only hurts innocent people. At best, it mildly annoys the terrorists.

To fix this, airport security has to make a choice. If something is dangerous, treat it as dangerous and treat anyone who tries to bring it on as potentially dangerous. If it's not dangerous, then stop trying to keep it off airplanes. Trying to have it both ways just distracts the screeners from actually making us safer.

Fixing Airport Security

Originally published in the New York Daily News, June 24, 2009

It's been months since the Transportation Security Administration has had a permanent director. If, during the job interview (no, I didn't get one), President Obama asked me how I'd fix airport security in one sentence, I would reply: "Get rid of the photo ID check, and return passenger screening to pre-9/11 levels."

Okay, that's a joke. While showing ID, taking your shoes off and throwing away your water bottles isn't making us much safer, I don't expect the Obama administration to roll back those security measures anytime soon. Airport security is more about CYA than anything else: defending against what the terrorists did last time.

But the administration can't risk appearing as if it facilitated a terrorist attack, no matter how remote the possibility, so those annoyances are probably here to stay.

This would be my real answer: "Establish accountability and transparency for airport screening." And if I had another sentence: "Airports are one of the places where Americans, and visitors to America, are most likely to interact with a law enforcement officer—and yet no one knows what rights travelers have or how to exercise those rights."

Obama has repeatedly talked about increasing openness and transparency in government, and it's time to bring transparency to the Transportation Security Administration (TSA).

Let's start with the no-fly and watch lists. Right now, everything about them is secret: You can't find out if you're on one, or who put you there and why, and you can't clear your name if you're innocent. This Kafkaesque scenario is so un-American it's embarrassing. Obama should make the no-fly list subject to judicial review.

Then, move on to the checkpoints themselves. What are our rights? What powers do the TSA officers have? If we're asked "friendly" questions by behavioral detection officers, are we allowed not to answer? If we object to the rough handling of ourselves or our belongings, can the TSA official retaliate against us by putting us on a watch list? Obama should make the rules clear and explicit, and allow people to bring legal action against the TSA for violating those rules; otherwise, airport checkpoints will remain a Constitution-free zone in our country.

Next, Obama should refuse to use unfunded mandates to sneak expensive security measures past Congress. The Secure Flight program is the worst offender. Airlines are being forced to spend billions of dollars redesigning their reservations systems to accommodate the TSA's demands to preapprove every passenger before he or she is allowed to board an airplane. These costs are borne by us, in the form of higher ticket prices, even though we never see them explicitly listed.

Maybe Secure Flight is a good use of our money; maybe it isn't. But let's have debates like that in the open, as part of the budget process, where it belongs.

And finally, Obama should mandate that airport security be solely about terrorism, and not a general-purpose security checkpoint to catch everyone from pot smokers to deadbeat dads.

The Constitution provides us, both Americans and visitors to America, with strong protections against invasive police searches. Two exceptions come into play at airport security checkpoints. The first is "implied consent," which means that you cannot refuse to be searched; your consent is implied when you purchased your ticket. And the second is "plain view," which means that if the TSA officer happens to see something unrelated to airport security while screening you, he is allowed to act on that.

Both of these principles are well established and make sense, but it's their combination that turns airport security checkpoints into police-state-like checkpoints.

The TSA should limit its searches to bombs and weapons and leave general policing to the police—where we know courts and the Constitution still apply.

None of these changes will make airports any less safe, but they will go a long way to de-ratcheting the culture of fear, restoring the presumption of innocence and reassuring Americans, and the rest of the world, that—as Obama said in his inauguration speech—"we reject as false the choice between our safety and our ideals."

Laptop Security while Crossing Borders —

Originally published in Wired News, July 15, 2009

Last year, I wrote about the increasing propensity for governments, including the US and Great Britain, to search the contents of people's laptops at customs. What we know is still based on anecdote, as no country has clarified the rules

about what their customs officers are and are not allowed to do, and what rights people have.

Companies and individuals have dealt with this problem in several ways, from keeping sensitive data off laptops traveling internationally, to storing the data—encrypted, of course—on websites and then downloading it at the destination. I have never liked either solution. I do a lot of work on the road, and need to carry all sorts of data with me all the time. It's a lot of data, and downloading it can take a long time. Also, I like to work on long international flights.

There's another solution, one that works with whole-disk encryption products like PGPDisk (I'm on PGP's advisory board), TrueCrypt, and BitLocker: Encrypt the data to a key you don't know.

It sounds crazy, but stay with me. Caveat: Don't try this at home if you're not very familiar with whatever encryption product you're using. Failure results in a bricked computer. Don't blame me.

Step One: Before you board your plane, add another key to your whole-disk encryption (it'll probably mean adding another "user")—and make it random. By "random," I mean *really* random: Pound the keyboard for a while, like a monkey trying to write Shakespeare. Don't make it memorable. Don't even try to memorize it.

Technically, this key doesn't directly encrypt your hard drive. Instead, it encrypts the key that is used to encrypt your hard drive—that's how the software allows multiple keys.

So now there are two different users named with two different keys: the one you normally use, and some random one you just invented.

Step Two: Send that new random key to someone you trust. Make sure the trusted recipient has it, and make sure it works. You won't be able to recover your hard drive without it.

Step Three: Burn, shred, delete or otherwise destroy all copies of that new random key. Forget it. If it was sufficiently random and non-memorable, this should be easy.

Step Four: Board your plane normally and use your computer for the whole flight.

Step Five: Before you land, delete the key you normally use.

At this point, you will not be able to boot your computer. The only key remaining is the one you forgot in Step Three. There's no need to lie to the customs official; you can even show him a copy of this article if he doesn't believe you.

Step Six: When you're safely through customs, get that random key back from your confidant, boot your computer and re-add the key you normally use to access your hard drive.

And that's it.

This is by no means a magic get-through-customs-easily card. Your computer might be impounded, and you might be taken to court and compelled to reveal who has the random key.

But the purpose of this protocol isn't to prevent all that; it's just to deny any possible access to your computer to customs. You might be delayed. You might have your computer seized. (This will cost you any work you did on the flight, but—honestly—at that point that's the least of your troubles.) You might be turned back or sent home. But when you're back home, you have access to your corporate management, your personal attorneys, your wits after a good night's sleep, and all the rights you normally have in whatever country you're now in.

This procedure not only protects you against the warrantless search of your data at the border, it also allows you to deny a customs official your data without having to lie or pretend—which itself is often a crime.

Now the big question: Who should you send that random key to?

Certainly it should be someone you trust, but—more importantly—it should be someone with whom you have a privileged relationship. Depending on the laws in your country, this could be your spouse, your attorney, your business partner or your priest. In a larger company, the IT department could institutionalize this as a policy, with the help desk acting as the key holder.

You could also send it to yourself, but be careful. You don't want to e-mail it to your webmail account, because then you'd be lying when you tell the customs official that there is no possible way you can decrypt the drive.

You could put the key on a USB drive and send it to your destination, but there are potential failure modes. It could fail to get there in time to be waiting for your arrival, or it might not get there at all. You could airmail the drive with the key on it to yourself a couple of times, in a couple of different ways, and also fax the key to yourself . . . but that's more work than I want to do when I'm traveling.

If you only care about the return trip, you can set it up before you return. Or you can set up an elaborate one-time pad system, with identical lists of keys with you and at home: Destroy each key on the list you have with you as you use it.

Remember that you'll need to have full-disk encryption, using a product such as PGPDisk, TrueCrypt or BitLocker, already installed and enabled to make this work.

I don't think we'll ever get to the point where our computer data is safe when crossing an international border. Even if countries like the US and Britain clarify their rules and institute privacy protections, there will always be other countries that will exercise greater latitude with their authority. And sometimes protecting your data means protecting your data from yourself.

Breaching the Secure Area in Airports

Originally published in Threatpost, January 5, 2010

An unidentified man breached airport security at Newark Airport on Sunday, walking into the secured area through the exit, prompting an evacuation of a terminal and flight delays that continued into the next day. This problem isn't common, but it happens regularly. The result is always the same, and it's not obvious that fixing the problem is the right solution.

This kind of security breach is inevitable, simply because human guards are not perfect. Sometimes it's someone going in through the out door, unnoticed by a bored guard. Sometimes it's someone running through the checkpoint and getting lost in the crowd. Sometimes it's an open door that should be locked. Amazing as it seems to frequent fliers, the perpetrator often doesn't even know he did anything wrong.

Basically, whenever there is—or could be—an unscreened person lost within the secure area of an airport, there are two things the TSA can do. They can say, "This isn't a big deal," and ignore it. Or they can evacuate *everyone* inside the secure area, search every nook and cranny—inside the large boxes of napkins at the fast food restaurant, above the false ceilings in the bathrooms, everywhere—looking for anyone hiding, and then rescreen everybody, causing delays of six, eight, twelve hours or more. That's it; those are his options. And there's no way he's choosing to ignore the risk; even if the odds are minuscule that it's a problem, it'll cost him his career if he's wrong.

Several European airports have their security screening organized differently. At Schiphol Airport in Amsterdam, for example, passengers are screened at the gates. This is more expensive and requires a substantially different

airport design, but it does mean that if there is a problem only the one gate has to be evacuated and searched and the people rescreened.

American airports can do more to secure against this risk, but I'm reasonably sure it's not worth it. We could double the guards to reduce the risk of inattentiveness, and redesign the airports to make this kind of thing less likely, but that's an expensive solution to an already rare problem. As much as I don't like saying it, the smartest thing is probably to live with this occasional but major inconvenience.

Stop the Panic on Air Security

Originally published in CNN, January 7, 2010

The Underwear Bomber failed. And our reaction to the failed plot is failing as well, by focusing on the specifics of this made-for-a-movie plot rather than the broad threat. While our reaction is predictable, it's not going to make us safer.

We're going to beef up airport security, because Umar Farouk AbdulMutallab allegedly snuck a bomb through a security checkpoint. We're going to intensively screen Nigerians, because he is Nigerian. We're going to field full body scanners, because they might have noticed the PETN that authorities say was hidden in his underwear. And so on.

We're doing these things even though security worked. The security checkpoints, even at their pre-9/11 levels, forced whoever made the bomb to construct a much worse bomb than he would have otherwise. Instead of using a timer or a plunger or another reliable detonation mechanism, as would any commercial user of PETN, he had to resort to an ad hoc homebrew—and a much more inefficient one, involving a syringe, and 20 minutes in the lavatory, and we don't know exactly what else—that didn't explode.

At that point, AbdulMutallab's fellow passengers quickly subdued him. Yes, the screeners didn't notice any PETN in his underwear, but the system was never intended to catch that particular tactic. There probably were intelligence failures—why wasn't his father's tip followed up on, and why wasn't his visa revoked?—but it's always easy to connect the dots in hindsight.

We're doing these things even though this particular plot was chosen precisely because we weren't screening for it; future al Qaeda attacks rarely look like past attacks; and the terrorist threat is far broader than attacks against airplanes.

We're doing these things even though airplane terrorism is incredibly rare, the risk is no greater today than it was in previous decades, the taxi to the airport is still more dangerous than the flight, and ten times as many Americans are killed by lightning as by terrorists.

In fact, we're focusing on the specifics of the plot, not despite these facts, but because of them.

The Underwear Bomber is precisely the sort of story we humans tend to overreact to. Our brains aren't very good at probability and risk analysis, especially when it comes to rare events. Our brains are much better at processing the simple risks we've had to deal with throughout most of our species' existence, and much poorer at evaluating the complex risks modern society forces us to face. We exaggerate spectacular rare events, and downplay familiar and common ones.

We can see the effects of this all the time. We fear being murdered, kidnapped, raped and assaulted by strangers, when it's far more likely that the perpetrator of such offenses is a relative or a friend. We fear school shootings, even though a school is almost always the safest place a child can be. We worry about shark attacks instead of fatal dog or pig attacks—both far more common. In the US, over 38,000 people die each year in car crashes; that's as many deaths as 9/11 each and every month, year after year.

Overreacting to the rare and spectacular is natural. We tend to base risk analysis on personal story rather than on data. If a friend gets mugged in a foreign country, that story is more likely to affect how safe you feel in that country than abstract crime statistics.

We give storytellers we have a relationship with more credibility than we give strangers, and stories that are close to us more weight than stories from foreign lands. And who is everyone's major storyteller these days? Television.

I tell people that if it's in the news, don't worry about it. The very definition of "news" is "something that hardly ever happens." It's when something isn't in the news, when it's so common that it's no longer news—car crashes, domestic violence—that you should start worrying.

But that's not the way we think. The more an event is talked about, the more probable we think it is. The more vivid our thoughts about the event are—again, think television—the more easily we remember it and the more convincing it is. So when faced with a very available and highly vivid event like the Underwear Bomber, 9/11, or a child kidnapping in a playground, we overreact. We get scared.

And once we're scared, we need to "do something"—even if that something doesn't make sense and is ineffective. We need to do something directly related to the story that's making us scared. We implement full body scanners at airports. We pass the Patriot Act. We don't let our children go to playgrounds unsupervised. Instead of implementing effective, but more general, security measures to reduce the overall risk, we concentrate on making the fearful story go away. Yes, it's security theater, but it makes us feel safer.

As circular as it sounds, rare events are rare primarily because they don't occur very often, and not because of any preventive security measures. If you want to do something that makes security sense, figure out what's common among a bunch of rare events, and concentrate your countermeasures there.

Focus on the general risk of terrorism, and not the specific threat of airplane bombings using PETN-filled underwear. Focus on the general risk of troubled teens, and not the specific threat of a lone gunman wandering around a school. Ignore the movie-plot threats, and concentrate on the real risks.

A Waste of Money and Time

Originally published in the New York Times Room for Debate blog, November 23, 2010

A short history of airport security: We screen for guns and bombs, so the terrorists use box cutters. We confiscate box cutters and corkscrews, so they put explosives in their sneakers. We screen footwear, so they try to use liquids. We confiscate liquids, so they put PETN bombs in their underwear. We roll out full-body scanners, even though they wouldn't have caught the Underwear Bomber, so they put a bomb in a printer cartridge. We ban printer cartridges over 16 ounces—the level of magical thinking here is amazing—and they're going to do something else.

This is a stupid game, and we should stop playing it.

It's not even a fair game. It's not that the terrorist picks an attack and we pick a defense, and we see who wins. It's that we pick a defense, and then the terrorists look at our defense and pick an attack designed to get around it. Our security measures only work if we happen to guess the plot correctly. If we get it wrong, we've wasted our money. This isn't security; it's security theater.

There are two basic kinds of terrorists. There are the sloppy planners, like the guy who crashed his plane into the Internal Revenue Service building in Austin. He's going to be sloppy and stupid, and even pre-9/11 airplane security is going to catch him. The second is the well-planned, well-financed, and much rarer sort of plot. Do you really expect the T.S.A. screeners, who are busy confiscating water bottles and making people take off their belts—and now doing uncomfortable pat-downs—to stop them?

Of course not. Airport security is the last line of defense, and it's not a very good one. What works is investigation and intelligence: security that works regardless of the terrorist tactic or target. Yes, the target matters too; all this airport security is only effective if the terrorists target airports. If they decide to bomb crowded shopping malls instead, we've wasted our money.

That being said, airplanes require a special level of security for several reasons: they're a favored terrorist target; their failure characteristics mean more deaths than a comparable bomb on a bus or train; they tend to be national symbols; and they often fly to foreign countries where terrorists can operate with more impunity.

But all that can be handled with pre-9/11 security. Exactly two things have made airplane travel safer since 9/11: reinforcing the cockpit door, and convincing passengers they need to fight back. Everything else has been a waste of money. Add screening of checked bags and airport workers and we're done. Take all the rest of the money and spend it on investigation and intelligence.

Immediately after the Christmas Day Underwear Bomber's plot failed, Homeland Security Secretary Janet Napolitano called airplane security a success. She was pilloried in the press and quickly backpedaled, but I think it was one of the most sensible things said on the subject. Plane lands safely, terrorist in custody, nobody injured except the terrorist: what more do people want out of a security success?

Look at what succeeded. Because even pre-9/11 security screened for obvious bombs, Abdulmutallab had to construct a far less reliable bomb than he would have otherwise. Instead of using a timer or a plunger or a reliable detonation mechanism, as would any commercial user of PETN, Abdulmutallab had to resort to an ad hoc and much more inefficient detonation mechanism involving a syringe, 20 minutes in the lavatory, and setting his pants on fire. As a result, his actions came to the notice of the other passengers, who subdued him.

Neither the full-body scanners or the enhanced pat-downs are making anyone safer. They're more a result of politicians and government appointees

capitulating to a public that demands that “something must be done,” even when nothing should be done; and a government bureaucracy that is more concerned about the security of their careers if they fail to secure against the last attack than what happens if they fail to anticipate the next one.

Why the TSA Can't Back Down

Originally published in the Atlantic, December 2, 2010

Organizers of National Opt Out Day, the Wednesday before Thanksgiving when air travelers were urged to opt out of the full-body scanners at security checkpoints and instead submit to full-body patdowns—were outfoxed by the TSA. The government pre-empted the protest by turning off the machines in most airports during the Thanksgiving weekend. Everyone went through the metal detectors, just as before.

Now that Thanksgiving is over, the machines are back on and the “enhanced” pat-downs have resumed. I suspect that more people would prefer to have naked images of themselves seen by TSA agents in another room, than have themselves intimately touched by a TSA agent right in front of them.

But now, the TSA is in a bind. Regardless of whatever lobbying came before, or whatever former DHS officials had a financial interest in these scanners, the TSA has spent billions on those scanners, claiming they're essential. But because people can opt out, the alternate manual method must be equally effective; otherwise, the terrorists could just opt out. If they make the pat-downs less invasive, it would be the same as admitting the scanners aren't essential. Senior officials would get fired over that.

So not counting inconsequential modifications to demonstrate they're “listening,” the pat-downs will continue. And they'll continue for everyone: children, abuse survivors, rape survivors, urostomy bag wearers, people in wheelchairs. It has to be that way; otherwise, the terrorists could simply adapt. They'd hide their explosives on their children or in their urostomy bags. They'd recruit rape survivors, abuse survivors, or seniors. They'd dress as pilots. They'd sneak their PETN through airport security using the very type of person who isn't being screened.

And PETN is what the TSA is looking for these days. That's pentaerythritol tetranitrate, the plastic explosive that both the Shoe Bomber and the Underwear Bomber attempted but failed to detonate. It's what was mailed from

Yemen. It's in Iraq and Afghanistan. Guns and traditional bombs are passé; PETN is the terrorist tool of the future.

The problem is that no scanners or puffers can detect PETN; only swabs and dogs work. What the TSA hopes is that they will detect the bulge if someone is hiding a wad of it on their person. But they won't catch PETN hidden in a body cavity. That doesn't have to be as gross as you're imagining; you can hide PETN in your mouth. A terrorist can go through the scanners a dozen times with bits in his mouth each time, and assemble a bigger bomb on the other side. Or he can roll it thin enough to be part of a garment, and sneak it through that way. These tricks aren't new. In the days after the Underwear Bomber was stopped, a scanner manufacturer admitted that the machines might not have caught him.

So what's next? Strip searches? Body cavity searches? TSA Administrator John Pistole said there would be no body cavity searches for now, but his reasons make no sense. He said that the case widely reported as being a body cavity bomb might not actually have been. While that appears to be true, what does that have to do with future bombs? He also said that even body cavity bombs would need "external initiators" that the TSA would be able to detect.

Do you think for a minute that the TSA can detect these "external initiators"? Do you think that if a terrorist took a laptop—or better yet, a less-common piece of electronics gear—and removed the insides and replaced them with a timer, a pressure sensor, a simple contact switch, or a radio frequency switch, the TSA guy behind the X-ray machine monitor would detect it? How about if those components were distributed over a few trips through airport security. On the other hand, if we believe the TSA can magically detect these "external initiators" so effectively that they make body-cavity searches unnecessary, why do we need the full-body scanners?

Either PETN is a danger that must be searched for, or it isn't. Pistole was being either ignorant or evasive.

Once again, the TSA is covering their own asses by implementing security-theater measures to prevent the previous attack while ignoring any threats of future attacks. It's the same thinking that caused them to ban box cutters after 9/11, screen shoes after Richard Reid, limit liquids after that London gang, and—I kid you not—ban printer cartridges over 16 ounces after they were used to house package bombs from Yemen. They act like the terrorists are incapable of thinking creatively, while the terrorists repeatedly demonstrate they can always come up with a new approach that circumvents the old measures.

On the plus side, PETN is very hard to get to explode. The pre-9/11 screening procedures, looking for obvious guns and bombs, forced the terrorists to

build inefficient fusing mechanisms. We saw this when Abdulmutallab, the Underwear Bomber, used bottles of liquid and a syringe and 20 minutes in the bathroom to assemble his device, then set his pants on fire—and still failed to ignite his PETN-filled underwear. And when he failed, the passengers quickly subdued him.

The truth is that exactly two things have made air travel safer since 9/11: reinforcing cockpit doors and convincing passengers they need to fight back. The TSA should continue to screen checked luggage. They should start screening airport workers. And then they should return airport security to pre-9/11 levels and let the rest of their budget be used for better purposes. Investigation and intelligence is how we're going to prevent terrorism, on airplanes and elsewhere. It's how we caught the liquid bombers. It's how we found the Yemeni printer-cartridge bombs. And it's our best chance at stopping the next serious plot.

Because if a group of well-planned and well-funded terrorist plotters makes it to the airport, the chance is pretty low that those blue-shirted crotch-groping water-bottle-confiscating TSA agents are going to catch them. The agents are trying to do a good job, but the deck is so stacked against them that their job is impossible. Airport security is the last line of defense, and it's not a very good one.

We have a job here, too, and it's to be indomitable in the face of terrorism. The goal of terrorism is to terrorize us: to make us afraid, and make our government do exactly what the TSA is doing. When we react out of fear, the terrorists succeed even when their plots fail. But if we carry on as before, the terrorists fail—even when their plots succeed.

The Trouble with Airport Profiling

Originally published in Forbes, May 9, 2012

Why do otherwise rational people think it's a good idea to profile people at airports? Recently, neuroscientist and best-selling author Sam Harris related a story of an elderly couple being given the twice-over by the TSA, pointed out how these two were obviously not a threat, and recommended that the TSA focus on the actual threat: "Muslims, or anyone who looks like he or she could conceivably be Muslim."

This is a bad idea. It doesn't make us any safer—and it actually puts us all at risk.

The right way to look at security is in terms of cost-benefit trade-offs. If adding profiling to airport checkpoints allowed us to detect more threats at a lower cost, than we should implement it. If it didn't, we'd be foolish to do so. Sometimes profiling works. Consider a sheep in a meadow, happily munching on grass. When he spies a wolf, he's going to judge that individual wolf based on a bunch of assumptions related to the past behavior of its species. In short, that sheep is going to profile . . . and then run away. This makes perfect sense, and is why evolution produced sheep—and other animals—that react this way. But this sort of profiling doesn't work with humans at airports, for several reasons.

First, in the sheep's case the profile is accurate, in that all wolves are out to eat sheep. Maybe a particular wolf isn't hungry at the moment, but enough wolves are hungry enough of the time to justify the occasional false alarm. However, it isn't true that almost all Muslims are out to blow up airplanes. In fact, almost none of them are. Post 9/11, we've had 2 Muslim terrorists on US airplanes: the shoe bomber and the underwear bomber. If you assume 0.8% (that's one estimate of the percentage of Muslim Americans) of the 630 million annual airplane fliers are Muslim and triple it to account for others who look Semitic, then the chances any profiled flier will be a Muslim terrorist is 1 in 80 million. Add the 19 9/11 terrorists—arguably a singular event—that number drops to 1 in 8 million. Either way, because the number of actual terrorists is so low, almost everyone selected by the profile will be innocent. This is called the “base rate fallacy,” and dooms any type of broad terrorist profiling, including the TSA's behavioral profiling.

Second, sheep can safely ignore animals that don't look like the few predators they know. On the other hand, to assume that only Arab-appearing people are terrorists is dangerously naive. Muslims are black, white, Asian, and everything else—most Muslims are not Arab. Recent terrorists have been European, Asian, African, Hispanic, and Middle Eastern; male and female; young and old. Underwear bomber Umar Farouk Abdul Mutallab was Nigerian. Shoe bomber Richard Reid was British with a Jamaican father. One of the London subway bombers, Germaine Lindsay, was Afro-Caribbean. Dirty bomb suspect Jose Padilla was Hispanic-American. The 2002 Bali terrorists were Indonesian. Both Timothy McVeigh and the Unabomber were white Americans. The Chechen terrorists who blew up two Russian planes in 2004 were female. Focusing on a profile increases the risk that TSA agents will miss those who don't match it.

Third, wolves can't deliberately try to evade the profile. A wolf in sheep's clothing is just a story, but humans are smart and adaptable enough to put

the concept into practice. Once the TSA establishes a profile, terrorists will take steps to avoid it. The Chechens deliberately chose female suicide bombers because Russian security was less thorough with women. Al Qaeda has tried to recruit non-Muslims. And terrorists have given bombs to innocent—and innocent-looking—travelers. Randomized secondary screening is more effective, especially since the goal isn't to catch every plot but to create enough uncertainty that terrorists don't even try.

And fourth, sheep don't care if they offend innocent wolves; the two species are never going to be friends. At airports, though, there is an enormous social and political cost to the millions of false alarms. Beyond the societal harms of deliberately harassing a minority group, singling out Muslims alienates the very people who are in the best position to discover and alert authorities about Muslim plots before the terrorists even get to the airport. This alone is reason enough not to profile.

I too am incensed—but not surprised—when the TSA singles out four-year old girls, children with cerebral palsy, pretty women, the elderly, and wheelchair users for humiliation, abuse, and sometimes theft. Any bureaucracy that processes 630 million people per year will generate stories like this. When people propose profiling, they are really asking for a security system that can apply judgment. Unfortunately, that's really hard. Rules are easier to explain and train. Zero tolerance is easier to justify and defend. Judgment requires better-educated, more expert, and much-higher-paid screeners. And the personal career risks to a TSA agent of being wrong when exercising judgment far outweigh any benefits from being sensible.

The proper reaction to screening horror stories isn't to subject only "those people" to it; it's to subject no one to it. (Can anyone even explain what hypothetical terrorist plot could successfully evade normal security, but would be discovered during secondary screening?) Invasive TSA screening is nothing more than security theater. It doesn't make us safer, and it's not worth the cost. Even more strongly, security isn't our society's only value. Do we really want the full power of government to act out our stereotypes and prejudices? Have we Americans ever done something like this and not been ashamed later? This is what we have a Constitution for: to help us live up to our values and not down to our fears.

8

Security, Policy, Liberty, and Law

Memo to Next President: How to Get Cybersecurity Right

Originally published in Wired News, August 7, 2008

Obama has a cybersecurity plan.

It's basically what you would expect: Appoint a national cybersecurity adviser, invest in math and science education, establish standards for critical infrastructure, spend money on enforcement, establish national standards for securing personal data and data-breach disclosure, and work with industry and academia to develop a bunch of needed technologies.

I could comment on the plan, but with security, the devil is always in the details—and, of course, at this point there are few details. But since he brought up the topic—McCain supposedly is “working on the issues” as well—I have three pieces of policy advice for the next president, whoever he is. They're too detailed for campaign speeches or even position papers, but they're essential for improving information security in our society. Actually, they apply to national security in general. And they're things only government can do.

One, use your immense buying power to improve the security of commercial products and services. One property of technological products is that most of the cost is in the development of the product rather than the production. Think software: The first copy costs millions, but the second copy is free.

You have to secure your own government networks, military and civilian. You have to buy computers for all your government employees. Consolidate those contracts, and start putting explicit security requirements into the RFPs. You have the buying power to get your vendors to make serious security

improvements in the products and services they sell to the government, and then we all benefit because they'll include those improvements in the same products and services they sell to the rest of us. We're all safer if information technology is more secure, even though the bad guys can use it, too.

Two, legislate results and not methodologies. There are a lot of areas in security where you need to pass laws, where the security externalities are such that the market fails to provide adequate security. For example, software companies who sell insecure products are exploiting an externality just as much as chemical plants that dump waste into the river. But a bad law is worse than no law. A law requiring companies to secure personal data is good; a law specifying what technologies they should use to do so is not. Mandating software liabilities for software failures is good; detailing how is not. Legislate for the results you want and implement the appropriate penalties; let the market figure out how—that's what markets are good at.

Three, broadly invest in research. Basic research is risky; it doesn't always pay off. That's why companies have stopped funding it. Bell Labs is gone because nobody could afford it after the AT&T breakup, but the root cause was a desire for higher efficiency and short-term profitability—not unreasonable in an unregulated business. Government research can be used to balance that by funding long-term research.

Spread those research dollars wide. Lately, most research money has been redirected through DARPA to near-term military-related projects; that's not good. Keep the earmark-happy Congress from dictating how the money is spent. Let the NSF, NIH and other funding agencies decide how to spend the money and don't try to micromanage. Give the national laboratories lots of freedom, too. Yes, some research will sound silly to a layman. But you can't predict what will be useful for what, and if funding is really peer-reviewed, the average results will be much better. Compared with corporate tax breaks and other subsidies, this is chump change.

If our research capability is to remain vibrant, we need more science and math students with decent elementary and high school preparation. The declining interest is partly from the perception that scientists don't get rich like lawyers and dentists and stockbrokers, but also because science isn't valued in a country full of creationists. One way the president can help is by trusting scientific advisers and not overruling them for political reasons.

Oh, and get rid of those post-9/11 restrictions on student visas that are causing so many top students to do their graduate work in Canada, Europe and

Asia instead of in the United States. Those restrictions will hurt us immensely in the long run.

Those are the three big ones; the rest is in the details. And it's the details that matter. There are lots of serious issues that you're going to have to tackle: data privacy, data sharing, data mining, government eavesdropping, government databases, use of Social Security numbers as identifiers, and so on. It's not enough to get the broad policy goals right. You can have good intentions and enact a good law, and have the whole thing completely gutted by two sentences sneaked in during rulemaking by some lobbyist.

Security is both subtle and complex, and—unfortunately—it doesn't readily lend itself to normal legislative processes. You're used to finding consensus, but security by consensus rarely works. On the Internet, security standards are much worse when they're developed by a consensus body, and much better when someone just does them. This doesn't always work—a lot of crap security has come from companies that have “just done it”—but nothing but mediocre standards come from consensus bodies. The point is that you won't get good security without pissing someone off: The information-broker industry, the voting-machine industry, the telcos. The normal legislative process makes it hard to get security right, which is why I don't have much optimism about what you can get done.

And if you're going to appoint a cybersecurity czar, you have to give him actual budgetary authority—otherwise he won't be able to get anything done, either.

CRB Checking

Originally published in Schneier on Security, November 8, 2008

Since the UK's Criminal Records Bureau (CRB) was established in 2002, an ever-increasing number of people are required to undergo a “CRB check” before they can interact with children. It's not only teachers and daycare providers, but football coaches, scoutmasters and Guiders, church volunteers, bus drivers, and school janitors—3.4 million checks in 2007, 15 million since 2002. In 2009, it will include anyone who works or volunteers in a position where he or she comes into contact with children: 11.3 million people, or a quarter of the adult population.

This might make sense if it worked, but it doesn't. CRB checks don't keep child predators away from children. And even worse, this bureaucratic process fosters an atmosphere of mistrust among parents, teaches parents to ignore their own intuitions about other adults, and limits children's activities as organizations find CRB checks too cumbersome.

Effectiveness first. CRB checking does not guarantee that only non-predatory adults interact with children. At best, it only protects children from recidivist predators. This is a real risk but less than first-time predators, predatory relatives, or predators they come into casual contact with.

The CRB cites statistics like "Since 2004 the CRB has stopped 80,000 unsuitable people working with vulnerable groups," but that's just false. In the first place, there are only 30,000 people on the list. Also, denying someone a CRB approval isn't the correct metric—protecting children is. Before CRB checks, there weren't 20,000 repeat-offender child-predatory crimes annually. The effectiveness of this program is the difference between the small handful in 2001 and the smaller handful today.

The Home Office admits that 9,000 potential predators who should be on the list are not, and their error rate means that 2,000 innocent people will be falsely labeled as child predators in 2009. But it's more than the errors; the list is filled with people who shouldn't be on it. For example, underage teenagers could be put on the list for having consensual sex—nothing that should prevent them from taking a summer job around kids.

CRB checking might not be effective at stopping child predators, but it's effective at instilling fear in parents. Sociology professor Frank Furedi wrote the book *License to Hug* about the issue, chronicles some examples: a mother who couldn't kiss her child goodbye on a school trip, another barred from taking her child to a school event, and a father who gets "filthy looks" when he takes his child swimming. The most horrific example is the story of a bricklayer who did not help a lost two-year-old because he feared being thought an abductor. The girl later drowned.

Pervasive CRB checking also teaches parents to ignore their own parenting instincts. As Furedi says, "If adults are not expected to respond to problems in accordance with their experience and intuition they will have little incentive to develop the kind of skills required to manage children and young people."

The assumption—contrary to all data—that everyone is a child predator unless checked by the police poisons the natural relationship between children and adults, and directly affects their welfare. Half of all adults fear being falsely accused of being abusers and 13% of men don't volunteer as a result.

At the same time, 50,000 girls can't join the Guides because of a shortage of CRB-checked adults, and kids' sports leagues are drastically cutting back.

A parent's natural reaction is "If it were your child, you'd do anything . . ." This is natural—we're all predisposed to exaggerate risks that are extreme, that are from strangers, and that are against our children. Over several generations, children have had less freedom than their parents, even though the data doesn't warrant it. When evaluating risks, we respond more to stories than to data.

This enormous government hydra, first proposed in 1996 and which Furedi estimates has cost half a billion pounds since it was implemented 2002, seems to have been sold primarily on two stories: the 1997 murder of ten-year-old Scott Simpson by convicted sex offender Stephen Leisk, and the 2002 murders of 10-year-olds Jessica Chapman and Holly Wells by school caretaker Ian Huntley. As gruesome as these stories are, they're don't make a basis for sound policy. Government database checking is no substitute for alert parenting.

State Data Breach Notification Laws: Have They Helped?

Originally published in Information Security, January 2009

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

There are three reasons for breach notification laws. One, it's common politeness that when you lose something of someone else's, you tell him. The prevailing corporate attitude before the law—"They won't notice, and if they do notice they won't know it's us, so we are better off keeping quiet about the whole thing"—is just wrong. Two, it provides statistics to security researchers as to how pervasive the problem really is. And three, it forces companies to improve their security.

That last point needs a bit of explanation. The problem with companies protecting your data is that it isn't in their financial best interest to do so. That is, the companies are responsible for protecting your data, but bear none of the costs if your data is compromised. You suffer the harm, but you have no control—or even knowledge—of the company's security practices. The idea behind such laws, and how they were sold to legislators, is that they would increase the cost—both in bad publicity and the actual notification—of

security breaches, motivating companies to spend more to prevent them. In economic terms, the law reduces the externalities and forces companies to deal with the true costs of these data breaches.

So how has it worked?

Earlier this year, three researchers at the Heinz School of Public Policy and Management at Carnegie Mellon University—Sasha Romanosky, Rahul Telang and Alessandro Acquisti—tried to answer that question. They looked at reported data breaches and rates of identity theft from 2002 to 2007, comparing states with a law to states without one. If these laws had their desired effects, people in states with notification laws should experience fewer incidences of identity theft. The result: not so much. The researchers found data breach notification laws reduced identity theft by just 2 percent on average.

I think there's a combination of things going on. Identity theft is being reported far more today than five years ago, so it's difficult to compare identity theft rates before and after the state laws were enacted. Most identity theft occurs when someone's home or work computer is compromised, not from theft of large corporate databases, so the effect of these laws is small. Most of the security improvements companies made didn't make much of a difference, reducing the effect of these laws.

The laws rely on public shaming. It's embarrassing to have to admit to a data breach, and companies should be willing to spend to avoid this PR expense. The problem is, in order for this to work well, public shaming needs the cooperation of the press. And there's an attenuation effect going on. The first major breach after the first state disclosure law was in February 2005 in California, when ChoicePoint sold personal data on 145,000 people to criminals. The event was big news, ChoicePoint's stock tanked, and it was shamed into improving its security.

Next, LexisNexis exposed personal data on 300,000 individuals, and then Citigroup lost data on 3.9 million. The law worked; the only reason we knew about these security breaches was because of the law. But the breaches came in increasing numbers, and in larger quantities. Data breach stories felt more like "crying wolf" and soon, data breaches were no longer news.

Today, the remaining cost is that of the direct mail campaign to notify customers, which often turns into a marketing opportunity.

I'm still a fan of these laws, if only for the first two reasons I listed. Disclosure is important, but it's not going to solve identity theft. As I've written previously, the reason theft of personal information is common is that the data is valuable

once stolen. The way to mitigate the risk of fraud due to impersonation is not to make personal information difficult to steal, it's to make it difficult to use.

Disclosure laws only deal with the economic externality of data owners protecting your personal information. What we really need are laws prohibiting financial institutions from granting credit to someone using your name with only a minimum of authentication.

How to Ensure Police Database Accuracy

Originally published in the Wall Street Journal, January 27, 2009

Earlier this month, the Supreme Court ruled that evidence gathered as a result of errors in a police database is admissible in court. Their narrow decision is wrong, and will only ensure that police databases remain error-filled in the future.

The specifics of the case are simple. A computer database said there was a felony arrest warrant pending for Bennie Herring when there actually wasn't. When the police came to arrest him, they searched his home and found illegal drugs and a gun. The Supreme Court was asked to rule whether the police had the right to arrest him for possessing those items, even though there was no legal basis for the search and arrest in the first place.

What's at issue here is the exclusionary rule, which basically says that unconstitutionally or illegally collected evidence is inadmissible in court. It might seem like a technicality, but excluding what is called "the fruit of the poisonous tree" is a security system designed to protect us all from police abuse.

We have a number of rules limiting what the police can do: rules governing arrest, search, interrogation, detention, prosecution, and so on. And one of the ways we ensure that the police follow these rules is by forbidding the police to receive any benefit from breaking them. In fact, we design the system so that the police actually harm their own interests by breaking them, because all evidence that stems from breaking the rules is inadmissible.

And that's what the exclusionary rule does. If the police search your home without a warrant and find drugs, they can't arrest you for possession. Since

the police have better things to do than waste their time, they have an incentive to get a warrant.

The Herring case is more complicated, because the police thought they did have a warrant. The error was not a police error, but a database error. And, in fact, Judge Roberts wrote for the majority: “The exclusionary rule serves to deter deliberate, reckless, or grossly negligent conduct, or in some circumstances recurring or systemic negligence. The error in this case does not rise to that level.”

Unfortunately, Roberts is wrong. Government databases are filled with errors. People often can't see data about themselves, and have no way to correct the errors if they do learn of any. And more and more databases are trying to exempt themselves from the Privacy Act of 1974, and specifically the provisions that require data accuracy. The legal argument for excluding this evidence was best made by an *amicus curiae* filed by the Electronic Privacy Information Center, but in short, the court should exclude the evidence because it's the only way to ensure police database accuracy.

We are protected from becoming a police state by limits on police power and authority. This is not a trade-off we make lightly: we deliberately hamper law enforcement's ability to do its job because we recognize that these limits make us safer. Without the exclusionary rule, your only remedy against an illegal search is to bring legal action against the police—and that can be very difficult. We, the people, would rather have you go free than motivate the police to ignore the rules that limit their power.

By not applying the exclusionary rule in the Herring case, the Supreme Court missed an important opportunity to motivate the police to purge errors from their databases. Constitutional lawyers have written many articles about this ruling, but the most interesting idea comes from George Washington University professor Daniel J. Solove, who proposes this compromise: “If a particular database has reasonable protections and deterrents against errors, then the Fourth Amendment exclusionary rule should not apply. If not, then the exclusionary rule should apply. Such a rule would create an incentive for law enforcement officials to maintain accurate databases, to avoid all errors, and would ensure that there would be a penalty or consequence for errors.”

Increasingly, we are being judged by the trail of data we leave behind us. Increasingly, data accuracy is vital to our personal safety and security. And if errors made by police databases aren't held to the same legal standard as errors made by policemen, then more and more innocent Americans will find themselves the victims of incorrect data.

How Perverse Incentives Drive Bad Security Decisions

Originally published in Wired News, February 26, 2009

An employee of Whole Foods in Ann Arbor, Michigan, was fired in 2007 for apprehending a shoplifter. More specifically, he was fired for touching a customer, even though that customer had a backpack filled with stolen groceries and was running away with them.

I regularly see security decisions that, like the Whole Foods incident, seem to make absolutely no sense. However, in every case, the decisions actually make perfect sense once you understand the underlying incentives driving the decision. All security decisions are trade-offs, but the motivations behind them are not always obvious: They're often subjective, and driven by external incentives. And often security trade-offs are made for nonsecurity reasons.

Almost certainly, Whole Foods has a no-touching-the-customer policy because its attorneys recommended it. "No touching" is a security measure as well, but it's security against customer lawsuits. The cost of these lawsuits would be much, much greater than the \$346 worth of groceries stolen in this instance. Even applied to suspected shoplifters, the policy makes sense: The cost of a lawsuit resulting from tackling an innocent shopper by mistake would be far greater than the cost of letting actual shoplifters get away. As perverse as it may seem, the result is completely reasonable given the corporate incentives—Whole Foods wrote a corporate policy that benefited itself.

At least, it works as long as the police and other factors keep society's shoplifter population down to a reasonable level.

Incentives explain much that is perplexing about security trade-offs. Why does King County, Washington, require one form of ID to get a concealed-carry permit, but two forms of ID to pay for the permit by check? Making a mistake on a gun permit is an abstract problem, but a bad check actually costs some department money.

In the decades before 9/11, why did the airlines fight every security measure except the photo-ID check? Increased security annoys their customers, but the photo-ID check solved a security problem of a different kind: the resale of nonrefundable tickets. So the airlines were on board for that one.

And why does the TSA confiscate liquids at airport security, on the off chance that a terrorist will try to make a liquid explosive instead of using the

more common solid ones? Because the officials in charge of the decision used CYA security measures to prevent specific, known tactics rather than broad, general ones.

The same misplaced incentives explain the ongoing problem of innocent prisoners spending years in places like Guantanamo and Abu Ghraib. The solution might seem obvious: Release the innocent ones, keep the guilty ones, and figure out whether the ones we aren't sure about are innocent or guilty. But the incentives are more perverse than that. Who is going to sign the order releasing one of those prisoners? Which military officer is going to accept the risk, no matter how small, of being wrong?

I read almost five years ago that prisoners were being held by the United States far longer than they should, because “no one wanted to be responsible for releasing the next Osama bin Laden.” That incentive to do nothing hasn't changed. It might have even gotten stronger, as these innocents languish in prison.

In all these cases, the best way to change the trade-off is to change the incentives. Look at why the Whole Foods case works. Store employees don't have to apprehend shoplifters, because society created a special organization specifically authorized to lay hands on people the grocery store points to as shoplifters: the police. If we want more rationality out of the TSA, there needs to be someone with a broader perspective willing to deal with general threats rather than specific targets or tactics.

For prisoners, society has created a special organization specifically entrusted with the role of judging the evidence against them and releasing them if appropriate: the judiciary. It's only because the George W. Bush administration decided to remove the Guantanamo prisoners from the legal system that we are now stuck with these perverse incentives. Our country would be smart to move as many of these people through the court system as we can.

It's Time to Drop the “Expectation of Privacy” Test

Originally published in Wired News, March 26, 2009

In the United States, the concept of “expectation of privacy” matters because it's the constitutional test, based on the Fourth Amendment, that governs when and how the government can invade your privacy.

Based on the 1967 *Katz v. United States* Supreme Court decision, this test actually has two parts. First, the government's action can't contravene an individual's subjective expectation of privacy; and second, that expectation of privacy must be one that society in general recognizes as reasonable. That second part isn't based on anything like polling data; it is more of a normative idea of what level of privacy people should be allowed to expect, given the competing importance of personal privacy on one hand and the government's interest in public safety on the other.

The problem is, in today's information society, that definition test will rapidly leave us with no privacy at all.

In *Katz*, the Court ruled that the police could not eavesdrop on a phone call without a warrant: Katz expected his phone conversations to be private and this expectation resulted from a reasonable balance between personal privacy and societal security. Given NSA's large-scale warrantless eavesdropping, and the previous administration's continual insistence that it was necessary to keep America safe from terrorism, is it still reasonable to expect that our phone conversations are private?

Between the NSA's massive Internet eavesdropping program and Gmail's content-dependent advertising, does anyone actually expect their e-mail to be private? Between calls for ISPs to retain user data and companies serving content-dependent web ads, does anyone expect their web browsing to be private? Between the various computer-infecting malware, and world governments increasingly demanding to see laptop data at borders, hard drives are barely private. I certainly don't believe that my SMSs, any of my telephone data, or anything I say on LiveJournal or Facebook—regardless of the privacy settings—is private.

Aerial surveillance, data mining, automatic face recognition, terahertz radar that can "see" through walls, wholesale surveillance, brain scans, RFID, "life recorders" that save everything: Even if society still has some small expectation of digital privacy, that will change as these and other technologies become ubiquitous. In short, the problem with a normative expectation of privacy is that it changes with perceived threats, technology and large-scale abuses.

Clearly, something has to change if we are to be left with any privacy at all. Three legal scholars have written law review articles that wrestle with the problems of applying the Fourth Amendment to cyberspace and to our computer-mediated world in general.

George Washington University's Daniel Solove, who blogs at *Concurring Opinions*, has tried to capture the byzantine complexities of modern privacy.

He points out, for example, that the following privacy violations—all real—are very different: A company markets a list of 5 million elderly incontinent women; reporters deceitfully gain entry to a person's home and secretly photograph and record the person; the government uses a thermal sensor device to detect heat patterns in a person's home; and a newspaper reports the name of a rape victim. Going beyond simple definitions such as the divulging of a secret, Solove has developed a taxonomy of privacy, and the harms that result from their violation.

His 16 categories are: surveillance, interrogation, aggregation, identification, insecurity, secondary use, exclusion, breach of confidentiality, disclosure, exposure, increased accessibility, blackmail, appropriation, distortion, intrusion and decisional interference. Solove's goal is to provide a coherent and comprehensive understanding of what is traditionally an elusive and hard-to-explain concept: privacy violations. (This taxonomy is also discussed in Solove's book, *Understanding Privacy*.)

Orin Kerr, also a law professor at George Washington University, and a blogger at Volokh Conspiracy, has attempted to lay out general principles for applying the Fourth Amendment to the Internet. First, he points out that the traditional inside/outside distinction—the police can watch you in a public place without a warrant, but not in your home—doesn't work very well with regard to cyberspace. Instead, he proposes a distinction between content and non-content information: the body of an e-mail versus the header information, for example. The police should be required to get a warrant for the former, but not for the latter. Second, he proposes that search warrants should be written for particular individuals and not for particular Internet accounts.

Meanwhile, Jed Rubenfeld of Yale Law School has tried to reinterpret the Fourth Amendment not in terms of privacy, but in terms of security. Pointing out that the whole "expectations" test is circular—what the government does affects what the government can do—he redefines everything in terms of security: the security that our private affairs are private.

This security is violated when, for example, the government makes widespread use of informants, or engages in widespread eavesdropping—even if no one's privacy is actually violated. This neatly bypasses the whole individual privacy versus societal security question—a balancing that the individual usually loses—by framing both sides in terms of personal security.

I have issues with all of these articles. Solove's taxonomy is excellent, but the sense of outrage that accompanies a privacy violation—"How could they know/do/say that!?"—is an important part of the harm resulting from a privacy violation. The non-content information that Kerr believes should be collectible

without a warrant can be very private and personal: URLs can be very personal, and it's possible to figure out browsed content just from the size of encrypted SSL traffic. Also, the ease with which the government can collect all of it—the calling and called party of every phone call in the country—makes the balance very different. I believe these need to be protected with a warrant requirement. Rubinfeld's reframing is interesting, but the devil is in the details. Reframing privacy in terms of security still results in a balancing of competing rights. I'd rather take the approach of stating the—obvious to me—individual and societal value of privacy, and giving privacy its rightful place as a fundamental human right. (There's additional commentary on Rubinfeld's thesis at ArsTechnica.)

The trick here is to realize that a normative definition of the expectation of privacy doesn't need to depend on threats or technology, but rather on what we—as society—decide it should be. Sure, today's technology makes it easier than ever to violate privacy. But it doesn't necessarily follow that we have to violate privacy. Today's guns make it easier than ever to shoot virtually anyone for any reason. That doesn't mean our laws have to change.

No one knows how this will shake out legally. These three articles are from law professors; they're not judicial opinions. But clearly something has to change, and ideas like these may someday form the basis of new Supreme Court decisions that brings legal notions of privacy into the 21st century.

Who Should Be in Charge of Cybersecurity?

Originally published in the Wall Street Journal, March 31, 2009

US government cybersecurity is an insecure mess, and fixing it is going to take considerable attention and resources. Trying to make sense of this, President Barack Obama ordered a 60-day review of government cybersecurity initiatives. Meanwhile, the US House Subcommittee on Emerging Threats, Cybersecurity, Science and Technology is holding hearings on the same topic.

One of the areas of contention is who should be in charge. The FBI, DHS and DoD—specifically, the NSA—all have interests here. Earlier this month, Rod Beckström resigned from his position as director of the DHS's National Cybersecurity Center, warning of a power grab by the NSA.

Putting national cybersecurity in the hands of the NSA is an incredibly bad idea. An entire parade of people, ranging from former FBI director Louis Freeh to Microsoft's Trusted Computing Group Vice President and former Justice Department computer crime chief Scott Charney, have told Congress the same thing at this month's hearings.

Cybersecurity isn't a military problem, or even a government problem—it's a universal problem. All networks, military, government, civilian and commercial, use the same computers, the same networking hardware, the same Internet protocols and the same software packages. We all are the targets of the same attack tools and tactics. It's not even that government targets are somehow more important; these days, most of our nation's critical IT infrastructure is in commercial hands. Government-sponsored Chinese hackers go after both military and civilian targets.

Some have said that the NSA should be in charge because it has specialized knowledge. Earlier this month, Director of National Intelligence Admiral Dennis Blair made this point, saying "There are some wizards out there at Ft. Meade who can do stuff." That's probably not true, but if it is, we'd better get them out of Ft. Meade as soon as possible—they're doing the nation little good where they are now.

Not that government cybersecurity failings require any specialized wizardry to fix. GAO reports indicate that government problems include insufficient access controls, a lack of encryption where necessary, poor network management, failure to install patches, inadequate audit procedures, and incomplete or ineffective information security programs. These aren't super-secret NSA-level security issues; these are the same managerial problems that every corporate CIO wrestles with.

We've all got the same problems, so solutions must be shared. If the government has any clever ideas to solve its cybersecurity problems, certainly a lot of us could benefit from those solutions. If it has an idea for improving network security, it should tell everyone. The best thing the government can do for cybersecurity world-wide is to use its buying power to improve the security of the IT products everyone uses. If it imposes significant security requirements on its IT vendors, those vendors will modify their products to meet those requirements. And those same products, now with improved security, will become available to all of us as the new standard.

Moreover, the NSA's dual mission of providing security and conducting surveillance means it has an inherent conflict of interest in cybersecurity. Inside the NSA, this is called the "equities issue." During the Cold War, it was easy; the NSA used its expertise to protect American military information and communications, and eavesdropped on Soviet information and communications. But what happens when both the good guys the NSA wants to protect, and the bad guys the NSA wants to eavesdrop on, use the same systems? They all use Microsoft Windows, Oracle databases, Internet email, and Skype. When the NSA finds a vulnerability in one of those systems, does it alert the manufacturer and fix it—making both the good guys and the bad guys more secure? Or does it keep quiet about the vulnerability and not tell anyone—making it easier to spy on the bad guys but also keeping the good guys insecure? Programs like the NSA's warrantless wiretapping program have created additional vulnerabilities in our domestic telephone networks.

Testifying before Congress earlier this month, former DHS National Cyber Security division head Amit Yoran said "the intelligence community has always and will always prioritize its own collection efforts over the defensive and protection mission of our government's and nation's digital systems."

Maybe the NSA could convince us that it's putting cybersecurity first, but its culture of secrecy will mean that any decisions it makes will be suspect. Under current law, extended by the Bush administration's extravagant invocation of the "state secrets" privilege when charged with statutory and constitutional violations, the NSA's activities are not subject to any meaningful public oversight. And the NSA's tradition of military secrecy makes it harder for it to coordinate with other government IT departments, most of which don't have clearances, let alone coordinate with local law enforcement or the commercial sector.

We need transparent and accountable government processes, using commercial security products. We need government cybersecurity programs that improve security for everyone. The NSA certainly has an advisory and a coordination role in national cybersecurity, and perhaps a more supervisory role in DoD cybersecurity—both offensive and defensive—but it should not be in charge.

Coordinate, but Distribute Responsibility

Originally published in NYTimes.com, May 29, 2009

This essay appeared as part of a round table about Obama's cybersecurity speech on the New York Times Room for Debate blog.

I am optimistic about President Obama's new cybersecurity policy and the appointment of a new "cybersecurity coordinator," though much depends on the details. What we do know is that the threats are real, from identity theft to Chinese hacking to cyberwar.

His principles were all welcome—securing government networks, coordinating responses, working to secure the infrastructure in private hands (the power grid, the communications networks, and so on), although I think he's overly optimistic that legislation won't be required. I was especially heartened to hear his commitment to funding research. Much of the technology we currently use to secure cyberspace was developed from university research, and the more of it we finance today the more secure we'll be in a decade.

Education is also vital, although sometimes I think my parents need more cybersecurity education than my grandchildren do. I also appreciate the president's commitment to transparency and privacy, both of which are vital for security.

But the details matter. Centralizing security responsibilities has the downside of making security more brittle by instituting a single approach and a uniformity of thinking. Unless the new coordinator distributes responsibility, cybersecurity won't improve.

As the administration moves forward on the plan, two principles should apply. One, security decisions need to be made as close to the problem as possible. Protecting networks should be done by people who understand those networks, and threats needs to be assessed by people close to the threats. But distributed responsibility has more risk, so oversight is vital.

Two, security coordination needs to happen at the highest level possible, whether that's evaluating information about different threats, responding to an Internet worm or establishing guidelines for protecting personal information. The whole picture is larger than any single agency.

“Zero Tolerance” Really Means Zero Discretion

Originally published in MPR NewsQ, November 4, 2009

Recent stories have documented the ridiculous effects of zero-tolerance weapons policies in a Delaware school district: a first-grader expelled for taking a camping utensil to school, a 13-year-old expelled after another student dropped a pocketknife in his lap, and a seventh-grader expelled for cutting paper with a utility knife for a class project. Where’s the common sense? the editorials cry.

These so-called zero-tolerance policies are actually zero-discretion policies. They’re policies that must be followed, no situational discretion allowed. We encounter them whenever we go through airport security: no liquids, gels or aerosols. Some workplaces have them for sexual harassment incidents; in some sports a banned substance found in a urine sample means suspension, even if it’s for a real medical condition. Judges have zero discretion when faced with mandatory sentencing laws: three strikes for drug offences and you go to jail, mandatory sentencing for statutory rape (underage sex), etc. A national restaurant chain won’t serve hamburgers rare, even if you offer to sign a waiver. Whenever you hear “that’s the rule, and I can’t do anything about it”—and they’re not lying to get rid of you—you’re butting against a zero discretion policy.

These policies enrage us because they are blind to circumstance. Editorial after editorial denounced the suspensions of elementary school children for offenses that anyone with any common sense would agree were accidental and harmless. The Internet is filled with essays demonstrating how the TSA’s rules are nonsensical and sometimes don’t even improve security. I’ve written some of them. What we want is for those involved in the situations to have discretion.

However, problems with discretion were the reason behind these mandatory policies in the first place. Discretion is often applied inconsistently. One school principal might deal with knives in the classroom one way, and another principal another way. Your drug sentence could depend considerably on how sympathetic your judge is, or on whether she’s having a bad day.

Even worse, discretion can lead to discrimination. Schools had weapons bans before zero-tolerance policies, but teachers and administrators enforced the rules disproportionately against African-American students. Criminal

sentences varied by race, too. The benefit of zero-discretion rules and laws is that they ensure that everyone is treated equally.

Zero-discretion rules also protect against lawsuits. If the rules are applied consistently, no parent, air traveler or defendant can claim he was unfairly discriminated against.

So that's the choice. Either we want the rules enforced fairly across the board, which means limiting the discretion of the enforcers at the scene at the time, or we want a more nuanced response to whatever the situation is, which means we give those involved in the situation more discretion.

Of course, there's more to it than that. The problem with the zero-tolerance weapons rules isn't that they're rigid, it's that they're poorly written.

What constitutes a weapon? Is it any knife, no matter how small? Should the penalties be the same for a first grader and a high school student? Does intent matter? When an aspirin carried for menstrual cramps becomes "drug possession," you know there's a badly written rule in effect.

It's the same with airport security and criminal sentencing. Broad and simple rules may be simpler to follow—and require less thinking on the part of those enforcing them—but they're almost always far less nuanced than our complex society requires. Unfortunately, the more complex the rules are, the more they're open to interpretation and the more discretion the interpreters have.

The solution is to combine the two, rules and discretion, with procedures to make sure they're not abused. Provide rules, but don't make them so rigid that there's no room for interpretation. Give the people in the situation—the teachers, the airport security agents, the policemen, the judges—discretion to apply the rules to the situation. But—and this is the important part—allow people to appeal the results if they feel they were treated unfairly. And regularly audit the results to ensure there is no discrimination or favoritism. It's the combination of the four that work: rules plus discretion plus appeal plus audit.

All systems need some form of redress, whether it be open and public like a courtroom or closed and secret like the TSA. Giving discretion to those at the scene just makes for a more efficient appeals process, since the first level of appeal can be handled on the spot.

Zachary, the Delaware first grader suspended for bringing a combination fork, spoon and knife camping utensil to eat his lunch with, had his punishment unanimously overturned by the school board. This was the right decision; but what about all the other students whose parents weren't as forceful or media-savvy enough to turn their child's plight into a national story? Common sense in applying rules is important, but so is equal access to that common sense.

US Enables Chinese Hacking of Google

Originally published in CNN, January 23, 2010

Google made headlines when it went public with the fact that Chinese hackers had penetrated some of its services, such as Gmail, in a politically motivated attempt at intelligence gathering. The news here isn't that Chinese hackers engage in these activities or that their attempts are technically sophisticated—we knew that already—it's that the US government inadvertently aided the hackers.

In order to comply with government search warrants on user data, Google created a backdoor access system into Gmail accounts. This feature is what the Chinese hackers exploited to gain access.

Google's system isn't unique. Democratic governments around the world—in Sweden, Canada and the UK, for example—are rushing to pass laws giving their police new powers of Internet surveillance, in many cases requiring communications system providers to redesign products and services they sell.

Many are also passing data retention laws, forcing companies to retain information on their customers. In the US, the 1994 Communications Assistance for Law Enforcement Act required phone companies to facilitate FBI eavesdropping, and since 2001, the National Security Agency has built substantial eavesdropping systems with the help of those phone companies.

Systems like these invite misuse: criminal appropriation, government abuse and stretching by everyone possible to apply to situations that are applicable only by the most tortuous logic. The FBI illegally wiretapped the phones of Americans, often falsely invoking terrorism emergencies, 3,500 times between 2002 and 2006 without a warrant. Internet surveillance and control will be no different.

Official misuses are bad enough, but it's the unofficial uses that worry me more. Any surveillance and control system must itself be secured. An infrastructure conducive to surveillance and control invites surveillance and control, both by the people you expect and by the people you don't.

China's hackers subverted the access system Google put in place to comply with US intercept orders. Why does anyone think criminals won't be able to use the same system to steal bank account and credit card information, use it to launch other attacks or turn it into a massive spam-sending network? Why does anyone think that only authorized law enforcement can mine collected Internet data or eavesdrop on phone and IM conversations?

These risks are not merely theoretical. After September 11, the NSA built a surveillance infrastructure to eavesdrop on telephone calls and e-mails within the US. Although procedural rules stated that only non-Americans and international phone calls were to be listened to, actual practice didn't match those rules. NSA analysts collected more data than they were authorized to and used the system to spy on wives, girlfriends and notables such as President Clinton.

But that's not the most serious misuse of a telecommunications surveillance infrastructure. In Greece, between June 2004 and March 2005, someone wiretapped more than 100 cell phones belonging to members of the Greek government: the prime minister and the ministers of defense, foreign affairs and justice.

Ericsson built this wiretapping capability into Vodafone's products and enabled it only for governments that requested it. Greece wasn't one of those governments, but someone still unknown—A rival political party? Organized crime? Foreign intelligence?—figured out how to surreptitiously turn the feature on.

And surveillance infrastructure can be exported, which also aids totalitarianism around the world. Western companies like Siemens and Nokia built Iran's surveillance. US companies helped build China's electronic police state. Just last year, Twitter's anonymity saved the lives of Iranian dissidents, anonymity that many governments want to eliminate.

In the aftermath of Google's announcement, some members of Congress are reviving a bill banning US tech companies from working with governments that digitally spy on their citizens. Presumably, those legislators don't understand that their own government is on the list.

This problem isn't going away. Every year brings more Internet censorship and control, not just in countries like China and Iran but in the US, the UK, Canada, and other free countries, egged on by both law enforcement trying to catch terrorists, child pornographers and other criminals, and by media companies trying to stop file sharers.

The problem is that such control makes us all less safe. Whether the eavesdroppers are the good guys or the bad guys, these systems put us all at greater risk. Communications systems that have no inherent eavesdropping capabilities are more secure than systems with those capabilities built in. And it's bad civic hygiene to build technologies that could someday be used to facilitate a police state.

Should the Government Stop Outsourcing Code Development?

Originally published in Information Security, March 2010

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

Information technology is increasingly everywhere, and it's the same technologies everywhere. The same operating systems are used in corporate and government computers. The same software controls critical infrastructure and home shopping. The same networking technologies are used in every country. The same digital infrastructure underpins the small and the large, the important and the trivial, the local and the global; the same vendors, the same standards, the same protocols, the same applications.

With all of this sameness, you'd think these technologies would be designed to the highest security standard, but they're not. They're designed to the lowest or, at best, somewhere in the middle. They're designed sloppily, in an ad hoc manner, with efficiency in mind. Security is a requirement, more or less, but it's a secondary priority. It's far less important than functionality, and security is what gets compromised when schedules get tight.

Should the government—ours, someone else's?—stop outsourcing code development? That's the wrong question to ask. Code isn't magically more secure when it's written by someone who receives a government paycheck than when it's written by someone who receives a corporate paycheck. It's not magically less secure when it's written by someone who speaks a foreign language, or is paid by the hour instead of by salary. Writing all your code in-house isn't even a viable option anymore; we're all stuck with software written by who-knows-whom in who-knows-which-country. And we need to figure out how to get security from that.

The traditional solution has been defense in depth: layering one mediocre security measure on top of another mediocre security measure. So we have the security embedded in our operating system and applications software, the security embedded in our networking protocols, and our additional security products such as antivirus and firewalls. We hope that whatever security flaws—either found and exploited, or deliberately inserted—there are in one layer are counteracted by the security in another layer, and that when they're not, we can patch our systems quickly enough to avoid serious long-term damage. That is a lousy solution when you think about it, but we've been more-or-less managing with it so far.

Bringing all software—and hardware, I suppose—development in-house under some misconception that proximity equals security is not a better solution. What we need is to improve the software development process, so we can have some assurance that our software is secure—regardless of what coder, employed by what company, and living in what country, writes it. The key word here is “assurance.”

Assurance is less about developing new security techniques than about using the ones we already have. It’s all the things described in books on secure coding practices. It’s what Microsoft is trying to do with its Security Development Lifecycle. It’s the Department of Homeland Security’s Build Security In program. It’s what every aircraft manufacturer goes through before it fields a piece of avionics software. It’s what the NSA demands before it purchases a piece of security equipment. As an industry, we know how to provide security assurance in software and systems. But most of the time, we don’t care; commercial software, as insecure as it is, is good enough for most purposes.

Assurance is expensive, in terms of money and time, for both the process and the documentation. But the NSA needs assurance for critical military systems and Boeing needs it for its avionics. And the government needs it more and more: for voting machines, for databases entrusted with our personal information, for electronic passports, for communications systems, for the computers and systems controlling our critical infrastructure. Assurance requirements should be more common in government IT contracts.

The software used to run our critical infrastructure—government, corporate, everything— isn’t very secure, and there’s no hope of fixing it anytime soon. Assurance is really our only option to improve this, but it’s expensive and the market doesn’t care. Government has to step in and spend the money where its requirements demand it, and then we’ll all benefit when we buy the same software.

Punishing Security Breaches

Originally published in Schneier on Security, April 26, 2009

The editor of the Freakonomics blog asked me to write about this topic. The idea was that they would get several opinions, and publish them all. They spiked the story, but I’d already written my piece. So here it is.

In deciding what to do with Gray Powell, the Apple employee who accidentally left a secret prototype 4G iPhone in a California bar, Apple needs to figure

out how much of the problem is due to an employee not following the rules, and how much of the problem is due to unclear, unrealistic, or just plain bad rules.

If Powell sneaked the phone out of the Apple building in a flagrant violation of the rules—maybe he wanted to show it to a friend—he should be disciplined, perhaps even fired. Some military installations have rules like that. If someone wants to take something classified out of a top secret military compound, he might have to secrete it on his person and deliberately sneak it past a guard who searches briefcases and purses. He might be committing a crime by doing so, by the way. Apple isn't the military, of course, but if its corporate security policy is that strict, it may very well have rules like that. And the only way to ensure rules are followed is by enforcing them, and that means severe disciplinary action against those who bypass the rules.

Even if Powell had authorization to take the phone out of Apple's labs—presumably someone has to test drive the new toys sooner or later—the corporate rules might have required him to pay attention to it at all times. We've all heard of military attachés who carry briefcases chained to their wrists. It's an extreme example, but demonstrates how a security policy can allow for objects to move around town—or around the world—without getting lost. Apple almost certainly doesn't have a policy as rigid as that, but its policy might explicitly prohibit Powell from taking that phone into a bar, putting it down on a counter, and participating in a beer tasting. Again, if Apple's rules and Powell's violation were both that clear, Apple should enforce them.

On the other hand, if Apple doesn't have clear-cut rules, if Powell wasn't prohibited from taking the phone out of his office, if engineers routinely ignore or bypass security rules and—as long as nothing bad happens—no one complains, then Apple needs to understand that the system is more to blame than the individual. Most corporate security policies have this sort of problem. Security is important, but it's quickly jettisoned when there's an important job to be done. A common example is passwords: people aren't supposed to share them, unless it's really important and they have to. Another example is guest accounts. And doors that are supposed to remain locked but rarely are. People routinely bypass security policies if they get in the way, and if no one complains, those policies are effectively meaningless.

Apple's unfortunately public security breach has given the company an opportunity to examine its policies and figure out how much of the problem is Powell and how much of it is the system he's a part of. Apple needs to fix its security problem, but only after it figures out where the problem is.

Three Reasons to Kill the Internet Kill Switch Idea

Originally published in AOL News, July 9, 2010

Last month, Sen. Joe Lieberman, I-Conn., introduced a bill that might—we're not really sure—give the president the authority to shut down all or portions of the Internet in the event of an emergency. It's not a new idea. Sens. Jay Rockefeller, D-W.Va., and Olympia Snowe, R-Maine, proposed the same thing last year, and some argue that the president can already do something like this. If this or a similar bill ever passes, the details will change considerably and repeatedly. So let's talk about the idea of an Internet kill switch in general.

It's a bad one.

Security is always a trade-off: costs versus benefits. So the first question to ask is: What are the benefits? There is only one possible use of this sort of capability, and that is in the face of a warfare-caliber enemy attack. It's the primary reason lawmakers are considering giving the president a kill switch. They know that shutting off the Internet, or even isolating the US from the rest of the world, would cause damage, but they envision a scenario where not doing so would cause even more.

That reasoning is based on several flawed assumptions.

Internet without Borders

The first flawed assumption is that cyberspace has traditional borders, and we could somehow isolate ourselves from the rest of the world using an electronic Maginot Line. We can't.

Yes, we can cut off almost all international connectivity, but there are lots of ways to get out onto the Internet: satellite phones, obscure ISPs in Canada and Mexico, long-distance phone calls to Asia.

The Internet is the largest communications system mankind has ever created, and it works because it is distributed. There is no central authority. No nation is in charge. Plugging all the holes isn't possible.

Even if the president ordered all US Internet companies to block, say, all packets coming from China, or restrict non-military communications, or just shut down access in the greater New York area, it wouldn't work. You can't

figure out what packets do just by looking at them; if you could, defending against worms and viruses would be much easier.

And packets that come with return addresses are easy to spoof. Remember the cyberattack July 4, 2009, that probably came from North Korea, but might have come from England, or maybe Florida? On the Internet, disguising traffic is easy. And foreign cyberattackers could always have dial-up accounts via US phone numbers and make long-distance calls to do their misdeeds.

Unpredictable Side Effects

The second flawed assumption is that we can predict the effects of such a shutdown. The Internet is the most complex machine mankind has ever built, and shutting down portions of it would have all sorts of unforeseen ancillary effects.

Would ATMs work? What about the stock exchanges? Which emergency services would fail? Would trucks and trains be able to route their cargo? Would airlines be able to route their passengers? How much of the military's logistical system would fail?

That's to say nothing of the variety of corporations that rely on the Internet to function, let alone the millions of Americans who would need to use it to communicate with their loved ones in a time of crisis.

Even worse, these effects would spill over internationally. The Internet is international in complex and surprising ways, and it would be impossible to ensure that the effects of a shutdown stayed domestic and didn't cause similar disasters in countries we're friendly with.

Security Flaws

The third flawed assumption is that we could build this capability securely. We can't.

Once we engineered a selective shutdown switch into the Internet, and implemented a way to do what Internet engineers have spent decades making sure never happens, we would have created an enormous security vulnerability. We would make the job of any would-be terrorist intent on bringing down the Internet much easier.

Computer and network security is hard, and every Internet system we've ever created has security vulnerabilities. It would be folly to think this one wouldn't as well. And given how unlikely the risk is, any actual shutdown would be far more likely to be a result of an unfortunate error or a malicious hacker than of a presidential order.

But the main problem with an Internet kill switch is that it's too coarse a hammer.

Yes, the bad guys use the Internet to communicate, and they can use it to attack us. But the good guys use it, too, and the good guys far outnumber the bad guys.

Shutting the Internet down, either the whole thing or just a part of it, even in the face of a foreign military attack would do far more damage than it could possibly prevent. And it would hurt others whom we don't want to hurt.

For years we've been bombarded with scare stories about terrorists wanting to shut the Internet down. They're mostly fairy tales, but they're scary precisely because the Internet is so critical to so many things.

Why would we want to terrorize our own population by doing exactly what we don't want anyone else to do? And a national emergency is precisely the worst time to do it.

Just implementing the capability would be very expensive; I would rather see that money going toward securing our nation's critical infrastructure from attack.

Defending his proposal, Sen. Lieberman pointed out that China has this capability. It's debatable whether or not it actually does, but it's actively pursuing the capability because the country cares less about its citizens.

Here in the US, it is both wrong and dangerous to give the president the power and ability to commit Internet suicide and terrorize Americans in this way.

Web Snooping Is a Dangerous Move

Originally published in CNN, September 29, 2010

On Monday, the *New York Times* reported that President Obama will seek sweeping laws enabling law enforcement to more easily eavesdrop on the Internet. Technologies are changing, the administration argues, and modern digital systems aren't as easy to monitor as traditional telephones.

The government wants to force companies to redesign their communications systems and information networks to facilitate surveillance, and to provide law enforcement with back doors that enable them to bypass any security measures.

The proposal may seem extreme, but—unfortunately—it's not unique. Just a few months ago, the governments of the United Arab Emirates and Saudi

Arabia threatened to ban BlackBerry devices unless the company made eavesdropping easier. China has already built a massive Internet surveillance system to better control its citizens.

Formerly reserved for totalitarian countries, this wholesale surveillance of citizens has moved into the democratic world as well. Governments like Sweden, Canada and the United Kingdom are debating or passing laws giving their police new powers of Internet surveillance, in many cases requiring communications system providers to redesign products and services they sell. More are passing data retention laws, forcing companies to retain customer data in case they might need to be investigated later.

Obama isn't the first US president to seek expanded digital eavesdropping. The 1994 CALEA law required phone companies to build ways to better facilitate FBI eavesdropping into their digital phone switches. Since 2001, the National Security Agency has built substantial eavesdropping systems within the United States.

These laws are dangerous, both for citizens of countries like China and citizens of Western democracies. Forcing companies to redesign their communications products and services to facilitate government eavesdropping reduces privacy and liberty; that's obvious. But the laws also make us less safe. Communications systems that have no inherent eavesdropping capabilities are more secure than systems with those capabilities built in.

Any surveillance system invites both criminal appropriation and government abuse. Function creep is the most obvious abuse: New police powers, enacted to fight terrorism, are already used in situations of conventional non-terrorist crime. Internet surveillance and control will be no different.

Official misuses are bad enough, but the unofficial uses are far more worrisome. An infrastructure conducive to surveillance and control invites surveillance and control, both by the people you expect and the people you don't. Any surveillance and control system must itself be secured, and we're not very good at that. Why does anyone think that only authorized law enforcement will mine collected Internet data or eavesdrop on Skype and IM conversations?

These risks are not theoretical. After 9/11, the National Security Agency built a surveillance infrastructure to eavesdrop on telephone calls and e-mails within the United States. Although procedural rules stated that only non-Americans and international phone calls were to be listened to, actual practice didn't always match those rules. NSA analysts collected more data than they

were authorized to and used the system to spy on wives, girlfriends and famous people like former President Bill Clinton.

The most serious known misuse of a telecommunications surveillance infrastructure took place in Greece. Between June 2004 and March 2005, someone wiretapped more than 100 cell phones belonging to members of the Greek government—the prime minister and the ministers of defense, foreign affairs and justice—and other prominent people. Ericsson built this wiretapping capability into Vodafone’s products, but enabled it only for governments that requested it. Greece wasn’t one of those governments, but some still unknown party—a rival political group? organized crime?—figured out how to surreptitiously turn the feature on.

Surveillance infrastructure is easy to export. Once surveillance capabilities are built into Skype or Gmail or your BlackBerry, it’s easy for more totalitarian countries to demand the same access; after all, the technical work has already been done.

Western companies such as Siemens, Nokia and Secure Computing built Iran’s surveillance infrastructure, and US companies like L-1 Identity Solutions helped build China’s electronic police state. The next generation of worldwide citizen control will be paid for by countries like the United States.

We should be embarrassed to export eavesdropping capabilities. Secure, surveillance-free systems protect the lives of people in totalitarian countries around the world. They allow people to exchange ideas even when the government wants to limit free exchange. They power citizen journalism, political movements and social change. For example, Twitter’s anonymity saved the lives of Iranian dissidents—anonymity that many governments want to eliminate.

Yes, communications technologies are used by both the good guys and the bad guys. But the good guys far outnumber the bad guys, and it’s far more valuable to make sure they’re secure than it is to cripple them on the off chance it might help catch a bad guy. It’s like the FBI demanding that no automobiles drive above 50 mph, so they can more easily pursue getaway cars. It might or might not work—but, regardless, the cost to society of the resulting slowdown would be enormous.

It’s bad civic hygiene to build technologies that could someday be used to facilitate a police state. No matter what the eavesdroppers say, these systems cost too much and put us all at greater risk.

The Plan to Quarantine Infected Computers

Originally published in Forbes, November 11, 2010

Last month, Scott Charney of Microsoft proposed that infected computers be quarantined from the Internet. Using a public health model for Internet security, the idea is that infected computers spreading worms and viruses are a risk to the greater community and thus need to be isolated. Internet service providers would administer the quarantine, and would also clean up and update users' computers so they could rejoin the greater Internet.

This isn't a new idea. Already there are products that test computers trying to join private networks, and only allow them access if their security patches are up-to-date and their antivirus software certifies them as clean. Computers denied access are sometimes shunned to a limited-capability sub-network where all they can do is download and install the updates they need to regain access. This sort of system has been used with great success at universities and end-user-device-friendly corporate networks. They're happy to let you log in with any device you want—this is the consumerization trend in action—as long as your security is up to snuff.

Charney's idea is to do that on a larger scale. To implement it we have to deal with two problems. There's the technical problem—making the quarantine work in the face of malware designed to evade it, and the social problem—ensuring that people don't have their computers unduly quarantined. Understanding the problems requires us to understand quarantines in general.

Quarantines have been used to contain disease for millennia. In general several things need to be true for them to work. One, the thing being quarantined needs to be easily recognized. It's easier to quarantine a disease if it has obvious physical characteristics: fever, boils, etc. If there aren't any obvious physical effects, or if those effects don't show up while the disease is contagious, a quarantine is much less effective.

Similarly, it's easier to quarantine an infected computer if that infection is detectable. As Charney points out, his plan is only effective against worms and viruses that our security products recognize, not against those that are new and still undetectable.

Two, the separation has to be effective. The leper colonies on Molokai and Spinalonga both worked because it was hard for the quarantined to leave. Quarantined medieval cities worked less well because it was too easy to leave, or—when the diseases spread via rats or mosquitoes—because the quarantine was targeted at the wrong thing.

Computer quarantines have been generally effective because the users whose computers are being quarantined aren't sophisticated enough to break out of the quarantine, and find it easier to update their software and rejoin the network legitimately.

Three, only a small section of the population must need to be quarantined. The solution works only if it's a minority of the population that's affected, either with physical diseases or computer diseases. If most people are infected, overall infection rates aren't going to be slowed much by quarantining. Similarly, a quarantine that tries to isolate most of the Internet simply won't work.

Fourth, the benefits must outweigh the costs. Medical quarantines are expensive to maintain, especially if people are being quarantined against their will. Determining who to quarantine is either expensive (if it's done correctly) or arbitrary, authoritative and abuse-prone (if it's done badly). It could even be both. The value to society must be worth it.

It's the last point that Charney and others emphasize. If Internet worms were only damaging to the infected, we wouldn't need a societally imposed quarantine like this. But they're damaging to everyone else on the Internet, spreading and infecting others. At the same time, we can implement systems that quarantine cheaply. The value to society far outweighs the cost.

That makes sense, but once you move quarantines from isolated private networks to the general Internet, the nature of the threat changes. Imagine an intelligent and malicious infectious disease: That's what malware is. The current crop of malware ignores quarantines; they're few and far enough between not to affect their effectiveness.

If we tried to implement Internet-wide—or even countrywide—quarantining, worm-writers would start building in ways to break the quarantine. So instead of nontechnical users not bothering to break quarantines because they don't know how, we'd have technically sophisticated virus-writers trying to break quarantines. Implementing the quarantine at the ISP level would help, and if the ISP monitored computer behavior, not just specific virus signatures, it would be somewhat effective even in the face of evasion tactics. But evasion would be possible, and we'd be stuck in another computer security arms race.

This isn't a reason to dismiss the proposal outright, but it is something we need to think about when weighing its potential effectiveness.

Additionally, there's the problem of who gets to decide which computers to quarantine. It's easy on a corporate or university network: the owners of the network get to decide. But the Internet doesn't have that sort of hierarchical control, and denying people access without due process is fraught with danger. What are the appeal mechanisms? The audit mechanisms? Charney proposes that ISPs administer the quarantines, but there would have to be some central authority that decided what degree of infection would be sufficient to impose the quarantine. Although this is being presented as a wholly technical solution, it's these social and political ramifications that are the most difficult to determine and the easiest to abuse.

Once we implement a mechanism for quarantining infected computers, we create the possibility of quarantining them in all sorts of other circumstances. Should we quarantine computers that don't have their patches up to date, even if they're uninfected? Might there be a legitimate reason for someone to avoid patching his computer? Should the government be able to quarantine someone for something he said in a chat room, or a series of search queries he made? I'm sure we don't think it should, but what if that chat and those queries revolved around terrorism? Where's the line?

Microsoft would certainly like to quarantine any computers it feels are not running legal copies of its operating system or applications software. The music and movie industry will want to quarantine anyone it decides is downloading or sharing pirated media files—they're already pushing similar proposals.

A security measure designed to keep malicious worms from spreading over the Internet can quickly become an enforcement tool for corporate business models. Charney addresses the need to limit this kind of function creep, but I don't think it will be easy to prevent; it's an enforcement mechanism just begging to be used.

Once you start thinking about implementation of quarantine, all sorts of other social issues emerge. What do we do about people who need the Internet? Maybe VoIP is their only phone service. Maybe they have an Internet-enabled medical device. Maybe their business requires the Internet to run. The effects of quarantining these people would be considerable, even potentially life-threatening. Again, where's the line?

What do we do if people feel they are quarantined unjustly? Or if they are using nonstandard software unfamiliar to the ISP? Is there an appeals process? Who administers it? Surely not a for-profit company.

Public health is the right way to look at this problem. This conversation—between the rights of the individual and the rights of society—is a valid one to have, and this solution is a good possibility to consider.

There are some applicable parallels. We require drivers to be licensed and cars to be inspected not because we worry about the danger of unlicensed drivers and uninspected cars to themselves, but because we worry about their danger to other drivers and pedestrians. The small number of parents who don't vaccinate their kids have already caused minor outbreaks of whooping cough and measles among the greater population. We all suffer when someone on the Internet allows his computer to get infected. How we balance that with individuals' rights to maintain their own computers as they see fit is a discussion we need to start having.

Close the Washington Monument

Originally published in New York Daily News, December 2, 2010

The published version of this essay was heavily edited. This is the complete version.

Securing the Washington Monument from terrorism has turned out to be a surprisingly difficult job. The concrete fence around the building protects it from attacking vehicles, but there's no visually appealing way to house the airport-level security mechanisms the National Park Service has decided are a must for visitors. It is considering several options, but I think we should close the monument entirely. Let it stand, empty and inaccessible, as a monument to our fears.

An empty Washington Monument would serve as a constant reminder to those on Capitol Hill that they are afraid of the terrorists and what they could do. They're afraid that by speaking honestly about the impossibility of attaining absolute security or the inevitability of terrorism—or that some American ideals are worth maintaining even in the face of adversity—they will be branded as “soft on terror.” And they're afraid that Americans would vote them out of office if another attack occurred. Perhaps they're right, but what has happened to leaders who aren't afraid? What has happened to “the only thing we have to fear is fear itself”?

An empty Washington Monument would symbolize our lawmakers' inability to take that kind of stand—and their inability to truly lead.

Some of them call terrorism an “existential threat” against our nation. It's not. Even the events of 9/11, as horrific as they were, didn't make an existential

dent in our nation. Automobile-related fatalities—at 42,000 per year, more deaths each month, on average, than 9/11—aren't, either. It's our reaction to terrorism that threatens our nation, not terrorism itself. The empty monument would symbolize the empty rhetoric of those leaders who preach fear and then use that fear for their own political ends.

The day after Umar Farouk Abdulmutallab failed to blow up a Northwest jet with a bomb hidden in his underwear, Homeland Security Secretary Janet Napolitano said “The system worked.” I agreed. Plane lands safely, terrorist in custody, nobody injured except the terrorist. Seems like a working system to me. The empty monument would represent the politicians and press who pilloried her for her comment, and Napolitano herself, for backing down.

The empty monument would symbolize our war on the unexpected,—our overreaction to anything different or unusual—our harassment of photographers, and our probing of airline passengers. It would symbolize our “show me your papers” society, rife with ID checks and security cameras. As long as we're willing to sacrifice essential liberties for a little temporary safety, we should keep the Washington Monument empty.

Terrorism isn't a crime against people or property. It's a crime against our minds, using the death of innocents and destruction of property to make us fearful. Terrorists use the media to magnify their actions and further spread fear. And when we react out of fear, when we change our policy to make our country less open, the terrorists succeed—even if their attacks fail. But when we refuse to be terrorized, when we're indomitable in the face of terror, the terrorists fail—even if their attacks succeed.

We can reopen the monument when every foiled or failed terrorist plot causes us to praise our security, instead of redoubling it. When the occasional terrorist attack succeeds, as it inevitably will, we accept it, as we accept the murder rate and automobile-related death rate; and redouble our efforts to remain a free and open society.

The grand reopening of the Washington Monument will not occur when we've won the war on terror, because that will never happen. It won't even occur when we've defeated al Qaeda. Militant Islamic terrorism has fractured into small, elusive groups. We can reopen the Washington Monument when we've defeated our fears, when we've come to accept that placing safety above all other virtues cedes too much power to government and that liberty is worth the risks, and that the price of freedom is accepting the possibility of crime.

I would proudly climb to the top of a monument to those ideals.

Whitelisting and Blacklisting

Originally published in Information Security, January 2011

This essay appeared as the second half of a point/counterpoint with Marcus Ranum.

The whitelist/blacklist debate is far older than computers, and it's instructive to recall what works where. Physical security works generally on a whitelist model: if you have a key, you can open the door; if you know the combination, you can open the lock. We do it this way not because it's easier—although it is generally much easier to make a list of people who should be allowed through your office door than a list of people who shouldn't—but because it's a security system that can be implemented automatically, without people.

To find blacklists in the real world, you have to start looking at environments where almost everyone is allowed. Casinos are a good example: everyone can come in and gamble except those few specifically listed in the casino's black book or the more general Griffin book. Some retail stores have the same model—a Google search on “banned from Wal-Mart” results in 1.5 million hits, including Megan Fox—although you have to wonder about enforcement. Does Wal-Mart have the same sort of security manpower as casinos?

National borders certainly have that kind of manpower, and Marcus is correct to point to passport control as a system with both a whitelist and a blacklist. There are people who are allowed in with minimal fuss, people who are summarily arrested with as minimal a fuss as possible, and people in the middle who receive some amount of fussing. Airport security works the same way: the no-fly list is a blacklist, and people with redress numbers are on the whitelist.

Computer networks share characteristics with your office and Wal-Mart: sometimes you only want a few people to have access, and sometimes you want almost everybody to have access. And you see whitelists and blacklists at work in computer networks. Access control is whitelisting: if you know the password, or have the token or biometric, you get access. Antivirus is blacklisting: everything coming into your computer from the Internet is assumed to be safe unless it appears on a list of bad stuff. On computers, unlike the real world, it takes no extra manpower to implement a blacklist—the software can do it largely for free.

Traditionally, execution control has been based on a blacklist. Computers are so complicated and applications so varied that it just doesn't make sense to limit users to a specific set of applications. The exception is constrained

environments, such as computers in hotel lobbies and airline club lounges. On those, you're often limited to an Internet browser and a few common business applications.

Lately, we're seeing more whitelisting on closed computing platforms. The iPhone works on a whitelist: if you want a program to run on the phone, you need to get it approved by Apple and put in the iPhone store. Your Wii game machine works the same way. This is done primarily because the manufacturers want to control the economic environment, but it's being sold partly as a security measure. But in this case, more security equals less liberty; do you really want your computing options limited by Apple, Microsoft, Google, Facebook, or whoever controls the particular system you're using?

Turns out that many people do. Apple's control over its apps hasn't seemed to hurt iPhone sales, and Facebook's control over its apps hasn't seemed to affect Facebook's user numbers. And honestly, quite a few of us would have had an easier time over the Christmas holidays if we could have implemented a whitelist on the computers of our less-technical relatives.

For these two reasons, I think the whitelist model will continue to make inroads into our general purpose computers. And those of us who want control over our own environments will fight back—perhaps with a whitelist we maintain personally, but more probably with a blacklist.

Securing Medical Research: a Cybersecurity Point of View

Originally published in Science, June 22, 2012

This article was based on a talk given at the meeting on H5N1 Research, Biosafety, Biosecurity, and Bioethics, Royal Society, London, April 3, 2012.

Science and Nature have each published papers on the H5N1 virus in humans after considerable debate about whether the research results in those papers could help terrorists create a bioweapon. This notion of “dual use” research is an important one for the community, and one that will sooner or later become critical. Perhaps these two papers are not dangerous in the wrong hands, but eventually there will be research results that are.

My background is in cryptography and computer security. I cannot comment on the potential value or harm from any particular piece of biological research, but I can discuss what works and what does not to keep research

data secure. The cryptography and computer security communities have been wrestling for decades now with dual-use research: for example, whether to publish new Windows (Microsoft Corporation) vulnerabilities that can be immediately used to attack computers but whose publication helps us make the operating system more secure in the long run. From this experience, I offer five points to the virology community.

First, security based on secrecy is inherently fragile. The more secrets a system has, the less secure it is. A door lock that has a secret but unchangeable locking mechanism is less secure than a commercially purchased door lock with an easily changeable key. In cryptography, this is known as Kerckhoffs' principle: Put all your secrecy into the key and none into the cryptographic algorithm. The key is unique and easily changeable; the algorithm is system-wide and much more likely to become public. In fact, algorithms are deliberately published so that they get analyzed broadly. The lesson for dual-use virology research is that it is risky to base your security on keeping research secret. Militaries spend an enormous amount of money trying to maintain secret research laboratories, and even they do not always get security right. Once secret data become public, there is no way to go back.

Second, omitting technical details from published research is a poor security measure. We tried this in computer security with regard to vulnerabilities, announcing general information but not publishing specifics. The problem is that once the general information is announced, it is much easier for another researcher to replicate the results and generate the details. This is probably even more true in virology research than in computer security research, where the very existence of a result can provide much of the road map to that result.

Third, technical difficulty as a security measure has only short-term value. Technology only gets better; it never gets worse. To believe that some research cannot be replicated by amateurs because it requires equipment only available to state-of-the-art research institutions is short-sighted at best. What is impossible today will be a Ph.D. thesis in 20 years, and what was a Ph.D. thesis 20 years ago is a high-school science fair project today.

Fourth, securing research data in computer networks is risky at best. If you read newspapers, you know the current state of the art in computer security: Everything gets hacked. Cyber criminals steal money from banks. Cyber spies steal data from military computers. Although people talk about H5N1 research in terms of securing the research papers, that is largely a red herring; even if no papers existed, the research data would still be on a network-connected computer somewhere.

Not all computers are hacked and not all data gets stolen, but the risks are there. There are two basic types of threats in cyberspace. There are the opportunists: for example, criminals who want to break into a retail merchant's system and steal a thousand credit card numbers. Against these attackers, relative security is what matters. Because the criminals do not care whom they attack, you are safe if you are more secure than other networks. The other type of threat is a targeted attack. These are attackers who, for whatever reason, want to attack a particular network. The buzzword in Internet security for this is "advanced persistent threat." It is almost impossible to secure a network against a sufficiently skilled and tenacious adversary. All we can do is make the attacker's job harder.

This does not mean that all virology data will be stolen via computer networks, but it does mean that, once the existence of that data becomes public knowledge, you should assume that the bad guys will be able to get their hands on it.

Lastly, national measures that prohibit publication will not work in an international community, especially in the Internet age. If either *Science* or *Nature* had refused to publish the H5N1 papers, they would have been published somewhere else. Even if some countries stop funding—or ban—this sort of research, it will still happen in another country.

The US cryptography community saw this in the 1970s and early 1980s. At that time, the National Security Agency (NSA) controlled cryptography research, which included denying funding for research, classifying results after the fact, and using export-control laws to limit what ended up in products. This was the pre-Internet world, and it worked for a while. In the 1980s they gave up on classifying research, because an international community arose. The limited ability for US researchers to get funding for block-cipher cryptanalysis merely moved that research to Europe and Asia. The NSA continued to limit the spread of cryptography via export-control laws; the US-centric nature of the computer industry meant that this was effective. In the 1990s they gave up on controlling software because the international online community became mainstream; this period was called "the Crypto Wars." Export-control laws did prevent Microsoft from embedding cryptography into Windows for over a decade, but it did nothing to prevent products made in other countries from filling the market gaps.

Today, there are no restrictions on cryptography, and many US government standards are the result of public international competitions. Right now the National Institute of Standards and Technology is working on a new Secure

Hash Algorithm standard. When it is announced next year, it will be the product of a public call for algorithms that resulted in 64 submissions from over a dozen countries and then years of international analysis. The practical effects of unrestricted research are seen in the computer security you use today: on your computer, as you browse the Internet and engage in commerce, and on your cell phone and other smart devices. Sure, the bad guys make use of this research, too, but the beneficial uses far outweigh the malicious ones.

The computer security community has also had to wrestle with these dual-use issues. In the early days of public computing, researchers who discovered vulnerabilities would quietly tell the product vendors so as to not also alert hackers. But all too often, the vendors would ignore the researchers. Because the vulnerability was not public, there was no urgency to fix it. Fixes might go into the next product release. Researchers, tired of this, started publishing the existence of vulnerabilities but not the details. Vendors, in response, tried to muzzle the researchers. They threatened them with lawsuits and belittled them in the press, calling the vulnerabilities only theoretical and not practical. The response from the researchers was predictable: They started publishing full details, and sometimes even code, demonstrating the vulnerabilities they found. This was called “full disclosure” and is the primary reason vendors now patch vulnerabilities quickly. Faced with published vulnerabilities that they could not pretend did not exist and that the hackers could use, they started building internal procedures to quickly issue patches. If you use Microsoft Windows, you know about “patch Tuesday,” the once-a-month automatic download and installation of security patches.

Once vendors started taking security patches seriously, the research community (university researchers, security consultants, and informal hackers) moved to something called “responsible disclosure.” Now it is common for researchers to alert vendors before publication, giving them a month or two head start to release a security patch. But without the threat of full disclosure, responsible disclosure would not work, and vendors would go back to ignoring security vulnerabilities.

Could a similar process work for viruses? That is, could the makers work in concert with people who develop vaccines so that vaccines become available at the same time as the original results are released? Certainly this is not easy in practice, but perhaps it is a goal to work toward.

Limiting research, either through government classification or legal threats from vendors, has a chilling effect. Why would professors or graduate students choose cryptography or computer security if they were going to be prevented from publishing their results? Once these sorts of research slow down, the increasing ignorance hurts us all.

On the other hand, the current vibrant fields of cryptography and computer security are a direct result of our willingness to publish methods of attack. Making and breaking systems are one and the same; you cannot learn one without the other. (Some universities even offer classes in computer virus writing.) Cryptography is better, and computers and networks are more secure, because our communities openly publish details on how to attack systems.

Virology is not computer science. A biological virus is not the same as a computer virus. A vulnerability that affects every individual copy of Windows is not as bad as a vulnerability that affects every individual person. Still, the lessons from computer security are valuable to anyone considering policies intended to encourage life-saving research in virology while at the same time prevent that research from being used to cause harm. This debate will not go away; it will only get more urgent.

Fear Pays the Bills, but Accounts Must Be Settled

*Originally published in the New York Times Room for Debate blog,
October 19, 2012*

A lot of the debate around President Obama's cybersecurity initiative centers on how much of a burden it would be on industry, and how that should be financed. As important as that debate is, it obscures some of the larger issues surrounding cyberwar, cyberterrorism, and cybersecurity in general.

It's difficult to have any serious policy discussion amongst the fear-mongering. Secretary Panetta's recent comments are just the latest; search the Internet for "cyber 9/11," "cyber Pearl-Harbor," "cyber Katrina," or—my favorite—"cyber Armageddon."

There's an enormous amount of money and power that results from pushing cyberwar and cyberterrorism: power within the military, the Department of Homeland Security, and the Justice Department; and lucrative government contracts supporting those organizations. As long as cyber remains a prefix that scares, it'll continue to be used as a bugaboo.

But while scare stories are more movie-plot than actual threat, there are real risks. The government is continually poked and probed in cyberspace, from attackers ranging from kids playing politics to sophisticated national intelligence gathering operations. Hackers can do damage, although nothing like the cyberterrorism rhetoric would lead you to believe. Cybercrime continues to rise, and still poses real risks to those of us who work, shop, and play on the Internet. And cyberdefense needs to be part of our military strategy.

Industry has definitely not done enough to protect our nation's critical infrastructure, and federal government may need more involvement. This should come as no surprise; the economic externalities in cybersecurity are so great that even the freest free market would fail.

For example, the owner of a chemical plant will protect that plant from cyberattack up to the value of that plant to the owner; the residual risk to the community around the plant will remain. Politics will color how government involvement looks: market incentives, regulation, or outright government takeover of some aspects of cybersecurity.

None of this requires heavy-handed regulation. Over the past few years we've heard calls for the military to better control Internet protocols; for the United States to be able to "kill" all or part of the Internet, or to cut itself off from the greater Internet; for increased government surveillance; and for limits on anonymity. All of those would be dangerous, and would make us less secure. The world's first military cyberweapon, Stuxnet, was used by the United States and Israel against Iran.

In all of this government posturing about cybersecurity, the biggest risk is a cyber-war arms race; and that's where remarks like Panetta's lead us. Increased government spending on cyberweapons and cyberdefense, and an increased militarization of cyberspace, are both expensive and destabilizing. Fears lead to weapons buildups, and weapons beg to be used.

I would like to see less fear-mongering, and more reasoned discussion about the actual threats and reasonable countermeasures. Pushing the fear button benefits no one.

Power and the Internet

Originally published in Edge, January 23, 2013

This essay appeared as a response to Edge's annual question, "What Should We Be Worried About?"

All disruptive technologies upset traditional power balances, and the Internet is no exception. The standard story is that it empowers the powerless, but that's only half the story. The Internet empowers everyone. Powerful institutions might be slow to make use of that new power, but since they are powerful, they can use it more effectively. Governments and corporations have woken up to the fact that not only can they use the Internet, they can control it for their interests. Unless we start deliberately debating the future we want to live in, and information technology in enabling that world, we will end up with an Internet that benefits existing power structures and not society in general.

We've all lived through the Internet's disruptive history. Entire industries, like travel agencies and video rental stores, disappeared. Traditional publishing—books, newspapers, encyclopedias, music—lost power, while Amazon and others gained. Advertising-based companies like Google and Facebook gained a lot of power. Microsoft lost power (as hard as that is to believe).

The Internet changed political power as well. Some governments lost power as citizens organized online. Political movements became easier, helping to topple governments. The Obama campaign made revolutionary use of the Internet, both in 2008 and 2012.

And the Internet changed social power, as we collected hundreds of "friends" on Facebook, tweeted our way to fame, and found communities for the most obscure hobbies and interests. And some crimes became easier: impersonation fraud became identity theft, copyright violation became file sharing, and accessing censored materials—political, sexual, cultural—became trivially easy.

Now powerful interests are looking to deliberately steer this influence to their advantage. Some corporations are creating Internet environments that maximize their profitability: Facebook and Google, among many others. Some industries are lobbying for laws that make their particular business models more profitable: telecom carriers want to be able to discriminate between different types of Internet traffic, entertainment companies want to crack down

on file sharing, advertisers want unfettered access to data about our habits and preferences.

On the government side, more countries censor the Internet—and do so more effectively—than ever before. Police forces around the world are using Internet data for surveillance, with less judicial oversight and sometimes in advance of any crime. Militaries are fomenting a cyberwar arms race. Internet surveillance—both governmental and commercial—is on the rise, not just in totalitarian states but in Western democracies as well. Both companies and governments rely more on propaganda to create false impressions of public opinion.

In 1996, cyber-libertarian John Perry Barlow issued his “Declaration of the Independence of Cyberspace.” He told governments: “You have no moral right to rule us, nor do you possess any methods of enforcement that we have true reason to fear.” It was a utopian ideal, and many of us believed him. We believed that the Internet generation, those quick to embrace the social changes this new technology brought, would swiftly outmaneuver the more ponderous institutions of the previous era.

Reality turned out to be much more complicated. What we forgot is that technology magnifies power in both directions. When the powerless found the Internet, suddenly they had power. But while the unorganized and nimble were the first to make use of the new technologies, eventually the powerful behemoths woke up to the potential—and they have more power to magnify. And not only does the Internet change power balances, but the powerful can also change the Internet. Does anyone else remember how incompetent the FBI was at investigating Internet crimes in the early 1990s? Or how Internet users ran rings around China’s censors and Middle Eastern secret police? Or how digital cash was going to make government currencies obsolete, and Internet organizing was going to make political parties obsolete? Now all that feels like ancient history.

It’s not all one-sided. The masses can occasionally organize around a specific issue—SOPA/PIPA, the Arab Spring, and so on—and can block some actions by the powerful. But it doesn’t last. The unorganized go back to being unorganized, and powerful interests take back the reins.

Debates over the future of the Internet are morally and politically complex. How do we balance personal privacy against what law enforcement needs to prevent copyright violations? Or child pornography? Is it acceptable to be judged by invisible computer algorithms when being served search results? When being served news articles? When being selected for additional scrutiny by airport security? Do we have a right to correct data about us? To

delete it? Do we want computer systems that forget things after some number of years? These are complicated issues that require meaningful debate, international cooperation, and iterative solutions. Does anyone believe we're up to the task?

We're not, and that's the worry. Because if we're not trying to understand how to shape the Internet so that its good effects outweigh the bad, powerful interests will do all the shaping. The Internet's design isn't fixed by natural laws. Its history is a fortuitous accident: an initial lack of commercial interests, governmental benign neglect, military requirements for survivability and resilience, and the natural inclination of computer engineers to build open systems that work simply and easily. This mix of forces that created yesterday's Internet will not be trusted to create tomorrow's. Battles over the future of the Internet are going on right now: in legislatures around the world, in international organizations like the International Telecommunications Union and the World Trade Organization, and in Internet standards bodies. The Internet is what we make it, and is constantly being re-created by organizations, companies, and countries with specific interests and agendas. Either we fight for a seat at the table, or the future of the Internet becomes something that is done to us.

Danger Lurks in Growing New Internet Nationalism

Originally published in MIT Technology Review, March 11, 2013

For technology that was supposed to ignore borders, bring the world closer together, and sidestep the influence of national governments, the Internet is fostering an awful lot of nationalism right now. We've started to see increased concern about the country of origin of IT products and services; US companies are worried about hardware from China; European companies are worried about cloud services in the US; no one is sure whether to trust hardware and software from Israel; Russia and China might each be building their own operating systems out of concern about using foreign ones.

I see this as an effect of all the cyberwar saber-rattling that's going on right now. The major nations of the world are in the early years of a cyberwar arms race, and we're all being hurt by the collateral damage.

Our nationalist worries have recently been fueled by a media frenzy surrounding attacks from China. These attacks aren't new—cyber-security experts have been writing about them for at least a decade, and the popular media reported about similar attacks in 2009 and again in 2010—and the current allegations aren't even very different than what came before. This isn't to say that the Chinese attacks aren't serious. The country's espionage campaign is sophisticated, and ongoing. And because they're in the news, people are understandably worried about them.

But it's not just China. International espionage works in both directions, and I'm sure we are giving just as good as we're getting. China is certainly worried about the US Cyber Command's recent announcement that it was expanding from 900 people to almost 5,000, and the NSA's massive new data center in Utah. The US even admits that it can spy on non-US citizens freely.

The fact is that governments and militaries have discovered the Internet; everyone is spying on everyone else, and countries are ratcheting up offensive actions against other countries.

At the same time, many nations are demanding more control over the Internet within their own borders. They reserve the right to spy and censor, and to limit the ability of others to do the same. This idea is now being called the "cyber sovereignty movement," and gained traction at the International Telecommunications Union meeting last December in Dubai. One analyst called that meeting the "Internet Yalta," where the Internet split between liberal-democratic and authoritarian countries. I don't think he's exaggerating.

Not that this is new, either. Remember 2010, when the governments of the UAE, Saudi Arabia, and India demanded that RIM give them the ability to spy on BlackBerry PDAs within their borders? Or last year, when Syria used the Internet to surveil its dissidents? Information technology is a surprisingly powerful tool for oppression: not just surveillance, but censorship and propaganda as well. And countries are getting better at using that tool.

But remember: none of this is cyberwar. It's all espionage, something that's been going on between countries ever since countries were invented. What moves public opinion is less the facts and more the rhetoric, and the rhetoric of war is what we're hearing.

The result of all this saber-rattling is a severe loss of trust, not just amongst nation-states but between people and nation-states. We know we're nothing more than pawns in this game, and we figure we'll be better off sticking with our own country.

Unfortunately, both the reality and the rhetoric play right into the hands of the military and corporate interests that are behind the cyberwar arms race in the first place. There is an enormous amount of power at stake here: not only power within governments and militaries, but power and profit amongst the corporations that supply the tools and infrastructure for cyber-attack and cyber-defense. The more we believe we are “at war” and believe the jingoistic rhetoric, the more willing we are to give up our privacy, freedoms, and control over how the Internet is run.

Arms races are fueled by two things: ignorance and fear. We don’t know the capabilities of the other side, and we fear that they are more capable than we are. So we spend more, just in case. The other side, of course, does the same. That spending will result in more cyber weapons for attack and more cyber-surveillance for defense. It will result in more government control over the protocols of the Internet, and less free-market innovation over the same. At its worst, we might be about to enter an information-age Cold War: one with more than two “superpowers.” Aside from this being a bad future for the Internet, this is inherently destabilizing. It’s just too easy for this amount of antagonistic power and advanced weaponry to get used: for a mistaken attribution to be reacted to with a counterattack, for a misunderstanding to become a cause for offensive action, or for a minor skirmish to escalate into a full-fledged cyberwar.

Nationalism is rife on the Internet, and it’s getting worse. We need to damp down the rhetoric and—more importantly—stop believing the propaganda from those who profit from this Internet nationalism. Those who are beating the drums of cyberwar don’t have the best interests of society, or the Internet, at heart.

IT for Oppression

*Originally published in IEEE Security & Privacy,
March/April 2013*

Whether it’s Syria using Facebook to help identify and arrest dissidents or China using its “Great Firewall” to limit access to international news throughout the country, repressive regimes all over the world are using the Internet to more efficiently implement surveillance, censorship, propaganda, and control. They’re getting really good at it, and the IT industry is helping. We’re helping by creating business applications—categories of applications, really—that are being repurposed by oppressive governments for their own use:

- What is called censorship when practiced by a government is content filtering when practiced by an organization. Many companies want to keep their employees from viewing porn or updating their Facebook pages while at work. In the other direction, data loss prevention software keeps employees from sending proprietary corporate information outside the network and also serves as a censorship tool. Governments can use these products for their own ends.
- Propaganda is really just another name for marketing. All sorts of companies offer social media-based marketing services designed to fool consumers into believing there is “buzz” around a product or brand. The only thing different in a government propaganda campaign is the content of the messages.
- Surveillance is necessary for personalized marketing, the primary profit stream of the Internet. Companies have built massive Internet surveillance systems designed to track users’ behavior all over the Internet and closely monitor their habits. These systems track not only individuals but also relationships between individuals, to deduce their interests so as to advertise to them more effectively. It’s a totalitarian’s dream.
- Control is how companies protect their business models by limiting what people can do with their computers. These same technologies can easily be co-opted by governments that want to ensure that only certain computer programs are run inside their countries or that their citizens never see particular news programs.

Technology magnifies power, and there’s no technical difference between a government and a corporation wielding it. This is how commercial security equipment from companies like BlueCoat and Sophos end up being used by the Syrian and other oppressive governments to surveil—in order to arrest—and censor their citizens. This is how the same face-recognition technology that Disney uses in its theme parks ends up identifying protesters in China and Occupy Wall Street protesters in New York.

There are no easy technical solutions, especially because these four applications—censorship, propaganda, surveillance, and control—are intertwined; it can be hard to affect one without also affecting the others. Anonymity helps prevent surveillance, but it also makes propaganda easier. Systems that block propaganda can facilitate censorship. And giving users the ability to run untrusted software on their computers makes it easier for governments—and criminals—to install spyware.

We need more research into how to circumvent these technologies, but it's a hard sell to both the corporations and governments that rely on them. For example, law enforcement in the US wants drones that can identify and track people, even as we decry China's use of the same technology. Indeed, the battleground is often economic and political rather than technical; sometimes circumvention research is itself illegal.

The social issues are large. Power is using the Internet to increase its power, and we haven't yet figured out how to correct the imbalances among government, corporate, and individual interests in our digital world. Cyberspace is still waiting for its Gandhi, its Martin Luther King, and a convincing path from the present to a better future.

The Public/Private Surveillance Partnership

Originally published in the Atlantic, April 30, 2013

Our government collects a lot of information about us. Tax records, legal records, license records, records of government services received—it's all in databases that are increasingly linked and correlated. Still, there's a lot of personal information the government can't collect. Either they're prohibited by law from asking without probable cause and a judicial order, or they simply have no cost-effective way to collect it. But the government has figured out how to get around the laws, and collect personal data that has been historically denied to them: ask corporate America for it.

It's no secret that we're monitored continuously on the Internet. Some of the company names you know, such as Google and Facebook. Others hide in the background as you move about the Internet. There are browser plugins that show you who is tracking you. One Atlantic editor found 105 companies tracking him during one 36-hour period. Add data from your cell phone (who you talk to, your location), your credit cards (what you buy, from whom you buy it), and the dozens of other times you interact with a computer daily, we live in a surveillance state beyond the dreams of Orwell.

It's all corporate data, compiled and correlated, bought and sold. And increasingly, the government is doing the buying. Some of this is collected using National Security Letters (NSLs). These give the government the ability

to demand an enormous amount of personal data about people for very speculative reasons, with neither probable cause nor judicial oversight. Data on these secretive orders is obviously scant, but we know that the FBI has issued hundreds of thousands of them in the past decade—for reasons that go far beyond terrorism.

NSLs aren't the only way the government can get at corporate data. Sometimes they simply purchase it, just as any other company might. Sometimes they can get it for free, from corporations that want to stay on the government's good side.

CISPA, a bill currently wending its way through Congress, codifies this sort of practice even further. If signed into law, CISPA will allow the government to collect all sorts of personal data from corporations, without any oversight at all, and will protect corporations from lawsuits based on their handing over that data. Without hyperbole, it's been called the death of the 4th Amendment. Right now, it's mainly the FBI and the NSA who are getting this data, but—all sorts of government agencies have administrative subpoena power.

Data on this scale has all sorts of applications. From finding tax cheaters by comparing data brokers' estimates of income and net worth with what's reported on tax returns, to compiling a list of gun owners from Web browsing habits, instant messaging conversations, and locations—did you have your iPhone turned on when you visited a gun store?—the possibilities are endless.

Government photograph databases form the basis of any police facial recognition system. They're not very good today, but they'll only get better. But the government no longer needs to collect photographs. Experiments demonstrate that the Facebook database of tagged photographs is surprisingly effective at identifying people. As more places follow Disney's lead in fingerprinting people at its theme parks, the government will be able to use that to identify people as well.

In a few years, the whole notion of a government-issued ID will seem quaint. Among facial recognition, the unique signature from your smart phone, the RFID chips in your clothing and other items you own, and whatever new technologies that will broadcast your identity, no one will have to ask to see ID. When you walk into a store, they'll already know who you are. When you interact with a policeman, she'll already have your personal information displayed on her Internet-enabled glasses.

Soon, governments won't have to bother collecting personal data. We're willingly giving it to a vast network of for-profit data collectors, and they're more than happy to pass it on to the government without our knowledge or consent.

Transparency and Accountability Don't Hurt Security—They're Crucial to It

Originally published in the Atlantic, May 8, 2013

As part of the fallout of the Boston bombings, we're probably going to get some new laws that give the FBI additional investigative powers. As with the Patriot Act after 9/11, the debate over whether these new laws are helpful will be minimal, but the effects on civil liberties could be large. Even though most people are skeptical about sacrificing personal freedoms for security, it's hard for politicians to say no to the FBI right now, and it's politically expedient to demand that *something* be done.

If our leaders can't say no—and there's no reason to believe they can—there are two concepts that need to be part of any new counterterrorism laws, and investigative laws in general: transparency and accountability.

Long ago, we realized that simply trusting people and government agencies to always do the right thing doesn't work, so we need to check up on them. In a democracy, transparency and accountability are how we do that. It's how we ensure that we get both effective and cost-effective government. It's how we prevent those we trust from abusing that trust, and protect ourselves when they do. And it's especially important when security is concerned.

First, we need to ensure that the stuff we're paying money for actually works and has a measureable impact. Law-enforcement organizations regularly invest in technologies that don't make us any safer. The TSA, for example, could devote an entire museum to expensive but ineffective systems: puffer machines, body scanners, FAST behavioral screening, and so on. Local police departments have been wasting lots of post-9/11 money on unnecessary high-tech weaponry and equipment. The occasional high-profile success aside, police surveillance cameras have been shown to be a largely ineffective police tool.

Sometimes honest mistakes led organizations to invest in these technologies. Sometimes there's self-deception and mismanagement—and far too often lobbyists are involved. Given the enormous amount of security money post-9/11, you inevitably end up with an enormous amount of waste. Transparency and accountability are how we keep all of this in check.

Second, we need to ensure that law enforcement does what we expect it to do and nothing more. Police powers are invariably abused. Mission creep is inevitable, and it results in laws designed to combat one particular type of

crime being used for an ever-widening array of crimes. Transparency is the only way we have of knowing when this is going on.

For example, that's how we learned that the FBI is abusing National Security Letters. Traditionally, we use the warrant process to protect ourselves from police overreach. It's not enough for the police to want to conduct a search; they also need to convince a neutral third party—a judge—that the search is in the public interest and will respect the rights of those searched. That's accountability, and it's the very mechanism that NSLs were exempted from.

When laws are broken, accountability is how we punish those who abused their power. It's how, for example, we correct racial profiling by police departments. And it's a lack of accountability that permits the FBI to get away with massive data collection until exposed by a whistleblower or noticed by a judge.

Third, transparency and accountability keep both law enforcement and politicians from lying to us. The Bush Administration lied about the extent of the NSA's warrantless wiretapping program. The TSA lied about the ability of full-body scanners to save naked images of people. We've been lied to about the lethality of Tasers, when and how the FBI eavesdrops on cell-phone calls, and about the existence of surveillance records. Without transparency, we would never know.

A decade ago, the FBI was heavily lobbying Congress for a law to give it new wiretapping powers: a law known as CALEA. One of its key justifications was that existing law didn't allow it to perform speedy wiretaps during kidnapping investigations. It sounded plausible—and who wouldn't feel sympathy for kidnapping victims?—but when civil-liberties organizations analyzed the actual data, they found that it was just a story; there were no instances of wiretapping in kidnapping investigations. Without transparency, we would never have known that the FBI was making up stories to scare Congress.

If we're going to give the government any new powers, we need to ensure that there's oversight. Sometimes this oversight is before action occurs. Warrants are a great example. Sometimes they're after action occurs: public reporting, audits by inspector generals, open hearings, notice to those affected, or some other mechanism. Too often, law enforcement tries to exempt itself from this principle by supporting laws that are specifically excused from oversight. . . or by establishing secret courts that just rubber-stamp government wiretapping requests.

Furthermore, we need to ensure that mechanisms for accountability have teeth and are used.

As we respond to the threat of terrorism, we must remember that there are other threats as well. A society without transparency and accountability is the very definition of a police state. And while a police state might have a low crime rate—especially if you don't define police corruption and other abuses of power as crime—and an even lower terrorism rate, it's not a society that most of us would willingly choose to live in.

We already give law enforcement enormous power to intrude into our lives. We do this because we know they need this power to catch criminals, and we're all safer thereby. But because we recognize that a powerful police force is itself a danger to society, we must temper this power with transparency and accountability.

It's Smart Politics to Exaggerate Terrorist Threats

Originally published in CNN, May 20, 2013

Terrorism causes fear, and we overreact to that fear. Our brains aren't very good at probability and risk analysis. We tend to exaggerate spectacular, strange and rare events, and downplay ordinary, familiar and common ones. We think rare risks are more common than they are, and we fear them more than probability indicates we should.

Our leaders are just as prone to this overreaction as we are. But aside from basic psychology, there are other reasons that it's smart politics to exaggerate terrorist threats, and security threats in general.

The first is that we respond to a strong leader. Bill Clinton famously said: "When people feel uncertain, they'd rather have somebody that's strong and wrong than somebody who's weak and right." He's right.

The second is that doing something—anything—is good politics. A politician wants to be seen as taking charge, demanding answers, fixing things. It just doesn't look as good to sit back and claim that there's nothing to do. The logic is along the lines of: "Something must be done. This is something. Therefore, we must do it."

The third is that the "fear preacher" wins, regardless of the outcome. Imagine two politicians today. One of them preaches fear and draconian security measures. The other is someone like me, who tells people that terrorism is

a negligible risk, that risk is part of life, and that while some security is necessary, we should mostly just refuse to be terrorized and get on with our lives.

Fast-forward 10 years. If I'm right and there have been no more terrorist attacks, the fear preacher takes credit for keeping us safe. But if a terrorist attack has occurred, my government career is over. Even if the incidence of terrorism is as ridiculously low as it is today, there's no benefit for a politician to take my side of that gamble.

The fourth and final reason is money. Every new security technology, from surveillance cameras to high-tech fusion centers to airport full-body scanners, has a for-profit corporation lobbying for its purchase and use. Given the three other reasons above, it's easy—and probably profitable—for a politician to make them happy and say yes.

For any given politician, the implications of these four reasons are straightforward. Overestimating the threat is better than underestimating it. Doing something about the threat is better than doing nothing. Doing something that is explicitly reactive is better than being proactive. (If you're proactive and you're wrong, you've wasted money. If you're proactive and you're right but no longer in power, whoever is in power is going to get the credit for what you did.) Visible is better than invisible. Creating something new is better than fixing something old.

Those last two maxims are why it's better for a politician to fund a terrorist fusion center than to pay for more Arabic translators for the National Security Agency. No one's going to see the additional appropriation in the NSA's secret budget. On the other hand, a high-tech computerized fusion center is going to make front page news, even if it doesn't actually do anything useful.

This leads to another phenomenon about security and government. Once a security system is in place, it can be very hard to dislodge it. Imagine a politician who objects to some aspect of airport security: the liquid ban, the shoe removal, something. If he pushes to relax security, he gets the blame if something bad happens as a result. No one wants to roll back a police power and have the lack of that power cause a well-publicized death, even if it's a one-in-a-billion fluke.

We're seeing this force at work in the bloated terrorist no-fly and watch lists; agents have lots of incentive to put someone on the list, but absolutely no incentive to take anyone off. We're also seeing this in the Transportation Security Administration's attempt to reverse the ban on small blades on airplanes. Twice it tried to make the change, and twice fearful politicians prevented it from going through with it.

Lots of unneeded and ineffective security measures are perpetrated by a government bureaucracy that is primarily concerned about the security of its members' careers. They know the voters are more likely to punish them more if they fail to secure against a repetition of the last attack, and less if they fail to anticipate the next one.

What can we do? Well, the first step toward solving a problem is recognizing that you have one. These are not iron-clad rules; they're tendencies. If we can keep these tendencies and their causes in mind, we're more likely to end up with sensible security measures that are commensurate with the threat, instead of a lot of security theater and draconian police powers that are not.

Our leaders' job is to resist these tendencies. Our job is to support politicians who do resist.

References

Consolidation: Plague or Progress

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1303850_idx2,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1303850_idx2,00.html)

Prediction: RSA Conference Will Shrink Like a Punctured Balloon

[http://www.wired.com/politics/security/news/2008/04/
securitymatters_0417](http://www.wired.com/politics/security/news/2008/04/securitymatters_0417)

How to Sell Security

<http://www.cio.com/article/367913>

Why Do We Accept Signatures by Fax?

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/05/securitymatters_0529](http://www.wired.com/politics/security/commentary/securitymatters/2008/05/securitymatters_0529)

The Pros and Cons of LifeLock

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/06/securitymatters_0612](http://www.wired.com/politics/security/commentary/securitymatters/2008/06/securitymatters_0612)

The Problem Is Information Insecurity

<https://www.schneier.com/essay-233.html>

Security ROI: Fact or Fiction?

[http://www.csoonline.com/article/446866/
Security_ROI_Fact_or_Fiction_](http://www.csoonline.com/article/446866/Security_ROI_Fact_or_Fiction_)

Social Networking Risks

[http://searchsecurity.techtarget.com/magazineContent/
Schneier-Ranum-debate-social-networking-risks](http://searchsecurity.techtarget.com/magazineContent/Schneier-Ranum-debate-social-networking-risks)

Do You Know Where Your Data Are?

<http://online.wsj.com/article/SB123997522418329223.html>

Be Careful When You Come to Put Your Trust in the Clouds

[http://www.guardian.co.uk/technology/2009/jun/04/
bruce-schneier-cloud-computing](http://www.guardian.co.uk/technology/2009/jun/04/bruce-schneier-cloud-computing)

Is Perfect Access Control Possible?

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1365957,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1365957,00.html)

News Media Strategies for Survival for Journalists

[http://www.tcdailyplanet.net/news/2009/11/11/
opinion-news-media-strategies-survival](http://www.tcdailyplanet.net/news/2009/11/11/opinion-news-media-strategies-survival)

Security and Function Creep

<https://www.schneier.com/essay-307.html>

Weighing the Risk of Hiring Hackers

[http://searchsecurity.techtarget.com/magazinePrintFriendly/
0,296905,sid14_gci1514250,00.html](http://searchsecurity.techtarget.com/magazinePrintFriendly/0,296905,sid14_gci1514250,00.html)

Should Enterprises Give In to IT Consumerization at the Expense of Security?

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1519679,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1519679,00.html)

The Vulnerabilities Market and the Future of Security

[http://www.forbes.com/sites/bruceschneier/2012/05/30/
the-vulnerabilities-market-and-the-future-of-security/](http://www.forbes.com/sites/bruceschneier/2012/05/30/the-vulnerabilities-market-and-the-future-of-security/)

So You Want to Be a Security Expert

[http://krebsonsecurity.com/2012/07/
how-to-break-into-security-schneier-edition/](http://krebsonsecurity.com/2012/07/how-to-break-into-security-schneier-edition/)

When It Comes to Security, We're Back to Feudalism

<http://www.wired.com/opinion/2012/11/feudal-security/>

You Have No Control Over Security on the Feudal Internet

http://blogs.hbr.org/cs/2013/06/you_have_no_control_over_s.html

America's Dilemma: Close Security Holes, or Exploit Them Ourselves

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/05/blog_securitymatters_0501](http://www.wired.com/politics/security/commentary/securitymatters/2008/05/blog_securitymatters_0501)

Are Photographers Really a Threat?

<http://www.guardian.co.uk/technology/2008/jun/05/news.terrorism>

CCTV Doesn't Keep Us Safe, yet the Cameras Are Everywhere

<http://www.guardian.co.uk/technology/2008/jun/26/politics.ukcrime>

Chinese Cyberattacks: Myth or Menace?

http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1321716,00.html

How a Classic Man-in-the-Middle Attack Saved Colombian Hostages

http://www.wired.com/politics/security/commentary/securitymatters/2008/07/securitymatters_0710

How to Create the Perfect Fake Identity

http://www.wired.com/politics/security/commentary/securitymatters/2008/09/securitymatters_0904

A Fetishistic Approach to Security Is a Perverse Way to Keep Us Safe

<http://www.guardian.co.uk/technology/2008/sep/04/terrorism.terrorismandtravel>

The Seven Habits of Highly Ineffective Terrorists

http://www.wired.com/politics/security/commentary/securitymatters/2008/10/securitymatters_1002

Why Society Should Pay the True Costs of Security

<http://www.guardian.co.uk/technology/2008/oct/02/chemicals.terrorism>

Why Technology Won't Prevent Identity Theft

<http://online.wsj.com/article/SB123125633551557469.html>

Terrorists May Use Google Earth, but Fear Is No Reason to Ban It

<http://www.guardian.co.uk/technology/2009/jan/29/read-me-first-google-earth>

Thwarting an Internal Hacker

<http://online.wsj.com/article/SB123447990459779609.html>

An Enterprising Criminal Has Spotted a Gap in the Market

<http://www.guardian.co.uk/technology/2009/apr/02/google-earth-censorship-crime-comodities>

We Shouldn't Poison Our Minds with Fear of Bioterrorism

<http://www.guardian.co.uk/technology/2009/may/14/bruce-schneier-bioterrorism>

Raising the Cost of Paperwork Errors Will Improve Accuracy

<http://www.guardian.co.uk/technology/2009/jun/24/read-me-first-identity-fraud>

So-Called Cyberattack Was Overblown

<http://minnesota.publicradio.org/display/web/2009/07/10/schneier/>

Why Framing Your Enemies Is Now Virtually Child's Play

[http://www.guardian.co.uk/technology/2009/oct/15/
bruce-schneier-internet-security](http://www.guardian.co.uk/technology/2009/oct/15/bruce-schneier-internet-security)

Beyond Security Theater

<http://www.schneier.com/essay-292.html>

Cold War Encryption Is Unrealistic in Today's Trenches

<http://search.japantimes.co.jp/cgi-bin/nc20091230a2.html>

Profiling Makes Us Less Safe

[http://roomfordebate.blogs.nytimes.com/2010/01/04/
will-profiling-make-a-difference/](http://roomfordebate.blogs.nytimes.com/2010/01/04/will-profiling-make-a-difference/)

Fixing Intelligence Failures

[http://www.sfgate.com/cgi-bin/article.cgi?f=/c/a/2010/01/14/
ED0L1BIFK6.DTL](http://www.sfgate.com/cgi-bin/article.cgi?f=/c/a/2010/01/14/ED0L1BIFK6.DTL)

Spy Cameras Won't Make Us Safer

[http://www.cnn.com/2010/OPINION/02/25/schneier.security.cameras/
index.html](http://www.cnn.com/2010/OPINION/02/25/schneier.security.cameras/index.html)

Scanners, Sensors Are Wrong Way to Secure the Subway

[http://www.nydailynews.com/opinions/2010/04/07/2010-04-07
_scanners_sensors_are_wrong_way_to_secure_the_subway.html](http://www.nydailynews.com/opinions/2010/04/07/2010-04-07_scanners_sensors_are_wrong_way_to_secure_the_subway.html)

Preventing Terrorist Attacks in Crowded Areas

[http://roomfordebate.blogs.nytimes.com/2010/05/03/
times-square-bombs-and-big-crowds/?src=tptw#bruce](http://roomfordebate.blogs.nytimes.com/2010/05/03/times-square-bombs-and-big-crowds/?src=tptw#bruce)

Where Are All the Terrorist Attacks?

[http://www.aolnews.com/2010/05/04/
opinion-where-are-all-the-terrorist-attacks/](http://www.aolnews.com/2010/05/04/opinion-where-are-all-the-terrorist-attacks/)

Worst-Case Thinking Makes Us Nuts, Not Safe

[http://www.cnn.com/2010/OPINION/05/12/schneier.worst.case
.thinking/](http://www.cnn.com/2010/OPINION/05/12/schneier.worst.case.thinking/)

Threat of "Cyberwar" Has Been Hugely Hyped

<http://www.cnn.com/2010/OPINION/07/07/schneier.cyberwar.hyped/>

Cyberwar and the Future of Cyber Conflict

[http://www.ft.com/cms/s/0/f863fb4c-fe53-11df-abac-00144feab49a
.html](http://www.ft.com/cms/s/0/f863fb4c-fe53-11df-abac-00144feab49a.html)

Why Terror Alert Codes Never Made Sense

[http://www.cnn.com/2011/OPINION/01/28/schneier.terror.threat
.level/](http://www.cnn.com/2011/OPINION/01/28/schneier.terror.threat.level/)

Debate Club: An International Cyberwar Treaty Is the Only Way to Stem the Threat

<http://www.usnews.com/debate-club/should-there-be-an-international-treaty-on-cyberwarfare/an-international-cyberwar-treaty-is-the-only-way-to-stem-the-threat>

Overreaction and Overly Specific Reactions to Rare Risks

<http://www.cnn.com/2012/07/31/opinion/schneier-aurora-aftermath/index.html>

Militarizing Cyberspace Will Do More Harm Than Good

<http://www.irishtimes.com/newspaper/finance/2012/1129/1224327257946.html>

Rhetoric of Cyber War Breeds Fear—And More Cyber War

<https://www.schneier.com/essay-421.html>

The Boston Marathon Bombing: Keep Calm and Carry On

<http://www.theatlantic.com/national/archive/2013/04/the-boston-marathon-bombing-keep-calm-and-carry-on/275014/>

Why FBI and CIA Didn't Connect the Dots

<http://www.cnn.com/2013/05/02/opinion/schneier-boston-bombing/index.html>

The FBI's New Wiretapping Plan Is Great News for Criminals

http://www.foreignpolicy.com/articles/2013/05/29/the_fbi_s_new_wiretapping_plan_is_great_news_for_criminals

US Offensive Cyberwar Policy

<http://www.cnn.com/2013/06/18/opinion/schneier-cyberwar-policy/index.html>

Secret Questions Blow a Hole in Security

<http://www.computerweekly.com/Articles/2008/04/04/230137/secret-questions-blow-a-hole-in-security.htm>

When You Lose a Piece of Kit, the Real Loss Is the Data It Contains

<http://www.guardian.co.uk/technology/2008/dec/04/computer-security-mobile-phones>

The Kindness of Strangers

<http://online.wsj.com/article/SB123567809587886053.html>

Blaming the User Is Easy—But It's Better to Bypass Them Altogether

<http://www.guardian.co.uk/technology/2009/mar/12/read-me-first>

The Value of Self-Enforcing Protocols

<http://threatpost.com/blogs/value-self-enforcing-protocols>

Reputation is Everything in IT Security

<http://www.guardian.co.uk/technology/2009/nov/11/schneier-reputation-it-security>

When to Change Passwords

<http://www.darkreading.com/hacked-off/when-to-change-passwords/228201016>

The Big Idea: Bruce Schneier

<http://whatever.scalzi.com/2012/02/16/the-big-idea-bruce-schneier/>

High-Tech Cheats in a World of Trust

<http://www.newscientist.com/article/mg21328530.400-foiling-the-cheats-in-a-world-of-hightech-trust.html>

Detecting Cheaters

<https://www.schneier.com/essay-337.html>

Lance Armstrong and the Prisoner's Dilemma of Doping in Professional Sports

<http://www.wired.com/opinion/2012/10/lance-armstrong-and-the-prisoners-dilemma-of-doping-in-professional-sports/>

Trust and Society

<http://www.themontrealreview.com/2009/Liars-and-Outliers-Bruce-Schneier.php>

How Secure Is the Papal Election?

<http://www.cnn.com/2013/02/20/opinion/schneier-papal-election-secure/index.html>

The Court of Public Opinion

<http://www.wired.com/opinion/2013/02/court-of-public-opinion/>

On Security Awareness Training

<http://www.darkreading.com/blog/240151108/on-security-awareness-training.html>

Our New Regimes of Trust

<https://www.schneier.com/essay-410.html>

<https://www.schneier.com/essay-253.html>

https://www.schneier.com/blog/archives/2013/02/our_new_regimes.html

The Myth of the “Transparent Society”

http://www.wired.com/politics/security/commentary/securitymatters/2008/03/securitymatters_0306

Our Data, Ourselves

http://www.wired.com/politics/security/commentary/securitymatters/2008/05/securitymatters_0515

The Future of Ephemeral Conversation

<http://online.wsj.com/article/SB122722381368945937.html>

How to Prevent Digital Snooping

<http://online.wsj.com/article/SB122877438178489235.html>

Architecture of Privacy

<https://www.schneier.com/essay-253.html>

Privacy in the Age of Persistence

<http://news.bbc.co.uk/1/hi/technology/7897892.stm>

Should We Have an Expectation of Online Privacy?

http://searchsecurity.techtarget.com/magazinePrintFriendly/0,296905,sid14_gc11354832,00.html

Offhand but On Record

<http://www.japantimes.co.jp/life/2009/08/19/life/offhand-but-on-record/>

Google’s and Facebook’s Privacy Illusion

<http://www.forbes.com/2010/04/05/google-facebook-twitter-technology-security-10-privacy.html>

The Internet: Anonymous Forever

<http://www.forbes.com/2010/05/12/privacy-hackers-internet-technology-security-anonymity.html>

A Taxonomy of Social Networking Data

<https://www.schneier.com/essay-322.html>

The Difficulty of Surveillance Crowdsourcing

http://threatpost.com/en_us/blogs/difficulty-surveillance-crowdsourcing-110810

The Internet Is a Surveillance State

<http://www.cnn.com/2013/03/16/opinion/schneier-internet-surveillance/index.html>

Surveillance and the Internet of Things

<http://www.guardian.co.uk/technology/2013/may/16/internet-of-things-privacy-google>

Government Secrets and the Need for Whistleblowers

<http://www.theatlantic.com/politics/archive/2013/06/what-we-dont-know-about-spying-on-citizens-scarier-than-what-we-know/276607/>

Before Prosecuting, Investigate the Government

<http://www.nytimes.com/roomfordebate/2013/06/11/in-nsa-leak-case-a-whistle-blower-or-a-criminal/before-prosecuting-snowden-investigate-the-government>

The Security Mindset

http://www.wired.com/politics/security/commentary/securitymatters/2008/03/securitymatters_0320

The Difference between Feeling and Reality in Security

http://www.wired.com/politics/security/commentary/securitymatters/2008/04/securitymatters_0403

How the Human Brain Buys Security

http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=4588238&isnumber=4588217&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxppls%2Fabs_all.jsp%3Farnumber%3D4588238%26isnumber%3D4588217

Does Risk Management Make Sense?

http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1332745_idx2,00.html

How the Great Conficker Panic Hacked into Human Credulity

<http://www.guardian.co.uk/technology/2009/apr/23/conficker-panic>

How Science Fiction Writers Can Help, or Hurt, Homeland Security

http://www.wired.com/politics/security/commentary/securitymatters/2009/06/securitymatters_0619

Privacy Salience and Social Networking Sites

<http://www.guardian.co.uk/technology/2009/jul/15/privacy-internet-facebook>

Security, Group Size, and the Human Brain

<https://www.schneier.com/essay-275.html> July/August 2009

People Understand Risks—But Do Security Staff Understand People?

[http://www.guardian.co.uk/technology/2009/aug/05/
bruce-schneier-risk-security](http://www.guardian.co.uk/technology/2009/aug/05/bruce-schneier-risk-security)

Nature's Fears Extend to Online Behavior

<http://www.japantimes.co.jp/text/nc20091118a1.html>

The Ethics of Vulnerability Research

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1313268,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1313268,00.html)

I've Seen the Future, and It Has a Kill Switch

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/06/securitymatters_0626](http://www.wired.com/politics/security/commentary/securitymatters/2008/06/securitymatters_0626)

Software Makers Should Take Responsibility

[http://www.guardian.co.uk/technology/2008/jul/17/internet
.security](http://www.guardian.co.uk/technology/2008/jul/17/internet.security)

Lesson from the DNS Bug: Patching Isn't Enough

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/07/securitymatters_0723](http://www.wired.com/politics/security/commentary/securitymatters/2008/07/securitymatters_0723)

Why Being Open about Security Makes Us All Safer in the Long Run

<http://www.guardian.co.uk/technology/2008/aug/07/hacking.security>

Boston Court's Meddling with "Full Disclosure" Is Unwelcome

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/08/securitymatters_0821](http://www.wired.com/politics/security/commentary/securitymatters/2008/08/securitymatters_0821)

Quantum Cryptography: As Awesome As It Is Pointless

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/10/securitymatters_1016](http://www.wired.com/politics/security/commentary/securitymatters/2008/10/securitymatters_1016)

Passwords Are Not Broken, but How We Choose Them Sure Is

[http://www.guardian.co.uk/technology/2008/nov/13/
internet-passwords](http://www.guardian.co.uk/technology/2008/nov/13/internet-passwords)

America's Next Top Hash Function Begins

[http://www.wired.com/politics/security/commentary/
securitymatters/2008/11/securitymatters_1120](http://www.wired.com/politics/security/commentary/securitymatters/2008/11/securitymatters_1120)

Tigers Use Scent, Birds Use Calls—Biometrics Are Just Animal Instinct

[http://www.guardian.co.uk/technology/2009/jan/08/
identity-fraud-security-biometrics-schneier-id](http://www.guardian.co.uk/technology/2009/jan/08/identity-fraud-security-biometrics-schneier-id)

The Secret Question Is: Why Do IT Systems Use Insecure Passwords?

[http://www.guardian.co.uk/technology/2009/feb/19/
insecure-passwords-conflickerb-worm](http://www.guardian.co.uk/technology/2009/feb/19/insecure-passwords-conflickerb-worm)

The Pros and Cons of Password Masking

http://www.schneier.com/blog/archives/2009/07/the_pros_and_co.html

Technology Shouldn't Give Big Brother a Head Start

<http://minnesota.publicradio.org/display/web/2009/07/30/schneier/>

Lockpicking and the Internet

[http://www.darkreading.com/hacked-off/
lockpicking-and-the-internet/227700647](http://www.darkreading.com/hacked-off/lockpicking-and-the-internet/227700647)

The Battle Is On against Facebook and Co. to Regain Control of Our Files

[http://www.guardian.co.uk/technology/2009/sep/09/
bruce-schneier-file-deletion](http://www.guardian.co.uk/technology/2009/sep/09/bruce-schneier-file-deletion)

The Difficulty of Un-Authentication

<http://threatpost.com/difficulty-un-authentication-092809/72267>

Is Antivirus Dead?

[http://searchsecurity.techtarget.com/magazinePrintFriendly/
0,296905,sid14_gci1373562,00.html](http://searchsecurity.techtarget.com/magazinePrintFriendly/0,296905,sid14_gci1373562,00.html)

Virus and Protocol Scares Happen Every Day—But Don't Let Them Worry You

[http://www.guardian.co.uk/technology/blog/2009/dec/09/
virus-protocol-scares-every-day](http://www.guardian.co.uk/technology/blog/2009/dec/09/virus-protocol-scares-every-day)

The Failure of Cryptography to Secure Modern Networks

[http://www.darkreading.com/blog/227700878/
the-failure-of-cryptography-to-secure-modern-networks.html](http://www.darkreading.com/blog/227700878/the-failure-of-cryptography-to-secure-modern-networks.html)

The Story behind The Stuxnet Virus

[http://www.forbes.com/2010/10/06/
iran-nuclear-computer-technology-security-stuxnet-worm
.html](http://www.forbes.com/2010/10/06/iran-nuclear-computer-technology-security-stuxnet-worm.html)

The Dangers of a Software Monoculture

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1522895,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1522895,00.html)

How Changing Technology Affects Security

<https://www.schneier.com/essay-392.html>

The Importance of Security Engineering

http://www.schneier.com/blog/archives/2012/08/the_importance_1.html

Technologies of Surveillance

http://www.schneier.com/blog/archives/2013/03/technologies_of.html

When Technology Overtakes Security

<http://www.wired.com/opinion/2013/03/security-when-the-bad-guys-have-technology-too-how-do-we-survive/>

Memo to Next President: How to Get Cybersecurity Right

http://www.wired.com/politics/security/commentary/securitymatters/2008/08/securitymatters_0807

CRB Checking

<https://www.schneier.com/essay-277.html>

State Data Breach Notification Laws: Have They Helped?

http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1344729_idx2,00.html

How to Ensure Police Database Accuracy

<http://online.wsj.com/article/SB123301316511017419.html>

How Perverse Incentives Drive Bad Security Decisions

http://www.wired.com/politics/security/commentary/securitymatters/2009/02/securitymatters_0226

It's Time to Drop the "Expectation of Privacy" Test

http://www.wired.com/politics/security/commentary/securitymatters/2009/03/securitymatters_0326

Who Should Be in Charge of Cybersecurity?

<http://online.wsj.com/article/SB123844579753370907.html>

Coordinate, but Distribute Responsibility

<http://roomfordebate.blogs.nytimes.com/2009/05/29/a-plan-of-attack-in-cyberspace/#bruce>

"Zero Tolerance" Really Means Zero Discretion

<http://minnesota.publicradio.org/display/web/2009/11/03/schneier/>

US Enables Chinese Hacking of Google

<http://www.cnn.com/2010/OPINION/01/23/schneier.google.hacking/>

Should the Government Stop Outsourcing Code Development?

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1507396,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1507396,00.html)

Punishing Security Breaches

https://www.schneier.com/blog/archives/2010/04/punishing_security_breaches.html

Three Reasons to Kill the Internet Kill Switch Idea

<https://www.schneier.com/essay-321.html>

Web Snooping Is a Dangerous Move

[http://www.cnn.com/2010/OPINION/09/29/schneier.web.surveillance/
index.html?iref=allsearch](http://www.cnn.com/2010/OPINION/09/29/schneier.web.surveillance/index.html?iref=allsearch)

The Plan to Quarantine Infected Computers

[http://www.forbes.com/2010/11/10/
microsoft-viruses-security-technology-quarantine.html](http://www.forbes.com/2010/11/10/microsoft-viruses-security-technology-quarantine.html)

Close the Washington Monument

<https://www.schneier.com/essay-332.html>
[http://www.nydailynews.com/opinions/2010/12/02/2010-12-02
_washington_monument_security_debate_is_the_height_of
_irrational_fearmongering.html](http://www.nydailynews.com/opinions/2010/12/02/2010-12-02_washington_monument_security_debate_is_the_height_of_irrational_fearmongering.html)

Whitelisting and Blacklisting

[http://searchsecurity.techtarget.com/
magazineFeature/0,296894,sid14_gci1526425,00.html](http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1526425,00.html)

Securing Medical Research: A Cybersecurity Point of View

<http://www.sciencemag.org/content/336/6088/1527.full>

Fear Pays the Bills, but Accounts Must Be Settled

[http://www.nytimes.com/roomfordebate/2012/10/17/
should-industry-face-more-cybersecurity-mandates/
fear-pays-the-bills-but-accounts-must-be-settled](http://www.nytimes.com/roomfordebate/2012/10/17/should-industry-face-more-cybersecurity-mandates/fear-pays-the-bills-but-accounts-must-be-settled)

Power And The Internet

<http://edge.org/response-detail/23818>

Danger Lurks in Growing New Internet Nationalism

[http://www.technologyreview.com/view/512386/
danger-lurks-in-growing-new-internet-nationalism/](http://www.technologyreview.com/view/512386/danger-lurks-in-growing-new-internet-nationalism/)

IT for Oppression

<https://www.schneier.com/essay-420.html>

The Public/Private Surveillance Partnership

<http://www.theatlantic.com/technology/archive/13/04/governments-wont-need-to-issue-ids-data-brokers-will-identify-you-for-them/275431/>

Transparency and Accountability Don't Hurt Security— They're Crucial to It

<http://www.theatlantic.com/politics/archive/2013/05/transparency-and-accountability-dont-hurt-security-theyre-crucial-to-it/275662/>

It's Smart Politics to Exaggerate Terrorist Threats

<http://www.cnn.com/2013/05/20/opinion/schneier-security-politics/index.html>

State Data Breach Notification Laws: Have They Helped?

http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1344729_idx2,00.html

How to Ensure Police Database Accuracy

<http://online.wsj.com/article/SB123301316511017419.html>

Crossing Borders with Laptops and PDAs

<http://www.guardian.co.uk/technology/2008/may/15/computing.security>

The TSA's Useless Photo ID Rules

<http://www.latimes.com/news/opinion/1a-oe-schneier28-2008aug28,0,3099808.story>

The Two Classes of Airport Contraband

http://www.wired.com/politics/security/commentary/securitymatters/2008/09/securitymatters_0918

Fixing Airport Security

http://www.nydailynews.com/opinions/2009/06/24/2009-06-24_clear_common_sense_for_takeoff_how_the_tsa_can_make_airport_security_work_for_pa.html

Laptop Security while Crossing Borders

[http://www.wired.com/politics/security/commentary/
securitymatters/2009/07/securitymatters_0715](http://www.wired.com/politics/security/commentary/securitymatters/2009/07/securitymatters_0715)

Breaching the Secure Area in Airports

[http://threatpost.com/en_us/blogs/
fixing-security-problem-isnt-always-right-answer-010510](http://threatpost.com/en_us/blogs/fixing-security-problem-isnt-always-right-answer-010510)

Stop the Panic on Air Security

<http://www.cnn.com/2010/OPINION/01/07/schneier.security/>

A Waste of Money and Time

[http://www.nytimes.com/roomfordebate/2010/11/22/
do-body-scanners-make-us-safer/a-waste-of-money-and-time](http://www.nytimes.com/roomfordebate/2010/11/22/do-body-scanners-make-us-safer/a-waste-of-money-and-time)

Why the TSA Can't Back Down

[http://www.theatlantic.com/national/archive/2010/12/
why-the-tsa-cant-back-down/67337/](http://www.theatlantic.com/national/archive/2010/12/why-the-tsa-cant-back-down/67337/)

The Trouble with Airport Profiling

[http://www.forbes.com/sites/bruce-schneier/2012/05/09/
the-trouble-with-airport-profiling/](http://www.forbes.com/sites/bruce-schneier/2012/05/09/the-trouble-with-airport-profiling/)

Index

Numbers

9/11, counterterrorism and, 6

A

Abrahms, Max, 55

access control, 22–24

accountability, security and, 327–329

actuarial tail, 15

airport

contraband, 262–263

full-body scanners, 254

National Opt Out Day, 273

panic over security, 269–275

profiling, 275–277

rarity of terrorism, 270–275

secure area breaches, 268–269

Secure Flight program, 265

security fixes, 264–265

ALE (annualized loss expectancy),

14–16

alert codes, 96–97

alibi clubs, 121

anonymity

Internet and, 173–175

Web browsing, 121

antivirus effectiveness, 238–240

Applied Cryptography, 2nd Edition

(Schneier), 220

architecture of privacy, 162–164

arms race for cyberwar, 98–99,

101–103

Armstrong, Lance, 137–141

Asian Disease Experiment, 5

audit, 161–162

authentication

decentralized systems, 60

expiring credentials, 127–129

identity theft and, 59–60

unauthentication, 237–238

AVG antivirus, 239

B

background checks, CRB (Criminal
Records Bureau), 281–283

backups, lost devices, 119

Baldwin, Bob, 233–235

behavioral data, 176

Bennett, Charles, 220–221

Berge, Tom, 65–66

biometrics, 225–227

bioterrorism, 66–68

BitLocker, 266

blacklisting/whitelisting, 312–313

border crossing, 259–260, 265–266

Boston Marathon bombing,

105–107

Boyd, Danah, 149

brain fight-or-flight reflex, 207

size and group size study, 203–204

Brassard, Giles, 220–221

breach notification laws, 283–285

breaches, punishing, 300–301

Brin, David, *The Transparent Society*, 155
 Broadwell, Paula, 179
 bypassing users, 122–123

C

cameras. *See also* CCTV (closed-circuit TV)
 photographers, 43–45
 safety and, 82–84
 surveillance crowdsourcing, 177–179

careers in security, 33–34

CCTV (closed-circuit TV), 45–46, 82–84

cell phones. *See* mobile phones

censorship
 China, 231–233
 content filtering, 324

certifications, security careers, 34

cheating
 detecting cheaters, 134–137
 trust and, 131–134

child safety, CRB (Criminal Records Bureau), 281–283

China
 censorship, 231–233
 cyberattacks, 47–48, 103–104, 179
 Google hacking, 297–298

CIA, 107–109

cloud computing, 19–22. *See also* Saas (software as a service)
 data availability, 169
 feudal Internet and, 37–40
 file control, 235–236

code as law, 163

code development, government outsourcing, 299–300

cognitive bias, 194–195, 199–200

Cold War, encryption, 78–79

collaborative spam filtering, 121

Colombian hostages, MITM (man-in-the-middle) attack and, 48–51

communications
 identity theft and, 59
 infrastructure, terrorism and, 61–62
 saved, 158–160

COMSEC (communications security), vulnerabilities and, 32

Concurring Opinions, 289–290

conferences, RSA, 2–4

Conficker, 197–198

Conficker.B, 227–228

consolidation, 1–2

consumerization of IT, 29–30

contamination of food, terrorism and, 66–68

content filtering, oppression and, 324

contraband in airports, 262–263

conversation
 digital footprints and, 164–166
 trails, 158–160

copper wiring theft, 65

correlation of digital footprint, 179–181

Cosmides, Leda, 137

counterterrorism, 6
 crowds, 86–87
 motivations of terrorists, 54–56

court of public opinion, 147–149

CRB (Criminal Records Bureau), 281–283

credit bureaus, fraud alert, 9–12

credit cards
 LifeLock and, 9–12
 signatures, 8

Crespo, Erik, 156

crowdsourcing surveillance, 177–179

cryptography, 34
 hash standards, 223–225
 modern networks, 242–244
 quantum cryptography, 220–221

customs
 devices and, 259–260
 encryption and, 266
 laptops and, 265–266

cyberattacks
 Chinese, 47–48, 103–104
 Estonia, 41–43, 70–72
 GhostNet, 104–105
 kids playing politics, 41–43, 70–72
 North Korea attack on US, 70–72
 technologies, 42–43

cybersecurity
 advice to president, 279–281
 financial issues, 317–318
 government oversight, 291–293
 medical research, 313–317
 Obama administration policy, 294

cyberspace, militarizing, 101–103

cyberwar
 arms race, 98–99, 101–103
 cyber conflict and, 94–95
 Estonia attacks, 92–93
 OCEO (Offensive Cyber Effect Operations), 112–113
 rhetoric and, 103–105
 Snowden, Edward, 112
 threat hype, 92–94
 treaty, 97–99
 US Cyber Command, 94, 113–114
 US policy, 112–115

D

data
 availability, 168–171
 behavioral, 176
 collection, digital footprints and, 164–165
 derived, 176
 disclosed, 175
 entrusted, 176
 government database accuracy, 285–286
 incidental, 176
 location, 18–20
 as pollution, 164–165
 service, 175
 social networking, taxonomy, 175–177
 trails, 168–171

data shadow, 157–158

Davis, Todd, 9, 10

death predictions of security industry, 3–4

default security practices, 241

DEFCON levels, 96

defection, scope of, 249–251

Department of Homeland Security.
See DHS (Department of Homeland Security)

derived data, 176

devices
 border agents and, 259–260
 border crossing and, 265–266
 lost, 118–120

DHS (Department of Homeland Security)
 open source software, 42
 science fiction writers and, 198–201
 terror alert codes, 96–97

digital footprints, 164–166
 conversation trails, 158–160
 correlation and, 179–181
 data shadow, 157–158
 trails of data, 157–158

Digital Manners Policies, 211

digital snooping, preventing, 160–162

disclosed data, 175

DNA fabrication, 72

DNS bug, 214–216

Doctorow, Cory, 252

Domain Awareness System (NYPD), 253–254

doping in professional sports, 137–141

driver's licenses, 26

dual use technology, 41–42
 NSA, 42–43

E

eavesdropping, 304–306
 wiretapping, 231–233
 Internet, 109–112

economics of security, 1–2
 elections, papal, 143–147
 emotion
 counterterrorism and, 6
 security sales, 6
 encryption, 77–79
 border agents and, 259–260
 customs and, 266
 keys, 242
 lost devices, 119
 passwords, 260
 Engle, Mary K., 20
 entrusted data, 176
 ephemeral conversation, 158–160
 EPIC (Electronic Privacy Information Center), 19
 equities issues, 42
 errors in paperwork, 68–70
 Estonia
 cyberattack on, 41–43, 70–72
 cyberwar and, 92–93
 evolutionary pressure on society,
 crime and, 65–66
 expectation of privacy, 288–291
 expiring passwords, 127–129

F

Facebook, data availability, 168
 failures in intelligence, 81–82
 false identity, 51–52
 Fannie Mae, 62
 fax signatures, 7–9
 FBI
 intelligence data, 107–109
 wiretapping the Internet, 109–112
 fear
 Conficker and, 197–198
 natural, online behavior and,
 206–208
 security, 194–195
 selling, 6–7
 fetishistic approach to security, 52–54

feudal Internet, 37–40
 feudal security, 34–37, 39
 fight-or-flight reflex, 207
 file control, 235–236
 financial issues of cybersecurity,
 317–318
 financial systems, MITM (man-in-the-
 middle) attacks, 49–50
 firewalls, 239
 food contamination (terrorism and),
 66–68
 forgery, 68–70
 Fourth Amendment, 288–291
 framing others, 72–73
 fraud, 72–73
 fraud alert, credit bureaus,
 9–12
 Freakonomics, 300–301
 full disclosure, 218–220
 full-body scanner in airports, 254
 National Opt Out Day, 273
 function creep, 26–27
 Furedi, Frank, 91

G

generation gap of Internet, 17–18, 82,
 162
 GhostNet, 104–105
 Google
 Chinese hacking, 297–298
 cloud computing and, 19–20
 Google Docs, 20
 Google Earth, 60–62
 Google Glass, 183
 government
 code development outsourcing,
 299–300
 cybersecurity oversight, 291–293
 whistleblowers and, 184–187
 Green Dam Youth Escort, 231–233
 group size, 203–204

H

hackers

- hiring risks, 27–28
- internal, 62–64

hacktivism, 102

hardware, lost, 118–120

Harris, Sam, 251

hash standards, cryptography,
223–225

helpful strangers, 120–122

hindsight bias, 108

HomeGuard, 178

I

ID checks

- security theater and, 74
- TSA, 261–262

identity, false, 51–52

identity farming, 51–52

identity theft

- authentication and, 59–60
- breach notification laws, 283–285
- communications and, 59
- forgery, 68–70
- LifeLock, 9–12
- signatures, 7–9
- technology in prevention, 58–60

incidental data, 176

information insecurity, 12–13

infrastructure

- IT as, 126–127
- outsourcing, 1–2

insecure passwords in IT, 227–229

intelligence

- CIA, 107–109
- failures, fixing, 81–82
- FBI, 107–109

internal hackers, 62–64

Internet

- anonymity and, 173–175
- code is law, 163

court of public opinion, 149

feudal, 37–40

generation gap, 17–18, 82, 162

locks and, 233–235

nationalism, 321–323

power and, 319–321

quarantine of infected computers,
307–310

as surveillance state, 179–181

wiretapping, 109–112

Internet of Things, 181–184

intuition and security, 251

Irish Potato Famine, 247–248

IT

consumerization, 29–30

as infrastructure, 126–127

oppression and, 323–325

passwords, 227–229

services, 125–126

J

Johnson, Eric, 23–24

journalists, 24–25

KKahneman, Daniel, 4–5,
108, 194

Kaminsky, 214–216

Kerr, Orin, 290

kids playing politics, 41–43,
70–72

kill switch, 211–212, 302–304

L

LifeLock, 9–12

liquid bombers, 75–76

airport contraband, 263

Litvinenko, Alexander, 67

logic bomb, 62

logical reasoning, 135–136

lost devices, 118–120

M

Malwarebytes' Anti-Malware, 239
 masking passwords, 229–231
 medical research, 313–317
 mental risk thermostat, 206
 Microsoft UAC (User Account Control), 122–123
 Mifare Classic Chip, 216–218
 military
 dual use technology, 41–42
 encryption, 77
 equities issues, 42
 militarization of cyberspace, 101–103
 Predators, 77–79
 MITM (man-in-the-middle) attacks
 Colombian hostages, 48–51
 financial systems, 49–50
 NSA and, 50
 mobile phones
 data availability, 169
 records to NSA, 184
 terrorists and, 61
modus ponens, 135, 136
modus tollens, 135, 136
 monoculture of software, 247–249
 Monsegur, Hector, 179
 Mook, Heather, 67
 movie-plot threat, 44–45. *See also*
 security theater
 false identity, 52
 not responding, 76–77
 subway security, 84–86
 terrorism, rarity of, 73–77
 myth of transparent society, 155–157

N

National Opt Out Day, 273
 nationalism, 321–323
 natural fears, 206–208
 networks, 242–244
 Nielson, Jakob, 229

nightmare scenarios, 89–92
 NIST (National Institute of Standards and Technology), 224–225
 no-fly list, 261–262
 North Korea, cyberattack on US, 70–72
 notary publics, 8
 NSA (National Security Agency)
 dual-use technology, 42–43
 equities issues, 42
 expectation of privacy and, 288–291
 legality, 188
 MITM (man-in-the-middle) attacks, 50
 mobile phone records, 184
 NYPD Domain Awareness System, 253–254

O

Obama, Barack
 advice on cybersecurity, 279–281
 BlackBerry surrender, 158
 cybersecurity policy, 294
 OCEO (Offensive Cyber Effect Operations), 112–113
 online behavior, 206–208
 online privacy, 167–168
 cloud computing, 169
 data availability, 168–171
 Facebook, 171–173
 Google, 171–173
 mobile phones, 169
 universal identification, 173–175
 outsourcing
 code development, government and, 299–300
 infrastructure, 1–2
 Oyster cards, 216

P

papal election, 143–147
 paperwork errors, 68–70
 Password Safe, 228

- passwords
 - changing, 127–129
 - encryption, 260
 - guessing, 222–223, 228
 - insecure in IT, 227–229
 - masking, 229–231
 - Password Safe, 228
 - SecurID, 228
 - security questions, 117–118
 - shoulder surfing, 229–230
 - social networking sites, 128
 - structure, 222
 - patches, 213
 - persistence, 164–166
 - Petraeus, David, 179
 - PGPDisk, 266
 - photo ID, 261–262
 - photographer rights document, 45
 - photographers
 - movie-plot threat, 44–45
 - terrorism and, 43–45
 - PLC (Programmable Logic Controller), 244–245
 - police database accuracy, 285–286
 - politics
 - cyberattacks, 70–72
 - privacy and, 163–164
 - terrorism and, 54–55
 - threat exaggeration, 329–331
 - pollution of data, 164–165
 - portable devices, 118–120
 - Powell, Gray, 300–301
 - Predators (military), 77–79
 - preventive measures, 161
 - Prisoner's Dilemma, 129–130
 - cheating and trust, 132
 - doping in professional sports, 137–141
 - privacy
 - architecture, 162–164
 - business, 163–164
 - conversation trails, 158–160
 - data availability, 168–171
 - data shadow, 157–158
 - digital snooping, 158–160
 - expectation of, 288–291
 - Facebook, 171–173
 - Google, 19–20, 171–173
 - Obama, Barack, 158
 - online, 167–168
 - persistence and, 164–166
 - politics and, 163–164
 - RealAge, 18–19
 - salience, 201–202
 - secrecy and, 171
 - surveillance crowdsourcing, 177–179
 - transparent society and, 155–157
 - Privacy International, 177
 - profiling, 80
 - airport, 275–277
 - Muslims, 251
 - propaganda, 324
 - propositional calculus, 135–137
 - Prospect Theory, 4–5
 - risk and, 194–195
 - protocols, self-enforcing, 123–125
 - psychology
 - risk analysis, 99–100
 - of security, 130–131
 - public opinion, 147–149
 - punishing security breaches, 300–301
 - purchasing security, 3–4
- Q**
- quantum cryptography, 220–221
 - quarantine of infected computers, 307–310
- R**
- RBAC. *See* role-based access control
 - RealAge, 18–19
 - regimes of trust, 152–154
 - reputational justice, 148
 - risk
 - human brain and, 194–195
 - mental risk thermostat, 206

Prospect Theory and, 194–195
 understanding of, 205–206
 risk analysis
 over-optimism and, 199
 psychology, 99–100
 risk management, 195–197
 ROI (return on investment), 14–16
 role-based access control, 23
 RSA conference, 2–4
 Rubinfeld, Jed, 290

S

SaaS (software as a service), 21–22
 Schneier's Law, 34
 science fiction writers, 198–201
 scope of defection, 249–251
 secret questions, 117–118
 secure area breaches, 268–269
 Secure Flight program, 265
 SecurID token, 228
 security
 awareness training, 150–152
 business of, 1
 costs paid by society, 56–58
 early model, 35
 economics of, 1
 as feeling, 74–75
 feeling vs. reality, 191–193
 industry, death predictions, 3–4
 intuition and, 251
 preventive measures, 161
 psychology of, 130–131
 as reality, 74–75
 selling, 3–4
 fear and, 6–7
 Prospect Theory, 4–5
 trade-offs, incentives, 287–288
 security creep, 26–27
 security engineering, 251–252
 security gap, 153–154
 security mindset, 189–191
 security theater, 74–77. *See also*
 movie-plot threat
 liquid bombers, 75–76
 self-enforcing protocols, 123–125
 September 11 attacks. *See* 9/11
 service data, 175
 SHA (Secure Hash Algorithm),
 224–225
 shadow data, 157–158
 Sidewiki, 73
 SIGINT (signals intelligence),
 vulnerabilities and, 32
 signatures
 credit cards, 8
 invalid documents, 8
 notary publics, 8
 via fax, 7–9
 SmartWater, 189
 Snowden, Edward, 112, 187–188
 social group size, 203–204
 social networking, 17–18
 data taxonomy, 175–177
 passwords, 128
 privacy salience and, 201–202
 societal pressures, 132, 141–143
 software
 insecurity, 12–13
 liabilities, 212–214
 producer responsibility, 212–214
 as a service, 3–4, 21–22
 software monoculture, 247–249
 Solove, Daniel, 289–290
 spam, 121
 SSL protocol security flaw, 240–242
 strangers, 120–122
 Stuxnet, 244–247
 subway security, 84–86
 surveillance, 231–233
 crowdsourcing, 177–179
 eavesdropping, 304–306
 Google Glass, 183

Internet of Things and, 181–184
 oppression and, 324
 public/private partnership, 325–326
 technologies, 253–254
 surveillance state, Internet as,
 179–181

T

taxonomy of social networking data,
 175–177
 technology
 attack technologies, 42–43
 changes effects on security, 249–251
 dual use, 41–42
 identity theft prevention and,
 58–60
 overtaking security, 255–257
 of surveillance, 253–254
 terror alert codes, 96–97
 terrorism
 Abrahms, Max, 55
 bioterrorism, 66–68
 communications infrastructure,
 61–62
 counterterrorism, 6
 crowded areas, 86–87
 difficulty of, 88
 driver's licenses, 26
 food contamination, 66–68
 goals, 76, 106
 Google Earth and, 60–62
 lack of attacks, 87–89
 liquid bombers, 75–76
 mobile phones, 61
 motivations of terrorists, 54–56
 Muslim profiling, 251
 no-fly list, 261–262
 number of terrorists, 88
 photographers, 43–45
 politics and, 54–55
 rarity of, 73–77, 100, 270–275

refusal to be terrorized, 76–77
 science fiction writers and, 198–201
 small attacks, 89
 social solidarity and, 55
 threat exaggeration, politics and,
 329–331
 truths, 53
 Underwear Bomber, 269
 Washington Monument closure,
 310–311
 worst-case scenarios, 89–92
 Texas Virtual Border Patrol, 177
 threats
 exaggeration, politics and, 329–331
 movie-plot threats, 44–45
 Tooby, John, 137
 Tor (anonymous Web browsing), 121
 trails of data, 157–158, 168–171
 transparency, security and, 327–329
The Transparent Society (Brin), 155
 transparent society myth, 155–157
 travel
 airport
 profiling, 275–277
 rarity of terrorism, 270–275
 secure area breaches, 268–269
 security fixes, 264–265
 border crossing, 259–260,
 265–266
 contraband in airports, 262–263
 no-fly list, 261–262
 panic over, 269–275
 Secure Flight program, 265
 TrueCrypt, 266
 trust, 130
 cheating and, 131–134
 regimes of, 152–154
 societal pressures and, 132,
 141–143
 trusted people, internal hackers and,
 62–64

TSA (Transportation Security Administration)
airport profiling, 275–277
photo ID rules, 261–262
transparency, 264–265
Tversky, Amos, 4–5, 194
hindsight bias, 108

U

UAC (User Account Control),
Microsoft, 122–123
unauthentication, 237–238
Underwear Bomber, 269
universal identification, 173–175
US Cyber Command, 94, 113–114
users, bypassing, 122–123
Utility Theory, 5

V

Vanish, 236
VAT (value added tax), 124–125
Virtual Border Patrol (Texas), 177
Volokh Conspiracy, 290
vulnerabilities

ethics of research, 209–210
full disclosure, 218–220
market, 30–32
SSL protocol security flaw,
240–242

W

Walker, Jay, 178
Wason selection task, 134–136
Web, anonymous browsing, 121
whistleblowers, 184–187
whitelisting/blacklisting, 312–313
Windows Personal Firewall, 239
wiretapping, 231–233
Internet, 109–112
witnesses, hindsight bias, 108
World Trade Center Attacks.
See 9/11
worst-case scenarios, 89–92

X–Y–Z

zero tolerance policies, 295–296
Zfone, 50