HAL ABELSON • KEN LEDEEN
HARRY LEWIS • WENDY SELTZER

BLOWN to BITS

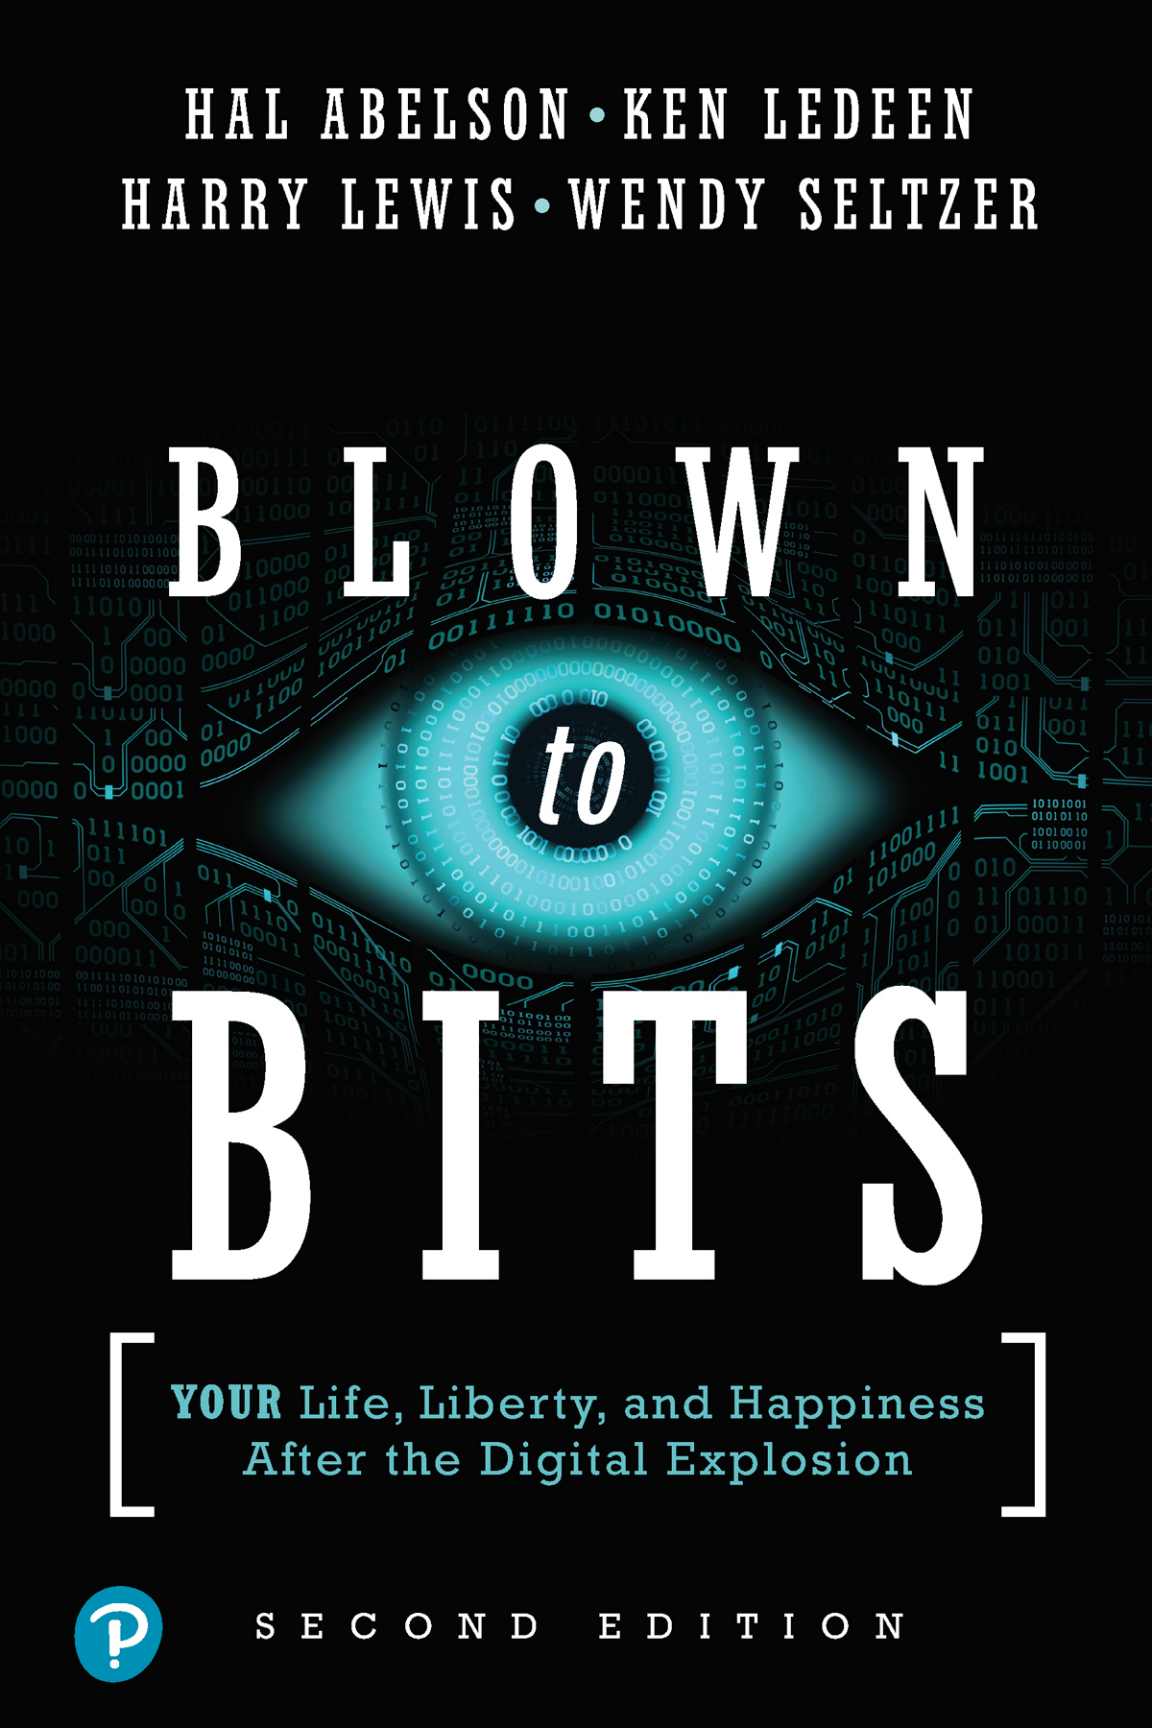

[YOUR Life, Liberty, and Happiness After the Digital Explosion]

SECOND EDITION

Blown To Bits

Second Edition

This page intentionally left blank

Blown To Bits

Your Life, Liberty, and Happiness After the Digital Explosion

Second Edition

Hal Abelson
Ken Ledeen
Harry Lewis
Wendy Seltzer

Visit us on the web: informit.com/business

Library of Congress Control Number: 2020941961

Published by Pearson Education, Inc.

ISBN-13: 978-0-13-485001-6
ISBN-10: 0-13-485001-7

ScoutAutomatedPrintCode

Editor-in-Chief
Mark L. Taub

Acquisitions Editor
Gregory Doench

Managing Editor
Sandra Schroeder

Senior Project Editor
Lori Lyons

Production Manager
Aswini Kumar/codeMantra

Copy Editor
Kitty Wilson

Indexer
Cheryl Ann Lenser

Proofreader
Abigail Manheim

Cover Designer
Chuti Prasertsith

Compositor
codeMantra

To our children, Amanda, Jennifer, Joshua, Elaheh, Annie, and Elizabeth, who will see the world changed yet again in ways we cannot imagine.

To our grandchildren, Connor, Rory, Abigail, Cameron, Juliet, Robert, Alexandra, and Stella, who are living with those changes.

And to our students, who always challenge us to think more deeply.

This page intentionally left blank

Contents at a Glance

Register Your Book

Register your copy of *Blown to Bits* at informit.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN 9780134850016 and click Submit. Once the process is complete, you will find any available bonus content under "Registered Products."

This page intentionally left blank

Contents

This page intentionally left blank

Acknowledgments

While we take full responsibility for any errors in the book, we owe thanks to a great many others for any enlightenment it may provide. Specifically, we are grateful to the following individuals, who commented on parts of the book while it was in draft or provided other valuable assistance: Lynn Abelson, Meg Ausman, Scott Bradner, Art Brodsky, Mike Carroll, Marcus Cohn, Frank Cornelius, Alex Curtis, Natasha Devroye, David Fahrenthold, Robert Faris, Johann-Christoph Freytag, Wendy Gordon, Tom Hemnes, Brian LaMacchia, Marshall Lerner, Anne Lewis, Elizabeth Lewis, Jessica Litman, Lory Lybeck, Fred vonLohmann, Marlyn McGrath, Michael Marcus, Michael Mitzenmacher, Steve Papa, Jonathan Pearce, Bradley Pell, Les Perelman, Thomas Roessler, Pamela Samuelson, Jeff Schiller, Katie Sluder, Gigi Sohn, Debora Spar, René Stein, Alex Tibbetts, Susannah Tobin, Salil Vadhan, David Warsh, Danny Weitzner, and Matt Welsh.

This page intentionally left blank

About the Authors

Hal Abelson is Class of 1922 Professor of Computer Science and Engineering at MIT, and an IEEE Fellow. He has helped drive innovative educational technology initiatives such MIT OpenCourseWare, co-founded Creative Commons and Public Knowledge, and was founding director of the Free Software Foundation.

Ken Ledeen, Chairman/CEO of Nevo Technologies, is a serial entrepreneur who has served on the boards of numerous technology companies.

Harry Lewis, former Dean of Harvard College and of Harvard's School of Engineering and Applied Sciences, is Gordon McKay Research Professor of Computer Science at Harvard and Faculty Associate of the Berkman Klein Center for Internet and Society. He is author of *Excellence Without a Soul: Does Liberal Education Have a Future?* and editor of *Ideas that Created the Future: Classic Papers of Computer Science*.

Wendy Seltzer is Counsel and Strategy Lead at the World Wide Web Consortium (W3C), based at MIT. She founded Lumen Database, the pioneering transparency report for online content removals.

This page intentionally left blank

Preface

More than a decade ago, we decided to write this book because the digital revolution was producing profound changes to nearly every aspect of our lives, and intelligent social choices required understanding of the underlying principles of digital technology and their implications for human institutions. What was lacking was not a book about how computers work, but rather one analyzing that human perspective.

This book has remained surprisingly current for more than a decade. The continuing progress of technology and its broad adoption, however, demanded a second edition. A few topics deserve particular mention.

It should come as no surprise that the impact of technology on privacy—or the lack thereof—has accelerated. Facial recognition technology was nascent when we first wrote this book, and is now pervasive. Cell phone apps track our every movement. Speaker identification software is in common use, both by governments and commercial organizations. (See Chapter 2, "Naked in the Sunlight: Privacy Lost, Privacy Abandoned.")

Artificial intelligence applications are now commonplace. No need for a CD collection (if you even remember what CDs are). Just tell your device what you'd like to hear. Or say "Siri, remind me when I get to Home Depot in Waltham to buy some light bulbs." Automated assistants are built into TV remotes and refrigerators. (See Chapter 9, "The Next Frontier: AI and the Bits World of the Future.")

In 2008, Facebook was popular for connecting with friends, and Twitter had barely begun. Today, these and other social media platforms have a profound impact on society, facilitating social uprisings, influencing elections, and offering soapboxes and megaphones to politicians. (See Chapter 3, "Who Owns Your Privacy? The Commercialization of Personal Data.")

None of us could have foreseen the spotlight that the coronavirus pandemic of 2020 would bring to the implications and impact of the digital revolution. In a matter of weeks, schools from Kindergarten through graduate programs went "virtual." Six year-old kids mastered video conferencing, and online ordering became the primary means of meeting daily needs. Telework became the norm, something that would have been hard to imagine a decade earlier, and impossible a decade before that.

We are optimistic that a deeper understanding of both the "what" and the "how" of the digital world will lead to wise choices as the revolution continues.

Every day, billions of photographs, news stories, songs, X-rays, TV shows, phone calls, and emails are being scattered around the world as sequences of zeros and ones: bits. Phone books, newspapers, CDs, handwritten letters—and privacy—are all relics of the pre-digital era.

We can't escape this explosion of digital information, and few of us want to—the benefits are too seductive. Digital technology has enabled unprecedented innovation, collaboration, entertainment, and democratic participation.

But the same engineering marvels are shattering centuries-old assumptions about privacy, identity, free expression, and personal control as more and more details of our lives are captured as digital data.

Can you control who sees all that personal information about you? Can anything be confidential, when nothing seems to be private? Should the Internet be censored the way radio and TV are? When you search for something, who decides what to show you? How do you know what is "true" when we live in a digital echo chamber with an unlimited array of information—and misinformation—sources? Do you still have free speech in the digital world? Do you have a voice in shaping government or corporate policies about any of this? In the world of Artificial Intelligence, how do we know why machines decide to do things? Does a small set of powerful corporations influence what we know, and how we perceive the world? Have we already lost control?

Blown to Bits guides you through the digital landscape, offering provocative answers to these questions through intriguing real-life stories. Understanding the potential and pitfalls of this transformed world is essential information for everyone.

This book is a wake-up call to the human consequences of the digital explosion.

CHAPTER 1

Digital Explosion

Why Is It Happening, and What Is at Stake?

This book isn't about computers. It's about your life and mine. It's about how the ground underneath us has shifted in fundamental ways. We all know it is happening. We see it all around us, every day. We all need to understand it more.

The digital explosion is changing everything. In this book we talk about both what is happening and how. We explain the technology itself—why it creates so many surprises and why things often don't work the way we expect them to. It is also about things the information explosion is destroying: old assumptions about our privacy, about our identity, and about who is in control of our lives. It's about how we got this way, what we are losing, and what remains that society still has a chance to put right.

The digital explosion is creating both opportunities and risks. Many of both will be gone in a decade, settled one way or another. Governments, corporations, and other authorities are taking advantage of the chaos, and most of us don't even see it happening. Yet we all have a stake in the outcome. Beyond the science, the history, the law, and the politics, this book is a wake-up call. The forces shaping your future are digital, and you need to understand them.

This book is about the stories we hear and read every day. Stories that are about the profound, often unexpected impact that digital technology is having on our lives. Let's begin with the story of Nicolette Vartuli.

Nicolette couldn't figure out why she didn't get the job. A college senior with a 3.5 GPA, she had prepared for her interview with the investment bank and stayed positive throughout. She kept her head up, smiled, and spoke with confidence. But when the company followed up, it was bad news. She would not be moving on in the hiring process.[1]

Nicolette wanted to know what she had done wrong, but no one could explain why she was rejected—because no one actually knew. She had been interviewed by a computer that used AI software from HireVue to assess her suitability. That software rejected her not because she didn't have some particular qualification but because, as it claimed, the software could detect patterns in people who were successful in the job—and what it observed in Nicolette didn't match. It is easy to understand being rejected because you don't have three years of required experience or some particular skill. This is different. And scary—especially because no explanation was offered for what the software was looking for. And it may be that no explanation could be offered, even if HireVue were willing to disclose its proprietary algorithms. (It is not.)

Companies like this new technology. It is cheaper and more efficient than human interviews. In fact, HireVue, just one of many providers, has completed more than 10 million interviews. Many applicants, by contrast, don't like these automated hiring assistants. It's not just that it feels dehumanizing to be judged by a machine. The companies that offer the service counter that by using technology, more people can get interviews now, and the likelihood of inherent bias on the part of interviewers is diminished. They claim the technology is opening up opportunities, not limiting them—but how do we know?

The instinctive antipathy to automated job screening can't really be because people don't want computers making life-critical decisions. Many such decisions are made by computers today; airplanes and radiation therapy machines are now largely automated systems, for example. Computers now beat highly trained radiologists at spotting cancer tumors in breast X-rays.[2] Would anyone prefer less accurate human screeners? But HireVue's judgments are of a different kind. The program made a decision about Nicolette's humanity. It decided that she was not the sort of person the company should hire, and it did so without explaining to her or anyone else what sort of person would be a good hire and how Nicolette fell short.

Many other systems are today making similar judgments in other human domains. Judges consult computers to assess the risk that criminal defendants will fail to show up for their trials—again by comparing the individuals with others who have been arrested in the past and have been given the benefit of avoiding pretrial detention.[3] Real-estate agents use computers to judge which prospective renters are likely to be deadbeats.[4]

Most of these systems are proprietary, and the companies that make them don't have to disclose how they work. And after all, they argue, human interviewers are no gold standard of impartial judgment. They are prone to all sorts of unfortunate biases and prejudices. That is why tryouts for instrumental musicians are now commonly held out of view of the listeners: When the performers could be seen, women were systematically judged more harshly than men.[5] By matching candidates' interview skills to those of existing workers, HireVue claims, it is

eliminating the most fallible part of the system. It's the human recruiters, HireVue says, who are the "ultimate black box." Maybe—except that HireVue says it is matching candidates to the profile of the best of the bank's current employees. How would anyone know if the software is simply replicating, now automatically, all the prejudices that gave the bank the workforce it now has?

What makes this whole story particularly important is not only that Nicolette was judged by a machine to be unsuitable but that no one—not a human resource manager, not even a programmer—told the HireVue software what criteria to use. It determined those all by itself. The software watched videos of existing employees and picked its own criteria.

The tale of Nicolette's rejected job application is what we call "a bits story." That is, it is not just a job search story; it is a story about the collection, storage, analysis, transmission, and use of trillions of trillions of trillions of individual 0s and 1s. By looking carefully at these stories, we can understand not only the technology behind them, but the implications and risks as well.

Bits represented Nicolette's image as it flowed from her own computer to HireVue's, over wires and cables and probably several kinds of radio waves. The bits were reassembled, taken apart, and analyzed by HireVue's programs. They were somehow compared to trillions of trillions of trillions of bits representing videos of other people, and then a single bit, a single yes or no, came out: continue to the next stage of the hiring process or reject immediately. That bit was a 0 for Nicolette, and that is all she heard back from the company. But HireVue kept all the bits of Nicolette's failed interview; she had to sign over her rights to them in order to get the interview in the first place.

"Algorithmic transparency" is the principle that we should know how computers are making decisions about us. In the words of EPIC (the Electronic Privacy Information Center), "The public has a right to know the data processes that impact their lives so they can correct errors and contest decisions made by algorithms."[6]

New technologies interact in odd ways with evolving standards of privacy, communications practices, and criminal law. Nicolette's story, while important to her, is just one of thousands of bits stories that could be told about any one of us. Every day we encounter unexpected consequences of data flows that could not have happened a few years ago.

When you have finished reading this book, you should see the world in a different way. You should hear a story from a friend or on a newscast and say to yourself, "That's really a bits story," even if no one mentions anything digital. The movements of physical objects and the actions of flesh-and-blood human beings are only the surface. To understand what is really going on, you have to see the virtual world, the eerie flow of bits steering the events of life.

This book is your guide to this new world.

The Explosion of Bits, and Everything Else

The world changed very suddenly. Almost everything is stored in a computer somewhere. Court records, grocery purchases, precious family photos and priceless Hollywood movies, pointless television shows....Computers contain a lot of stuff that isn't useful today but somebody thinks might someday come in handy. It is all being reduced to 0s and 1s—"bits." The bits are stashed on disks of home computers, in the data centers of big corporations and government agencies. Many of the disks aren't even round, spinning things—they are a different kind of storage media, called "disks" for historical reasons. Most of the disks these days are "in the cloud"—just a fancy name for disks owned by a big company such as Amazon and rented out to whoever needs space to store stuff. The disks can hold so many bits that there is no need to pick and choose what gets remembered.

"Bit" is shorthand for "binary digit." The binary number system uses just two digits, 0 and 1, instead of the ten digits 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 used in the decimal number system. The first clear statement of the principles of binary notation was given by Gottfried Wilhelm Leibniz in 1679.

So much digital information, misinformation, data, and garbage is being squirreled away that most of it will be seen only by computers, never by human eyes. And computers are getting better and better at extracting meaning from all those bits—finding patterns that sometimes solve crimes, diagnose diseases, and make useful suggestions—and sometimes reveal things about us we did not expect others to know.

The tale of Edward Snowden, who leaked thousands of highly secret government documents in 2013, is a bits story. He brought the documents out of the United States on his laptop; only a few years earlier, he would have needed to carry hundreds of pounds of paper. And everything he disclosed was about government electronic surveillance, raising fundamental questions about trade-offs between privacy and security.

The grounding of the 737 Max in 2019 was not just an airplane story. It was also a bits story. The engines of earlier model 737s had been moved, changing the airplane's weight distribution; software written to process

sensor data and automatically control the airplane's movements did not work as intended.[7]

But it is not just events of global significance that are bits stories; it's the day-to-day stories of ordinary life. The creepy experience of recreational runner Rosie Spinks is a bits story. Spinks used an app on her phone to keep track of her routes and times, and she thought her whereabouts were being kept secret because she had the app's so-called "Enhanced Privacy" setting on. Only when strangers started "liking" her workouts while she was traveling abroad did she realize that "Enhanced Privacy" actually meant "tell random men about my runs if I'm on the leader board." The fitness app was also a social network app, and Rosie's data was being commercialized.[8]

Once something is on a computer, it can replicate and move around the world in a heartbeat. Making a million perfect copies takes but an instant—copies of things we want everyone in the world to see and also copies of things that weren't meant to be copied at all.

The digital explosion is changing the world as much as printing once did—and some of the changes are catching us unaware, blowing to bits our assumptions about the way the world works.

The digital explosion can seem benign, amusing, or even utopian. Instead of sending prints through the mail to Grandma, we share pictures of our children on Instagram. Then not only can Grandma see them but so can Grandma's friends and anyone else. We enjoy the benefits, but what are the risks? The photos are cute and harmless. But suppose a tourist takes a vacation snapshot, and you just happen to appear in the background, at a restaurant where no one knew you were dining. If the tourist uploads his photo and makes it public, the whole world could know where you were and when you were there. Face recognition, which only a few years ago was beyond the capabilities of computers, is now good enough that the tourist photo could even get tagged with your name. It might not happen because a policy or a law prohibits it, but technological limitations won't. Identifying faces in crowds automatically is now a solved problem, and software to do this is being used in China and other authoritarian regimes to discourage public protests and generally to keep track of citizens' whereabouts. And this technology is being used in the United States, too: With the aid of billions of labeled photos gathered from Facebook and other social media, a small company named Clearview AI suddenly became a tool of many law enforcement agencies and even security-conscious private companies.[9] It wasn't even very hard to do—for an entrepreneurial company willing to stretch the limits of the appropriate use

of the massive photographic databases that Facebook and other companies had gathered.

And before we leave the topic of family photos—remember when they were all printed on paper and lasted for decades? Not so anymore. The wondrous benefits of digital images also make them inaccessible. You can't put digital images in a box under your bed for your grandchildren to find. All those family memories may well be lost in the future. There is a good and bad side to pretty much everything in the digital world.

Data leaks. Credit card records are supposed to stay locked up in a data warehouse, but they escape into the hands of identity thieves. And we give information away just because we get something back for doing so. A company will give you free phone calls to anywhere in the world—if you don't mind watching ads for the products its computers hear you talking about. Google will suggest restaurants you might like—if you will leave location tracking on so Google knows what restaurants you already frequent. If you have a meal, Google will ask you if you enjoyed it—no software is yet able to figure *that* out by itself—and into the data maw goes your answer to help Google make recommendations to you and others (and make a little money along the way).

And these are merely some of the things that are happening today. The explosion, and the social disruption it will create, have barely begun.

We already live in a world in which there is enough memory *just in cell phones* to store every word of every book in the Library of Congress billions of times over. Every day, enough video is uploaded to YouTube to record every moment of an entire human lifetime. The explosive growth is still happening. Every year we can store more information, move it more quickly, and do far more ingenious things with it than we could the year before. Now that refrigerators and vacuum cleaners create data, the increasing rate at which data is created is almost unimaginable. Most of the data that has ever existed was created in the past year—and that will be true again next year and the year after that.

So much disk storage is being produced every year that it could be used to record a page of information, every few seconds, about you *and every other human being on earth*. A remark made long ago can come back to haunt a political candidate, and a letter jotted quickly can be a key discovery for a biographer. Imagine what it would mean to record every word every human being speaks or writes in a lifetime. The technological barrier to that has already been removed: There is enough storage to remember it all. YouTube says that 500 hours of video are uploaded every minute.[10] Should any social barrier stand in the way?

Sometimes things seem to work both better and worse than they used to. A "public record" is now *very* public. Before you get hired in Nashville,

Tennessee, your employer can figure out if you were caught ten years ago taking an illegal left turn in Lubbock, Texas. The old notion of a "sealed court record" is mostly a fantasy in a world where tidbits of information are duplicated, cataloged, and moved around endlessly. In Europe a new "right to be forgotten" has been added to the list of human rights, intended to protect people from having to carry every youthful indiscretion with them forever; but in the United States the right to free speech remains dominant, and the collision between these conflicting rights is inevitable. In the bits world, the Atlantic ocean can be crossed in microseconds.

With hundreds of TV and radio stations and millions of websites, Americans love the variety of news sources, but they have adjusted uncomfortably to the displacement of more authoritative sources. In China, the situation is reversed: The technology creates greater government control of the information its citizens receive, as well as better tools for monitoring their behavior.

The Koans of Bits

Bits behave strangely. They travel almost instantaneously, and they take almost no space to store. We have to use physical metaphors to make them understandable. We liken them to dynamite exploding or water flowing. We even use social metaphors for bits. We talk about two computers agreeing on some bits and about people using burglary tools to steal bits. Getting the right metaphor is important, but so is knowing the limitations of our metaphors. An imperfect metaphor can mislead as much as an apt metaphor can illuminate.

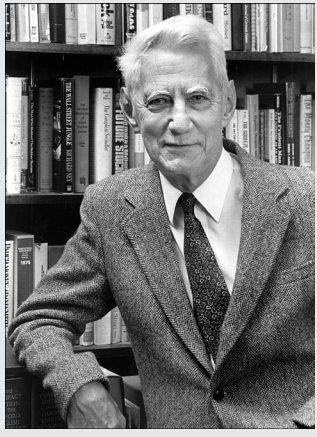

CLAUDE SHANNON

Claude Shannon (1916–2001) is the undisputed founding figure of information and communication theory. While working at Bell Telephone Laboratories after the Second World War, he wrote the seminal paper "A Mathematical Theory of Communication," which foreshadowed much of the subsequent development of digital technologies. Published in 1948, this paper gave birth to the now-universal realization that the bit is the natural unit of information and to the use of the term.

Reused with permission of Nokia Corporation and AT&T Archives. http:www.bell-labs.com/news/2001/february/26/shannon2_lg.jpeg.

We offer seven truths about bits. We call them "koans" because they are paradoxes, like the Zen verbal puzzles that provoke meditation and enlightenment. These koans are oversimplifications and overgeneralizations. They describe a world that is developing but hasn't yet fully emerged. But even today they are truer than we often realize. These themes will echo through our tales of the digital explosion.

Koan 1: It's All Just Bits

Your computer and your smartphone (really just another computer) successfully create the illusion that they contain photographs, letters, songs, and movies. All they really contain is bits—lots of them—patterned in ways you can't see. Your computer was designed to store just bits; all the files and folders and different kinds of data are illusions created by computer programmers. When you send a message containing a photograph, the computers that handle your message as it flows through the Internet have no idea that what they are handling is part text and part graphic. Telephone calls are also just bits, and that has helped create competition: Traditional phone companies, cell phone companies, cable TV companies, and voice over IP (VoIP) service providers can shuffle bits around to each other to complete calls. The Internet was designed to handle just bits, not emails or attachments, which are inventions of software engineers. We couldn't live without those more intuitive concepts, but they are artifices. Underneath, it's all just bits.

This koan is more consequential than you might think. Consider the story of Naral Pro-Choice America and Verizon Wireless. Naral wanted to form a text messaging group to send alerts to its members. Verizon decided not to allow it, citing the "controversial or unsavory" things the messages might contain.[11] Text message alert groups for political candidates it would allow, but not for political causes it deemed controversial. Had Naral simply wanted telephone service or an 800 number, Verizon would have had no choice. Telephone companies were long ago declared "common carriers." Like railroads, phone companies are legally prohibited from picking and choosing customers from among those who want their services. In the bits world, there is no difference between a text message and a wireless phone call. It's all just bits, traveling through the air by radio waves. But the law hasn't caught up to the technology. It doesn't treat all bits the same, and the common carriage rules for voice bits don't apply to text message bits.

Exclusive and Rivalrous

Economists would say that bits, unless controlled somehow, tend to be non-exclusive (once a few people have them, it is hard to keep them from others) and non-rivalrous (when someone gets them from me, I don't have any less). In a letter he wrote about the nature of ideas, Thomas Jefferson eloquently stated both properties: "If nature has made any one thing less susceptible than all others of exclusive property, it is the action of the thinking power called an idea, which an individual may exclusively possess as long as he keeps it to himself; but the moment it is divulged, it forces itself into the possession of every one, and the receiver cannot dispossess himself of it. Its peculiar character, too, is that no one possesses the less, because every other possesses the whole of it."[12]

Verizon backed down in the case of Naral, but not on the principle. A phone company can do whatever it thinks will maximize its profits in deciding whose messages to distribute. Yet no sensible engineering distinction can be drawn between text messages, phone calls, and any other bits traveling through the digital airwaves.

Koan 2: Perfection Is Normal

To err is human. When books were laboriously transcribed by hand in ancient scriptoria and medieval monasteries, errors crept in with every copy. Computers and networks work differently. Every copy is perfect. If you email a photograph to a friend, the friend won't receive a fuzzier version than the original. The copy will be identical, down to levels of detail too small for the eye to see.

Computers do fail, of course. Networks break down, too. If the power goes out with no battery backup, nothing works at all. So the statement that copies are normally perfect is only relatively true. Digital copies are perfect only to the extent that they can be communicated at all. And yes, it is possible in theory that a single bit of a big message will arrive incorrectly—but it's also possible that a volcano will erupt under you, and you won't get the message at all. The odds of an erroneous bit are lower than the odds of a physical catastrophe, and that is good enough for all practical purposes.

Networks don't just pass bits from one place to another. They check to see if the bits seem to have been damaged in transit and correct them or retransmit them if they seem incorrect. As a result of these error detection and correction mechanisms, the odds of an actual error—a character being wrong in an email, for example—are so low that we would be wiser to worry instead about a meteor hitting our computer, improbable though precision meteor strikes may be.

The phenomenon of perfect copies has drastically changed the law, a story told in Chapter 6, "Balance Toppled." In the days when music was distributed on audio tape, teenagers were not prosecuted for making copies of songs because the copies weren't as good as the originals, and copies of copies would be even worse. The reason that people today more often subscribe to music services than own their own copies of recordings is that copies are perfect—not just as good as the original but identical to the original so that even the notion of "original" is meaningless. The consequences of digital disruption of "intellectual property" are not over yet. Bits are an odd kind of property. Once I release them, everybody has them. And if I give you my bits, I don't have any fewer.

Koan 3: There Is Want in the Midst of Plenty

Vast as worldwide data storage is today, two years from now, it will be twice as large. Yet the information explosion means, paradoxically, the loss of information that is not online. One of us saw a new doctor at a clinic he had been using for decades. She showed him dense charts of his blood chemistry, data transferred from his home medical device to the clinic's computer—more data than any specialist could have had at her disposal five years ago. The doctor then asked whether he had ever had a stress test and what the test had shown. Those records should be all there, the patient explained, in the medical file. But the information was in the *paper* file, to which the doctor did not have access. It wasn't in the *computer's* memory, and the patient's memory was being used as a poor substitute. The old data might as well not have existed at all since it wasn't digital.

Even information that exists in digital form is useless if there are no devices to read it. The rapid progress of storage engineering has meant that data stored on obsolete devices effectively ceases to exist. A twentieth-century digital update of the eleventh-century British *Domesday Book* shown in Figure 1.1 was useless by the time it was only one-sixtieth the age of the original.

FIGURE 1.1 The Domesday Book of 1086. A 900th anniversary digital update was unreadable 15 years later.[13]

Or consider search, among the subjects of Chapter 4, "Gatekeepers." At first, search engines such as Google were interesting conveniences that a few people used for special purposes. With the growth of the World Wide Web and the explosion of online information, search engines became the first place many people look for information—even before they look in books or ask friends. Appearing prominently in search results has become a matter of life or death for businesses. We may move on to purchase from a competitor if we can't find the site we wanted in the first page or two of results. We may assume something didn't happen if we can't find it quickly in an online news source. If it can't be found—quickly—it's just as though it doesn't exist at all.

And some information isn't true. All of the mechanisms that enable the communication and storage of facts also work for falsehoods. Ugliness and cruelty are as easily captured in bits as beauty and kindness. The market economics of information change when everyone can be a publisher and no one needs an editor. Floods of misinformation, disinformation, and garbage can overwhelm truth and beauty. Authoritarian societies may be able to manage the flow of bits more efficiently than free societies, which risk being undercut by their own principles of information freedom.

Koan 4: Processing Is Power

Moore's Law

Gordon Moore, founder of Intel Corporation, observed that the density of integrated circuits seemed to double every couple of years. This observation is referred to as "Moore's Law." Of course, it is not a natural law, like the law of gravity. Instead, it is an empirical observation of the progress of engineering and a challenge to engineers to continue their innovation. In 1965, Moore predicted that this exponential growth would continue for quite some time.[14] That it has continued for more than 40 years is one of the great marvels of engineering. No other effort in history has sustained a growth rate anywhere close to this.

The speed of a computer is usually measured by the number of basic operations, such as additions, that can be performed in one second. The fastest computers available in the early 1940s could perform about five operations per second. The fastest today can perform about a trillion. Buyers of personal computers know that a machine that seems fast today will seem slow in a year or two.

For at least three decades, the increase in processor speeds was exponential. Computers became twice as fast every couple of years. These increases were one consequence of Moore's Law (see sidebar).

Since 2001, processor speed has not followed Moore's Law; in fact, processors have hardly grown faster at all. But that doesn't mean that computers

won't continue to get faster. New chip designs include multiple processors on the same chip so the work can be split up and performed in parallel. Such design innovations promise to achieve the same effect as continued increases in raw processor speed. And the same technology improvements that make computers faster also make them less expensive.

The rapid increase in processing power means that inventions move out of labs and into consumer goods very quickly. Robot vacuum cleaners and self-parking vehicles were possible in theory a decade ago, but now they have become consumer items. Tasks that today seem to require uniquely human skills are no longer just the subject of research projects in corporate or academic laboratories; they are incorporated in consumer products. Face recognition and voice recognition are here and now; telephones know who is calling, and surveillance cameras don't need humans to watch them. The power comes not just from the bits but from being able to do things with the bits.

Koan 5: More of the Same Can Be a Whole New Thing

The explosive growth is exponential growth—doubling at a steady rate. Imagine earning 100% annual interest on your savings account: In 10 years, your money would have increased more than a thousandfold, and in 20 years, more than a millionfold. A more reasonable interest rate of 5% will hit the same growth points, but it will do so 14 times more slowly. Epidemics initially spread exponentially, as each infected individual infects several others.

When something grows exponentially, for a long time it may seem not to be changing at all. If we don't watch it steadily, it will seem as though something discontinuous and radical occurred while we weren't looking.

That is why epidemics at first go unnoticed, no matter how catastrophic they may be when full-blown. Imagine 1 sick person infecting 2 healthy people, and the next day each of those 2 infects 2 others, and the next day after that each of those 4 infects 2 others, and so on. The number of newly infected each day grows from 2 to 4 to 8. In a week, 128 people come down with the disease in a single day, and twice that number are now sick, but in a population of 10 million, no one notices. Even after two weeks, barely 3 people in a 1,000 are sick. But after another week, 40% of the population is sick, and society collapses. The 2019–2020 coronavirus pandemic followed pretty much this pattern in parts of the world where societies did not react quickly. At the start of the epidemic in Wuhan, the number of cases doubled about every three days.[15]

Exponential growth is actually smooth and steady; it just takes very little time to pass from unnoticeable change to highly visible. Exponential growth of anything can suddenly make the world look utterly different than it had been. When that threshold is passed, changes that are "just" quantitative can look qualitative.

Another way of looking at the apparent abruptness of exponential growth—its explosive force—is to think about how little lead time we have to respond to it. Our hypothetical epidemic took three weeks to overwhelm the population. At what point was it only half as devastating? The answer is *not* "a week and a half." The answer is *on the next-to-last day*. Suppose it took a week to develop and administer a vaccine. Then noticing the epidemic after a week and a half would have left ample time to prevent the disaster. But that would have required understanding that there *was* an epidemic when only 2,000 people out of 10 million were infected.

The information story is full of examples of unperceived changes followed by dislocating explosions. Those with the foresight to notice the explosion just a little earlier than everyone else can reap huge benefits. Those who move a little too slowly may be overwhelmed by the time they try to respond. Take the case of digital photography.

In 1983, Christmas shoppers could buy digital cameras to hook up to their IBM PC and Apple II home computers. The potential was there for anyone to see; it was not hidden in secret corporate laboratories. But digital photography did not take off. Economically and practically, it couldn't. Cameras were too bulky to put in your pocket, and digital memories were too small to hold many images. Even 14 years later, film photography was still a robust industry. In early 1997, Kodak stock hit a record price, with a 22% increase in quarterly profit, "fueled by healthy film and paper sales...[and] its motion picture film business," according to a news report.[16] The company raised its dividend for the first time in eight years. But by 2007, digital memories had become huge, digital processors had become fast and compact, and both were cheap. As a result, cameras had become little computers. The company that was once synonymous with photography was a shadow of its former self. Kodak announced that its employee force would be cut to 30,000, barely a fifth the size it was during the good times of the late 1980s.[17] By 2018, that number was down to about 5,400. Bits took away 90% of the jobs. Moore's Law moved faster than Kodak did.

In the rapidly changing world of bits, it pays to notice even small changes—and to do something about them.

Koan 6: Nothing Goes Away

25,000,000,000,000,000,000.

That is the number of bits that were created and stored every day of 2019, according to one industry estimate. The capacity of disks has followed its own version of Moore's Law, doubling every two or three years. Far more data is created by every manner of device but not stored.

In financial industries, federal laws now *require* massive data retention to assist in audits and investigations of corruption. In many other businesses, economic competitiveness drives companies to save everything they collect and to seek out new data to retain. Tens of millions of transactions occur in Walmart stores every day, and every one of them is saved: date, time, item, store, price, who made the purchase, and how—credit, debit, cash, or gift card. Such data is so valuable to planning the supply chain that stores will pay money to get more of it from their customers. That is really what loyalty programs at supermarkets and other stores provide: Shoppers are supposed to think that the store is granting them a discount in appreciation for their steady business, but actually the store is paying them for information about their buying patterns. We might better think of a privacy tax: We pay the regular price *unless* we want to keep information about our food, alcohol, and pharmaceutical purchases from the market; to keep our habits to ourselves, we pay extra.

The massive databases challenge our expectations about what will happen to the data about us. Take something as simple as a stay in a hotel. When you check in, you are given a keycard, not a mechanical key. In fact, some hotels have gone one step further, having you use your own cell phone as the room key. Because the keycards can be deactivated instantly, there is no longer any great risk associated with losing your key, as long as you report it missing quickly. On the other hand, the hotel now has a record, accurate to the second, of every time you entered your room, used the gym or the business center, or went in the back door after-hours. The same database could identify every cocktail and steak you charged to the room, which other rooms you phoned and when, and the brands of tampons and laxatives you charged at the hotel's gift shop. This data might be merged with billions of bits' worth of other data, analyzed, and transferred to the parent company, which owns restaurants and fitness centers as well as hotels. It might also be lost, or stolen, or subpoenaed in a court case.

The ease of storing information has meant asking for more of it. Birth certificates used to include just the information about the child's and parents' names, birthplaces, and birthdates, plus the parents' occupations. Now the electronic birth record includes how much the mother drank and smoked during her pregnancy, whether she had genital herpes or a variety of other medical conditions, and both parents' Social Security numbers. Opportunities for research are plentiful, and so are opportunities for mischief and catastrophic accidental data loss.

The data will all be kept forever, unless there are policies to get rid of it.

And the data will all be kept forever unless there are policies to get rid of it. For the time being at least, the data sticks around. And because databases are intentionally

duplicated—backed up for security or shared while pursuing useful analyses—it is far from certain that data can ever be permanently expunged, even if we wish that to happen. The Internet consists of millions of interconnected computers; once data gets out, there is no getting it back. Victims of identity theft experience daily the distress of having to remove misinformation from the record. It seems never to go away.

Koan 7: Bits Move Faster Than Thought

The Internet existed before there were personal computers. It predates the fiber-optic communication cables that now hold it together. When it started around 1970, ARPANET, as it was called, was designed to connect a handful of university and military computers. No one imagined a network connecting tens of millions of computers and shipping information around the world in the blink of an eye. (In fact, no one imagined that so many computers would even exist.) Even the engineer who was charged with designing the gateways that would connect the computers together remembers his reaction to the idea of a computer network: "Looks like a straightforward engineering job; we could certainly do it, but I can't imagine why anyone would want such a thing."[18] Along with processing power and storage capacity, networking has experienced its own exponential growth in the number of computers interconnected and the rate at which data can be shipped over long distances, from space to earth and from service providers into private homes.

The Internet has caused drastic shifts in business practice. Customer service calls are outsourced to India today not just because labor costs are low there. Labor costs have *always* been low in India, but international telephone calls used to be expensive. Calls about airline reservations and lingerie returns are answered in India today because it now takes almost no time and costs almost no money to send to India the bits representing your voice. The same principle holds for professional services. When you are X-rayed at your local hospital in Iowa, the radiologist who reads the X-ray may be half a world away. The digital X-ray moves around the world and back faster than a physical X-ray could be moved between floors of a hospital. When you place an order at a drive-through station at a fast food restaurant, the person taking the order may be in another state. She keys the order so it appears on a computer screen in the kitchen, a few feet from your car, and you are none the wiser. Such developments are causing massive changes to the global economy, as industries figure out how to keep their workers in one place and ship their business as bits.

In the bits world, in which messages flow instantaneously, it sometimes seems that distance doesn't matter at all. The consequences can be startling. One of us, while dean of an American college, witnessed the shock of a father

receiving condolences on his daughter's death. The story was sad but familiar, except that this version had a startling twist. Father and daughter were both in Massachusetts, but the condolences arrived from halfway around the world before the father had learned that his daughter had died. News, even the most intimate news, travels fast in the bits world once it gets out.

When everyone is connected all the time, people can organize themselves as never before. Those afflicted with rare diseases or inspired by idiosyncratic interests can stroke a few keys and share their experiences, though they are separated by oceans and will never meet in person. And those united by a common cause can organize to air their grievances, as the digital-savvy youth of Hong Kong did with great skill during the pro-democracy protests of 2014. But in the hands of the authorities, the bits the protesters exchanged became evidence against them. By the time of the 2019 Hong Kong protests, organizers had abandoned Facebook and were resorting to less convenient encrypted texting apps—and were wearing face masks to confuse the government's face surveillance systems.[19]

And if surveillance fails, governments can simply shut down the Internet. That happened in the Muslim-majority Indian state of Kashmir in 2019; there was no Internet for 7 months, in the interest of "public safety."[20] Similar shutdowns happened in 2019 in Iran, Congo, Bangladesh, and more than a dozen other countries.[21] And in the United States, Section 706 of the Communications Act of 1934 authorizes the president to shut down "a facility for wire communication" in case of "a state or threat of war"—a very sweeping authorization, thus far never invoked to gain control of the Internet.

The instantaneous communication of massive amounts of information has created the misimpression that there is a place called "cyberspace," a land without frontiers where all the world's people can be interconnected as though they were residents of the same small town. That concept has been decisively refuted by the actions of the world's courts. National and state borders still count—and count a lot. If a book is bought online in England, the publisher and author are subject to British libel laws rather than those of the homeland of the author or publisher. Under British law, defendants have to prove their innocence; in the United States, plaintiffs have to prove the guilt of the defendants. An ugly downside to the explosion of digital information and its movement around the world is that information may become less available even where it would be legally protected. (We return to this subject in Chapter 7, "You Can't Say That on the Internet.") "Right to be forgotten" laws may require information to disappear—not just in the country where an individual has asked for some past misdeed to be struck from the electronic record but everywhere. Such a law might seem to be unenforceable, but the companies making the information available—Google, for example—operate internationally, and if they violate the law anywhere, they risk having

employees harassed or arrested any time they are within a jurisdiction where the law has been violated or ignored. Similarly, the publishing world has been blown to bits. It used to be possible to publish a bowdlerized edition of a book or an edited edition of a newspaper in countries with strict speech codes, but now the bits can readily flow from less censorious regions to more. It may prove simpler to publish only a single version of a work for sale everywhere, an edition omitting information that might somewhere excite a lawsuit.

Good and Ill, Promise and Peril

The digital explosion has thrown a lot of things up for grabs, and we all have a stake in who does the grabbing. The way the technology is offered to us, the way we use it, and the consequences of the vast dissemination of digital information are matters not in the hands of technology experts alone. Governments and corporations and universities and other social institutions have a say. And ordinary citizens, to whom these institutions are accountable, can influence their decisions. Important choices are made every year, in government offices and legislatures, in town meetings and police stations, in the corporate offices of banks and insurance companies, in the purchasing departments of chain stores and pharmacies. We can all help raise the level of discourse and understanding. We can all help ensure that technical decisions are made in a context of ethical standards.

We offer two basic morals. First, information technology is inherently neither good nor bad; it can be used for good or ill, to free us or to shackle us. Second, new technologies bring social change, and change comes with both risks and opportunities. All of us, and all of our public agencies and private institutions, have a say in whether technology will be used for good or ill and whether we will fall prey to its risks or prosper from the opportunities it creates.

Technology Is Neither Good nor Bad

Any technology can be used for good or ill; digital technologies, in particular, can be simultaneously good and bad. Nuclear reactions create electric power and weapons of mass destruction. These two uses share a common core but are otherwise quite distinct. Not so in the world after the digital explosion.

The same encryption technology that makes it possible for you to email your friends with confidence that no eavesdropper will be able to decipher your message also makes it possible for terrorists to plan their attacks undiscovered. The same Internet technology that facilitates the widespread distribution of educational works to impoverished students in remote locations

also enables massive copyright infringement. The photomanipulation tools that enhance your snapshots are used by child pornographers to escape prosecution.

The same technologies can be used to monitor individuals, to track their behaviors, and to control what information they receive. Search engines need not return unbiased results. Many users of web browsers do not realize that the sites they visit archive their actions. There is probably a record of exactly what you have been accessing and when, as you browse a clothing or book store catalog, a site selling pharmaceuticals, or a service offering advice on contraception or drug overdose. There are vast opportunities to use this information for invasive but relatively benign purposes, such as marketing, and also for more questionable purposes, such as blacklisting and blackmail.

The key to managing the ethical and moral consequences of technology while nourishing economic growth is to *regulate the use* of technology without *banning or restricting its creation*.

Few regulations mandate disclosure that the information is being collected or restrict the use to which the data can be put. The USA PATRIOT Act and other federal laws give government agencies sweeping authority to sift through mostly innocent data looking for signs of "suspicious activity" by potential terrorists—and to notice lesser transgressions in the process. Although the World Wide Web reaches into millions of households, the rules and regulations governing it are not much better than those of a lawless frontier town of the Old West.

New Technologies Bring Both Risks and Opportunities

The same large storage media that enable anyone to analyze millions of baseball statistics also allow anyone with access to confidential information to jeopardize its security. Access to aerial maps via the Internet makes it possible for criminals to plan burglaries of upscale houses, but technologically sophisticated police know that records of such queries can also be used to solve crimes.

Social networking tools such as Facebook and Twitter have made their founders quite wealthy and have given birth to many thousands of new friendships, marriages, and other ventures. But interconnectivity has unexpected side effects, too. A woman in England discovered that her fiancé was married when Facebook suggested his wife as someone she might want as a friend.[22] And in 2019 a Massachusetts college student committed suicide by jumping from the fourth floor of a parking garage, having received some 47,000 text messages, many allegedly abusive, from his girlfriend in the previous two months. She was charged with involuntary manslaughter—the same crime she might have been charged with had she instead struck him with her car while driving and texting.[23] In a nation deeply committed to free expression as a legal right, which Internet evils should be crimes, and which are just wrong?

Vast data networks have made it possible to move work to where the people are, not people to the work. The results are enormous business opportunities for entrepreneurs who take advantage of these technologies and new enterprises around the globe, and also the other side of the coin: jobs lost to outsourcing.

The difference every one of us can make, to our workplace or to another institution, can be to ask a question at the right time about the risks of some new technological innovation—or to point out the possibility of doing something in the near future that a few years ago would have been utterly impossible.

We begin our tour of the digital landscape with a look at our privacy, a social structure that the explosion has left in shambles. While we enjoy the benefits of ubiquitous information, we also sense the loss of the shelter that privacy once gave us. And we don't know what we want to build in its place. The good and ill of technology, and its promise and peril, are all thrown together when information about us is spread everywhere. In the post-privacy world, we stand exposed to the glare of noonday sunlight—and sometimes it feels strangely pleasant.

Endnotes

1 Drew Harwell, "A Face-Scanning Algorithm Increasingly Decides Whether You Deserve the Job," *Washington Post*, November 6, 2019, https://www.washingtonpost.com/technology/2019/10/22/ai-hiring-face-scanning-algorithm-increasingly-decides-whether-you-deserve-job/.

2 Scott Mayer McKinney et al., "International Evaluation of an AI System for Breast Cancer Screening," *Nature* 577, no. 7788 (January 2020): 89–94, https://doi.org/10.1038/s41586-019-1799-6.

3 Julia Angwin et al., "Machine Bias," ProPublica, May 23, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

4 Elizabeth Fernandez, "Will Machine Learning Algorithms Erase the Progress of the Fair Housing Act?" *Forbes*, November 17, 2019, https://www.forbes.com/sites/fernandezelizabeth/2019/11/17/will-machine-learning-algorithms-erase-the-progress-of-the-fair-housing-act/.

5 Claudia Goldin and Cecilia Rouse, "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians," *The American Economic Review* 90, no. 4 (September 2000), https://pubs.aeaweb.org/doi/pdfplus/10.1257/aer.90.4.715.

6 "Algorithmic Transparency: End Secret Profiling," Electronic Privacy Information Center, March 1, 2020, https://epic.org/algorithmic-transparency/.

7 Benjamin Zhang, "The Boeing 737 Max Is Likely to Be the Last Version of the Best-Selling Airliner of All Time," *Business Insider*, March 19, 2019, https://www.businessinsider.com/boeing-737-max-design-pushed-to-limit-2019-3.

8 Rosie Spinks, "Confused About How to Use Strava Safely? You Are Not Alone," *Quartz*, January 29, 2018, https://qz.com/1191431/strava-privacy-concerns-here-is-how-to-safely-use-the-app/.

9 Kashmir Hill, "The Secretive Company That Might End Privacy as We Know It," *The New York Times*, January 18, 2020, https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html.

10 J. Clement, "Hours of Video Uploaded to YouTube Every Minute, 2007–2019," *Statista*, August 9, 2019, https://www.statista.com/statistics/259477/hours-of-video-uploaded-to-youtube-every-minute/.

11 Adam Liptak, "Verizon Blocks Messages of Abortion Rights Group," *The New York Times*, September 27, 2007, https://www.nytimes.com/2007/09/27/us/27verizon.html.

12 "Article 1, Section 8, Clause 8: Thomas Jefferson to Isaac McPherson," in Andrew A. Lipscomb and Albert Ellery Bergh, eds., *The Writings of Thomas Jefferson* (Thomas Jefferson Memorial Association, 1905), http://press-pubs.uchicago.edu/founders/documents/a1_8_8s12.html.

13 Robin McKie and Vanessa Thorpe, "Digital Domesday Book lasts 15 years not 1000," *Guardian Unlimited*, March 3, 2002.

14 G. E. Moore, "Cramming More Components onto Integrated Circuits," *Proceedings of the IEEE* 86, no. 1 (January 1998): 82–85, https://doi.org/10.1109/JPROC.1998.658762.

15 Steven Sanche et al., "High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2," *Emerging Infectious Diseases* 26, no. 7 (July 2020): 1470–1477, https://dx.doi.org/10.3201/eid2607.200282.

16 "Kodak, GE, Digital Report Strong Quarterly Results," *Atlanta Constitution*, January 17, 1997.

17 Claudia H. Deutsch, "Shrinking Pains at Kodak," *The New York Times*, February 9, 2007.

18 Harry R. Lewis, "A Science Is Born," *Harvard Magazine*, September-October 2020: 42, https://harvardmagazine.com/2020/09/features-a-science-is-born/

19 Lily Kuo, "Hong Kong's Digital Battle: Tech That Helped Protesters Now Used Against Them," *The Guardian*, June 14, 2019, https://www.theguardian.com/world/2019/jun/14/hong-kongs-digital-battle-technology-that-helped-protesters-now-used-against-them.

20 Billy Perrigo, "India's Supreme Court Orders Review of Internet Shutdown in Kashmir. But for Now, It Continues," *Time*, January 10, 2020, https://time.com/5762751/internet-kashmir-supreme-court/.

21 Samuel Woodhams and Simon Migliano, "The Global Cost of Internet Shutdowns in 2019," Top10VPN, January 7, 2020, https://www.top10vpn.com/cost-of-internet-shutdowns/.

22 "Mum-of-Three Uncovered her Cheating Fiance's Double Life After His Wife Came Up as a Friend Suggestion on Facebook," *The Sun*, September 7, 2017, https://www.thesun.co.uk/fabulous/4411305/mum-of-three-uncovered-her-cheating-fiances-double-life-after-his-wife-came-up-as-a-friend-suggestion-on-facebook/.

23 Julia Jones, "Girlfriend Charged in Boston College Student's Death After Telling Him 'Hundreds of Times' to Kill Himself, prosecutors say," CNN, October 29, 2019, https://www.cnn.com/2019/10/28/us/boston-college-student-suicide-charges/index.html.

CHAPTER 2

Naked in the Sunlight

Privacy Lost, Privacy Abandoned

1984 Is Here, and We Like It

Fans attending Taylor Swift's packed Rose Bowl concert in the spring of 2018 saw her take the stage in a cloud of fog to sing hits from *Reputation*. As they entered or mingled between sets, some of those fans visited video kiosks to watch clips of the star's earlier performances and rehearsals, to get a behind-the-scenes glimpse of a favorite artist. What they didn't know was that the kiosk was watching them, too. The video booth was fitted with a camera that sent its visitors' images back to a "command post" in Nashville, where facial recognition software scanned them, reportedly looking for matches against a database of people who had stalked Swift in the past.[1] Were these images kept, or were they deleted securely? We don't know, just as we don't know how many other cameras capture us every day. Scanners like Swift's have been spotted at entrances to sports arenas, concert halls, and other entertainment venues. The public is often in the dark about their existence—and about policies related to how the images and other captured data are to be used, stored, or shared.

George Orwell's *1984* was published in 1948. Over subsequent years, the book became synonymous with a world of permanent surveillance, a society devoid of both privacy and freedom:

> ...there seemed to be no color in anything except the posters that were plastered everywhere. The black-mustachio'd face gazed down from every commanding corner. There was one on the house front immediately opposite. BIG BROTHER IS WATCHING YOU.[2]

The real 1984 came and went decades ago. Today, Big Brother's two-way telescreens would be amateurish toys today. Orwell's imagined London had cameras everywhere. His actual city now has at least half a million. Across the United Kingdom, there is one surveillance camera for every ten people.[3] The average Londoner is photographed hundreds of times a day by electronic eyes on the sides of buildings and on utility poles.

Yet there is much about the digital world that Orwell did not imagine. He did not anticipate that cameras are far from the most pervasive of today's tracking technologies. There are dozens of other kinds of data sources, and the data they produce is retained and analyzed. Cell phone companies know not only what numbers you call but where you have carried your phone. Credit card companies know not only where you spent your money but what you spent it on. Your friendly bank keeps electronic records of your transactions not only to keep your balance right but because it has to tell the government if you make huge withdrawals. When you go to a restaurant or a store, an app that has been quietly tracking your location asks you how you liked it, to feed your response into its recommendation-making engine.

The digital explosion has scattered the bits of our lives everywhere: records of the clothes we wear, the soaps we wash with, the streets we walk, and the cars we drive and where we drive them. And although Orwell's Big Brother had cameras, he didn't have search engines to piece the bits together, to find the needles in the haystacks. Wherever we go, we leave digital footprints, and computers of staggering capacity reconstruct our movements from those tracks. Computers reassemble the clues to form a comprehensive image of who we are, what we do, where we are doing it, and whom we are discussing it with.

Perhaps none of this would have surprised Orwell. Had he known about electronic miniaturization, he might have guessed that we would develop an astonishing array of tracking technologies. But there is something more fundamental that distinguishes the world of *1984* from the actual world of today. We have fallen in love with this always-on world. We accept our loss of privacy in exchange for efficiency, convenience, and small price discounts.

Attitudes have changed in the past decade. In a 2007 Pew/Internet Project report, 60% of Internet users were "not worried about how much information is available about them online," but by 2018, the ratio had flipped, and more than 60% "would like to do more to protect their privacy"; just 9% believe they have "a lot of control" over the information that is collected about them.[4] Although we're getting more worried about the loss of control over personal information, we're not sure there's much we can do about it.

In the world of bits, Big Brother has gotten both bigger and smaller. Technologically sophisticated nations like the United States and China have unprecedented ability to watch us, and they exercise that ability more often

than we might like. Companies do, too. They've built new businesses around ubiquitous data collection, much of it geared toward marketing directly to us. Commercial data also forms a rich lode for government to mine, a public–private surveillance partnership.

We, too, are a part of the surveillance networks, keeping tabs on ourselves and one another. We invite apps to track our movements and smart assistants to listen in to our conversations. We record our changes of mood and chitchat with friends, and we snap photos of friends and strangers. About seven in ten adults have created profiles on social networking websites. Yet most are dissatisfied with the level of control they have over what happens to the data they post there.[5]

There are hints that the privacy tide may be changing, that we're not willing to trade privacy for the benefits of the digital world. Regulators are giving us new protections (although often not from government surveillance), and companies are now marketing privacy as a feature.

Bits Cubed: The Snowden Files

When a 29-year-old Edward Snowden met with journalists in the lobby of Hong Kong's Mira Hotel in June 2013, he told them to look for the guy with a Rubik's Cube.[6] They eventually did so and got a trove of classified documents and PowerPoint presentations describing massive U.S. government communications surveillance: a series of front-page stories for the journalists. Snowden, as a systems administrator for the National Security Agency (NSA), extracted gigabytes of material, copying it to micro-SD cards smaller than the stickers on his Rubik's Cube.

The Snowden revelations fueled a series of front-page stories in the *New York Times*, *Washington Post*, and *Guardian* in 2013.[7] They showed the NSA engaged in pervasive communications surveillance—not just of foreigners and suspected terrorists but of law-abiding American citizens. If you used Yahoo! Mail or Google Search or dozens of other popular services, you were swept up in the dragnet. While the U.S. Constitution and laws make a sharp distinction between U.S. citizens and "foreign persons" that limit the government's ability to spy on its citizens, the bits carried no such distinction, and citizens ended up in the same buckets.

After the September 11, 2001, terrorist attacks, Congress passed new laws increasing spying powers. Notably, the USA PATRIOT Act authorized national security letters, which are secret demands for communications records; warrantless wiretaps of foreigners suspected of terrorist activity; and increasing ability to collect information on citizens any time obtaining foreign intelligence information is "a significant purpose" of the surveillance. Civil liberties groups expressed concern at the time that the act

eliminated judicial checks and balances on surveillance,[8] but the act passed the Senate 98:1. Snowden's documents showed how far the NSA was pushing these new authorities.

The NSA exploited several properties of electronic communications. The popularity of centralized services for phone, email, search, and storage meant that taps at these corporate networks captured significant activity. The global nature of the Internet meant these taps could reach around the world from a few implant sites. A single request to Verizon for "business records" enabled the collection of millions of Americans' telephone call activity.[9] The Upstream program made full copies of everything carried along major domestic fiber-optic cable networks. Other top-secret warrantless data collection tools included XKEYSCORE[10] and EGOTISTICALGIRAFFE.[11]

U.S. officials defending the programs said they were only collecting metadata, not the contents of communications—the envelopes and addresses, not the letters inside. However, the web of contacts itself is tremendously informative. "We kill people based on metadata," said General Michael Hayden, former director of the NSA and the CIA.[12] William Binney, another ex-NSA whistleblower, left after the agency cut a program to conduct privacy-preserving searches.

The way we leave fingerprints and footprints is only part of what is new. We have always left a trail of information behind us—in our tax records, hotel reservations, and long-distance telephone bills. True, the footprints are far clearer and more complete today than ever before. But something else has changed: the harnessing of computing power to correlate data, to connect the dots, to put pieces together, and to create cohesive, detailed pictures from what would otherwise be meaningless fragments. The digital explosion does not just blow things apart. Like the explosion at the core of an atomic bomb, it blows things together as well. Gather up the details, connect the dots, and assemble the parts of the puzzle, and a clear picture will emerge.

Computers can sort through databases too massive and too boring to be examined with human eyes. They can assemble colorful pointillist paintings out of millions of tiny dots when any few dots would reveal nothing. When a federal court released half a million Enron emails obtained during the corruption trial, computer scientists quickly identified the subcommunities, and perhaps conspiracies, among Enron employees, using no data other than the pattern of who was emailing whom (see Figure 2.1). The same kinds of clustering algorithms work on patterns of telephone calls. You can learn a lot by knowing who is calling or emailing whom, even if you don't know what they are saying to each other—especially if you know the times of the communications and can correlate them with the times of other events.

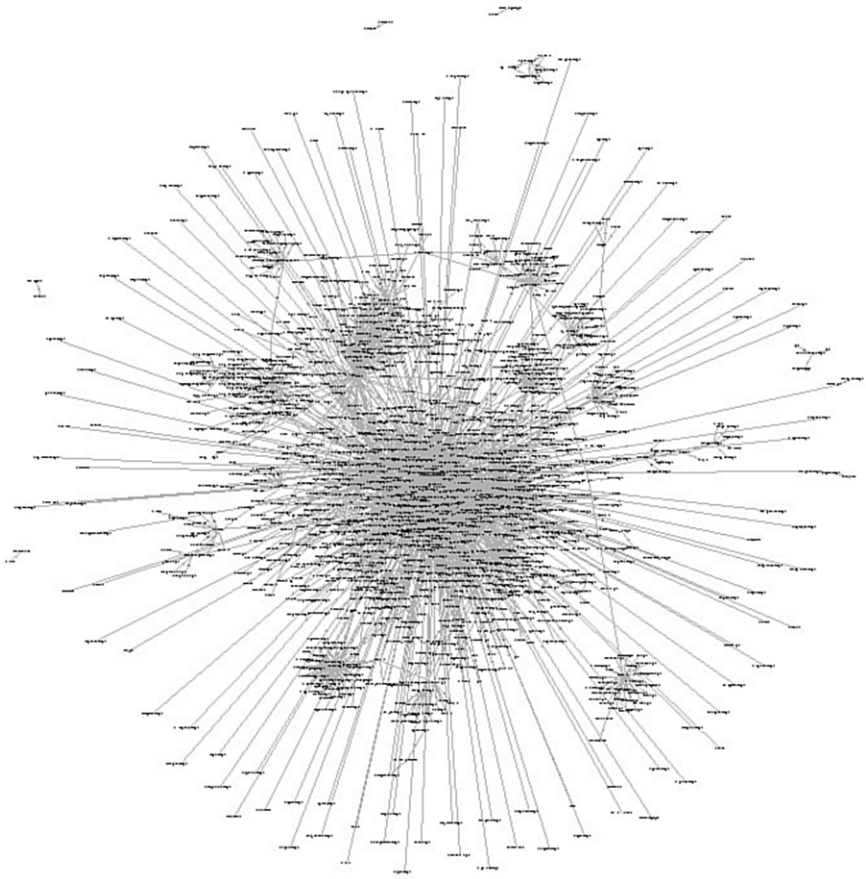

Source: Enron, Jeffrey Heer, Figure 3 from http://jheer.org/enron/v1/

FIGURE 2.1 Diagram showing clusters of Enron emailers, indicating which employees carried on heavy correspondence with which others. The evident "blobs" may be the outlines of conspiratorial cliques.

The tale of Snowden and the NSA is two bits stories at once. Digital communication made it possible for the NSA to collect vast quantities of information, millions and millions of calls and emails, from just a few locations—something that would have been impossible if we were still communicating with regular phones and paper letters. And when Snowden took copies of everything, he could fit the equivalent of thousands of file cabinets of information into his pocket.

What can we do in the face of such government-directed surveillance? Snowden chose exposure, aiming for his disclosures to help support lawsuits against the programs and public pressure on lawmakers to rein in the NSA.

When he finished opening his Pandora's memory card of documents, he left us with a cause for hope: Math works. The NSA may have the world's best cryptographers and cryptanalysts, but the fundamental mathematics of encryption are still effective. The years since Snowden's disclosures have seen a dramatic increase in the use of encryption in basic Internet and web protocols[13] and in the applications that run on them.[14] End-to-end encryption enables us to reclaim some of the privacy that pervasive monitoring of unencrypted traffic unraveled.

"Reasonable Expectations of Privacy" Technology and the Fourth Amendment

Technological change has stood in tension with privacy before. When the Supreme Court first encountered the telephone wiretap in 1928, the president did not yet have a phone on his desk, although traffickers in illegal liquor (this was during Prohibition) had found the technology, and law enforcement wanted to listen in.[15] When the bootleggers challenged the tapping of their phone lines—alligator clips on physical wires outside homes and offices—the Court's majority put the telephone, which was high-tech at the time, in a frame they recognized, of physical intrusion and trespass. Without trespass, the Court held, there was no "search" or "seizure" and therefore no need for a warrant:

> The reasonable view is that one who installs in his house a telephone instrument with connecting wires intends to project his voice to those quite outside, that the wires beyond his house, and messages while passing over them, are not within the protection of the Fourth Amendment. Here those who intercepted the projected voices were not in the house of either party to the conversation.

Justice Brandeis, who did not agree, wrote in his dissent:

> Whenever a telephone line is tapped, the privacy of the persons at both ends of the line is invaded, and all conversations between them upon any subject, and although proper, confidential, and privileged, may be overheard. Moreover, the tapping of one man's telephone line involves the tapping of the telephone of every other person whom he may call, or who may call him."

But he was in the minority; for decades, warrantless wiretapping was lawful.

The Court's ruling in *Olmstead v. United States* increased the vulnerability of telephonic communications to police snooping, but it also publicly exposed

that lack of privacy. Criminals, judges, and the general public learned that their conversations were liable to be tapped. As the telephone itself became more widely used, the legal rule triggered responses. States passed wiretap acts to protect by statute what the Constitution would not, and in 1934, Congress included anti-interception prohibitions in The Communications Act, section 605.

When Charles Katz came before the Court in 1967 to challenge the wiretapping of his (illegal wagering) conversation from a public telephone booth, the times, technologies, and legal norms had all changed. The telephone was part of everyday life, for personal and intimate communications as well as businesses both lawful and unlawful. The public and the justices themselves had experience to color their views of the technology. Asked again "whether a public telephone booth is a constitutionally protected area," the Court said that was the wrong formulation: It wasn't place but context. Telephone calls now demanded greater protection, even when conducted from the relative publicity of a glass-walled "public" phone booth. Justice Harlan, concurring in the judgment throwing out Katz's wiretap, articulated the test that still defines the Fourth Amendment's privacy protection: a "twofold requirement, first that a person have exhibited an actual (subjective) expectation of privacy and, second, that the expectation be one that society is prepared to recognize as 'reasonable.'"[16]

Location, Location, Location

Buy a navigation-equipped car, and it will listen to precisely timed signals from satellites reporting their positions in space. The Global Positioning System (GPS) calculates locations based on the satellites' locations and the times their signals are received. The 24 satellites spinning 12,500 miles above the earth enable your car to locate itself within 25 feet, at a price so low that most new cars have it as a standard feature. What was once a military secret now comes free in every smartphone.

If you carry a GPS-enabled cell phone, your friends can find you if that is what you want. If your GPS-enabled rental car has a radio transmitter, you can be found whether you want to be or not. Car leasing companies are adding transponders, including auto-immobilizers, to enable remote repossession, without even sending a repo man to the site. Those who fall behind in their car payments may suddenly find themselves unable to get to or from work.

GPS enables you to determine your location anywhere on earth, and even a low-end cell phone serves as a rudimentary positioning system. If you are traveling in settled territory—anyplace where you can get cell coverage—you

move from the range of one cell tower to the range others, pinging the towers as you go. Triangulation among these signals can be used to locate you. The location is less precise than that supplied by GPS—only within ten city blocks or so—but the fact that it is possible at all means that a pattern of your behavior can be built, or photos can be stamped with identifying information about where they were shot, as well as when and with what camera.

Timothy Carpenter was given away by the bits from his cell phone, and then he was given a second chance by the law. A string of robberies of Radio Shack and T-Mobile stores in the Detroit area led to the arrest of four men in April 2011.[17] One of them confessed and gave the FBI the cell phone numbers of his accomplices; he also let law enforcement collect recently called numbers from his phone. With that evidence, prosecutors obtained an order directing wireless phone carriers to disclose information and cell-site location histories on the called numbers. They concluded that a phone registered to Timothy Carpenter had been near four store locations at the times when the stores were robbed. Carpenter was taken to court, and at trial, several confederates testified that he had been the leader of the robbery operation. With the corroboration of the cell-site mapping data, he was sentenced to more than 100 years in prison.

Carpenter appealed his case to the U.S. Supreme Court, arguing that the use of cell-site location data amounted to a "search,"[18] which could be conducted only with a warrant based on probable cause—not the mere order prosecutors had used to obtain records from the wireless carriers.

In 2018, the Court agreed with Carpenter: Because "cell phone location information is detailed, encyclopedic, and effortlessly compiled," the equivalent of a long-term digital "tail," individuals should have a reasonable expectation of privacy that their location history will not be exposed without a search warrant. Just because we're all carrying detailed location trackers, and those devices locate us to third parties in order to function, doesn't mean law enforcement gets automatic access to our map history. As it had earlier in *Katz*, the Court said that new technological capabilities shouldn't upend the balance between law enforcement and the public. Bits might trace our every move, but police need the judicial oversight of a search warrant to see them.

Lawyers and technologists argue about the balance between their disciplines. Post-Snowden revelations, they wonder whether we can trust the government to check its own power, or if records supposed to be available only by application to a neutral magistrate for a search warrant will instead be handed over or gathered in bulk. In 2013, after Edward Snowden revealed the existence of a secret, and extensive, data collection program code-named "PRISM," the NSA argued that records aren't effectively "collected" until they are searched, even once they are gathered in data banks. But while you can

encrypt your conversations, it's much harder to hide the metadata of your digital footprints. (Tor Project's onion routing, https://www.torproject.org, is the best option.) To protect the privacy of those activities that must be public to be effective or that depend on interactions with others we don't necessarily trust to keep our secrets, we need the force of law and social norms.

Black Boxes: Not Just for Airplanes Anymore

On April 12, 2007, John Corzine, governor of New Jersey, was heading back to the governor's mansion in Princeton to mediate a discussion between Don Imus, the controversial radio personality, and the Rutgers University women's basketball team.[19]

His driver, 34-year-old state trooper Robert Rasinski, headed north on the Garden State Parkway. He swerved to avoid another car and flipped the governor's Chevy Suburban. Governor Corzine, who had not fastened his seatbelt, broke 12 ribs, a femur, his collarbone, and his sternum. The details of exactly what happened were unclear. When questioned, Trooper Rasinski said he was not sure how fast they were going—but we *do* know. He was going 91 in a 65-mile-per-hour zone. There were no police with radar guns around; no human was tracking his speed. We know his exact speed at the moment of impact because his car, like 30 million other cars in America, had a black box—an event data recorder (EDR) that captured every detail about what was going on just before the crash. An EDR is an automotive "black box" like the ones recovered from airplane crashes.

EDRs started appearing in cars around 1995, and they now appear in almost all models. Your insurance company is probably entitled to its data if you have an accident. Yet most people do not realize that EDRs exist, unless they've gotten an offer from their insurance company to give up real-time data rather than pay higher premiums.

EDRs capture information about speed, braking time, turn signal status, seat belts: information needed for accident reconstruction, to establish responsibility, or to prove innocence. CSX Railroad was exonerated of all liability in the death of the occupants of a car when its EDR showed that the car was stopped on the train tracks when it was hit. Police generally obtain search warrants before downloading EDR data—but not always; in some cases, they do not have to. When Robert Christmann struck and killed a pedestrian on October 18, 2003, Trooper Robert Frost of the New York State Police downloaded data from the car at the accident scene. The EDR revealed that Christmann had been going 38 miles per hour in an area where the speed limit was 30. When the data was introduced at trial, Christmann claimed that the state had violated his Fourth Amendment rights because it had not asked his permission or obtained a search warrant before retrieving the data. That

It is almost as hard to avoid leaving digital footprints as it is to avoid touching the ground when we walk.

was not necessary, ruled a New York court. Taking bits from the car was not like taking something out of a house, and no search warrant was necessary.[20]

Bits mediate our daily lives. It is almost as hard to avoid leaving digital footprints as it is to avoid touching the ground when we walk, and even if we lived our lives without walking, we would unsuspectingly be leaving fingerprints.

SEARCHING LOCATION

Download your location history from Google or Facebook and look at the picture it paints. Does anything there make you nervous (whether it should or not)? What would you have difficulty explaining? Have you ever changed the settings from their account defaults? Should you?

Saving Time: Electronic Tolling and License Plate Readers

For commuters who use toll roads or bridges, the risk–reward calculation is not even close. Time is money, and time spent waiting in a car also means anxiety and frustration. If there is an option to get a toll booth transponder, many commuters will get one, even if the device costs a few dollars up front. Cruising past the cars waiting to pay with dollar bills is not just a relief; it actually brings the driver a certain satisfied glow.

The transponder, which the driver attaches to the windshield inside the car, is an RFID-enabled device powered with a battery that sends information to a sensor several feet away as the driver whizzes past. The sensor can be mounted in a constricted travel lane, where a toll booth for a human toll taker might have been. Or it can be mounted on a boom above traffic so the driver doesn't even need to change lanes or slow down. And what is the possible harm? Of course, the state is recording the fact that the car has passed the sensor; that is how the proper account balance can be debited to pay the toll. When the balance gets too low, the driver's credit card may get billed automatically to replenish the balance, which only makes the system better—no fumbling for change or doing anything else to pay for your travels.

The monthly bill—for the Massachusetts Fast Lane, for example—shows where and when you got on the highway, accurate to the second. It also shows how far you traveled on the highway and where you got off. Fast Lane also informs you of the mileage, which is another useful service because Massachusetts drivers can get a refund on certain fuel taxes if the fuel was used on the state toll road. Of course, you do not need a PhD to

figure out that the state also knows when you got off the road, to the second, and that with one subtraction and one division, its computers could figure out if you were speeding. Technically, in fact, it would be trivial for the state to print the appropriate speeding fine at the bottom of the statement and to bill your credit card for that amount at the same time it charges you for the tolls. That would be taking convenience a bit too far, and no state does it—yet.

What does happen right now, however, is that toll transponder records are introduced into divorce and child custody cases. You've never been within 5 miles of that lady's house? Really? Why have you gotten off the highway at the exit near it so many times? You say you can be the better custodial parent for your children, but the facts suggest otherwise. As one lawyer put it, "When a guy says, 'Oh, I'm home every day at 5, and I have dinner with my kids every single night,' you subpoena his E-ZPass and you find out he's crossing that bridge every night at 8:30. Oops!" Such records have been subpoenaed hundreds of times in family law cases. They have also been used in employment cases, to prove that the car of a worker who said he was working was actually far from the workplace.

But most of us aren't planning to cheat on our spouses or our bosses, so the loss of privacy seems like no loss at all—at least compared to the time saved. Of course, if we actually *were* cheating, we *would* be in a big hurry and might take some risks to save a few minutes!

Massachusetts toll roads eliminated toll takers in 2017. Drivers can save some money by equipping their cars with transponders, but if they don't have a transponder, never fear: "We will bill you," the state announces in billboards along the highway. There are no cash lanes now. Gantries equipped with both transponder antennas and automated license plate readers record every car or truck that passes through. To stay anonymous, you need to take the slow road.

The License Plate Tells More Than You Think

In June 2018, southern California mall operator Irvine Company was found to be collecting the license plate numbers of vehicles entering its parking areas. When a visitor, 14-year-old Zoe Wheatcroft, dug deeper into the company's "privacy policy," she found that Irvine was not only collecting license plate information but sharing it with law enforcement, in a database that might be accessed by agents from Immigration and Customs Enforcement (ICE).[21] When word got out, Irvine and Vigilant, the database company, responded that their policy was in fact narrower and more restrictive but gave customers no way to know that a shopping trip wouldn't put them in surveillance crosshairs.

Automated license plate recognition is a form of mass surveillance enabled by cheaper and more sophisticated cameras, software, and network capabilities. Automatic cameras capture images of license plates, convert the plate numbers into plaintext characters, and annotate the images with time, date, and GPS-derived location before transmitting and storing each instance. The data stream may be queried in real time, as in a search for a wanted criminal or stolen vehicle, or it may be retrieved later to give a picture of shoppers' demographics or a particular shopper's travel pattern.

Loose Fitbits Sink Ships?

The Strava fitness-mapping application offers a connection to users' GPS-enabled smartphones, watches, and Fitbit devices in order to enable athletes to track their runs, cycle routes, and other activities. Strava combined the data into a "heatmap" visualization, aggregating more than a billion activity logs into colored streaks across a map. While the Strava team highlighted a few recreation images on their blog—the Ironman triathlon swim off Hawaii, mountain biking in Whistler, British Columbia—a researcher, noting what appeared to be the outlines of military bases in Afghanistan, posted screenshots to Twitter and reminded people "turning off data sharing is an option."[22] Strava's CEO followed up with a blog post pointing to explanations of the privacy settings and promising to work with military and government officials "to address potentially sensitive data."[23]

Of course, one can say that soldiers in sensitive locations should turn off their location reporting—which means they need to know that their devices and applications have that setting and consider its consequences. But the Strava heatmap may be only the most visible and most easily changed of the places we leave these trails. Cell phones build location maps as they ping nearby towers; frequently accessed websites have logs of the IP addresses from which they are viewed (from which the site operator can map corresponding geolocation); and many mobile apps collect location information to target advertising. Individual data points may seem harmless, but points gathered over time and space can paint a detailed picture of travel patterns or home life—and even secret military strategy.

Big Brother, Abroad and in the United States

Big Brother really is watching today, and his job has gotten much easier, thanks to the digital explosion. In China, which has a long history of tracking individuals as a mechanism of social control, the millions of residents of Shenzhen are being issued identity cards, which record far more than the

bearer's name and address. According to a report in the *New York Times*,[24] the cards document the individual's work history, educational background, religion, ethnicity, police record, medical insurance status, landlord's phone number, and reproductive history. Touted as a crime-fighting measure, the new technology—developed by an American company—will come in handy in dealing with cases of street protests and individual activities deemed suspicious by the authorities. The sort of record keeping that used to be the responsibility of local authorities is becoming automated and nationalized as the country prospers and its citizens become increasingly mobile. The technology makes it easier to know where everyone is, and the government is taking advantage of that opportunity. In Xinjiang, where the Uighur minority faces especially strict scrutiny, police have an app that can flag when someone has stopped using a smartphone or avoids the front door. Facial recognition is targeted at Uighurs, who are made to pass through checkpoints that Han (the ethnic majority elsewhere in China) are permitted to avoid. Chinese tracking is far more detailed and pervasive than Britain's system of ubiquitous surveillance cameras.

Identifying Citizens—Without ID Cards

In the age of global terrorism, democratic nations are resorting to digital surveillance to protect themselves, creating hotly contested conflicts with traditions of individual liberty. In the United States, the idea of a national identification card prompts a furious libertarian reaction from parties not usually outspoken in defense of individual freedom. Under the REAL ID Act of 2005, uniform federal standards were to be implemented for state-issued driver's licenses. Although it passed through Congress without debate, the law is opposed by at least 18 states. Resistance pushed back the implementation timetable multiple times. In 2018, 13 years later, only 37 states met the REAL ID rules. Finally, in 2019, states were told their final extension would expire, and only REAL ID–compliant documents would be accepted for federal identification by October 2020. Then COVID-19 hit, and the deadline was extended yet again. Yet even fully implemented, REAL ID would fall far short of the true national ID preferred by those charged with fighting crime and preventing terrorism.

As the national ID card debate continues in the United States, the FBI is making it irrelevant by exploiting emerging technologies. There would be no need for anyone to carry an ID card if the government had enough biometric data on Americans—that is, detailed records of their fingerprints, irises, voices, walking gaits, facial features, scars, and earlobe shapes. Gather a combination of measurements on individuals walking in public places, consult the databases, connect the dots, and—bingo!—their names pop up on the computer screen. No need for them to carry ID cards; the combination of biometric data would pin them down perfectly.

As the national ID card debate continues in the United States, the FBI is making it irrelevant by exploiting emerging technologies.

Well, only imperfectly at this point, but the technology is improving. And the data is already being gathered and deposited in the data vault of the FBI's Criminal Justice Information Services database in Clarksburg, West Virginia. The database already holds some 75 million sets of fingerprints, and the FBI processes 100,000 requests for matches every day. Any of 900,000 federal, state, and local law enforcement officers can send a set of prints and ask the FBI to identify it. If a match comes up, the individual's criminal history can be accessed in the database, too.

But fingerprint data is hard to gather; mostly it is obtained when people are arrested. The goal of the project is to get identifying information on nearly everyone—and to get it without bothering people too much. For example, a simple notice at airport security could advise travelers that, as they pass through airport security, a detailed "snapshot" will be taken as they enter the secure area. The traveler would then know what is happening and could have refused (and stayed home). As an electronic identification researcher puts it, "That's the key. You've chosen it. You have chosen to say, 'Yeah, I want this place to recognize me.'"[25] The project eliminates the issue of REAL ID controversies, as all the data being gathered is, in some sense at least, offered voluntarily.

San Francisco, California, the epicenter of the technology boom, moved in the other direction, banning law enforcement use of facial recognition technology.[26] The Board of Supervisors heard concerns that the technologies were biased, lacked transparency, and could be abused by government. At the same time, however, data-based identification flourishes in private hands. The company PatronScan asserts that its database of IDs swiped at bars contains more than 60 million IDs across 200 cities. PatronScan not only checks that bar-goers are of legal drinking age but maintains a blacklist of patrons flagged for "bad behavior."[27]

Friendly Cooperation Between Big Siblings

In fact, there are two Big Brothers, and they often work together. We are, by and large, glad they are watching—if we are aware of it at all. Only occasionally are we alarmed about their partnership.

The first Big Brother is Orwell's—the government. And the other Big Brother is the industry about which most of us know very little: the business of aggregating, consolidating, analyzing, and reporting on the billions of individual transactions, financial and otherwise, that take place electronically every day. Of course, the commercial data aggregation companies are not in the spying

business; none of their data reaches them illicitly. But they do know a lot about us, and what they know can be extremely valuable, both to businesses and to the government.

The new threat to privacy is that computers can extract significant information from billions of apparently uninteresting pieces of data, in the way that mining technology has made it economically feasible to extract precious metals from low-grade ore. Computers can correlate databases on a massive level, linking government data sources together with private and commercial ones to create comprehensive digital dossiers on millions of people. With their massive data storage and processing power, they can make connections in the data, by using brute force rather than ingenuity. And the computers can discern even very faint traces in the data—traces that may help track payments to terrorists, set insurance rates, or simply help us make sure our new babysitter is not a sex offender.

And so we turn to the story of the government and the aggregators.

Acxiom is the country's biggest customer data company. Its business is to aggregate transaction data from all those swipes of cards in card readers all over the world. This amounted to more than a hundred billion transactions in 2018.[28] The company uses its massive amounts of data about financial activity to support the credit card industry, banks, insurers, and other consumers of information about how people spend money. Unsurprisingly, after the War on Terror began, the Pentagon also got interested in Acxiom's data and the ways the company gathers and analyzes it. Tracking how money gets to terrorists might help find the terrorists and prevent some of their attacks.

ChoicePoint is the other major U.S. data aggregator. ChoicePoint has more than 100,000 clients, which call on it for help in screening employment candidates, for example, or determining whether individuals are good insurance risks.

Acxiom and ChoicePoint are different from older data analysis operations in the scale of their operations. Quantitative differences have qualitative effects, as we said in Chapter 1; what has changed is not the technology but rather the existence of rich data sources. Forty years ago, credit cards had no magnetic stripes. Charging a purchase was a mechanical operation; the raised numerals on the card made an impression through carbon paper so you could have a receipt, and the top copy went to the company that issued the card. Today, if you charge something using your CapitalOne card, the bits go instantly not only to CapitalOne but to Acxiom and other aggregators. The ability to search through huge commercial data sources—including not just credit card transaction data but phone call records, travel tickets, and banking transactions, for example—is another illustration that more of the same can create something new.

Privacy laws do exist, of course. For a bank, or a data aggregator, to post your financial data on its website would be illegal. But privacy is still developing as an area of the law, and it is connected to commercial and government interests in uncertain and surprising ways.

A critical development in privacy law was precipitated by the presidency of Richard Nixon. In what is generally agreed to be an egregious abuse of presidential power, Nixon used his authority as president to gather information on those who opposed him—in the words of his White House counsel at the time, to "use the available federal machinery to screw our political enemies." Among the tactics Nixon used was to have the Internal Revenue Service audit the tax returns of individuals on an "enemies list," which included members of Congress, journalists, and major contributors to Democratic causes. Outrageous as it was to use the IRS for this purpose, it was not illegal, so Congress moved to ban it in the future.

The Privacy Act of 1974 established broad guidelines for when and how the federal government can assemble dossiers on citizens it is not investigating for crimes. The government has to give public notice about what information it wants to collect and why, and it has to use what it collects only for those reasons.

The Privacy Act limits what the government can do to gather information about individuals and what it can do with records it holds. Specifically, it states, "No agency shall disclose any record which is contained in a system of records by any means of communication to any person, or to another agency, except pursuant to a written request by, or with the prior written consent of, the individual to whom the record pertains, unless...." If the government releases information inappropriately, even to another government agency, the affected citizen can sue for damages in civil court. The protections provided by the Privacy Act are sweeping—although not as sweeping as they may seem. Not every government office is in an "agency"; the courts are not, for example. The act requires agencies to give public notice of the uses to which they will put the information, but the notice can be buried in the Federal Register, where the public probably won't see it unless news media happen to report it. Then there is the "unless" clause, which includes significant exclusions. For example, the law does not apply to disclosures for statistical, archival, or historical purposes; civil or criminal law enforcement activities; Congressional investigations; or valid Freedom of Information Act requests.

Despite the act's exclusions, government practices changed significantly because of this law. Then, a quarter century later, came 9/11. *Law enforcement should have seen it all coming*, was the constant refrain as investigations revealed how many unconnected dots were in the hands of different government agencies. *It all could have been prevented if the investigative fiefdoms*

had been talking to each other. They should have been able to connect the dots. But they could not—in part because the Privacy Act restricted interagency data transfers. A response was badly needed. The Department of Homeland Security was created to ease some of the interagency communication problems, but that government reorganization was only a start.

In January 2002, just a few months after the World Trade Center attack, the Defense Advanced Research Projects Agency (DARPA) established the Information Awareness Office (IAO) with a mission to:

> imagine, develop, apply, integrate, demonstrate, and transition information technologies, components and prototype, closed-loop, information systems that will counter asymmetric threats by achieving total information awareness useful for preemption; national security warning; and national security decision making. The most serious asymmetric threat facing the United States is terrorism, a threat characterized by collections of people loosely organized in shadowy networks that are difficult to identify and define. IAO plans to develop technology that will allow understanding of the intent of these networks, their plans, and potentially define opportunities for disrupting or eliminating the threats. To effectively and efficiently carry this out, we must promote sharing, collaborating, and reasoning to convert nebulous data to knowledge and actionable options.

Vice Admiral John Poindexter directed the effort that came to be known as "Total Information Awareness" (TIA). The growth of enormous private data repositories provided a convenient way to avoid many of the prohibitions of the Privacy Act. The Department of Defense can't get data from the Internal Revenue Service because of the 1974 Privacy Act. *But the government can buy the very same data it is barred from collecting from private data aggregators!* In a May 2002 email to Adm. Poindexter, Lt. Col Doug Dyer discussed negotiations with Acxiom:

> Acxiom's Jennifer Barrett is a lawyer and chief privacy officer. She's testified before Congress and offered to provide help. One of the key suggestions she made is that people will object to Big Brother, wide-coverage databases, but they don't object to use of relevant data for specific purposes that we can all agree on. Rather than getting all the data for any purpose, we should start with the goal, tracking terrorists to avoid attacks, and then identify the data needed (although we can't define all of this, we can say that our templates and models of terrorists are good places to start). Already, this guidance has shaped my thinking.

> Ultimately, the U.S. may need huge databases of commercial transactions that cover the world or certain areas outside the U.S. This information provides economic utility, and thus provides two reasons why foreign countries would be interested. Acxiom could build this mega-scale database.

The *New York Times* broke the story in October 2002. As Poindexter had explained in speeches, the government had to "break down the stovepipes" separating agencies and get more sophisticated about how to create a big picture out of a million details, no one of which might be meaningful in itself. The *Times* story set off a sequence of reactions from the Electronic Privacy Information Center and civil libertarians. Congress defunded the office in 2003—but that was not the end of the idea.

The key to TIA was data mining: looking for connections across disparate data repositories, finding patterns, or "signatures," that might identify terrorists or other undesirables. The General Accountability Office report on Data Mining (GAO-04-548) reported on a survey of 128 federal departments.[29] It described 199 separate data mining efforts, of which 122 used personal information.

Although IAO and TIA went away, Project ADVISE at the Department of Homeland Security continued with large-scale profiling system development. Eventually, Congress demanded that the privacy issues concerning this program be reviewed as well. In his June 2007 report (OIG-07-56), Richard Skinner, the DHS inspector general, stated that "program managers did not address privacy impacts before implementing three pilot initiatives," and a few weeks later, the project was shut down. But ADVISE was only one of a dozen data-mining projects going on in DHS at the time.

Similar privacy concerns led to the cancellation of the Pentagon's TALON database project. That project sought to compile a database of reports of suspected threats to defense facilities as part of a larger program of domestic counterintelligence.

Despite these privacy concerns, as Edward Snowden revealed, many surveillance and data mining programs simply carried on under the radar.

The government creates projects, the media and civil liberties groups raise serious privacy concerns, the projects are canceled, and new projects arise to take their place. The cycle seems to be endless. In spite of Americans' traditional suspicions about government surveillance of their private lives, the cycle seems to be almost an inevitable consequence of Americans' concerns about their security and the responsibility that government officials feel to use the best available technologies to protect the nation. Corporate databases often contain the best information on the people about whom the government is curious.

Data Collection, Data Breach

Storage is cheap, but security is difficult. One of the depressingly common events in the digital era is data breach. A customer database is exposed, and user accounts or credit cards are misused until the breach is rectified. Data breach notification laws in many states now provide some transparency, as well as incentive for companies storing data to clean up to avoid class action lawsuits.

Amid numerous breaches, Equifax and OPM stand out. Equifax, one of the big credit-reporting companies, stores records of credit card account payment histories. If you go to take out a car loan or a mortgage, the lender will check your credit score with Equifax. In September 2017, Equifax announced a data breach that exposed the personal information—names and dates of birth, Social Security numbers, physical addresses, and other personal information that could lead to identity theft and fraud—of 147 million people, more than half the adult population of the United States.[30] The Federal Trade Commission complaint alleged that Equifax failed to take basic network security measures, including failing to update database software when notified of an access control vulnerability. This occurred even as the company had a privacy policy promising it implemented "reasonable physical, technical and procedural safeguards" to protect consumer data. Equifax settled the FTC complaint with an agreement to pay at least $575 million and potentially up to $700 million. As part of the settlement, affected consumers were offered free credit monitoring services. Those trying to exclude themselves from future databases, however, were told "You cannot opt out of this data collection."[31]

As the human resources arm of the U.S. government, the Office of Personnel Management collects a great deal of sensitive information: identification, background checks, and fingerprints.[32] Over 21 million of these records were stolen when OPM's data stores were breached in 2014. When people's credit cards are stolen, they get new cards. When their Social Security numbers are taken, they can be enrolled in credit monitoring services. But you can't be issued a new set of fingerprints.

The number of new data sources—and the proliferation and interconnection of old data sources—is part of the story of how the digital explosion shattered privacy. But the other part of the technology story is about how all that data is put together.

Exponential growth—in storage size, processing speed, and communication speed—have changed the same old thing into something new. Blundering, stupidity, curiosity, malice, and thievery are not new. The fact that sensitive data about everyone in a nation could fit on a laptop *is* new. The ability to search for a needle in the haystack of the Internet *is* new. Easily connecting "public" data sources that used to be stored in file drawers in Albuquerque and Atlanta but are now both electronically accessible from Algeria—*that* is new, too.

Training, laws, and software all can help. But the truth of the matter is that, as a society, we don't really know how to deal with these consequences of the digital explosion. The technology revolution is outstripping society's capacity to adjust to the changes in what can be taken for granted.

Sometimes even public information is revealing. In Massachusetts, the Group Insurance Commission (GIC) is responsible for purchasing health insurance for state employees. When the premiums it was paying jumped one year, the GIC asked for detailed information on every patient encounter. And for good reason: All kinds of health care costs had been growing at prodigious rates. In the public interest, the state had a responsibility to understand how it was spending taxpayer money. The GIC did not want to know patients' names; it did not want to track individuals, and it did not want people to *think* they were being tracked. Indeed, tracking the medical visits of individuals would have been illegal.

So, the GIC data had no names, no addresses, no Social Security numbers, no telephone numbers—nothing that would be a "unique identifier" enabling a mischievous junior staffer in the GIC office to see who exactly had a particular ailment or complaint. To use the official lingo, the data was "deidentified"—that is, stripped of identifying information. The data did include the gender, birth date, zip code, and similar facts about individuals making medical claims, along with some information about why they had sought medical attention. That information was gathered not to challenge any particular person but to learn about patterns; if the truckers in Worcester are having lots of back injuries, for example, maybe workers in that region need better training on how to lift heavy items. Most states do pretty much the same kind of analysis of deidentified data about state workers.

Now this was a valuable data set not just for the Insurance Commission but for others studying public health and the medical industry in Massachusetts. Academic researchers, for example, could use such a large inventory of medical data for epidemiological studies. Because it was all deidentified, there was no harm in letting others see it, the GIC figured. In fact, it was such good data that private industry—for example, businesses in the health management sector—might pay money for it. And so the GIC sold the data to businesses. The taxpayers might even benefit doubly from this decision: The data sale would provide a new revenue source to the state, and in the long run, a more informed health care industry might run more efficiently.

But how deidentified was the material—really?

Latanya Sweeney was at the time a researcher at MIT. (She went on to become a professor at Carnegie Mellon University and then Harvard University.) She wondered how hard it would be for those who had received the deidentified data to "reidentify" the records and learn the medical problems of a particular state employee—for example, the governor of the Commonwealth.

Governor Weld lived, at that time, in Cambridge, Massachusetts. Cambridge, like many other municipalities, makes its voter lists publicly available for a charge of $15—and free for candidates and political organizations. For a particular precinct, you can obtain the records for only $.75. Sweeney spent a few dollars and got the voter lists for Cambridge. Anyone else could have done the same.

According to the Cambridge voter registration list, there were only six people in Cambridge with Governor Weld's birth date, only three of those were men, and only one of those lived in Governor Weld's five-digit zip code. Sweeney could use that combination of factors—birth date, gender, and zip code—to recover the Governor's medical records; she could therefore also recover the records of members of his family, since the data was organized by employee. This type of reidentification is straightforward. In Cambridge, in fact, birth date alone was sufficient to identify more than 10% of the population. Nationally, gender, zip code, and date of birth are all it takes to identify 87% of the U.S. population uniquely.

The data set contained far more than gender, zip code, and birth date. In fact, any of the 58 individuals who received the data in 1997 could have identified any of the 135,000 people in the database. "There is no patient confidentiality," said Dr. Joseph Heyman, president of the Massachusetts Medical Society. "It's gone."[33]

It is easy to read a story like this and scream, "Heads should roll!" But it is actually quite hard to figure out *who, if anyone, made a mistake*. Certainly collecting the information was the right thing to do, given that health costs are a major expense for all businesses and institutions. The GIC made an honest effort to deidentify the data before releasing it. Arguably the GIC might not have released the data to other state agencies. Data is a valuable resource, and once someone has collected it, the government is entirely correct in wanting it used for the public good. Forbidding such sharing would be like saying that every department of government should acquire its heating oil independently. Some might object to selling the data to an outside business—but only in retrospect; had the data really been better deidentified, whoever made the decision to sell the data might well have been rewarded for helping to hold down the cost of government.

Perhaps the mistake was the ease with which voter lists can be obtained. However, it is a tradition deeply ingrained in our system of open elections that the public may know who is eligible to vote and, indeed, who has voted. And voter lists are only one source of public data about the U.S. population. How many

It is easy to read a story like this and scream, "Heads should roll!" But it is actually quite hard to figure out who, if anyone, made a mistake.

21-year-old male Native Hawaiians live in Middlesex County, Massachusetts? In the year 2000, there were four. Anyone can browse the U.S. Census data, and sometimes it can help fill in pieces of a personal picture: Just go to factfinder.census.gov.

The mistake was thinking that the GIC data was truly deidentified when it was not. But with so many data sources available, and so much computing power that could be put to work connecting the dots, it is very hard to know just how much information has to be discarded from a database to make it truly anonymous. Aggregating data into larger units certainly helps; releasing data by five-digit zip codes reveals less than releasing it by nine-digit zip codes. But the coarser the data, the less it reveals also of the valuable information for which it was made available.

The Internet of Things

We have already observed that even privacy-conscious people surrender their privacy in exchange for convenience and small cost savings. Nowhere is this principle more evident than in the networking of light switches, refrigerators, and doorbells known as the Internet of Things (IoT). And it turns out that privacy is not the only thing we sacrifice when we let the Internet grow into everything we touch (and everything we no longer need to touch). The security of everything in the network can be compromised.

On October 21, 2016, the U.S. East Coast woke up to a massive Internet outage. Many popular websites for work and play, such as Twitter, Netflix, GitHub, and Reddit, wouldn't load.[34] It turned out that critical Internet services were under attack by hordes of machines elsewhere on the Internet. These machines were sending so many simultaneous requests that nameservers, key components of the Internet's traffic management infrastructure, couldn't keep up with the load. Trying to respond to malicious requests left them unable to answer legitimate users. Without nameservers to give directions, requesting computers couldn't find the sites, and so Twitter was "down" for users even while the service itself was still functioning.

Investigating researchers and engineers found that the requests were coming from an army of "smart" home devices: Internet-connected baby monitors, light bulbs, and routers. The devices' owners hadn't intended this activity and were mostly unaware of it. The gadgets had been infected by malicious software and enlisted in the Mirai botnet. Together, a group of devices operating with weak computing power on home Internet connections

Unfortunately, this kind of mass attack has become common enough to get a name and an acronym: distributed denial of service (DDoS).

combined into a force strong enough to disrupt global Internet services. The malicious software, or *malware*, exploited some common security flaws—default administrative passwords that hadn't been changed and unpatched and outdated software on devices exposed directly to the Internet—to infect one device and then copy itself to other devices (an infection pattern known as a *worm*). Once installed, the malware turned each device into a waiting "bot," listening for commands.

On October 21, the controller directed the cohort to send a rapid stream of requests for domain names, which caused a burst of traffic to publicly accessible nameservers, including those of major nameservice provider Dyn DNS. Dyn reported that, under the attack, it was getting 10 to 20 times the normal volume of requests, which it estimated came from 100,000 malicious or infected end devices.[35] These requests, along with retry efforts from real end users who couldn't get through, overwhelmed Dyn's defenses and left the company's servers unable to respond to legitimate lookups.

What's New Here? Scale, Control, Connectedness, and Interoperability

The Internet of Things promises to connect the physical world much as the Internet of bits connects computers and data. Sometimes that means putting general-purpose connected computers into previously "dumb" devices like refrigerators. A smart refrigerator might not only warn you when you are running out of milk but contact your grocery store and have it delivered to your home and charged to your credit card. Other times, being "smart" means opening an interface by which one can remotely read from and control sensors (devices that see, hear, or otherwise perceive their environment) and actuators (devices that do something, like shut off the dryer). Smart thermostats, for example, can be triggered by motion detectors to turn on the heat or air conditioning when someone is in the room. These connected things enable a vision for automated factories and supply chains, smart homes and cities, and self-driving car fleets.

As sensors, actuators, and chips get cheaper, they grow in number. They also propagate down the value chain. When chips were expensive, it may have made sense to put them in costly equipment like airplanes, but today they are in doorbells. Even low-end smartphones are now smart enough to be at the center of a home appliance network. Capabilities that once were purchased only by factories and run by experts are now available to the general public. Sometimes smart gadgets are capable of doing much more than the purchaser realizes because the functionality is simplified for marketing reasons. At the same time, safety, reliability, and other less marketable features

are given short shrift—and the manufacturers try to excuse their misplaced priorities by saying that devices must be kept small and operate on low power.

A light bulb or thermostat is often "set and forget": Once the device is functioning, its owner thinks of it as an appliance rather than as a small computer in need of security monitoring and software updates. Moreover, budget sellers of the devices might view them as one-time sales with no follow-on support, and even users who want to update software may find themselves with no option to do so. Another alternative is a suite of centrally managed devices, but this option may be more costly in terms of both dollars and customer privacy. Not everyone wants to share their lighting and temperature preferences, much less the audio and video streams from their baby monitors, with a company storing that data somewhere unknown and using it for who knows what.

Many IoT devices are always on, awaiting the moment when their owners will throw the switch to light a room. That makes them attractive hosts for writers of malware—programs designed with evil intent. The cleverest malware doesn't interfere with the devices' normal function; rather, it lurks invisibly, waiting for the "attack" command.

Threats: One-to-One Versus One-to-Many

When one home has a smart refrigerator, its behavior is interesting and important to only a handful of people.[36] An attacker could spoil a gallon of milk and make a mess of the kitchen, drain a bank account by ordering caviar instead of milk (if the right limits aren't set), or use the machine to cause local damage, possibly shorting the whole house's circuitry. Multiply the devices, however, and they can be leveraged to do damage beyond their neighborhood. The first *D* in DDoS, is for *distributed*. Replicating an attack from thousands of distributed devices can have an overwhelming cumulative effect. Denial of service can take many forms: requests for service that look legitimate but that are sent at a high volume; requests that take an unusually long time to fulfill; or requests that are malformed in such a way as to disable or crash the server to which they are sent. For example, consider what happens when all of a town's high schoolers call the local pizza parlor at the same Friday lunch time to ask the price of a slice with the works. A customer genuinely calling to order a pie will probably give up after a few busy signals.

What makes the IoT an "Internet" is standard protocols that enable the devices to communicate with one another and their controllers, Internet protocols, or special-purpose Wi-Fi or Bluetooth-based communications standards. Standard interfaces enable users to address multiple devices together, such as to plug a new light bulb in to an existing setup or add a freezer to a smart refrigerator. Designers might have anticipated that a connected baby

monitor could share updates with grandparents or enable caregivers to listen from the far reaches of their homes; connectivity could enable a refrigerator to consult a weather report and order ice cream when the temperature hits 80 degrees. Connectivity can enable devices to get smarter over time, with software updates and new possibilities for interaction. However, unguarded connectivity can leave openings for intrusions like the Mirai worm, and the common interfaces and underlying software enable malware writers to "break once, run anywhere."

In December 2017, more than a year after attacks took Dyn offline, three men pleaded guilty to charges of computer fraud and abuse, admitting to having written the software behind Mirai: Paras Jha, an undergraduate studying computer science at Rutgers University in New Jersey, and two friends or associates. According to their plea, they first targeted their attack against online gaming servers for the popular Minecraft game, where they were attempting to overwhelm the servers to gain an advantage. Later, Jha started a business selling computer-protection services and launched attacks against Rutgers while taunting the university that it should be buying DDoS protection. Jha and his associates didn't necessarily intend to disrupt Dyn or the Internet at large, but after they posted the software's source code online, others modified and redeployed the malware, pointing it at new targets.

Who's Responsible for IoT Security?

When the Tesla Model 3 electric car was first reviewed by *Consumer Reports*, it got poor marks for braking.[37] "The Tesla's stopping distance of 152 feet from 60 mph was far worse than any contemporary car we've tested," wrote the reviewer. A week after publication of the *Consumer Reports* review, however, the car manufacturer sent an over-the-air software update to cars across the country, including those that had already been sold. The car's braking distance improved by 19 feet, performance comparable to that of other compact cars, prompting *Consumer Reports* to upgrade its review.[38] Tesla told *Consumer Reports* that it had updated software controlling the Model 3's antilock braking system.

This wasn't the first time an over-the-air update changed vehicle performance. While Hurricane Irma was heading for Florida, Tesla overrode software-defined range limitations for cars in the storm's path, enabling owners to escape further.[39] Both cases illustrate the blurred line between software and hardware and murky outlines of product boundaries. Physical features of the car were changed by a remote software update, and a car's owner might not have even been aware of the change or given an opportunity to accept or reject it. Few owners would reject longer range (a feature that was a costly upgrade outside Irma's wake), but what if the more consistent braking came

at the cost of other performance features? Some owners found themselves thinking their cars were slower after the upgrade. What if public safety came at the price of speed? Should owners get to refuse updates?

Driving poses an externality problem. It's not just Tesla drivers who take risks if their cars don't stop in time; their vehicles are more dangerous to everyone who shares the road with them. We impose safety standards and inspection requirements on automobiles to reduce such risks and make roads safer. We might similarly impose a duty to upgrade on software and hardware users. If your software-enabled product is causing risks to others, if a safer alternative is found, you could be required to update, even if doing so would cause you some inconvenience or cost. Yet it's not just obviously dangerous and expensive objects like cars that require this caution; some of the connected devices taken over by the Mirai botnet were cheap toys. Some of their vendors might no longer be in business. Would this requirement change the nature of ownership?

Bruce Schneier speaks of the Internet of Things as a "world-sized robot," with sensors and actuators spanning the globe.[40] As the capabilities of this robot to cause harm—and actual examples of harm—multiply, he predicts that demands for regulation and liability will follow. Unless those who are building the technologies also build safeguards, the political and regulatory responses are likely to be blunt and may include prohibitions on connecting or using devices or broad restrictions on their use. Worse, regulations that do not account for the architectures and incentives of connectedness may fail to protect us.

Smart Cities: Efficiency, Individual Choice, Privacy, and Systemic Risks

An older man in a New York apartment complained that he was virtually imprisoned in his own home after the landlord installed app-controlled "smart locks" at the building's lobby entrance. Tony Mysak, 93 and blind in one eye, was unable to use the smartphone app required to open the lock. Mary Beth McKenzie, Tony's wife, objected to giving a record of her entries to the building and to Latch, the lock provider. The Latch app's privacy policy (since changed) noted that the app collected GPS location information that Latch might use for marketing purposes, as well as providing a record of door accesses and photographs to building management. When she asked for a physical key, the landlord laughed and offered only a smartcard. McKenzie and Mysak and a group of tenants had to sue their landlord to win the right to use keys instead of apps.[41]

The tenants had several complaints about the digitization of their front door. For some, it was the change in usability, from a familiar physical key to

a new application. For others, it was the privacy, the sense that their entries and exits—and perhaps even their travels—would be tracked and compiled, not by a human and fallible doorman but in an impersonal corporate database that wasn't even visible to them. The new affordances of this system, such as the ability to let a guest or super into the building without leaving a key under the doormat, weren't enough to compensate tenants for their loss of control.

Scale this "smart" building up a few orders of magnitude, and you get the "smart city," full of embedded and networked sensors. Traffic lights might coordinate with cars and buses for efficiency; power meters might communicate with the electrical grid in real time to smooth demand.

The city of Toronto planned a revitalization of its industrial waterfront "from the Internet up." The new Quayside would be built as a "smart city" in a partnership between the city and Google/Alphabet's Sidewalk Labs. But as they engaged in giddy futuristic speculation, planners were surprised by the opposition their announcement sparked. People complained about privacy, security, and loss of control. Who gets to see the data generated by the digital infrastructure; who gets to make decisions based on it? Sadly, we won't learn the answer. In May 2020, Toronto and Google scrapped the project. The decision was taken amid the Coronavirus pandemic—but privacy advocates claimed credit.

> "This is a major victory for the responsible citizens who fought to protect Canada's democracy, civil and digital rights," said one opponent of the project. "Toronto will go down in history as one of the more disturbing planned experiments in surveillance capitalism"[42]—referring to the title of a best-selling business book.

Interconnection brings new privacy and security concerns. Who can learn when you're out of town by watching for changes in power usage patterns? Who can learn when you have company or take a hot shower? By monitoring the power signatures of home devices, a watcher could even see when you start the morning coffee pot or turn on the evening news.

The flow of bits, storage capacity, and processing power needed for analysis all tend to heighten the power disadvantage of individuals against governments and corporations. Privacy serves as a way of taking back some control, a zone of autonomy. In Orwell's imagined London, only O'Brien and other members of the Inner Party could escape the gaze of the telescreen. For now, individuals can employ a mix of mathematical and legal protections to shut out the watching eyes of Big Brother—at least most of the time.

Endnotes

1 Sopan Deb and Natasha Singer, "Taylor Swift Keeping An Eye Out For Stalkers," *New York Times*, December 15, 2018, C6, https://www.nytimes.com/2018/12/13/arts/music/taylor-swift-facial-recognition.html

2 George Orwell, *1984* (Signet Classic, 1977), p. 2.

3 Silkie Carlo, "Britain Has More Surveillance Cameras per Person Than Any Country Except China. That's a Massive Risk to Our Free Society," *Time*, May 17, 2019, https://news.yahoo.com/britain-more-surveillance-cameras-per-151641361.html.

4 Lee Rainie, "Americans' Complicated Feelings About Social Media in an Era of Privacy Concerns," Pew Research Center, March 27, 2018, https://www.pewresearch.org/fact-tank/2018/03/27/americans-complicated-feelings-about-social-media-in-an-era-of-privacy-concerns/.

5 Lee Raine, Americans' complicated feelings about social media in an era of privacy concerns, Pew Research Center, March 27, 2018, https://www.pewresearch.org/fact-tank/2018/03/27/americans-complicated-feelings-about-social-media-in-an-era-of-privacy-concerns/.

6 Edward Snowden, *Permanent Record* (Metropolitan Books, 2019).

7 https://www.washingtonpost.com/investigations/us-intelligence-mining-data-from-nine-us-internet-companies-in-broad-secret-program/2013/06/06/3a0c0da8-cebf-11e2-8845-d970ccb04497_story.html.

8 Kevin Bankston, "EFF Analysis of the Provisions of the USA PATRIOT Act," Electronic Frontier Foundation, October 27, 2003, https://www.eff.org/deeplinks/2003/10/eff-analysis-provisions-usa-patriot-act.

9 Glenn Greenwald, "NSA Collecting Phone Records of Millions of Verizon Customers Daily," *The Guardian*, June 6, 2013, https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order.

10 Micah Lee et al., "A Look at the Inner Workings of NSA's XKEYSCORE," *The Intercept*, July 2, 2015, https://theintercept.com/2015/07/02/look-under-hood-xkeyscore/.

11 Tom Bowman, "Why Does the NSA Keep an EGOTISTICALGIRAFFE? It's Top Secret," NPR, November 10, 2003, https://www.npr.org/2013/11/10/244240199/why-does-the-nsa-keep-an-egotisticalgiraffeits-top-secret.

12 David Cole, "We Kill People Based on Metadata," *The New York Review of Books*, May 10, 2014, https://www.nybooks.com/daily/2014/05/10/we-kill-people-based-metadata/.

13 Stephen Farrell and Hannes Tschofenig, "Pervasive Monitoring Is an Attack," Internet Engineering Task Force, RFC 7258, May 2014, https://tools.ietf.org/html/rfc7258.

14 "HTTPS Encryption on the Web," Google Transparency Report, accessed May 18, 2020, https://transparencyreport.google.com/https/overview;%20https:=.

15 *Olmstead v. United States*, 277 U.S. 438 (1928), https://supreme.justia.com/cases/federal/us/277/438/.

16 *Katz v. United States*, 389 U.S. 347 (1967), https://supreme.justia.com/cases/federal/us/389/347/.

17 Eric Sandy, "Supreme Court Case Has Roots in Radio Shack Robberies in Michigan and Ohio," *Detroit Metro Times*, November 28, 2017, https://www.metrotimes.com/news-hits/archives/2017/11/28/supreme-court-case-has-roots-in-radio-shack-robberies-in-michigan-and-ohio.

18 *Carpenter v. United States*, 138 S. Ct. 2206 (2018), https://www.oyez.org/cases/2017/16-402.

19 Ellen Messmer, "Black Hat: Researcher Claims Hack of Processor Used to Secure Xbox 360, Other Products," *Network World,* February 2, 2010, https://www.networkworld.com/article/2243700/black-hat--researcher-claims-hack-of-processor-used-to-secure-xbox-360--other-products.html.

20 *People v. Christmann*, 861 N.W.2d 18 (2015), https://caselaw.findlaw.com/ny-justice-court/1143124.html.

21 Taylor Hatmaker, "California Malls Are Sharing License Plate Tracking Data with an ICE-Linked Database," TechCrunch, July 10, 2018, https://social.techcrunch.com/2018/07/10/alpr-license-plate-recognition-ice-irvine-company/.

22 Thomas Brewster, "Why Strava's Fitness Tracking Should Really Worry You," *Forbes*, January 29, 2018, https://www.forbes.com/sites/thomasbrewster/2018/01/29/strava-fitness-data-location-privacys-care/#46e488aa55c3.

23 James Quarles, "A Letter to the Strava Community," Strava, January 29, 2018 https://blog.strava.com/press/a-letter-to-the-strava-community/.

24 Chris Buckley and Paul Mozur, "How China Uses High-Tech Surveillance to Subdue Minorities," *The New York Times*, May 22, 2019, https://www.nytimes.com/2019/05/22/world/asia/china-surveillance-xinjiang. html.

25 Ellen Nakashima, "FBI Prepares Vast Database of Biometrics," *The Washington Post*, December 22, 2007, http://www.washingtonpost.com/wp-dyn/content/article/2007/12/21/AR2007122102544.html.

26 Kate Conger et al., "San Francisco Bans Facial Recognition Technology," *The New York Times*, May 14, 2019, https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html.

27 Susie Cagle, "This ID Scanner Company Is Collecting Sensitive Data on Millions of Bargoers," Medium, May 29, 2019, https://onezero.medium.com/id-at-the-door-meet-the-security-company-building-an-international-database-of-banned-bar-patrons-7c6d4b236fc3.

28 "The 2019 Federal Reserve Payments Study," Board of Governors of the Federal Reserve System, January 6, 2020, https://www.federalreserve.gov/paymentsystems/2019-December-The-Federal-Reserve-Payments-Study.htm.

29 GAO, U.S. Government Accountability Office, https://www.gao.gov/products/GAO-04-548.

30 "Equifax to Pay $575 Million as Part of Settlement with FTC, CFPB, and States Related to 2017 Data Breach," Federal Trade Commission, July 22, 2019, https://www.ftc.gov/news-events/press-releases/2019/07/equifax-pay-575-million-part-settlement-ftc-cfpb-states-related.

31 "Equifax Data Breach Settlement," Federal Trade Commission, July 11, 2019, https://www.ftc.gov/enforcement/cases-proceedings/refunds/equifax-data-breach-settlement.

32 OPM.GOV, Cybersecurity Resource Center, https://www.opm.gov/cybersecurity/cybersecurity-incidents/

33 Michael Lasalandra, "Panel told releases of med records hurt privacy," *Boston Herald*, March 20, 1997.

34 Manos Antonakakis et al., "Understanding the Mirai Botnet," *Proceedings of 26th USENIX Security Symposium*, April 16, 2017; "Mirai IoT Botnet Co-Authors Plead Guilty," Krebs on Security, December 13, 2017, https://krebsonsecurity.com/2017/12/mirai-iot-botnet-co-authors-plead-guilty/.

35 Scott Hilton, "Dyn Analysis Summary of Friday October 21 Attack," Oracle, 2016. http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/.

36 C. J. Hughes, "The Latest in Apartment Technology: Fridge Cams and Robotic Valets," *The New York Times*, December 15, 2017, https://www.nytimes.com/2017/12/15/realestate/apartment-technology-fridge-cams-robotic-valets.html.

37 Patrick Olsen, "Tesla Model 3 Falls Short of a CR Recommendation," *Consumer Reports*, May 30, 2018, https://www.consumerreports.org/hybrids-evs/tesla-model-3-review-falls-short-of-consumer-reports-recommendation/.

38 Patrick Olsen, "Tesla Model 3 Gets CR Recommendation After Braking Update," *Consumer Reports*, May 30, 2018, https://www.consumerreports.org/car-safety/tesla-model-3-gets-cr-recommendation-after-braking-update/.

39 Andrew Liptak, "Tesla Extended the Range of Some Florida Vehicles for Drivers to Escape Hurricane Irma," *The Verge*, September 10, 2017, https://www.theverge.com/2017/9/10/16283330/tesla-hurricane-irma-update-florida-extend-range-model-s-x-60-60d.

40 Bruce Schneier, *Click Here to Kill Everybody: Security and Survival in a Hyper-Connected World* (WW Norton & Company, 2018).

41 Alfred Ng, "Tenants Win as Settlement Orders Landlords Give Physical Keys over Smart Locks," CNET, May 7, 2019, https://www.cnet.com/news/tenants-win-rights-to-physical-keys-over-smart-locks-from-landlords/.

42 Rob Gillies, Google Affiliate scraps plan for Toronto Smart City Project, *US News and World Report*, May 7, 2020, https://www.usnews.com/news/business/articles/2020-05-07/google-affiliate-scraps-plan-for-toronto-smart-city-project and Shoshana Zuboff, *The Age of Surveillance Capitalism,* PublicAffairs, 2019.

CHAPTER 3

Who Owns Your Privacy?

The Commercialization of Personal Data

What Kind of Vegetable Are You?

It didn't raise eyebrows when Aleksandr Kogan offered "This is Your Digital Life" as a Facebook quiz app. Quiz apps are a staple of Facebook marketing, enticing users to participate and then harvesting marketing data. These apps—which are enticing, seductive, and highly effective—have spawned an entire subindustry of quiz-marketing tools and specialists.

About 270,000 Facebook users installed Kogan's app and took its personality test, in the process giving the app access to their contacts to invite them to follow suit. Kogan's ostensible motivation was academic research—studying how emojis are used to convey emotion. But what he did with all the data he collected was quite different. Through Kogan's app, the firm Cambridge Analytica harvested data on more than 50 million people. Cambridge Analytica used that information to help presidential candidate Donald Trump's campaign target audiences for digital advertising and fundraising, model voter turnout, identify markets to air television ads, and even plan Trump's travel. Cambridge Analytica asserted that its "psychographic profiles" helped to identify likely voters and the kinds of messages that would sway them to vote Trump.[1]

But how did a quarter million people downloading an app turn into data spillage from 50 million? Through the porous privacy model of Facebook apps. Each of the 270,000 users who installed the app was connected to an average of 200 friends. "This is Your Digital Life" based its assessment not so much on the quiz as on the history of pages "liked." The quiz was a pretext to

obtain access to users' likes and those of their contacts. Facebook permitted that data shoveling in 2015—although it says Kogan violated the program's terms by sharing profile data with Cambridge Analytica.

Your privacy is not your own. Even if you rejected "This is Your Digital Life," any of your friends—or the apps they installed—could have compromised *your* data. This has parallels in the non-digital world as well, of course. (Consider the old saying "Two people can keep a secret if one of them is dead.") But offline, you may have better intuitions about it. You know not to share a story with the gossipy neighbor until you're ready to be asked questions by strangers in the supermarket. Online, it took a long time for Facebook's privacy settings to gain simple audience controls, and not until after the Cambridge Analytica scandal did the social network stop allowing apps to traverse the social graph, slurping up the network of friend connections.

Leave Me Alone

More than a century ago, two lawyers raised the alarm about the impact technology and the media were having on personal privacy:

> Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life; and numerous mechanical devices threaten to make good the prediction that "what is whispered in the closet shall be proclaimed from the house-tops."

This statement is from the seminal *Harvard Law Review* article on privacy published in 1890 by Boston attorney Samuel Warren and his law partner, Louis Brandeis, later to be a justice of the U.S. Supreme Court (where, as we saw, he dissented in defense of privacy in *Olmstead v. U.S.*).[2] Warren and Brandeis went on to say,

> Gossip is no longer the resource of the idle and of the vicious, but has become a trade, which is pursued with industry as well as effrontery. To satisfy a prurient taste the details of sexual relations are spread broadcast in the columns of the daily papers. To occupy the indolent, column upon column is filled with idle gossip, which can only be procured by intrusion upon the domestic circle.

New technologies made this garbage easy to produce, and then the supply created the demand. And those candid photographs and gossip columns were not merely tasteless; they were bad. Sounding like modern critics of mindless

reality TV, Warren and Brandeis raged that society was going to hell in a handbasket because of all that stuff that was being spread about:

> Even gossip apparently harmless, when widely and persistently circulated, is potent for evil. It both belittles and perverts. It belittles by inverting the relative importance of things, thus dwarfing the thoughts and aspirations of a people. When personal gossip attains the dignity of print, and crowds the space available for matters of real interest to the community, what wonder that the ignorant and thoughtless mistake its relative importance. Easy of comprehension, appealing to that weak side of human nature which is never wholly cast down by the misfortunes and frailties of our neighbors, no one can be surprised that it usurps the place of interest in brains capable of other things. Triviality destroys at once robustness of thought and delicacy of feeling. No enthusiasm can flourish, no generous impulse can survive under its blighting influence.

The problem Warren and Brandeis perceived was that it was hard to say just why such invasions of privacy should be unlawful. In individual cases, you could say something sensible, but the individual legal decisions were not part of a general regime. The courts had certainly applied legal sanctions for defamation—publishing malicious gossip that was false—but then what about malicious gossip that was true? Other courts had imposed penalties for publishing an individual's private letters—but on the basis of property law, just as though the individual's horse had been stolen rather than the words in his letters. That did not seem to be the right analogy either. No, they concluded, such rationales didn't get to the nub. When something private is published about you, something has been taken from you, you are a victim of theft—but the thing stolen from you is part of your identity as a person. In fact, privacy was a right, they said, a "general right of the individual to be let alone." That right had long been in the background of court decisions, but the new technologies had brought this matter to a head. In articulating this new right, Warren and Brandeis were, they asserted, grounding it in the principle of "inviolate personhood," the sanctity of individual identity.

Privacy and Freedom

The Warren–Brandeis articulation of privacy as a right to be left alone was influential, but it was never really complete. Throughout the twentieth century, there were simply too many good reasons for *not* leaving people alone, and there were too many ways in which people *preferred* not to be left alone. And in the United States, First Amendment rights stood in tension with privacy rights. As a general rule, the government cannot stop me from saying

anything truthful. In particular, it usually cannot stop me from saying what I lawfully discover about your private affairs. Yet the Warren–Brandeis definition worked well enough for a long time because, as Robert Fano put it, "The pace of technological progress was for a long time sufficiently slow as to enable society to learn pragmatically how to exploit new technology and prevent its abuse, with society maintaining its equilibrium most of the time."[3] By the late 1950s, the emerging electronic technologies, both computers and communication, had destroyed that balance. Society could no longer adjust pragmatically because surveillance technologies were developing too quickly.

The result was a landmark study of privacy by the Association of the Bar of the City of New York, which culminated in the publication, in 1967, of a book by Alan Westin, titled *Privacy and Freedom*.[4] (Fano was reviewing Westin's book when he painted the picture of social disequilibrium caused by rapid technological change.) Westin proposed a crucial shift of focus.

Brandeis and Warren had seen a loss of privacy as a form of personal injury, which might be so severe as to cause "mental pain and distress, far greater than could be inflicted by mere bodily injury." Individuals had to take responsibility for protecting themselves. "Each man is responsible for his own acts and omissions only." But the law had to provide the weapons with which to resist invasions of privacy.

Westin recognized that the Brandeis–Warren formulation was too absolute, in the face of the speech rights of other individuals and society's legitimate data-gathering practices. Protection might come not from protective shields but from control over the uses to which personal information could be put. "Privacy," wrote Westin, "is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others." Westin proposed:

> ...what is needed is a structured and rational weighing process, with definite criteria that public and private authorities can apply in comparing the claim for disclosure or surveillance through new devices with the claim to privacy. The following are suggested as the basic steps of such a process: measuring the seriousness of the need to conduct surveillance; deciding whether there are alternative methods to meet the need; deciding what degree of reliability will be required of the surveillance instrument; determining whether true consent to surveillance has been given; and measuring the capacity for limitation and control of the surveillance if it is allowed.[5]

So even if there were a legitimate reason why the government, or some other party, might know something about you, your right to privacy might limit what the knowing party could do with that information.

This more nuanced understanding of privacy emerged from the important social roles that privacy plays. Privacy is not, as Warren and Brandeis had it, the right to be isolated from society; privacy is a right that makes society work.

Fano mentioned three social roles of privacy. First, "the right to maintain the privacy of one's personality can be regarded as part of the right of self-preservation"—the right to keep your adolescent misjudgments and personal conflicts to yourself, as long as they are of no lasting significance to your ultimate position in society. Second, privacy is the way society allows deviations from prevailing social norms, given that no one set of social norms is universally and permanently satisfactory—and, indeed, given that social progress requires social experimentation. And third, privacy is essential to the development of independent thought; it enables some decoupling of the individual from society so that thoughts can be shared in limited circles and rehearsed before public exposure.

Philosopher Helen Nissenbaum similarly grounds privacy in social being, describing privacy as "contextual integrity."[6] Privacy depends on a match between data flows and the expectations and norms of the setting in which information was generated and shared. When Facebook invites you to friend your therapist or a fellow patient, that's a context violation. Online spaces offer the opportunity to multiply contexts: You can be one persona on your Instagram feed and another in the classroom. But online spaces also threaten context collapse, as Stacy Snyder found way back in the days of Myspace, when her photograph captioned "drunken pirate" on what she thought was a merely social post cost her a teaching degree.[7]

Privacy is the way society allows deviations from prevailing social norms, given that social progress requires social experimentation.

The explosive growth in digital technologies has radically altered our expectations about what will be private and shifted our thinking about what *should* be private. It has made privacy violations easier and potentially more numerous. Indeed, it is remarkable that we no longer blink at intrusions that a decade ago would have seemed shocking. Unlike with the story of secrecy, there was no single technological event that caused the change, no privacy-shattering breakthrough—only a steady advance on several technological fronts that ultimately passed a tipping point.

Sensor devices got cheaper, better, and smaller. Tiny cameras, GPS units, and microphones have gone from the stuff of spy museums to the banality of everyday carry. Once they became useful consumer goods, we seemingly worried less about their uses as surveillance devices. Instead of trying to come up with a unifying theory of privacy and its value, we find ourselves piecing together privacy from feelings of discomfort and regret amid the abundance. It's that much harder when we're the ones bringing spies into our own homes and those of our friends, when we trade privacy against conviviality and convenience.

Smile While We Snap!

Big Brother had his legions of cameras, and the City of London has theirs today. But for sheer photographic pervasiveness, nothing beats the cameras in the cell phones in the hands of everyday people. Flying out before the Fourth of July, Helen was asked to switch seats with another woman who wanted to be seated with her boyfriend. She took her seat a row up and struck up a conversation with her new seatmate, unaware that the row behind was filming them as romance. The pair she had helped were tweeting the flight, hashtagged #PlaneBae, and the story soon made the rounds of television morning shows. Innocent fun, it might seem, but not for Helen, who stated (through lawyers),

> Without my knowledge or consent, other passengers photographed me and recorded my conversation with a seatmate. They posted images and recordings to social media, and speculated unfairly about my private conduct.
>
> Since then, my personal information has been widely distributed online. Strangers publicly discussed my private life based on patently false information.
>
> I have been doxxed, shamed, insulted and harassed. Voyeurs have come looking for me online and in the real world.[8]

The massive dissemination of cheap cameras coupled with universal access to the Web enables a kind of vigilante justice—a ubiquitous Little-Brotherism, in which we can all be detectives, judges, and corrections officers. Bloggers can bring global attention to ordinary citizens.

For every lens aimed deliberately, there are also scores more watching unattended: public and private observation and surveillance. Main Street is lined with security cameras peeping from store windows and police surveillance cameras, some of which even offer public viewing. Leafy Lane may be watching, too, thanks to networks of Ring doorbells and vigilant neighbors in Nextdoor groups. Coupled with automated facial recognition, the wired streets could be building dossiers on us all.

Looking at images on the Web is now a leisure activity that anyone can do at any time, anywhere in the world. Using Google Street View, you can sit in a café in Tajikistan and identify a car that was parked in my driveway when Google's camera came by (perhaps months ago). From Seoul, you can see what's happening right now, updated every few seconds, in Piccadilly Circus or on the strip in Las Vegas. These views were always available to the public, but cameras plus the Web change the meaning of "public."

Some of the intrusions into our privacy come because of the unexpected, unseen side effects of things we do quite voluntarily. While the Fourth Amendment protects us from overreach of government surveillance, there is only patchwork legal consideration of private information gathering in the United States. Companies routinely gather and infer information about individuals and use it to customize product offerings and advertisements. As the saying goes, if you're not paying, you're the product.

Footprints and Fingerprints

As we do our daily business and lead our private lives, we leave footprints and fingerprints. We can see our footprints in mud on the floor and in the sand and snow outdoors. We would not be surprised that anyone who went to the trouble to match our shoes to our footprints could determine, or guess, where we had been. Fingerprints are different. It doesn't even occur to us that we are leaving them as we open doors and drink out of tumblers. Those who have guilty consciences may think about fingerprints and worry about where they are leaving them, but the rest of us don't.

THE UNWANTED GAZE

The Unwanted Gaze by Jeffrey Rosen (Vintage, 2000) details many ways in which the legal system has contributed to our loss of privacy.

In the digital world, we all leave both electronic footprints and electronic fingerprints—data trails we leave intentionally, and data trails of which we are unaware or unconscious. The identifying data may be useful for forensic purposes. Because most of us don't consider ourselves criminals, however, we tend not to worry about that. What we don't think about is that the various small smudges we leave on the digital landscape may be useful to someone else—someone who wants to use the data we left behind to make money or to get something from us. It is therefore important to understand how and where we leave these digital footprints and fingerprints.

Tracing Paper

If I send an email or download a web page, it should come as no surprise that I've left some digital footprints. After all, the bits have to get to me, so some part of the system knows where I am. In the old days, if I wanted to be anonymous, I could write a note, but my handwriting might be recognizable, and I might leave fingerprints (the oily kind) on the paper. I might have typed, but Perry Mason regularly solved crimes by matching a typewritten note with the unique signature of the suspect's typewriter. More fingerprints.

So, today I would laser print the letter and wear gloves. But even that might not suffice to disguise me. Researchers at Purdue have developed techniques for matching laser-printed output to a particular printer.[9] They analyze printed sheets and detect unique characteristics of each manufacturer and each individual printer—fingerprints that can be used, like the smudges of old typewriter hammers, to match output with source. It may be unnecessary to put the microscope on individual letters to identify what printer produced a page.

The Electronic Frontier Foundation has demonstrated that many color printers nearly invisibly encode the printer serial number, date, and time on every page they print (see Figure 3.1). Therefore, when you print a report, you should not assume that no one can tell who printed it.

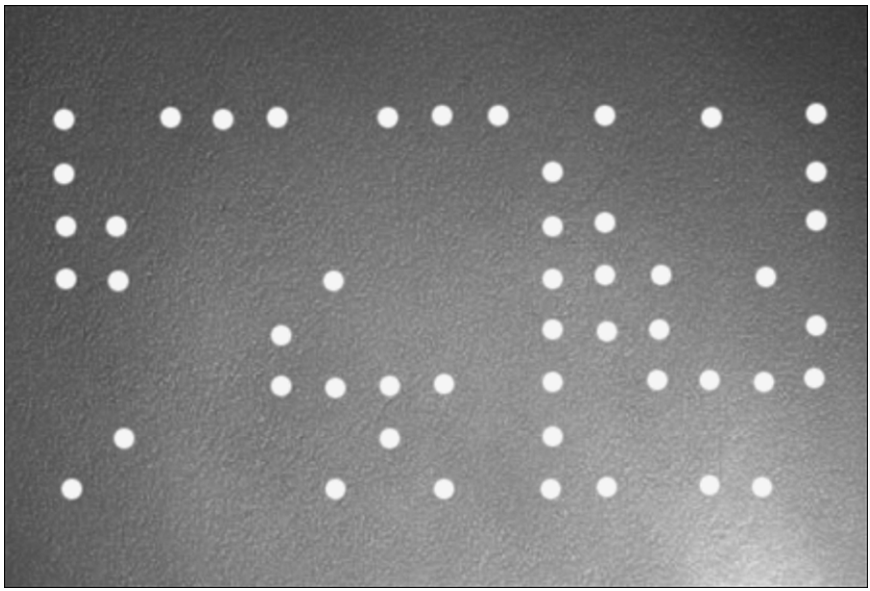

Source: Electronic Frontier Foundation, http://w2.eff.org/Privacy/printers/docucolor/

FIGURE 3.1 Fingerprint left by a Xerox DocuColor 12 color laser printer. The dots are very hard to see with the naked eye; the photograph was taken under blue light. The dot pattern encodes the date (2005-05-21), the time (12:50), and the serial number of the printer (21052857).

There was a sensible rationale behind this technology. The government wanted to make sure that office printers could not be used to turn out sets of hundred-dollar bills. The technology that was intended to frustrate counterfeiters makes it possible to trace every page printed on color laser printers back to the source. Useful technologies often have unintended consequences.

Many people, for perfectly legal and valid reasons, would like to protect their anonymity. They might be whistleblowers or dissidents. Perhaps they are merely railing against injustice in their workplace. Will technologies that undermine anonymity in political discourse also stifle free expression? A measure of anonymity is essential in a healthy democracy—and in the United States, anonymity has been a weapon used to advance free speech since the time of the Revolution. We may regret a complete abandonment of anonymity in favor of communication technologies that leave fingerprints.

The problem is not just the existence of fingerprints but that no one told us that we are creating them.

The problem is not just the existence of fingerprints but that no one told us that we are creating them.

When NSA contractor Reality Winner leaked classified information to The Intercept, she might have thought that sending a paper copy would thwart attempts to trace the leaks.[10] The Intercept had shared the document with NSA to verify its authenticity, and Winner was arrested a few days later. Initial reports speculated that she was traced through printer microdots, but the truth appears to have been even more mundane: NSA logs showed that only six accounts, including Winner's, had accessed the document, and Winner had used a personal account to contact The Intercept shortly beforehand.[11]

Advertising

If you ride the T in Boston, you'll see lots of advertisements for college and graduate programs. They all have phone numbers and URLs, and many direct you places like college.edu/recruiting/redline. That web address isn't saying the college has a special program on the Red Line, but it does have a special *advertising* program there. The "redline" at the end of the URL lets the college know that you were referred there by its subway ad. It might use that to direct you to the particular programs advertised on the poster and to track the effectiveness of this ad campaign.

Ads on the Web use the referring page as just one of many signifiers; others are less visible than the URL decoration visible on the subway poster. When you follow a link to open a web page in your browser, that click kicks off a series of events that starts with an electronic request for the web page and a request for any cookies the site may have set previously. All but the simplest of pages will then trigger requests for more subresources: images, fonts, scripts to make the page dynamic. A commercial site may have dozens of advertisements and tracking pixels, or "web bugs"—invisible elements that make your computer call out to yet another source for the purpose of tracking your activity.

How Sites Know Who You Are (An Incomplete List)

1. **You tell them.** Log in to Gmail, Amazon, or eBay, and you are letting them know exactly who you are.
2. **They've left cookies on one of your previous visits.** A *cookie* is a small text file stored on your local hard drive that contains information that a particular website wants to have available during your current session (about your shopping cart, for example) or from one session to the next. Cookies give sites persistent information for tracking and personalization. Your browser has a command for showing cookies; if you use it, you may be surprised how many websites have left them!
3. **They have your IP address.** The web server has to know where you are so that it can ship its web pages to you. Your IP address is a number like 66.82.9.88 that locates your computer in the Internet. That address may change from one day to the next. But in a residential setting, your Internet service provider (ISP; typically your phone or cable company) knows who was assigned each IP address at any time. Those records are often subpoenaed in court cases.
4. **You look like someone they already recognize.** Users who log in to Facebook often share a lot of detail about their lives and networks: friends and family connections, favorite bands and restaurants, political leanings—and that's just things they deliberately connect or "like." Facebook also creates shadow audiences, matching people on whom they have little information with others they already know, who share these characteristics.
5. **They've fingerprinted your browser and linked it to profiles from previous visits.** Websites can access lots of seemingly innocent details about your browser (which type, version, graphics encoding, language, and much more). These tend to remain fairly static, and often will uniquely identify a particular browser instance. This technique is simple, and remarkably accurate and effective.

If you are curious about who is using a particular IP address, you can check the American Registry of Internet Numbers (www.arin.net). Services such as whatismyip.com, whatismyip.org, and ipchicken.com also allow you to check your own IP address. And www.whois.net allows you to check who owns a domain name such as harvard.com—which turns out to be the Harvard Bookstore, a privately owned bookstore right across the street from the university.

Unfortunately, IP address information won't reveal who is sending you spam, since spammers routinely forge the source of email they send you. In addition,

between the time you request a web page and its ads are displayed in your browser, there's often a real-time auction, in which your eyeballs (or at least the ad spaces in the web page your browser is about to display) are sold to the highest bidder. Ad networks collect the information from tracking pixels and page context to determine what ads to offer and how much to bid to place them in these auctions.

Why are these shoes following me around? Maybe you saw them on Instagram, tagged them on Pinterest, or searched for a new pair of sneakers on your favorite retailer's website. Maybe you even put them into a shopping cart before deciding they weren't in your budget at this time. Now, you can't seem to escape the shoes: whether you're reading the news or Facebooking with friends, there are the shoes, stalking you from the ad banners, urging you to click "buy."

Known in the trade as "retargeting," these ads are some of the products of real-time bidding. The marketer who dropped a tracking cookie in your browser during an earlier browsing session or the shopping visit you cut short is using it to identify you as a shoe-interested shopper and bidding to show you those ads in the hopes of luring you back to purchase. If you clicked through any of the ads, the marketer would register a "conversion" and factor this data further into your profile for future ad opportunities.

Web browsing users haven't taken all of this sitting quietly. The *Economist* calls data "the new oil," and browsers who are unwilling to be seen as gushers download ad blockers. As of early 2020, all of the major web browsers have incorporated tracker-blocking features or announced plans to limit third-party cookies.

Arvind Narayanan and his team at Princeton University have set up a laboratory for web measurement[12] and discovered new techniques for browser tracking. Through web "crawls," they find tracking techniques used in the wild to identify users and reidentify those who think they've cleared all previous interactions. One of the paradoxes of privacy on the Web is that browsers can be fingerprinted by their unique features, including features the user might enable with the goal of securing greater privacy. That means turning on such protections can make the privacy-seeking user stand out. In such cases, privacy depends on the actions of many to provide a crowd in which the privacy-seeking browser can blend. Standardized processes and well-thought-out default settings are necessary to preserve the opportunities for privacy.

Target Knows You're Pregnant

In 2012, as Charles Duhigg reported in the *New York Times*,[13] a man walked into a Minneapolis-area Target store, furiously asking to speak with the manager: "My daughter got this in the mail!" he said. "She's still in high school, and you're sending her coupons for baby clothes and cribs? Are you trying to encourage her to get pregnant?"

The store manager apologized to the Minneapolis man for their apparent mistake, but he returned a few weeks later with an apology of his own: His daughter was, in fact, pregnant. The store's predictive models had recognized the young woman's pregnancy even before her father had. Target's models didn't have access to her private information. They had the power of analytical tools and readily available data.

Like many other stores with loyalty cards or user accounts, Target built statistical models of shopper behavior to predict hot products for inventory and pricing and to make recommendations. Target correlated shopper purchase history based on an internal guest ID and purchased external data to supplement its logs. From those records, the company's statistician could derive patterns, noticing, for instance, that women in the second trimester of pregnancy would often purchase unscented moisturizing lotions and supplements. After watching this pattern play out many times, the store could anticipate future purchases of baby clothes and diapers from the earlier unscented lotion—and advertise to the mother-to-be at a time when her shopping habits were in flux—responding to a signal she didn't even know she was sending.

How can we solve a privacy problem that results from many developments, no one of which is really a problem in itself?

You Pay for the Mic, We'll Just Listen In

Planting tiny microphones where they might pick up conversations of underworld figures used to be risky work for federal authorities. There are much safer alternatives, now that most people carry their own radio-equipped microphones with them all the time or invite Alexa, Siri, Cortana, or Google into their homes.

Many cell phones can be reprogrammed remotely so that the microphone is always on and the phone is transmitting, even if you think you have powered it off. The FBI used this technique in 2004 to listen to John Tomero's conversations with other members of his organized crime family. A federal court ruled that this "roving bug," installed after due authorization, constituted a legal form of wiretapping. Tomero could have prevented it by removing the battery, and now some nervous business executives routinely do exactly that.

The microphone in a General Motors car equipped with the OnStar system can also be activated remotely, a feature that can save lives when OnStar operators contact the driver after receiving a crash signal. OnStar warns, "OnStar will cooperate with official court orders regarding criminal investigations from law enforcement and other agencies," and indeed, the FBI has used this method to eavesdrop on conversations held inside cars. In one case, a federal court ruled against this way of collecting evidence—but not on privacy grounds. The roving bug disabled the normal operation of OnStar, and

the court simply thought that the FBI had interfered with the vehicle owner's contractual right to chat with the OnStar operators!

Danielle, an Amazon Echo customer in Portland, Oregon, was alarmed by a call from one of her husband's colleagues, who said, "Unplug your Alexa devices right now; you're being hacked."[14] The gadget, which was supposed to record only when triggered by the wake word "Alexa," must have heard both that and a "send message" command in Danielle's conversation. Her chat about hardwood floors turned into a voice message to a business acquaintance. A freak occurrence, perhaps, but one that may be repeated as we invite tiny networked recorders into more corners of our lives. German authorities banned "My Friend Cayla,"[15] a talking doll, over concerns about its spying and data-collecting abilities. To engage in conversation with children, Cayla uploaded the sounds she heard over the Internet. German parents were told to destroy the "illegal espionage apparatus." Meanwhile, here in the United States, your smart TV may be watching your viewing habits to tailor advertising. Vizio's CTO told the Consumer Electronics Show that TVs would cost more if it weren't for this revenue stream.[16]

Venmo: It All Adds Up

Earlier we discussed the tracking that credit cards enable in credit reporting bureaus and data analysis firms. Newer payment technologies bring the reporting directly to you. Venmo lets you send someone money or split a bill by entering the person's phone number. It's so easy that as you send money to friends or roommates using the Venmo app, you might not notice that these payment transactions are public, including any memo you write along with the payment. A researcher who found the feed correlated just a few of the threads among millions of transactions into "Venmo stories":[17] a student's fast food habit, a cannabis vendor's sales, a budding relationship? You might not mind sharing your passion for elote (seasoned corn) but might feel differently about recreational marijuana purchases, even in states where those are legal. The researcher, Hang Do Thi Duc, anonymized the details but notes that the feed, which includes everything except dollar values, remained accessible to any visitor to Venmo's public API. (Every page of the site Duc developed, publicbydefault.fyi, encourages Venmo users to change their privacy settings from the default to make transactions private between sender and recipient.)

DNA: The Ultimate Digital Fingerprint

In April 2018, the state of California arraigned Joseph James DeAngelo on a series of decades-old murder and rape charges. The Golden State Killer had been a cold case until an investigator uploaded DNA from a crime scene to a

public genealogy website, GEDmatch. The investigator created a fake profile for the unknown person whose recovered DNA he uploaded. After GEDmatch compared this person's DNA against its existing database to identify partial genetic matches, it showed profiles of people who were likely distant relatives of the suspected killer. Those names led to family trees and to genealogy that could be traced further through census records, obituaries, gravesites, and commercial and law enforcement databases. After these searches put a name to their suspect, investigators confirmed their suspicions by tracking him down and obtaining another DNA sample, from skin cells he left on the car door when he parked in a Hobby Lobby parking lot. That DNA matched the original crime scene samples.[18]

DeAngelo had not posted to the ancestry site, but because a parent passes roughly half of his or her genes to a child (notwithstanding a few mutations along the way), much of DeAngelo's genetic record could be read or revealed by relatives. If your family members explore their genetic profiles and family trees on GEDmatch, they are also exposing information about traits you might share. Your privacy can be invaded through no actions of your own. While the Genetic Information Nondiscrimination Act prohibits employers or health insurers from discriminating based on DNA, the law doesn't restrict numerous other ways DNA can be used.

The Golden State Killer case started a boom in DNA forensic genealogy. By the end of 2018, more than a dozen violent criminals and perpetrators of sexual assault had been identified through GEDmatch. But the site also heard privacy alarm and changed its terms of service to prohibit law enforcement matching of DNA profiles unless users opted in for their own records.

Fair Information Practice Principles

An earlier information revolution, set in rooms full of disk drives that sprouted in government and corporate buildings in the 1960s, set off a round of soul searching about the operational significance of privacy rights. What, in practice, should those holding a big data bank think about when collecting the data, handling it, and giving it to others?

In 1973, the Department of Health, Education, and Welfare issued "Fair Information Practice Principles" (FIPP), as follows:

> **Openness.** There must be no personal data record-keeping systems whose very existence is secret.
>
> **Disclosure.** There must be a way for a person to find out what information about the person is in a record and how it is used.

Secondary use. There must be a way for a person to prevent information about the person that was obtained for one purpose from being used or made available for other purposes without the person's consent.

Correction. There must be a way for a person to correct or amend a record of identifiable information about the person.

Security. Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for its intended use and must take precautions to prevent misuses of the data.

These principles were proposed for U.S. medical data but were never adopted. Nevertheless, they have been the foundation for many corporate privacy policies. Variations on these principles were codified in international trade agreements by the Organisation for Economic Co-operation and Development (OECD) in 1980 and within the European Union (EU) in 1995. In the United States, echoes of these principles can be found in some state laws, but federal laws generally treat privacy on a case-by-case, or "sectorial," basis. The 1974 Privacy Act applies to interagency data transfers within the federal government but places no limitations on data handling in the private sector. The Fair Credit Reporting Act applies only to consumer credit data but does not apply to medical data. The Video Privacy Act applies only to videotape rentals but not to on-demand movie downloads, which did not exist when the act was passed. Finally, few federal or state laws apply to the huge data banks in the file cabinets and computer systems of cities and towns. American government is decentralized, and authority over government data is decentralized as well.

The United States is not lacking in privacy laws. But privacy has been legislated inconsistently and confusingly and in terms dependent on technological contingencies. There is no national consensus on what should be protected and how protections should be enforced. Without a more deeply informed collective judgment on the benefits and costs of privacy, the current legislative hodgepodge may well get worse in the United States.

The discrepancy between American and European data privacy standards threatened U.S. involvement in international trade because an EU directive would prohibit data transfers to nations, such as the United States, that do not meet the European "adequacy" standard for privacy protection. In 2000 the European Commission created a "safe harbor" for American businesses with multinational operations, but the European Court of Justice declared it inadequate to protect the rights of European data subjects. In 2016, the FTC developed an alternative, Privacy Shield, with a salient enforcement difference: "While joining the Privacy Shield Framework will be voluntary, once an eligible company makes the public commitment to comply with the Framework's requirements, the commitment will become enforceable under U.S. law."[19]

In 2020, The Court of Justice of the European Union (CJEU) ruled that even Privacy Shield was inadequate, because European citizens' data in the United States would be subject to U.S. government surveillance.[20]

Privacy as a Basic Right

Browse the Web on a visit to Europe, and you may notice a profusion of pop-ups and banners. Every site, it seems, wants you to consent to the use of cookies and the "processing of your data," assertedly to improve your browsing experience. While European law takes a stronger view of personal privacy as a fundamental right, European advertisers are just as eager to gather personal data as those in the United States. These banners are the means of asking for "consent to data processing," as the E-Privacy Directive required.

In 2018, the General Data Protection Regulation (GDPR) established specific individual rights in personal data and obliged businesses to give individuals ("data subjects") the ability to control the use of that data. Those who collect or process personal data must be able to justify the privacy intrusion based on consent or another "legitimate purpose"; for example, an email provider needs the email addresses of your contacts in order to send emails to their destination, but it doesn't need their home addresses. Individuals even have the right to withdraw consent, demanding that providers erase the data collected about them. Because the GDPR asserts extraterritorial reach, applying to European citizens wherever they are physically located, many providers outside Europe have also adopted cookie-consent requests and adapted their data handling to be able to respond to data deletion requests.

Despite the paper promise of European law, as of 2020, enforcement has been limited. Only one major fine has been issued, against Google for 50 million euros (roughly $54 million), or about one-tenth of what Google generates in a single day's ad sales. Without investigation of the hundreds of complaints raised by citizens to their national data protection authorities, it is difficult to say whether Europeans have more privacy online or just more pop-ups to click through.

It is, unfortunately, too easy to debate whether the European omnibus approach is more principled than the U.S. piecemeal approach, when the real question is whether either approach accomplishes what we want it to achieve. The Privacy Act of 1974 assured us that obscure statements would be buried deep in the Federal Register, providing the required official notice about massive governmental data collection plans; it was better than nothing but provided "openness" only in a narrow and technical sense. Most large corporations doing business with the public have privacy notices, and virtually no one reads them. Only 0.3% of Yahoo! users read its privacy notice in 2002, for example. In the midst of massive negative publicity that year, when Yahoo! changed its privacy policy to allow advertising messages, the number of users

who accessed the privacy policy rose only to 1%. None of the many U.S. privacy laws prevented the warrantless wiretapping program instituted by the Bush administration, nor the cooperation with it by major U.S. telecommunications companies.

Indeed, cooperation between the federal government and private industry seems more essential than ever before for gathering information about drug trafficking and international terrorism—because of yet another technological development. Twenty years ago, most long-distance telephone calls spent at least part of their time in the air, traveling by radio waves between microwave antenna towers or between the ground and a communication satellite. Government eavesdroppers could simply listen in. Now many phone calls travel through fiber-optic cables instead, and the government is tapping this privately owned infrastructure.

High privacy standards have a cost. They can limit the public usefulness of data. Public alarm about the release of personal medical information has led to major legislative remedies. The Health Insurance Portability and Accountability Act (HIPAA) was intended both to encourage the use of electronic data interchange for health information and to impose severe penalties for the disclosure of "protected health information," a very broad category including not just medical histories but, for example, medical payments. The bill mandates the removal of anything that could be used to reconnect medical records to their source. HIPAA is fraught with problems in an environment of ubiquitous data and powerful computing. Connecting the dots by assembling disparate data sources makes it

Ever Read Those "I Agree" Documents?

Companies can do almost anything they want with your information, as long as you agree. It seems hard to argue with this principle, but the deck can be stacked against the consumer who is "agreeing" to the company's terms. Sears Holding Corporation (SHC), the parent of Sears, Roebuck and Kmart, gave consumers an opportunity to join "My Sears Holding Community," which the company describes as "something new, something different…a dynamic and highly interactive online community…where your voice is heard and your opinion matters." When you went online to sign up, the terms appeared in a window on the screen.

The scroll box held only 10 lines of text, and the agreement was 54 boxfuls long. Deep in the terms was a detail: You were allowing Sears to install software on your PC that "monitors all of the Internet behavior that occurs on the computer…, including…filling a shopping basket, completing an application form, or checking your…personal financial or health information." So your computer might send your credit history and AIDS test results to SHC, and you had said that was fine!

extremely difficult to achieve the level of anonymity that HIPAA sought to guarantee. But help is available, for a price, from a whole industry of HIPAA compliance advisors. If you search for HIPAA online, you will likely see advertisements for services that will help you protect your data and also keep you out of jail.

At the same time as HIPAA and other privacy laws have safeguarded our personal information, they are making medical research costly and sometimes impossible to conduct. It is likely that classic studies such as the Framingham Heart Study, on which much public policy about heart disease was founded, could not be repeated in today's environment of strengthened privacy rules. Dr. Roberta Ness, president of the American College of Epidemiology, reported that "there is a perception that HIPAA may even be having a negative effect on public health surveillance practices."[21]

The five FIPP principles, and the spirit of transparency and personal control that lay behind them, have doubtless led to better privacy practices. But they have been overwhelmed by the digital explosion, along with the insecurity of the world and all the social and cultural changes that have occurred in daily life. Fred H. Cate, a privacy scholar at Indiana University, characterizes the FIPP principles as almost a complete bust:

> Modern privacy law is often expensive, bureaucratic, burdensome, and offers surprisingly little protection for privacy. It has substituted individual control of information, which it in fact rarely achieves, for privacy protection. In a world rapidly becoming more global through information technologies, multinational commerce, and rapid travel, data protection laws have grown more fractured and protectionist. Those laws have become unmoored from their principled basis, and the principles on which they are based have become so varied and procedural, that our continued intonation of the FIPP mantra no longer obscures the fact that this emperor indeed has few if any clothes left.[22]

Only sects such as the Amish still live without electricity. It is almost that unusual to live without Internet connectivity, with all the fingerprints it leaves of your daily searches and logins and downloads. Even the old "over-the-air" TV is rapidly disappearing in favor of digital communications.[23]

Digital TV brings the advantages of video on demand, but with a steep privacy cost. Your television service provider records everything you watch, and when. It is so attractive to be able to watch what we want when we want to watch it that we don't miss either the inconvenience or the anonymity of the days when all the TV stations washed your house with their airwaves. You couldn't pick the broadcast times, but at least no one knew which waves you were grabbing out of the air.

Privacy as a Right to Control Information

Privacy is complex and under attack from our peers, our own devices, and governments and corporate marketers. The bits are everywhere; there is simply no locking them down, and no one really wants to do that anymore. The meaning of privacy has changed, and we do not have a good way of describing it. It is not the right to be left alone, because not even the most extreme measures will disconnect our digital selves from the rest of the world. It is not the right to keep our private information to ourselves because the billions of atomic factoids don't lend themselves to being simply and uniquely classified as either private or public.

The bits are everywhere; there is simply no locking them down, and no one really wants to do that anymore.

Which would we prefer: the new world with digital fingerprints everywhere and the constant awareness that we are being tracked, or the old world with few digital footprints and a stronger sense of security from prying eyes? And what is the point of even asking the question when the world cannot be restored to its old information lockdown?

In a world that has moved beyond the old notion of privacy as a wall around the individual, we could instead regulate those who would inappropriately *use* information about us. If I post a YouTube video of myself dancing in the nude, I should expect to suffer some personal consequences. Ultimately, as Warren and Brandeis said, individuals have to take responsibility for their actions. But society has drawn lines in the past around which facts are relevant to certain decisions and which are not. Perhaps, the border of privacy having become so porous, the border of relevancy could be stronger. As Daniel Weitzner explains:

> New privacy laws should emphasize usage restrictions to guard against unfair discrimination based on personal information, even if it's publicly available. For instance, a prospective employer might be able to find a video of a job applicant entering an AIDS clinic or a mosque. Although the individual might have already made such facts public, new privacy protections would preclude the employer from making a hiring decision based on that information and attach real penalties for such abuse.[24]

There can still be principles of accountability for the *misuse* of information. Some ongoing research is outlining a possible new web technology to help ensure that information is used appropriately when it is known. Perhaps automated classification and reasoning tools, developed to help connect the dots in networked information systems, can be retargeted to limit inappropriate

use of networked information. A continuing border war is likely to be waged, however, along an existing free speech front: the line separating my right to tell the truth about you from your right not to have that information used against you. In the realm of privacy, the digital explosion has left matters deeply unsettled.

Paul Ohm posits a "database of ruin":

> Almost every person in the developed world can be linked to at least one fact in a computer database that an adversary could use for blackmail, discrimination, harassment, or financial or identity theft.[25]

We must, through a combination of law, technology, and norms of behavior, find ways to avoid a mutually assured privacy destruction.

A few beacons of hope come from state lawmakers, most notably in California, and a growing culture of privacy among engineers. Some corporate privacy notices are still boilerplate, but others give the impression that privacy is a product feature, designed to add value for users and respond to their needs.

Always On

In *1984,* the pervasive, intrusive technology could be turned off:

> As O'Brien passed the telescreen a thought seemed to strike him. He stopped, turned aside and pressed a switch on the wall. There was a sharp snap. The voice had stopped.
>
> Julia uttered a tiny sound, a sort of squeak of surprise. Even in the midst of his panic, Winston was too much taken aback to be able to hold his tongue.
>
> "You can turn it off!" he said.
>
> "Yes," said O'Brien, "we can turn it off. We have that privilege....Yes, everything is turned off. We are alone."

Sometimes we can still turn it off today—and should. But mostly we don't want to. We don't want to be alone; we want to be connected. We find it convenient to leave it on, to leave our footprints and fingerprints everywhere, so we will be recognized when we come back. We don't want to have to keep retyping our name and address when we return to a website. We like it when the restaurant remembers our name, perhaps because our phone number showed up on caller ID and is linked to our record in their database. We appreciate buying grapes for $1.95/lb instead of $3.49, just by letting the

store know that we bought them. We may want to leave it on for ourselves because we know it is on for criminals. Being watched reminds us that they are watched as well. Being watched also means we are being watched over.

And perhaps we don't care that so much is known about us because that is the way human society used to be: In kinship groups and small settlements, knowing everything about everyone else was a matter of survival. Having it on all the time may resonate with inborn preferences we acquired millennia ago, before urban life made anonymity possible. Still, today, privacy means something very different in a small rural town than it does on the Upper East Side of Manhattan.

We cannot know what the cost will be of having it on all the time. Just as troubling as the threat of authoritarian measures to restrict personal liberty is the threat of voluntary conformity. As Fano astutely observed, privacy allows limited social experimentation—the deviations from social norms that are much riskier to the individual in the glare of public exposure, but which can be, and often have been in the past, the leading edges of progressive social changes. With it always on, we may prefer not to try anything unconventional and stagnate socially by collective inaction.

For the most part, it is too late, realistically, ever to turn it off. We may once have had the privilege of turning it off, but we have that privilege no more. We have to solve our privacy problems another way.

The digital explosion is shattering old assumptions about who knows what. Bits move quickly, cheaply, and in multiple perfect copies. Information that used to be public in principle—for example, records in a courthouse, the price you paid for your house, or stories in a small-town newspaper—is now available to everyone in the world. Information that used to be private and available to almost no one—medical records and personal snapshots, for example—can become equally widespread through carelessness or malice. The norms and business practices and laws of society have not caught up to the change.

Endnotes

1 Matthew Rosenberg, Nicholas Confessore, and Carole Cadwalladr, "Firm That Assisted Trump Exploited Data of Millions," *New York Times*, March 18, 2018: A1, https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html.

2 Samuel A. Warren and Louis D. Brandeis, "The Right to Privacy," *Harvard Law Review* 4, no. 5 (December 15, 1890), https://groups.csail.mit.edu/mac/classes/6.805/articles/privacy/Privacy_brand_warr2.html.

3 Robert Fano, "Review of Alan Westin's *Privacy and Freedom*," *Scientific American* (May 1968): 148–152.

4 Alan F. Westin, *Privacy and Freedom* (Atheneum, 1967).

5 Ibid.

6 Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford Law Books, 2009).

7 "Judge Sides with University Against Student-Teacher with 'Drunken Pirate' Photo," *The Chronicle of Higher Education*, December 4, 2008, https://www.chronicle.com/article/Judge-Sides-With-University/42066.

8 Taylor Lorenz, "Unidentified Plane-Bae Woman's Statement Confirms the Worst," *The Atlantic*, July 13, 2018, https://www.theatlantic.com/technology/archive/2018/07/unidentified-plane-bae-womansstatement-confirms-the-worst/565139/.

9 Emil Venere, "Printer Forensics to Aid Homeland Security, Tracing Counterfeiters," Purdue University, October 12, 2004, https://www.purdue.edu/uns/html4ever/2004/041011.Delp.forensics.html.

10 Michael M. Grynbaum and John Koblin, "Journalists Fear Effects of Arrest," *New York Times*, June 7, 2017: A19, https://www.nytimes.com/2017/06/06/business/media/intercept-reality-winner-russia-trump-leak.html.

11 Jake Swearingen, "Did the Intercept Betray Its NSA Source?", New York Magazine, June 6, 2017. https://nymag.com/intelligencer/2017/06/intercept-nsa-leaker-reality-winner.html.

12 "Web Privacy—Arvind Narayanan," accessed May 18, 2020, https://www.cs.princeton.edu/~arvindn/web-privacy/.

13 Charles Duhugg, "How Companies Learn Your Secrets," *The New York Times*, February 16, 2012, https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.

14 "Amazon Alexa Can Accidentally Record and Share Your Conversations," *Vanity Fair*, May 24, 2018, https://www.vanityfair.com/news/2018/05/yes-amazons-alexa-can-secretly-record-and-share-conversations.

15 Katie Collins, "That Smart Doll Could be a Spy. Parents, Smash!", *CNET*, February 17, 2018, https://www.cnet.com/news/parents-told-to-destroy-connected-dolls-over-hacking-fears/.

16 Ben Gilbert, "There's a simple reason your new smart TV was so affordable: It's collecting and selling your data, and serving you ads," *Business Insider*, April 5, 2019, https://www.businessinsider.com/smart-tv-data-collection-advertising-2019-1.

17 Hang Do Thi Duc, Public By Default, Venmo Stories of 2017, HYPERLINK \h https://publicbydefault.fyi/.

18 Avi Selk, "The ingenious and 'dystopian' DNA technique police used to hunt the 'Golden State Killer' suspect," *Washington Post*, April 28, 2018, https://www.washingtonpost.com/news/true-crime/wp/2018/04/27/golden-state-killer-dna-website-gedmatch-was-used-to-identify-joseph-deangelo-as-suspect-police-say/.

19 "Fact Sheet: Overview of the EU–U.S. Privacy Shield Framework for Interested Participants," U.S. Department of Commerce, July 12, 2016, https://2014-2017.commerce.gov/sites/commerce.gov/files/media/files/2016/fact_sheet-_eu-us_privacy_shield_7-16_sc_cmts.pdf.

20 Court of Justice of the European Union, Press Release No 91/20, Luxembourg, July 16, 2020, Judgment in Case C-311/18: Data Protection Commissioner v Facebook Ireland and Maximillian Schrems. https://curia.europa.eu/jcms/upload/docs/application/pdf/2020-07/cp200091en.pdf.

21 "Privacy Rule Slows Scientific Discovery and Adds Cost To Research, Scientists Say," University of Pittsburgh Schools of the Health Sciences. https://www.sciencedaily.com/releases/2007/11/071113165648.htm. Also: Roberta B. Ness, MD, MPH, "Influence of the HIPAA Privacy Rule on Health Research," JAMA. 2007;298(18):2164-2170. doi:10.1001/jama.298.18.2164

22 Fred H. Cate, "The failure of Fair Information Practice Principles," in Jane K. Winn, ed., *Consumer Protection in the Age of the "Information Economy"* (Ashgate, 2006).

23 https://www.nielsen.com/us/en/insights/report/2019/nielsen-local-watch-report-the-evolving-ota-home/

24 Daniel. J. Weitzner, "Beyond Secrecy: New Privacy Protection Strategies for Open Information Spaces," in *IEEE Internet Computing* 11, no. 5 (September–October 2007): 96–95, https://dl.acm.org/doi/10.1109/MIC.2007.101

25 Paul Ohm, "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization," SSRN Scholarly Paper (Social Science Research Network, August 13, 2009), https://papers.ssrn.com/abstract=1450006.

This page intentionally left blank

CHAPTER 4

Gatekeepers

Who's in Charge Here?

Who Controls the Flow of Bits?

When the Telecommunications Workers Union went on strike against Telus, the leading telecommunications company in western Canada,[1] a discussion about strike-breaking sprung up on a pro-union website operated by a Telus employee. Then suddenly the site became inaccessible to anyone who was using Telus for Internet service. Telus subscribers could get bits originating in places from Afghanistan to Zimbabwe. They could get bits representing symphony orchestras and pornography. But if they wanted to see the discussion about resisting management's efforts to break the strike, they couldn't. Telus had taken the position that because the cables delivering the bits belonged to the company, it could choose to deliver or not deliver bits.

The union was enraged, and the legal experts were confused. It seemed to be clear enough that Telus couldn't cut off phone service to the union or its supporters if it wanted to, but the laws had been written in pre-Internet days. Was Telus within its rights to cut off Internet service in this case? The company noted that it had also blocked telusscabs.ca, which showed pictures of managers and employees who were going to work in spite of the strike. Telus said that it had a responsibility for their safety and felt compelled to protect them. But it turned out that Telus had blocked many more than these 2 sites. The web server hosting these 2 sites also hosted 766 other sites, including an alternative medicine site and a fundraising site for breast cancer research. In successfully blocking the 2 offending sites, Telus had blocked all the others, too.

Telus backtracked after making sure that no threatening material would be posted. But the incident—and others like it—raised questions that even today have no clear and generally accepted answers. Who controls what people can use the Internet to do?

The Open Internet?

Nobody was supposed to be in charge of the Internet. It wasn't even supposed to be the sort of thing that could be owned or controlled. It was meant to be more like a language that different people could use in any way they could imagine, to talk to each other or recite poetry or sing songs. It was to be like the "luminiferous ether," the invisible space-filling substance that physicists used to think must exist because light could not get from one place to another without it. The Internet was supposed to be a medium that could make communication possible by anyone to anyone and from anywhere to anywhere, but control of communication was assumed to be impossible because there would be no place to throttle it. Anyone who wanted to join in a conversation could—just by speaking the language of Internet protocols.

Ask Alex Jones if it has worked out that way. A leading American conspiracy theorist—or, as he considers himself, a "thought criminal against Big Brother"—Jones developed huge followings on YouTube, Facebook, LinkedIn, and other social networking sites. Millions of people followed his every word—and had every kooky rumor he promoted pushed to their mobile phones so they would see it instantly. And then suddenly many sites banned him. Apple stopped providing Jones's app to users. You can still find his website if you look for it, but Pinterest won't suggest it to you.

If this doesn't convince you that the Internet is not a participatory paradise, use the Internet to ask anyone in China what happened on June 4. Whether you use email, text messaging, or Weibo (China's version of Twitter), it is unlikely that your message will reach anyone because mention of June 4 has been censored thoroughly. Ask anyone in Hong Kong, and you won't even have to say which year; the Tiananmen massacre of June 4, 1989, is still vividly remembered more than 30 years on. But on the Internet of the mainland, the billions of users never talk about June 4. Mention it, and the conversation is snuffed out immediately rather than spreading like wildfire. And the government is not the only gatekeeper controlling the electronic sharing of information in China. When protesters in Hong Kong used an app called HKmap.live to organize themselves, the Chinese government became enraged, and Apple removed the app from its App Store. Google responded similarly to a request from the Hong Kong police to take down a game that enabled users to play the roles of protesters.[2]

Or ask Hasan Minhaj, an American comedian whose "Patriot Act" show is distributed via Netflix. His critical comments about Saudi Crown Prince Mohammed bin Salman can be watched almost everywhere in the world—but not where they would have the most meaning, in Saudi Arabia. The Saudi government demanded that the episode be taken down, citing a law that criminalizes the "production, preparation, transmission, or storage of material impinging on public order, religious values, public morals, and privacy, through the information network or computers." Netflix responded by saying it supported artistic freedom—but then took the video down anyway.[3] Artistic freedom for Minhaj proved to be incompatible with the personal freedom of Netflix employees, who might be subject to ten years' imprisonment or more for the vaguely defined crime of transmitting material impinging on public order.

Or ask anyone who sells anything. Google's search engine is used for more than 90% of Internet searches; Microsoft's Bing is in second place, at 3%. If you want to sell widgets and you're not on the first page of results when people search for "widget," it may be difficult to attract attention. "Google is the gatekeeper for the World Wide Web, for the internet as we know it," as lawyer Gary Reback put it. "It is every bit as important today as petroleum was when John D. Rockefeller was monopolizing that."[4] Google defended itself by responding that it is not a monopoly because Amazon search results also influence shopping outcomes. (Google might have noted that this is especially so because Amazon awards an "Amazon Choice" badge to products it favors.) But Google's defense only drives home the point that that the number of gatekeepers is tiny, even if it is modestly greater than one. There may be 1,000 small businesses making and selling widgets, but only the handful that appear on the first search results page are likely to get any Internet business—and they are competing for visibility against much bigger companies.

Or ask anyone in Browning, Montana, if anyone can use the Internet to communicate with anyone. On the Blackfoot Indian Reservation there, only 0.1% of the population has high-speed Internet. The cheapest Internet connectivity of any kind is at 10 Mbps and costs almost $780 annually[5]—in a place with a 35% poverty rate and a median annual household income of less than $22,000.[6] In most American cities, getting bits over the Internet is like turning on the tap to get water, and the same is true all over Finland and Japan. Internet service is abundant in places where it is profitable for private suppliers or where it is provided as a matter of public policy. Neither is true in Browning, where Internet service is all but unavailable.

Whatever the history, the theory, and the potential, the reality is that a few corporations and a few governments exert enormous control over what most people actually see over the Internet and what they can do with the Internet. If these enterprises and institutions don't provide the infrastructure to deliver

your message, or if they otherwise deprioritize delivery of your ad, your news, or your political barbs, you might as well have shouted in the wilderness. Someone might hear what you say, but probably not many people will. That is exactly the opposite of the way the Internet was designed to work. The Internet has evolved from its publicly funded design as an open system with unlimited potential uses to a system in which a few private businesses hold near-monopoly control over each of its major aspects.

This chapter focuses on three kinds of Internet gatekeepers. The first are the controllers of the data pipes through which the bits flow. We'll call these the *links gatekeepers*. The second are the controllers of the tools we use to find things on the Web. We'll call these the *search gatekeepers*. And the third are the controllers of the social connections that are, for many of us, our most important use of the Internet. We'll call these the *social gatekeepers*.

The links gatekeepers control the physical media through which the bits flow, while the search gatekeepers and social gatekeepers control what those bits express—that is, they are content gatekeepers. But such distinctions are not as sharp as they might seem. Links gatekeepers may be able, for example, to censor or to favor certain content over others or certain customers over others. Content gatekeepers may enter the links market if they think it will be to their advantage to resist the near-monopoly control of the links gatekeepers—whether or not consolidating links and content control is in the more broadly construed public interest. Social gatekeepers have added search within their social platforms to undercut the near-monopoly control of the search gatekeepers.

In the United States, all three gatekeeping functions are largely in private hands. In other parts of the world, governments have assumed some of these gatekeeping roles. Familiar debates about private versus public services have played out as part of the Internet almost from its beginning. Competition among private parties drives down costs and improves quality, goes one familiar argument; but consolidation, argue others, results in efficiencies of scale that more than outweigh the negative effects of reduced competition. According to another narrative, the government should provide infrastructure of general benefit to the people, paid for through general taxation rather than private purchase; it should provide the ether, in the same way it provides roadways and mail service, equally to all. But such analogies only invite the question of whether the Internet really is more like the public roadways, on which anyone can drive, or like cable television or movie theaters, which are more accessible in urban areas than rural and not available at all to people who are unwilling to pay the fees.

The results of the various possible answers to such questions have been mixed and depend to some degree on fundamental questions of civic and economic goals. In authoritarian regimes, committed to social "harmony" at the expense of individual liberty, control content may be even more centralized than in the United States. On the other hand, substantial government

investment in the grid itself outside the United States has resulted in far better connectivity in certain democratic and undemocratic countries alike. The debates about the right level of government investment and oversight of the Internet are no simpler than the story of government involvement in the delivery of postal mail, electricity, telephone service, education, or medical care. After quickly telling the story of how the open Internet fell under the control of oligopolies of gatekeepers, we'll raise the questions with which society is left about what, if anything, to do.

Let's start with how it all works.

Connecting the Dots: Designed for Sharing and Survival

The Internet grew out of the ARPANET, a U.S. Department of Defense (DoD) computer networking project from the 1970s. Through its Advanced Research Projects Agency (ARPA, later DARPA), the DoD was directly or indirectly paying for state-of-the-art machines in many academic and national research laboratories. ARPA had two worries.

One worry was pedestrian: The agency was paying for big, expensive computers across the country, but there was no way for underutilized machines at one location to be put to work on problems researchers needed solved at other locations. So every researcher wanted the biggest possible computer, and a lot of computer time was going unused. Scientists could put their data on tapes and send it across the country by air freight, but there was no way to ship the bits without shipping atoms. So ARPA wanted improve utilization by networking the research computers it was funding.

ARPA's other concern cut to the heart of the military mission. The DoD had for some time been worried that its far-flung bases and ships might not be able to communicate if critical sites were destroyed in a nuclear war. In the early 1960s, the worry was whether the telephone network would survive an attack that knocked out a few key switching centers, where many long-distance telephone lines interconnected.

At that time, the researcher Paul Baran studied the properties of a decentralized network, one in which there were many junction points, each connected to only a few others. (The telephone network, by contrast, consisted of a small number of central switching stations connected to customers like spokes joining the hub of a wheel to its rim.) In Baran's proposed mesh-like network, there would be many paths between any two points, so knocking out any one of the junction points would not prevent other points from communicating. Baran imagined an irregular connection of switching points, like the one shown below in an illustration from his 1962 paper.[7]

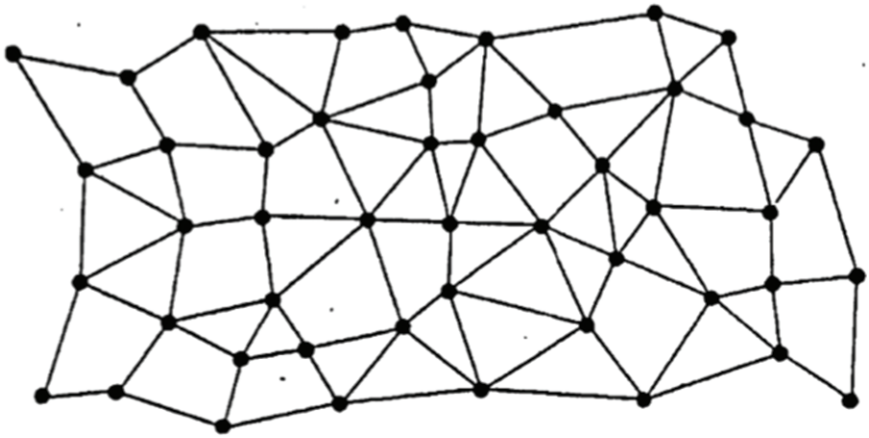

Baran contributed a second important idea: If a switchpoint went down, another route could be found that did not go through it, as long as the switchpoint itself was neither end of the communication. By setting the switches the right way, communication between two points could be established along a particular path. But knocking out any of the points along the path would then interrupt that communication. So would an ordinary hardware failure at any of those intermediate points. It was important to protect the integrity of individual communications even as the network components failed in unpredictable ways.

Baran proposed to chop communications into small chunks of bits, what we today refer to as "packets." In addition to the "payload," a fragment of the communication itself, the packets would contain information identifying the source and destination (much like the address information on a postal letter), and also a serial number so that the destination node could reassemble them in the right order if they happened to arrive out of sequence. With this much information on the "envelope," the packets comprising a single communication did not need to follow the same path. If a portion of the network was unavailable, the network nodes could direct packets along a different path. Making this all work was not simple—how would the network nodes know in what direction to forward a packet?—but in principle, Baran's idea of a mesh-like interconnection and packetized communication would meet the military requirements for survivability.

Protocols: How to Shake Hands with Strangers

Once ARPANET was operational and connected a few dozen computers, it started to become clear that what needed to be connected were not individual computers but existing computer networks. Different ways of networking computers together could coexist, as long as the networks used some common

language for communicating with each other. And in the 1970s and 1980s, different kinds of computer networks did exist, each using the standards of a different computer company. IBM had its SNA (Systems Network Architecture). Digital Equipment Corporation had DECnet. Apollo Computer connected its machines in a ring rather than a branching tree or mesh. Each company touted the advantages of its networking scheme, and some of the claims were valid for particular use cases. But none of the manufacturers had any incentive to make their machines interoperable with those of other manufacturers—until ARPA declared that it would pay for no more computers unless they could be interconnected. Starting from the success of the ARPANET, Vinton Cerf and Robert Kahn designed the protocols for interconnecting computer networks.[8] That is, they designed the Internet.

The Internet *is* its protocols. The Internet is not a machine or even a collection of machines. It's not some piece of software. It is a set of rules. Any person or organization can build hardware or write software that abides by those rules and become a functioning part of the Internet.

Protocols are communication conventions, like the convention that people shake their right hands. Having everyone greet each other by shaking their left hands would work equally well, but the established convention of right-hand shaking makes it possible for strangers to greet each other with no prior mediation. Internet protocols are the conventions by which different networks shake hands in order to pass information from one network to the other. Each network can operate as it wishes internally; only at the points where networks are connected together do the Internet protocols become relevant.

The decision to make ARPANET a packet-switched network simplified the Internet design considerably. Networks were connected to the Internet via connection points called *gateways*. If a gateway behaved as it should, information would flow through it. If it didn't behave properly, that caused no harm except to cut off that network from the rest of the Internet. No computer or network of computers needed permission to join the Internet. If it adhered to Internet standards, it could be understood by others and could interpret messages directed to it.

As we look at the Internet today, it seems varied and complicated: so many different kinds of content, so many different kinds of devices, and so many different kinds of connections. But it's all built on top of a single protocol, known simply as Internet Protocol (IP). It's the job of IP to get a single packet of around a thousand bits from one end of a communications network link to the other. The bits, as delivered, may contain errors; nothing physical ever works perfectly all the time. But errors can be recognized and, if necessary, dealt with. To get packets across the network, IP is used repeatedly, bucket brigade style, with each switching point receiving packets, checking them, and then dispatching them toward their intended destination.

The simplicity of the Internet design made it possible to build protocols on top of protocols to expand the Internet's utility. The earliest uses of the Internet were for logging in to time-shared computers remotely, for moving files from place to place, and for electronic mail. All these services required the data to arrive error free but not necessarily instantaneously. No one would notice if a file transfer or an email delivery took an extra fraction of a second, but having a single bit turn from a 0 to a 1 in transit could have catastrophic consequences. For such transfers, a protocol was developed to make sure that packets sent by the source were received correctly and reassembled in the correct order. Given the unreliability of the intermediate nodes of the network, this requires some bookkeeping at both source and destination. A packet, once received, is acknowledged by sending a special packet back from the destination to the source. The source runs a timer; if a packet sent is not acknowledged before the timer runs down to zero, the source figures that the packet has gotten lost somehow and retransmits it.

The details are tricky, but they are not important to the big picture. The result is that as long as the switches are making their best effort to pass packets along toward the destination, any message sent will be received in perfect order. The protocol that ensures such perfect transmissions is called Transmission Control Protocol (TCP). Because the underlying protocol for moving packets along single links of the network is IP, TCP/IP is the everyday name for the pair of conventions that make reliable communications possible across an unreliable network.

Since there are no rules for joining the Internet, it is fair to wonder about the "best effort" assumption. Couldn't a rogue actor try to sabotage the network by adding switches that would discard or misdirect packets rather than send them toward their intended destination? Indeed, that could happen, but neighboring switch points would eventually realize that the packets were not being delivered and would start avoiding the rogues. Internet routing heals itself by learning to avoid trouble spots—not just in case of hardware failures but also in case of malice. The Internet becomes more reliable the larger it becomes and the more interconnected it becomes.

The Internet worked because once a large enough number of parties agreed to use it in the intended way, bad actors could in effect be frozen out since they were few in number.

In addition to routing information and payload, packets also include some redundant bits to aid error detection. For example, a single extra bit might be added to every packet so that all transmitted packets have an odd number of 1 bits. If a packet arrives with an even number of 1 bits, it can be recognized as having been corrupted in transit and discarded so that the sender will retransmit it. Such extra bits can't guarantee that every packet received is correct. But they do guarantee correct transmission with overwhelming likelihood,

and from a practical standpoint, this process suffices to make the likelihood of an undetected error less than the likelihood of a catastrophe, such as a meteor strike, at the source of the transmission.

IP, the best-effort packet forwarding protocol, can also be used for delivering messages imperfectly but quickly. For example, think about how the Internet might be used to convey voice communications, such as telephone calls. The voice signal can be chopped up into small time slices, each digitized and sent over the Internet. But instead of using TCP, which guarantees delivery but not timeliness, a different protocol, called UDP, is used. UDP accepts some packet loss in exchange for speedy delivery. Voice tones change slowly enough from one instant to the next that packets of a telephone conversation can be scrambled a bit, and some could be omitted entirely, without causing the conversation to become hard to understand—as long as the packets that make it arrive at about the right rate.

Many other protocols have been designed for other purposes and to layer on top of these, using TCP and UDP to carry out more complex communications tasks. For example, Hypertext Transfer Protocol (HTTP) is designed for communication between a web browser on a user's computer and a web server anywhere else in the world. HTTP relies on TCP to retrieve web pages on the basis of location information such as lewis.seas.harvard.edu. So without knowing the details of how TCP operates, anyone could set up a web server that would deliver web pages in response to incoming requests.

Who's in Charge?

There are no Internet cops to force anyone to format their packets as TCP, IP, UDP, or other protocols stipulated. No one will throw you off the Internet if you put your source address where the destination belongs and vice versa. If your packets don't conform to the standards, they just won't be delivered, or they will be ignored if they are delivered.

The Internet does, however, have some governing authorities. One is the Internet Engineering Task Force (IETF), which establishes the standards for Internet protocols. The IETF is a remarkable organization. It is open to anyone who wants to join, and it makes decisions on the basis of "rough consensus and running code." In earlier years, the IETF would meet in a room and determine "rough consensus" by having members hum. Substantial majorities were evident to everyone, and individual preferences enjoyed a level of anonymity because in a large group it is hard to tell who is humming and who isn't. Because most changes to the Internet protocols are enhancements and additions that do not change anything that is already working, there is rarely a need to make a positive decision under time pressure; the IETF can defer decisions, let people talk more while tweaking their proposals, and wait for true consensus to develop.

So the Internet is *open* by design. Anyone can join the decision-making process. You would not be wrong to be reminded of the communal utopianism of the 1960s. Early IETF member David Clark famously said, "We reject kings, presidents, and voting. We believe in rough consensus and running code"—the last phrase indicating the engineer's preference for proofs of concept over concepts alone.[9] Of course, once the Internet became widely adopted, you would need to do a lot of persuading to develop a consensus to change anything that had become important to lots of people. But if you and I were halfway around the world from each other and decided to develop our own secret protocol for (say) trans-Pacific xylophone duets, we could happily program our computers to exchange IP packets that no one else would know what to do with. The IETF explains its role this way in its mission statement:

> When the IETF takes ownership of a protocol or function, it accepts the responsibility for all aspects of the protocol, even though some aspects may rarely or never be seen on the Internet. Conversely, when the IETF is not responsible for a protocol or function, it does not attempt to exert control over it, even though it may at times touch or affect the Internet.[10]

This is a remarkable statement, and it shows how badly the "Information Superhighway" metaphor breaks down when applied to the Internet. If the Internet is a highway, it is one in which motor vehicles voluntarily adhere to certain conventions so they can share the highway safely, but bicyclists and skateboarders are welcome to use the roadways, too—though at their own risk.

The Internet is open in another direction as well. Just as IP serves as the base layer for a hierarchy of protocols, IP itself is a logical, not physical, protocol. Internet packets can be transmitted on copper wire, through fiber-optic cables, or by radio waves. If you are an ordinary personal computer user buying something on Amazon, it is likely that the packets going back and forth between you and Amazon pass through all three and more, as they move from your computer to your wireless router, to your ISP, through the Internet, into Amazon's corporate network, and to one of Amazon's computers. Whenever engineers come up with a new way to move bits through physical media, they can also develop an implementation of IP that runs on that physical medium. There is even a carrier-pigeon protocol that could, in principle, be used to implement IP.

IP, the format in which all packets pass through the Internet, plays a role like the design of the ubiquitous 120V electric outlet, with three holes of specified shape and dimensions. The electric source on one side of the outlet may ultimately be a hydroelectric dam hundreds of miles away, solar panels only a few feet away, or battery storage. As long as the electricity conforms to the standards, the outlet is doing its job. The devices that get plugged into the outlet can be refrigerators, toothbrushes, vacuum cleaners, or dental drills.

As long as a device is fitted with the right plug and is designed to run on standard alternating current, it will work. In the same way, Internet Protocol acts as a universal mediator between applications and physical media.

In fact, the standardization of IP is the reason the Internet has so many uses that were not initially anticipated. Zoom and Facetime—Internet applications for connecting people via live audio and video links—were built on IP, even though there was absolutely nothing about such services in the original Internet design. The inventors of the Internet telephone system Skype—a small group of Scandinavian and Estonian engineers—just needed to adapt the Internet protocols to their purposes. And they did not need to ask the permission of the IETF or any other authority to start using Skype or to encourage others to start paying them to use it.

The Internet Has No Gatekeepers?

Now it has always been an overstatement to say that the Internet has no gatekeepers, but it is less true now than it used to be. As we will soon see, in some countries, governments are the primary gatekeepers, and in others, such as the United States, private corporations assume gatekeeper roles. Let's start with the forms of gatekeeping that have long existed.

Names to Numbers: What's Your Address?

The first fact of Internet life is that it does no good to be "on" the Internet if no one can find you. Packets flowing through the Internet have numeric addresses. Some entity has to translate the symbolic names—like cornell.edu and Skype.com—into numbers and keep track of which connecting points have which numbers.

The Internet Corporation for Assigned Names and Numbers (ICANN) is the body that decrees which numeric addresses are assigned to Cornell University or the nation of Australia. It oversees publication of electronic directories in such a way that anyone sending an email to the address president@cornell.edu or retrieving a web page from an address such http://anu.edu.au (the home page of the Australian National University) is directed to the correct place on the Internet. The translation tables, from letters to numbers, are held on Domain Name System (DNS) servers, which other computers consult in order to look up the numeric addresses to insert in the "destination" field of IP packets before they are launched into the Internet. If the Internet has a single vulnerability, it is control over the DNS servers. Does the island nation of Tuvalu get its own Internet top-level domain, like .au for Australia? (It does—and a very valuable one at that. It's .tv, and the nation, which used to make money from selling postage

stamps, now gets some revenue by letting video sites like twitch.tv use the .tv extension.) Who decides whether Coca-Cola is entitled to cocacola.com or, for that matter, cocacola.sucks? It should be no surprise that such high-stakes questions trigger strong interest from governments and multinational corporations. Such questions generally cannot be settled with anonymous humming or any similarly inclusive process.

Nonetheless, such territorial disputes get resolved without force and without fracturing the network. Indeed, from the first few machines connected as ARPANET in 1969, the number of connected devices has grown into the billions today. Any of them can, in principle, connect to any of the others.[11] The serious gatekeeping problems of the Internet lie elsewhere.

Links Gatekeepers: Getting Connected

The Internet isn't very useful if you can't get connected.

If you drive west from Boston across the northern United States, the offerings at supermarket deli counters change around the time you hit Iowa. Suddenly, they feature gelatin salads in great variety, incorporating various chopped fruits, colorful layers, and creamed toppings. Sometimes the gelatin is molded around fish or meat. The fashion persists across the Great Plains and up into the Rockies, but it disappears on the downward slopes. By the time you reach the Pacific Ocean, Jell-O is again mostly for children and hospital patients. The love of Jell-O concoctions in the rural heartland is so prevalent that these dishes are commonly featured on the covers of food magazines available at the checkout counters of coastal and urban supermarkets even though no one shopping there would dream of serving such a thing. Among the coastal and urban elites, gelatin salads are considered unsophisticated.

The midwestern fondness for gelatin salads is fading now, as food culture, like the rest of American culture, is becoming geographically homogenized. In some areas, grandma's gelatin salad recipes are remembered like handmade straw hats and calico dresses—as artifacts of a rural past out of place in more advanced times. But gelatin salads were not taken West in covered wagons; that would have been impossible since creating them requires refrigeration. They are instead a byproduct of a twentieth century technological revolution: rural electrification. Gelatin salads were a delicacy on farms because you couldn't make them if your farm wasn't electrified. People who served gelatin salads also had electric well pumps and electric lights. Serving Jell-O proved you were technologically advanced.

The economics of diffusing electricity and of diffusing bits through wires or cables are similar. It is expensive to lay cables over long distances, and if there are few customers at the end of the line, it's not cost-effective to lay the

cable. It's also expensive to pull wires to many individual residences if they are far apart. The profits are much higher for wiring up a city because once the line is brought down a street, every housing unit on the street becomes a customer, and the distance from the main line to the customers tends to be short. Indeed, cities were electrified first, and primarily for street lighting, which did not require installation of wiring inside private buildings. All the other uses of electricity—for interior lighting, refrigerators, washing machines, dishwashers, and radios—were derivative of the use for street lighting. To this day, a stroll down Beacon Street near Fenway Park in Boston takes you past a structure that reads "The Edison Electric Illuminating Company."

It would have made no sense to diffuse electricity nationally just so people could have Jell-O. In fact, there was no such thing as an electric home refrigerator when the first electric streetlights lit up. Household electricity was what Jonathan Zittrain calls a *generative* technology.[12] Once the infrastructure was in place, creative people began dreaming up uses for it, and whole new technologies developed that could not have existed without it. In the process, some old industries died, at enormous costs to those who had profited from them. "Ice box" is at best a nostalgic phrase for a refrigerator today, but a century ago there was an enormous industry devoted to cutting the frozen surface of lakes into blocks, shipping them long distances, and distributing them to American homes. Generative technologies are also destructive technologies.

The diffusion of the Internet has followed much the same stages as the diffusion of electricity once did. Lighting was the killer app for electricity; gelatin salads were the cat videos of the electric grid. And yet the U.S. experience with the Internet has been very different than it was with the electric grid.

The United States was electrified quickly and ubiquitously, but the fast build-out would not have happened without an impetus from the federal government. Wiring remote areas was unprofitable, and it would have been even more unprofitable for a competitor to lay a second cable to serve the same customers. So electricity was readily available in cities but extremely costly in remote areas, if it was available at all.

Franklin Roosevelt could see the difference in the price of electricity when he retreated from New York and Washington, DC, to Warm Springs, Georgia, where he sought respite and spa treatments for his paralysis. He conceived the Rural Electrification Act in Warm Springs and signed it into law in 1936. The initiative stimulated not only the diffusion of electricity but the invention of new ways for ordinary consumers to use it.

In the early 1920s, fewer than 1% of U.S. households had electricity. Six years after the signing of the Rural Electrification Act, on the heels of a terrible economic recession, half of U.S. households were electrified. By 1960, virtually the whole country had electricity.

To make a sensible comparison between diffusion of the Internet and diffusion of electricity, we need to define our terms. Electricity as delivered to the home became standardized: In the United States, electricity is alternating current, 60 Hz frequency and 120V. These standards are analogous to IP for the Internet. They guarantee that the same appliances can be plugged into outlets anywhere in the country.

But there is another important parameter: the *amount* of electric energy used by an appliance or a household. The rate at which electricity is used is called *power* and is measured in watts or kilowatts (for thousands of watts). The amount of energy used is measured in kilowatt-hours; a kilowatt-hour (kwh) is the amount of energy consumed by using it for an hour at a rate of 1,000 watts. Electric codes were standardized so that household circuits can handle around 2,000 watts; if you use much more than that, a circuit breaker will trip, and the wire might melt if there were no fuse or breaker. The wiring in an old house may have to be upgraded when an owner wants to use more electric equipment or more powerful electric equipment, such as air conditioning or a hot tub. On the other hand, new equipment tends to use less power than old, so average consumption per household has increased only slowly over time. The electric utilities that actually supply the power may have to upgrade their distribution grid to keep up with demand, but a combination of consumer pressure and federal standards generally make it rare in the United States to have "brownouts," when a whole city or neighborhood has insufficient power. In the United States, electricity is, for the most part, a successfully regulated industry.

The analog of power for the Internet is bit rate, and here the experience of the Internet and the electric grid have diverged significantly. Provisioning of Internet connectivity has been left almost entirely to the private sector, with minimal government regulation and minimal government support. Almost nowhere does serious competition exist, so consumers cannot switch to better providers. The monopoly provider of Internet services may offer a choice of speeds, but the higher speeds are likely to be exorbitantly expensive. In a word, rather than providing high-speed Internet, most suppliers of Internet services try to convince us that we have it already, and the government is assisting in their deceit rather than prodding them to improve their services.

Where Are the Bottlenecks, and What Counts as "High Speed"?

The rate at which bits complete their transit—for example, from a web server somewhere to the browser running in your home computer or from your office computer to the video chat room at your London headquarters—is the slowest of the rates of any of the links along the way. The rate at which

bits flow through a link is affected by some physical parameters—the electrical or electromagnetic properties of the copper or glass of which the link is composed—and how heavy the traffic is. If your communication has the link to itself, it can utilize all the bits-per-second of which the communication medium is capable, but if it has to share that capacity with a million other transmissions, yours may get only a millionth.

Think again of your home computer retrieving a web page from a server belonging to some big company. Your request has to get to the corporate server, and the packets comprising the web page have to get back to your home computer. Simplifying greatly, you can think of the bits in your request making three hops. They have to get from your computer to the outside wall of your house; then from your house to the "backbone," the long-haul cables that crisscross the country; and then across the hundreds or thousands of miles of the backbone. The connection from your house to the backbone is commonly referred to as the "last mile." The same hierarchy is, in principle, traversed at the other end, except that Amazon and Google are connected directly to the backbone because of their enormous capacity needs. If you were communicating with someone else sitting at a home computer, the bits would have to traverse the "last mile" to that person's house.

Inside a house, most people use Wi-Fi, a kind of short-distance radio communication. Newer Wi-Fi technologies can reach gigabit speeds, but in practice connections are likely to be slower because of interference or obstructions. Power users may still have their houses wired so they can connect their computers using Ethernet cable rather than wireless.

But slow wireless may be fast enough if the last mile is slow anyway. And in the United States, it almost certainly is slow. By global standards, what is described in the United States as "high-speed Internet" just isn't.

The backbone of the Internet is fiber-optic cable. Fiber is amazing stuff; the glass itself has almost limitless information-carrying capacity. Its actual capacity is limited not by the glass but by the electronics that connect the network at the switching points. The electronics are constantly being improved; fiber, once installed, is never replaced (unless it breaks—for example, because a fishing trawler snags it).[13]

In some parts of the world, the last mile is also fiber, so those amazing information capacities go right to the doors of homes and businesses. In Singapore and Sweden, virtually everyone has access to Internet speeds in the billions of bits per second. By contrast, perhaps 15% of Americans have fiber connections to their homes, and the percentage is not increasing. Most of us are connected by legacy telephone wiring, using so-called DSL service, or by the coaxial cable that was installed to bring cable TV to our houses. Even DSL service is being phased out in some places as unprofitable. And the cable and telecommunications providers have effectively divided up the map between

them; in few places can a consumer choose between cable and DSL service, much less between multiple cable services.

The Internet itself does not connect to your home or office or cell phone. Ordinary people connect their devices not to the Internet but to an Internet service provider (ISP). Since Internet packets can travel over a variety of physical media, there is in principle no limit to the number of ISPs that might carry packets into and out of your house, for a price. The reality is very different. In the United States, it is likely that your home ISP is AT&T, Time Warner, Comcast/Xfinity, Verizon, or Charter. The reason there are so few is that each of these is either a telecommunications or a cable television company, and they are using the wiring or fiber-optic cable that they have already brought to your house to provide telephone or television service. Wireless Internet access is also possible—by which is meant not the Wi-Fi connection between your computer and a router, which is then connected to your ISP, but a wireless connection directly to the ISP. Connection via satellite is possible in rural areas where no form of wired connection has been installed, but satellite Internet is both slow and expensive. Cell phones connect to the Internet via the cellular telephone network, but the cellular network is not a viable option for home use. And so called 5G radio signals, which are being billed as the future of Internet connectivity, travel only short distances. So a 5G infrastructure is realistic only in densely populated areas, where it is possible to install many hubs economically.

Americans are bombarded with advertisements for "high-speed Internet," but in reality, even the government definition of "high speed" is deceptive. The last time it was updated was in 2015, when—over the objections of Internet service providers—the FCC raised the standard from 4 Mbps to 25 Mbps. That is one-fortieth of the gigabit speeds that are now standard in Japan and Sweden and which even China is diffusing in rural areas. And the U.S. 25 Mbps standard applies only to download speeds, as though the Internet were basically a broadcast medium for consumers to receive Netflix movies. A great many applications, from video chats to transfer of medical imagery, require high speeds in both directions. Businesses of every kind are dependent on Internet connectivity for both uploads and downloads; they use the Internet to get information about their products and services, and indeed the products and services themselves, to their customers. So it is very difficult to start a business in an area where connectivity is poor or limited to fast downloads and slow uploads. And yet in the United States, the Internet has been optimized as a replacement for television—as a way for the few to supply content to the many.

A second form of creative semantics distorts government statistics about the availability of "high-speed Internet" in the United States. The government considers a census tract—one of the some 75,000 geographic areas into which the United States is divided for census purposes—to have high-speed Internet if even a single household has access to it, regardless of the price and regardless

of whether any household has actually signed up for a connection. Thus government estimates of the diffusion of high-speed Internet are wildly inflated.

And price matters. In large parts of the world, gigabit speeds are available for less than $50 per month. In the Boston area as we write, it's $70 per month in the restricted areas where it is available at all—and requires a 24-month contract.

The typical American has only one or two realistic ISP choices. More than 30% of U.S. households have no providers of Internet service at 25 Mbps or better, and fewer than a quarter of U.S. households have more than one choice.[14]

For the most part, diffusion of Internet connectivity has been left to the private sector in the United States. In fact, rules in 26 states hinder or prohibit local governments from offering Internet connectivity, forcing individuals to whom no affordable household service is available to connect at public libraries or fast food restaurants.[15] Montana's code[16] is typical: "Except as provided in subsection (2)(a) or (2)(b), an agency or political subdivision of the state may not directly or through another agency or political subdivision be an internet services provider." So unless all the private ISPs pull out, the people of Browning, Montana, are stuck with the poor-quality but expensive private offerings. No amount of municipal entrepreneurship can help its people out. The lobbyists from the private telecommunications firms got to the state legislature first.

Now, of course, there are reasons to keep governments from competing with private companies. The arguments are familiar. Competition drives down prices and improves quality. Taxpayer dollars should not be used to undercut private vendors. Governments should not interfere with free markets.

But there just aren't enough people at the other end to pay for the long-haul connection. This is no different from the situation of rural free delivery of postal mail—which became the law in 1893 and reached Billings in 1902—or the electric grid in the 1930s and 1940s. Connecting the country requires viewing electronic communications like electricity or postal mail: It would have to be available to everyone at an affordable price. That principle is in fact not generally accepted. Instead, the operative metaphor for the Internet is television or a multi-screen movie theater. The dominant Internet service providers see the Internet as a way of connecting active content providers to passive content consumers. That is why ISPs offer packages that encourage downloads and limit uploads.

In the absence of the kind of push from the government that has resulted in far better Internet service in South Korea, Switzerland, and even largely rural Finland than in the United States, why hasn't competition driven down prices and quality up?

Some would argue that corporate greed, collusion, and corruption are to blame, and while that perspective may have some validity in some cases, the

reality is that communications networks grow and consolidate almost organically for reasons of efficiency. Paul Baran himself anticipated that this phenomenon would affect computer networks several years before he designed one of the earliest networks. In testimony before Congress on electronic privacy in 1966, he said,[17]

> Our first railroads in the 1830's were short routes connecting local population centers. No one sat down and laid out a master plan for a network of railroad rails. With time, an increasing number of such separate local systems were built. A network gradually grew as economic pressure caused the new links to be built to span the gaps between the individual routes.
>
> We didn't start to build a nationwide telegraph network in the late 1840's; only independent telegraph links. But it was not long before we had an integrated nationwide network. Even the name, Western Union, recalls the pattern of independent links joined together to provide a more useful system.
>
> We didn't start to build a nationwide telephone system in the early days of the telephone in the 1890's. Yet, today we have a highly integrated telephone network.
>
> Such patterns of growth are not accidents. Communications and transportation are services that historically tend to form "natural monopolies." The reason is well understood. It's cheaper to share use of a large entity than to build your own facilities. Hence, if you were to look at the earth, say, from the far-off vantage point of the moon, it would appear that the growth of these integrated networks out of individual pieces is almost biological.

So it's not very complicated. It is more valuable to be part of a big network than a small one, and the bigger the network, the more valuable it is to be part of it. In the absence of pushback from some social structure with the authority to resist mergers, consolidations, buyouts, and strategic corporate decisions to gain control over network traffic, communication networks will grow larger and fewer with time. Such monopolization is not necessarily against the public interest—as long as the public interest is at the table when distribution and pricing decisions are being made. Today, they rarely are.

Can the Letter Carrier Decide What Mail to Deliver?

The story of Telus and its striking workers with which this chapter began demonstrates that the dichotomy between links and content is unhelpful when

the links gatekeeper takes on the role of content gatekeeper. The notion that Internet service providers should not be deciding which bits to deliver is known as *net neutrality*. In principle, it can seem simple and inarguable; after all, we don't want the phone company deciding which conversations it is going to allow to happen over its voice lines. True, when customers don't pay their bills, their phone service can be cut off, but even that is rare because society generally recognizes—or used to recognize—that phone service is important to daily life. But the Internet is not exactly like the telephone network.

About the same time as Telus was shutting down pro-union websites in Canada, a small North Carolina ISP by the name of Madison River Communications shut down Vonage, which offered Voice over IP (VoIP) service. Using the Internet to deliver live voice conversations would have seemed crazy at the birth of the Internet because the network was too slow and the computers connected to it couldn't keep up with the flood of packets in order to reassemble them into comprehensible speech. But times change. As link speeds improved and new protocols—based on IP—were optimized for voice communications, a systematic difference between telephone and Internet services intervened. Telephone companies charged extra for long-distance service; Internet service providers didn't care where packets were coming from or going to. They might charge more for higher data rates but not for more distant destinations. Inevitably, VoIP software—Skype was the earliest commercial success in this space—was developed to replace telephony by Internet communications and make long-distance "calling" virtually free for anyone with an Internet connection. Vonage was using Madison River's data service to undercut Madison River's phone service.

When Vonage was blocked, the company complained to the Federal Communications Commission, which has jurisdiction over telephone services. The case was resolved when Madison River agreed to pay a fine and not block VoIP for three years, but this resolution left more questions than answers in its wake. What if Madison River had been a cable company offering Internet services rather than a phone company? On the other hand, what if Madison River had been big enough to fight the FCC in court? It was not at all clear that the FCC had the congressional authority to back up its strong-arm tactics on the way even telephone ISPs were picking and choosing what bits to deliver.[18]

Matters came to a head in 2008, when the FCC ordered the ISP Comcast to stop "throttling"—that is, slowing down—BitTorrent, a peer-to-peer file-sharing service heavily used for delivering movies to the home. Comcast was profitably delivering movies over the same cable it was using to deliver Internet service, so BitTorrent was undercutting its video delivery business. Comcast successfully sued the FCC, establishing that indeed the FCC lacked the authority to regulate its Internet service business. This decision kicked off a net neutrality debate that has raged for more than a decade.

The details are complex. In a nutshell, pro-neutrality voices have argued for consumer choice and freedom; opponents have argued that market forces would resolve any tensions, an argument greeted skeptically by those noting how little competition exists in the ISP space. In the United States, net neutrality rules were instituted during the Obama administration and repealed during the Trump administration. Many other nations have adopted net neutrality in principle, but some of them allow usage-based billing, which may have the effect of rendering certain applications, such as watching movies, unacceptably expensive, thus achieving the same result—prioritizing other means of delivering movies to the home—that Comcast had achieved by throttling peer-to-peer services in 2008.

Search Gatekeepers: If You Can't Find It, Does It Exist?

Prescient as Baran was, he could not have anticipated the extent to which, as communication networks became accessible to everyone, control over the information they carry would also tend to fall into a small number of private hands. Search technology was a surprising development of the 1990s; it is now hard to imagine a world without it. And yet it is *not* hard to imagine a world in which Google does not control most of the searches in the Western world. It has just turned out that way, with troubling consequences.

Found After 70 Years

Rosalie Polotsky was 10 years old when she waved goodbye to her cousins, Sophia and Ossie, at the Moscow train station in 1937. The two sisters were fleeing the oppression of Soviet Russia to start a new life. Rosalie's family stayed behind. She grew up in Moscow, taught French, married Nariman Berkovich, and raised a family. In 1990, she emigrated to the United States and settled near her son, Sasha, in Massachusetts. Rosalie, Nariman, and Sasha always wondered about the fate of Sophia and Ossie. The Iron Curtain had utterly severed communication among Jewish relatives. By the time Rosalie left for the United States, her ties to Sophia and Ossie had been broken for so long that she had little hope of reconnecting with them—and, as the years wore on, she had less reason for optimism that her cousins were still alive. Although his grandfather dreamed of finding them, Sasha's search of immigrant records at Ellis Island and the International Red Cross provided no clues. Perhaps, traveling across wartime Europe, the little girls had never even made it to the United States.

Then one day, Sasha's cousin typed "Polotsky" into Google's search window and found a clue. An entry on a genealogical website mentioned "Minacker," the name of Sophia's and Ossie's father. In short order, Rosalie, Sophia, and Ossie were reunited in Florida, after 70 years apart. "All the time when he was alive, he asked me to do something to find them," said Sasha, recalling his grandfather's wish. "It's something magic."[19]

The World Wide Web has put vast amounts of information within reach of millions of ordinary people. But you can't reach for something if you don't know where it is. Most of that vast store of digital information might as well not exist without a way to find it. Indeed, the "dark web" exists as a kind of parallel universe, with troves of information invisible to search engines and to users who don't know where to look for it.

Search both fulfills dreams and shapes human knowledge. The search tools that help us find needles in the digital haystack are the lenses through which we view the digital landscape. But the "lenses" are not passive. They actively color what we see by their selection of what to show us on the first page of results and by the order in which the results are presented to us. Whoever controls the search engine shapes—and distorts—the reality we see through it. Google, which is used for more than 90% of the world's searches,[20] is supported by advertising, so questions inevitably arise about whether results optimize Google's profits or users' satisfaction. Microsoft's Bing is no less good at producing results, but it has less than 5% of the market. DuckDuckGo, which offers much stronger privacy protections than Google or Bing but produces less targeted results, has a negligible share of the market.[21] Baidu is the dominant search engine in China but for a reason: It censors heavily, as any search engine in the Chinese market must do. How did these lopsided statistics arise, and what are their consequences?

The Fall of Hierarchy

From the dawn of writing until about 1994, there were only two ways to organize information so it could be retrieved quickly. You could put it in a hierarchy, or you could create an index.

A hierarchy enables you to put things into categories and divide those categories into subcategories. Aristotle tried to classify everything. Living things, for example, were either plants or animals. Animals either had red blood or did not; red-blooded animals were either live-bearers or egg-bearers; live-bearers were either humans or other mammals; egg-bearers either swam or flew; and so on. Sponges, bats, and whales all presented classification enigmas, on which Aristotle did not think he had the last word. At the dawn of the Enlightenment, Linnaeus provided a more useful way of classifying

living things, using an approach that gained intrinsic scientific validity once it reflected evolutionary lines of descent.

Our traditions of hierarchical classification are evident everywhere. We just love outline structures. The law against cracking copyright protection is Title 17, Section 1201, paragraph (a), part (1), subpart (A). In the Library of Congress system, every book is in one of 26 major categories, designated by an uppercase letter, and these major categories are internally divided in a similar way; for example, in category B, philosophy, you find BQ, Buddhism.

If the categories are clear, it may be possible to use an organizing hierarchy to locate what you are looking for. This requires that the person doing the searching not only know the classification system but be skilled at making all the necessary decisions. For example, if knowledge about living things was organized as Aristotle had it, anyone wanting to know about whales would have to know already whether a whale is a fish or a mammal in order to go down the proper branch of the classification tree. As more and more knowledge has to be stuffed into the tree, the tree grows and sprouts twigs, which over time become branches sprouting more twigs. The classification problem becomes unwieldy, and the retrieval problem becomes practically impossible.

In 1991, when the Internet was barely known outside academic and government circles, some academic researchers offered a program called Gopher. This program provided a hierarchical directory of many websites, by organizing the directories provided by the individual sites into one big outline. Finding things using Gopher was tedious by today's standards, and it was dependent on the organizational skills of the contributors. Yahoo! was founded in 1994 as an online Internet directory, with human editors placing products and services in categories, making recommendations, and generally trying to make the Internet accessible to non-techies. Although it is today a search and news site, the name "Yahoo" was originally said to be an acronym for "Yet Another Hierarchical Organized Oracle."

The practical limitations of hierarchical organization trees were foreseen 60 years ago, long before the explosive growth of the World Wide Web and the countless daily changes to it. During World War II, President Franklin Roosevelt appointed Vannevar Bush of MIT to serve as director of the Office of Strategic Research and Development (OSRD). The OSRD coordinated scientific research in support of the war effort. It was a large effort, with 30,000 people and hundreds of projects covering the spectrum of science and engineering. The Manhattan Project, which produced the atomic bomb, was just a small piece of it.

From his vantage point, Bush saw a major obstacle to continued scientific progress. We were producing information faster than it could be consumed—or even classified. Decades before computers became commonplace, he wrote about this problem in a visionary article, "As We May Think."[22] It appeared in the *Atlantic Monthly*—a popular magazine, not a technical journal. As Bush saw it,

> The difficulty seems to be, not so much that we publish unduly...but rather that publication has been extended far beyond our present ability to make real use of the record. The summation of human experience is being expanded at a prodigious rate, and the means we use for threading through the consequent maze to the momentarily important item is the same as was used in the days of square-rigged ships....Our ineptitude in getting at the record is largely caused by the artificiality of systems of indexing.

The dawn of the digital era was at this time barely a glimmer on the horizon. But Bush imagined a machine, which he called a "memex," that would augment human memory by storing and retrieving all the information needed. It would be an "enlarged intimate supplement" to human memory, which could be "consulted with exceeding speed and flexibility."

Bush clearly perceived the problem, but the technologies available at the time—microfilm and vacuum tubes—could not solve it. He understood that the problem of finding information would eventually overwhelm the progress of science in creating and recording knowledge, and he anticipated that it would be possible to search using multiple terms to isolate special kinds of information:

> Wholly new forms of encyclopedias will appear, ready made with a mesh of associative trails running through them, ready to be dropped into the memex and there amplified....
>
> The historian, with a vast chronological account of a people, parallels it with a skip trail which stops only on the salient items, and can follow at any time contemporary trails which lead him all over civilization at a particular epoch. There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes, not only his additions to the world's record, but for his disciples the entire scaffolding by which they were erected.

Bush was intensely aware that civilization itself had been imperiled in the war, but he thought we must proceed with optimism about what the record of our vast knowledge might bring us:

> Presumably man's spirit should be elevated if he can better review his shady past and analyze more completely and objectively his present problems. He has built a civilization so complex that he needs to mechanize his records more fully if he is to push his experiment to its logical conclusion and not merely become bogged down part way

there by overtaxing his limited memory. His excursions may be more enjoyable if he can reacquire the privilege of forgetting the manifold things he does not need to have immediately at hand, with some assurance that he can find them again if they prove important.

...He may perish in conflict before he learns to wield that record for his true good. Yet, in the application of science to the needs and desires of man, it would seem to be a singularly unfortunate stage at which to terminate the process, or to lose hope as to the outcome.

A Futurist Precedent

In 1937, H. G. Wells anticipated Vannevar Bush's 1945 vision of a memex. Wells wrote even more clearly about the possibility of indexing everything and what that would mean for civilization:

> There is no practical obstacle whatever now to the creation of an efficient index to all human knowledge, ideas and achievements, to the creation, that is, of a complete planetary memory for all mankind. And not simply an index; the direct reproduction of the thing itself can be summoned to any properly prepared spot.... This in itself is a fact of tremendous significance. It foreshadows a real intellectual unification of our race. The whole human memory can be, and probably in a short time will be, made accessible to every individual....This is no remote dream, no fantasy.[23]

Capabilities Bush could not have seen clearly are commonplace now. Digital computers, vast storage, and high-speed networks make information search and retrieval necessary. They also make it possible. The Web is a realization of Bush's memex, and search is key to making it useful.

Search Histories

Bush did not imagine that everyone would have a memex, but he did foresee that "associative trails" would endure. It's worth looking a little more closely at what that implies about the way search engines work. What Bush saw as an important new knowledge structure has turned out to be something more like digital exhaust—a mostly harmless side effect of the fancy digital engines we use to get things done.

"Search" is something of a misnomer for what Google and other search engines actually do. When you type something into a search engine, the engine does not go check the entire World Wide Web, looking for it. It looks up your search term in an index that has already been created. It is a very large index

and very cleverly organized so that it can be updated constantly and so that searches can focus on multiple terms, but it is fundamentally no different from the index in the back of a book...except that when you look up something in the index of a book, nobody but you knows that you did so. If you ask Google to look something up in its index, it remembers that you did so.

There are good reasons for Google to remember what you searched for. The information might be useful in helping Google respond to future searches more appropriately. It would certainly be useful in helping Google target advertising to you. You can use a privacy-preserving search engine (such as DuckDuckGo, mentioned earlier), but you may not be as happy with the quality of the results. The overwhelming dominance of Google suggests that people are happy to trade their privacy for quality—or are just going with the household name and don't realize what exchange they are making.

Casey Anthony may not have been thinking about the endurance of search histories when she or someone using her computer Googled "neck breaking" and "how to make chloroform" before the mysterious death of her daughter Caylee in 2012.[24] The disclosure of this search history in her trial didn't result in her conviction; later it came out that the same computer had been used to search for "foolproof suffocation" on the very day the girl went missing.[25] (The detectives had missed this because the search was done using the Firefox browser rather than the Internet Explorer browser that gave up the information about the other searches.) A few years before, James Petrick had been convicted of killing his wife in part on the basis of searches he had done for terms such as "neck" and "snap break," and for topographic information about the lake where her body was found.[26] An appeals court in Illinois upheld the murder and solicitation of murder convictions of Steven Louis Zirko[27] in part on the basis of searches on Zirko's computer for terms like "mercenary for hire" and the hours when the child of one of his victims would be at school.

These cases all involved police searches of someone's home computer. But there is another way to get information about search history: by asking Google. Google won't just give that information out to anyone who asks, but you can see for yourself what Google is remembering about your searches and other activity that was conducted while you were logged into Google. (Or at least what Google tells you it is remembering; it probably is remembering a great deal more.) Under your Google account page, there is a "Data and personalization" screen, where you can turn off the recording of your search history, for example. You can even edit the history without getting rid of it entirely, if you want to—as Dr. Brent Dennis apparently did in order to mislead law enforcement. He told the police that his wife had died from drinking antifreeze, but it was he himself who searched for "antifreeze" and then hired someone to clean up the search history.[28]

Whenever you go to a new website that asks you to create an account or to "sign in using Google" or "sign in using Facebook," you might be pleased to save time and have one fewer password to remember. But what you are really doing if you choose to use your existing Google or Facebook login credentials is giving Google or Facebook permission to add to the enormous store of information it already holds about you the new information it gleans from your activity on the new site.

And the government can force Google to turn over what it knows about you—your search history, for example. It's not even that complicated. In 2018 Google received about 130,000 requests from courts and other government agencies, and complied, at least in part, with about two-thirds of them.[29] Google says it will inform you when an agency seeks your records, but it is under no obligation to do so or to comply with your wishes if you object.

To get that information from your laptop, law enforcement would need a search warrant; that is, it would

Can the Police Search Your Searches?

Can law enforcement get information about your searches even if you have done nothing wrong? It seems to have happened in Edina, Minnesota, in early 2017.[30] Someone impersonating a man named Douglas called Spire Credit Union and persuaded the clerk to transfer $28,500 of Douglas's money to another bank. The fake Douglas duly supplied Douglas's name, birth date, Social Security number, and a faxed copy of his passport—or at least a passport that had Douglas's photo on it, which matched the bank's records and completed its authentication.

When Douglas realized that his money was gone, he contacted law enforcement. Detective David Lindman used Google's image search to find a matching photo of Douglas online. Lindman asked the Hennepin County judge to issue a search warrant to Google for records of anyone in Edina who had searched for Douglas's image during a five-week period prior to the incident. So at least in some cases, a search warrant can now be issued not against a particular individual but for the set of individuals who have performed a certain kind of search.

Google the gatekeeper turns out to have gates swinging both ways. As you search, it is determining what information to show you, and it is also collecting information about you at the same time. It can use that information for its own advertising purposes, and under court order, it can open the gate to outsiders.

have to make a case to a judge that your Fourth Amendment rights against unreasonable searches were not being violated. Why can the police more easily get the same information directly from Google?

The underlying legal principle is simple: In the absence of any more specific legislation, if law enforcement asks Google what you searched for, it may spill the beans on you because of the "Third Party Doctrine." You can tell me a secret, and the government can't make either of us disclose it. But if you use Gmail—essentially asking Google to pass your secret to me—then Google is a third party and is not bound under the Fourth Amendment to respect your desire and mine to keep the information secret, any more than if you shared your secret with a stranger on the street. The same applies to other information you have entrusted to Google—for example, the terms you've asked it to search for. Those searches are Google's property, not yours. It can use them to generate targeted advertising and for other purposes.

In 2017 Google announced that it would not scan users' email in order to improve its advertisement targeting, but that was a policy decision, not a response to any U.S. law. In fact, the United States lacks comprehensive privacy laws, and corporate policies need not be unchangeable or consistent with users' expectations. When Gmail was launched in 2004, Google explained its practice of scanning users' email as being helpful in targeting advertisements to offset the cost of offering a free service. A decade of mounting criticism and some litigation resulted in Google's decision to stop scanning email. What Google did not explain at the time was that it was enabling certain corporate partners to scan emails—and sometimes have humans read them.

Navideh Forghani of Phoenix, Arizona, seems to have been unaware of such practices when she signed up for Earny, a money-saving service.[31] Earny checks the customer's inbox for receipts of items she has purchased, searches the Web to see if it can locate the same items at cheaper prices, files with her credit card company for a refund of the price difference, and then splits the proceeds with the customer, all quietly in the background. Once Navideh had signed up, the only thing she had to do was to watch the credits appear on her credit card bills.

Earny is a private company separate from Google, and Google is not scanning her email. But when she clicked a button to give Earny access to her inbox, she was authorizing Earny to do exactly that. Earny, in turn, was sharing her email with Return Path, a company with which Earny had partnered to do the actual scanning.

Google, Earny, and Return Path all explained that they had done nothing wrong because these practices were authorized under the companies' privacy policies. Navideh acknowledges that she did not read Earny's privacy policy[32]. No surprise: It was almost 3,000 words long when I checked it in 2019. Earny's policy is sufficiently dense as to be difficult for many readers; and, to make matters worse, it links to a host of other privacy policies making it nearly

impossible to understand what you are relinquishing when you sign up for their service. Return Path's privacy policy is almost 6,000 words.

The bottom line is that your data is valuable, and every convenience you accept comes at a price. As Marc Rotenberg, president of EPIC, a major privacy advocacy organization, stated, "The privacy policy model is simply broken beyond repair. There is simply no way that Gmail users could imagine that their personal data would be transferred to third parties."[33] When a product has so little competition and is so useful to everyday life and to the conduct of business, the "notice and consent" protocol does not realistically protect users from having their data used in unexpected ways.

How Did Google Get So Big?

As the Web grew in the early 1990s, hierarchical structures, never satisfactory for finding unclassifiable information, quickly failed to keep up with the size of the Web. Several search engines based on an automatically constructed index began to appear, and some were modestly successful. But very shortly Google became dominant, to the extent that the name of the search engine, and the company, has become a verb synonymous with "web search."

In 1996, the Google founders, Larry Brin and Sergey Brin, had a good idea while they were in graduate school. An important web page is one referenced—that is, linked to—by a lot of important pages. That sounds like a circular definition, but if the entire structure of the Web can be captured and analyzed, some fairly simple mathematics can be used to get a consistent measure of the importance of every web page. That mathematics, plus some solid engineering to get all the data organized and processed in the limited storage available at the time, got the company off the ground. Its dramatically simple interface—just type in something and get answers back, no options, bells, or whistles—comforted even the most naïve users and lured them into using it more.

Google's search engine was good, but it was not ten times better than others available when the company was founded in 1998. For example, by this time AltaVista had been operational for three years[34] and was processing hundreds of millions of search queries as a free service to the public.

Digital Equipment Corporation, which developed AltaVista, never figured out how to make it profitable. Digital was primarily a hardware company, and it sold AltaVista to another company. (Digital itself was bought by Compaq soon after.) AltaVista changed hands again and was finally quietly shut down. Microsoft didn't launch its Bing search engine until 2009. By then Google had a seemingly unsurmountable head start, in spite of the fact that users can switch search engines from Google to Bing with minimal effort.

Google gained its advantage by carrying advertising from the very beginning. The ads are generated in response to search terms; search for "cell phones," and you are likely to see ads for products and services related to cell

phones. Which ads come up, among all the advertisers who might want to tout their goods to people interested in cell phones, is determined by a continuous auction. Advertisers willing to pay more for their ads are more likely to have their ads appear. The auctions run automatically and invisibly, and the result is a system of unprecedented efficiency. An advertiser in a newspaper, magazine, or radio station has to hope that among the undifferentiated mass of people exposed to the ad, a few will be interested in the product advertised. Advertisers can try to tilt the odds in their favor by, for example, putting sports-related advertisements on radio stations that cater to sports fans. But associating advertising with search directs the ads only to those individuals who have shown at least enough interest in a topic to search for it.

The Google founders themselves recognized[35] the downside of mixing advertising with search, which was already being done with a few of the other search engines then in use. It would lower confidence in search results, for example, if users suspected that the search results themselves were biased to favor the advertisers:

> For example, in our prototype search engine one of the top results for cellular phone is "The Effect of Cellular Phone Use Upon Driver Attention," a study which explains in great detail the distractions and risk associated with conversing on a cell phone while driving. This search result came up first because of its high importance as judged by the PageRank algorithm, an approximation of citation importance on the web. It is clear that a search engine which was taking money for showing cellular phone ads would have difficulty justifying the page that our system returned to its paying advertisers.

After mentioning a few other examples of conflicts of interest between returning useful search results and gaining advertising revenue, Page and Brin concluded, "we believe the issue of advertising causes enough mixed incentives that it is crucial to have a competitive search engine that is transparent and in the academic realm." Be that as it may, no such search engine is widely used today. Google's enormous revenues are largely derived from exactly this kind of advertising, and in 2017 the European Union fined Google 2.4 billion euros for biasing its search results in favor of its advertisers.[36] And without knowing exactly what is going on inside Google's code, it is hard to know whether results are being biased. Brin and Page anticipated this, too: "For example, a search engine could add a small factor to search results from 'friendly' companies and subtract a factor from results from competitors. This type of bias is very difficult to detect but could still have a significant effect on the market." How the conflicting interests and lack of transparency will be resolved remains unknown, but the stakes are extremely high.

Social Gatekeepers: Known by the Company You Keep

When it was created, the Internet was a means of connecting one computer to another and, ultimately, one network of computers to another (hence the name "Internet"). It expanded from connecting machines to connecting users to information. The complexity of these connections led to the role of the search gatekeeper. In its latest phase of connectivity, the Internet has facilitated the connection of people to each other at a level, and with implications, that were not imagined even by the creators of the dominant solutions.

The Social Network: Facebook and More

The digital explosion has never been so powerful as in the growth of Facebook. As glamorized in the movie *The Social Network*, Facebook's success would seem to be the result of adolescent dumb luck and capitalist ruthlessness. The full story is more interesting—and more telling in terms of the ways people share bits.

There were online social networks before Facebook. The earliest was Sixdegrees.com in 1997. The name was derived from the early 1990s play and movie *Six Degrees of Separation*. The site grew to include millions of users, but it eventually stalled and died in 2000 for lack of a sustainable business model and because there wasn't much for people to do on it once they had connected to each other.[37]

Friendster launched in 2002 and quickly grew to be among the most popular sites on the Web. Originally it billed itself as a place where users could meet new people to date, make new friends, or help friends meet new people.[38] At its peak, it allowed easy search through the entire membership database of tens of millions of people.

Friendster collapsed in 2006, a victim of its own success. Growth was so explosive that the site was plagued with technical problems. People have very limited tolerance for delay while waiting for pages to load; they will abandon a site completely if it doesn't work and they have no strong reason to keep trying. Friendster also made too much of its main merit, which was making it easy to connect to people you didn't already know but might be interested in meeting. Moreover, unanticipated social problems arose because it was so easy to view the profiles of other users. For example, it turned out that not everything users put in their profiles to stimulate social connections was something they wanted their bosses to know.

Myspace was started in 2003 as a competitor to Friendster. It drew Friendster users who were either dissatisfied or had been thrown off the site. Myspace had a loyal following among indie rock bands and their followers who had refused to play by Friendster's rules. Soon Myspace had more visitors than Google or any other website, though culturally it retained the naughty and creative feel of its rebellious origins. But within a few years, users were abandoning it in droves after a number of highly publicized meetups between adults and children who found each other online. The moral panic caused the U.S. government to consider legislation to control online social networking (see box).[39]

Do You Know Where Your Child Is on the Web Tonight?

It was every parent's worst nightmare. Katherine Lester, a 16-year-old honors student from Fairgrove, Michigan, went missing in June 2006. Her parents had no idea what had happened to her; she had never given them a moment's worry. They called the police. Then federal authorities got involved.

After three days of terrifying absence, she was found, safe—in Amman, Jordan.

Fairgrove is too small to have a post office, and the Lesters lived in the last house on a dead-end street. In another time, Katherine's school, 6 miles away, might have been the outer limit of her universe. But through the Internet, her universe was the whole world. Katherine met a Palestinian man, Abdullah Jinzawi, from Jericho on the West Bank. She found his profile on the social networking website Myspace and sent him a message: "u r cute." They quickly learned everything about each other through online messages. Lester tricked her mother into getting her a passport and then took off for the Middle East. When U.S. authorities met her plane in Amman, she agreed to return home and apologized to her parents for the distress she had caused them.

A month later, Representative Judy Biggert of Illinois rose in the House to co-sponsor the Deleting Online Predators Act (DOPA). "MySpace.com and other networking websites have become new hunting grounds for child predators," she said, noting that "we were all horrified" by the story of Katherine Lester. "At least let's give parents some comfort that their children won't fall prey while using the Internet at schools and libraries that receive federal funding for Internet services." The law would require those institutions to prevent children from using on-location computers to access chat rooms and social networking websites without adult supervision.

Speaker after speaker rose in the House to stress the importance of protecting children from online predators, but not all supported the bill. The language was "overbroad and ambiguous," said one. As originally drafted, it seemed to cover not just Myspace but sites such as Amazon and Wikipedia. These sites possess some of the same characteristics as Myspace: Users can create personal profiles and continually share information with each other by using the Web. Although the law might block children in schools and libraries from "places" where they meet friends (and sometimes predators), it would also prevent access to online encyclopedias and bookstores, which rely on content posted by users.

Instead of taking the time to develop a sharper definition of what exactly was to be prohibited, DOPA's sponsors hastily redrafted the law to omit the definition, leaving it to the Federal Communications Commission to decide later just what the law would cover. Some murmured that the upcoming midterm elections were motivating the sponsors to put forward an ill-considered and showy effort to protect children—an effort that would likely be ineffective and so vague as to be unconstitutional.

Children use computers in lots of places; restricting what happens in schools and libraries would hardly discourage determined teenagers from sneaking onto Myspace. Only the most overbearing parents could honestly answer the question *USA Today* asked in its article about "cyber-predators": "It's 11 p.m. Do you know where your child is on the Web tonight?"

The statistics about what can go wrong were surely terrifying. The Justice Department has made thousands of arrests for "cyber enticement"—almost always older men using social networking websites to lure teenagers into meetings, some of which end very badly. Yet, as the American Library Association stated in opposition to DOPA, education, not prohibition, is the "key to safe use of the Internet." Students have to learn to cooperate online because network use and all the human interactions it enables are basic tools of the new, globally interconnected world of business, education, and citizenship—and perhaps even the globally interconnected world of true love.

The tale of Katherine Lester took an unexpected turn. From the moment she was found in Jordan, Lester steadily insisted that she intended to marry Jinzawi. Jinzawi, who was 20 when he and Lester first made contact, claimed to be in love with her—and his mother loved her, too. Jinzawi begged Lester to tell her parents the truth before she headed off to meet him, but she refused. Upon her return, authorities charged Lester as a runaway child and took her passport away from her. But on September 12, 2007, having attained legal independence by turning 18, she again boarded a plane to the Middle East, finally to meet her beloved face to face. The affair finally ended a few weeks later in an exchange of accusations and denials, as well as a hint that a third party had attracted Lester's attentions. There was no high-tech drama to the breakup—except that it was televised on *Dr. Phil*.

This was the environment in which Facebook was launched in a Harvard dorm room in 2004. Zuckerberg had been studying sociology and psychology and computer networking, and he hacked together a simple website that he proposed to call "Six Degrees to Harry Lewis." As he wrote to Lewis,

> Professor, I've been interested in graph theory and its applications to social networks for a while now, so I did some research...that has to do with linking people through articles they appear in from the Crimson. I thought people would find this interesting, so I've set up a preliminary site that allows people to find the connection (through people and articles) from any person to the most frequently mentioned person in the time frame I looked at. This person is you. I wanted to ask your permission to put this site up though, since it has your name in its title.

Lewis briefly demurred. "Can I see it before I say yes? It's all public information, but there is somehow a point at which aggregation of public information feels like an invasion of privacy." Shortly thereafter, thinking that the project sounded educational, Lewis replied, "Sure, what the hell. Seems harmless." thefacebook, as it was originally known, launched a week later.

Within two years, Facebook had overtaken Myspace. It has succeeded through a combination of good decisions:

- Good engineering (it was generally reliable even in its most explosive growth phase)
- Design that balanced, more successfully than its competitors, the opposing imperatives to connect the world and to provide spaces for more intimate conversation between birds of a feather

- An interface that was calmer and more standardized than that of Myspace, perhaps reflecting that Facebook had originated as a community of college students rather than indie rock music lovers
- An advertising model that was largely palatable to its users (in part because they were unaware of the extent to which their data was repurposed)
- A large number of strategic acquisitions that had the combined effect of making Facebook a single-stop platform not just for social networking but for text messaging, video search, photo storing and viewing, shopping, and gaming, among other things

The result was growth at an astonishing rate.[40] Facebook was launched on February 4, 2004, as a site just for Harvard students, a replacement for the "face books" printed by the Harvard student residences to familiarize students with each other. A month later the network was extended to Stanford, Yale, and Columbia, and by the end of the year it had more than 1 million users. The site may have become more popular because it began with an aura of exclusivity. During 2005 it added hundreds more colleges as well as high schools, and by the end of 2006 anyone could join, and the user base was over 12 million. A year later it was up to 60 million, and it hit 500 million by mid-2010. At the time of this writing, Facebook puts the number of users at 1.59 billion daily users, with 2.41 billion who use it at least once a month. That's about a third of the population of the earth, including infants. The number is still rising, in spite of adverse publicity due to misuse of the company's data.

Indeed, Facebook has a lot of data about its users, and it has done a lot of "learning on the job" about how to handle it. It was aware, early on, that privacy would be important to users and stated unequivocally in 2007, "We do not and will not use cookies to collect private information from any user."[41] But only a few months later, Facebook user Sean Lane bought a diamond eternity ring online—and his wife learned about it instantly, from Facebook. Facebook had recently launched a new feature called Beacon. In an attempt to keep Facebook friends up to date about what users were doing—and also to expand the opportunities for advertising on Facebook—Beacon posted information about what users were buying from non-Facebook sites on friends' news feeds. Lane's wife not only learned about the ring he had purchased but that he had gotten a 51% discount on it. "Who is the ring for?" she asked.[42] It was for her, so only the surprise was spoiled, and not the marriage!

Facebook had partnered with other sites in an information-sharing scheme. When users made a purchase on a partner site, Facebook was informed, and it would sometimes insert the information in friends' news feeds. Users could opt out—if they noticed the tiny box on the partner site and understood what they were being invited to opt out of. Tens of thousands of users were furious

and petitioned the company to remove the feature. Matters got worse when a researcher discovered that under certain circumstances, the information was sent to Facebook even when the user was logged out and no opt-out box was shown. Denials by company spokespeople proved to be inaccurate. Lawsuits ensued. Zuckerberg first apologized,[43] had to pay millions of dollars to settle, and then shut down the Beacon feature entirely.

But the financial penalty didn't prevent further privacy flubs. In late 2009, when the network had grown to 350 million users, its privacy policy was updated without prior notice. Wrapped in an announcement[44] touting that users would now be expected to "personalize their privacy," the company noted that the default privacy settings had changed. It had been the case that only a user's name and "network" were visible to the outside world. ("Networks" were a vestige of the days when members of only certain groups could join Facebook—a user's network was his or her college or high school, for example.) According to the new policy,[45]

> Certain categories of information such as your name, profile photo, list of friends and pages you are a fan of, gender, geographic region, and networks you belong to are considered publicly available to everyone, including Facebook-enhanced applications, and therefore do not have privacy settings. You can, however, limit the ability of others to find this information through search using your search privacy settings.

The announcement noted that most people make this information public anyway. Perhaps so, but there was a big difference between what some people chose to do and what others were expecting. In short order it became apparent how revealing friends lists and fan pages might be. MIT researchers, for example, found that it was not difficult to figure out, with high accuracy, who was gay, even in the absence of any explicit information about sexual orientation:

> Public information about one's coworkers, friends, family, and acquaintances, as well as one's associations with them, implicitly reveals private information....Our research demonstrates a method for accurately predicting the sexual orientation of Facebook users by analyzing friendship associations. After analyzing 4,080 Facebook profiles from the MIT network, we determined that the percentage of a given user's friends who self-identify as gay male is strongly correlated with the sexual orientation of that user, and we developed a logistic regression classifier with strong predictive power.[46]

The aggregation of enough public information does indeed constitute an invasion of privacy—something Zuckerberg seems not to have thought

through. After the Facebook defaults changed, he discovered that his own photographs had become public, and he quickly reset his privacy settings.[47] Another Facebook executive suggested that users who did not want their hometown to be made public should lie about it—apparently forgetting that Facebook requires such information to be truthful.[48]

The reaction was vocal, not just from users but from government officials and privacy organizations. On April 27, 2010, Zuckerberg received a polite but ominous letter from four U.S. senators that concluded,

> We look forward to the FTC examining this issue, but in the meantime we believe Facebook can take swift and productive steps to alleviate the concerns of its users. Providing opt-in mechanisms for information sharing instead of expecting users to go through long and complicated opt-out processes is a critical step towards maintaining clarity and transparency.[49]

In May of 2010, Facebook reversed the defaults, so that only name, photo, gender, and networks would automatically be public.[50]

In spite of the outcry from users and the hint of forthcoming involvement by a federal agency, in the few months between the ill-considered change of privacy defaults and the decision to revert to the previous assumptions, Facebook had added 50 million users. People were complaining but were finding Facebook too useful to give up. Every new member of the network made it that much more valuable to join. People went to Facebook not to make new friends but because all their friends were already there. This phenomenon is known as a *network* effect: As Paul Baran had anticipated (see page 92), the value to an individual of being in the network increases as the size of the network as a whole increases. After 2010, Facebook had no serious social network competition in the United States, although in some parts of the world, Facebook was not even allowed to sign up users.

As most of the United States and a large number of non-U.S. individuals joined, the network effect was boosted by product diversification. Facebook added a text messaging service in 2008 and acquired Instagram in 2012 and WhatsApp in 2014. Facebook had become a full-service platform for all kinds of communication—good, bad, and fraudulent. In 2019 it was reported that Facebook was used for 90% of reported instances of sharing of child pornography.[51] The bulk of that communication was via Facebook's Messenger app, and when Facebook announced that it would be adding end-to-end encryption to Messenger, which would make it impossible for anyone but the recipient to decipher Messenger communications en route, the U.S. attorney general and his counterparts in other countries firmly asked Zuckerberg not to follow

through.[52] The letter cited one of thousands of examples of how law enforcement had used electronic surveillance to catch a criminal:

> To take one example, Facebook sent a priority report to NCMEC, having identified a child who had sent self-produced child sexual abuse material to an adult male. Facebook located multiple chats between the two that indicated historical and ongoing sexual abuse. When investigators were able to locate and interview the child, she reported that the adult had sexually abused her hundreds of times over the course of four years, starting when she was 11. He also regularly demanded that she send him sexually explicit imagery of herself. The offender, who had held a position of trust with the child, was sentenced to 18 years in prison. Without the information from Facebook, abuse of this girl might be continuing to this day.

Encryption would make it impossible to catch such offenders. "We therefore call on Facebook and other companies," the letter continued, "whatever form of encryption they use, to enable law enforcement to obtain lawful access to content in a readable and usable format."

Chapter 5, "Secret Bits," traces the history of encryption. Most privacy and security experts agree that allowing law enforcement access to encrypted communications greatly increases the risk that others will be able to gain access to them as well. But an even larger issue is at stake: Facebook and Google have more information about most human beings and every form of human activity than any government. The mistakes they make—data leaks and service interruptions, for example—can affect significant portions of the population of the earth. In the case of Facebook, Mark Zuckerberg alone can, in principle, make the decisions by himself.

As a vehicle for many forms of communication, Facebook is under no obligation to allow every lawful form of speech—and given the protections that Internet companies enjoy under Section 230 (see Chapter 7, "You Can't Say That on the Internet") of the Communications Decency Act, their responsibility even for unlawful forms of speech is limited in the United States. They recognize, however, both a business interest and a political interest in limiting some forms of speech even in the United States, and abroad their obligation to censor is much more explicit. Fraudulent political ads may have affected voters in the 2016 U.S. presidential election. Several mass shooters were steeped in violent social media content.

Yet these companies can largely set policies—what and whether to censor and whether and how to encrypt messages, for example—as they see fit. They will doubtless tend to make those decisions with their shareholders' interests in mind; it would actually be improper for them to do otherwise routinely. Of course, those interests are best served by following policies of which the public will generally

approve. But does the public itself have an interest in intervening, as Attorney General Barr suggests, and as Senator Schumer had suggested earlier?

And having a policy is not the same as executing it perfectly. Human review of every comment or advertisement or video posted to Facebook is impossible; even responding in a timely manner to those that users complain about would be incredibly difficult. Inevitably the gatekeepers have turned to artificial intelligence (AI) software to do some of the screening for them. But where intent and context are important, AI is not yet a match for human readers.

What if a word or phrase in a post trips Facebook's "hate speech" prohibition, but it is taken out of context? An algorithm cannot make that judgement. Reluctant to be accused of political bias, Facebook has announced that it won't generally remove political advertising, even when it is known to be blatantly false. Its decision, it says, "is grounded in Facebook's fundamental belief in free expression, respect for the democratic process, and the belief that, in mature democracies with a free press, political speech is arguably the most scrutinized speech there is."[53] Is this what will best serve the public interest? It is not surprising that not everyone agreed. What to do instead is far less clear.

Or are these companies just too big? That notion became popular among some of the presidential candidates during the 2020 election cycle. Elizabeth Warren thinks Facebook should be broken up—perhaps by unwinding some of its acquisitions. Whether that would be either legal or helpful will excite intense discussion. Another idea would be to leave them intact but to regulate them more tightly—though the devil would be in the details. It would not be a small step on the part of the government to treat a private company whose product is bits as though it were a public utility whose product is water.

One part of the government has particular concerns about a specific aspect of the services offered by technology companies. As a result of one of the most remarkable discoveries of the twentieth century—just a little bit of arithmetic on bits—private citizens can and do now exchange over the public Internet encrypted messages that law enforcement can intercept but cannot decode. How this happened and what it portends is the subject of the next chapter.

Endnotes

1 Ian Austen, "A Canadian Telecom's Labor Dispute Leads to Blocked Web Sites and Questions of Censorship," *The New York Times*, August 1, 2005, https://www.nytimes.com/2005/08/01/business/worldbusiness/a-canadian-telecoms-labor-dispute-leads-to-blocked.html.

2 Tripp Mickle et al., "Apple, Google Pull Hong Kong Protest Apps Amid China Uproar," *Wall Street Journal*, October 10, 2019, https://www.wsj.com/articles/apple-pulls-hong-kong-cop-tracking-map-app-after-china-uproar-11570681464.

3 Jim Rutenberg, "Netflix's Bow to Saudi Censors Comes at a Cost to Free Speech," *The New York Times*, January 6, 2019, https://www.nytimes.com/2019/01/06/business/media/netflix-saudi-arabia-censorship-hasan-minhaj.html.

4 Steve Kroft, "How Did Google Get so Big?" CBS News, May 21, 2018, https://www.cbsnews.com/news/how-did-google-get-so-big/.

5 "Internet Providers in Browning, Montana," Broadband Now, accessed April 27, 2020, https://broadbandnow.com/Montana/Browning.

6 "Browning, MT," Data USA, accessed April 27, 2020, https://datausa.io/profile/geo/browning-mt/.

7 Paul Baran, "On Distributed Communications Networks," (RAND Corporation, Santa Monica, CA, September 1962), Reprinted with permission. https://www.rand.org/pubs/papers/P2626.html.

8 Vinton G. Cerf and Robert E Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Transactions on Communications*, no. 5 (1974): 13.

9 Pete Resnick, "On Consensus and Humming in the IETF," Internet Engineering Task Force, June 2014, https://tools.ietf.org/html/rfc7282.

10 Harald Tveit Alvestrand, "A Mission Statement for the IETF," Internet Engineering Task Force, October 2004, https://tools.ietf.org/html/rfc3935.

11 The Internet protocol packet design was set before the days of massive miniaturization and at a time when computer memory was limited and costly. Nobody then imagined the need to connect more than 4 billion computers to the Internet, so only 32 bits were reserved for the address fields. But now wristwatches and refrigerators have their own IP addresses, and the total number of connected computers is more than can be distinguished using 32 bits. Various workarounds have been developed, and a new protocol, IPv6, which has 128-bit addresses, is slowly being rolled out.

12 Jonathan Zittrain, *The Future of the Internet—and How to Stop It* (Yale University Press, 2008).

13 For a full account of the story of fiber, see Susan Crawford, Fiber: *The Coming Tech Revolution—and Why America Might Miss It*, Yale University Press, 2018.

14 Jon Brodkin, "US Broadband: Still No ISP Choice for Many, Especially at Higher Speeds," *Ars Technica*, August 10, 2016, https://arstechnica.com/information-technology/2016/08/us-broadband-still-no-isp-choice-for-many-especially-at-higher-speeds/.

15 Kendra Chamberlain, "Municipal Broadband Is Roadblocked or Outlawed in 25 States," *Broadband Now*, May 13, 2020, https://broadbandnow.com/report/municipal-broadband-roadblocks/.

16 "Government Competition with Private Internet Services Providers Prohibited—Exceptions," Montana Code Annotated 2019, https://leg.mt.gov/bills/mca/title_0020/chapter_0170/part_0060/section_0030/0020-0170-0060-0030.html.

17 Paul Baran, "Full Text of 'The Computer and Invasion of Privacy,'" July 26, 1966, https://archive.org/stream/U.S.House1966TheComputerAndInvasionOfPrivacy/U.S.%20House%20%281966%29%20-%20The%20Computer%20and%20Invasion%20of%20Privacy_djvu.txt.

18 Scott Bradner, "The Internet: Unblocking Pipes," Network World, March 14, 2005, https://www.networkworld.com/article/2319666/the-internet--unblocking-pipes.html.

19 Eva Wolchover, "Web Reconnects Cousins Cut off by Iron Curtain," *Boston Herald*, December 18, 2007.

20 "Search Engine Market Share Worldwide 2019," Statista, accessed April 27, 2020, https://www.statista.com/statistics/216573/worldwide-market-share-of-search-engines/.

21 Nathaniel Popper, "A Feisty Google Adversary Tests How Much People Care About Privacy," *The New York Times*, July 15, 2019, https://www.nytimes.com/2019/07/15/technology/duckduckgo-private-search. html.

22 Vannevar Bush, "As We May Think," *The Atlantic*, July 1, 1945, https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/.

23 H. G. Wells, *World Brain* (Methuen, 1938), pp. 60–61.

24 "Shady Web Searches in Missing Girl Case," CBS News, November 26, 2008, https://www.cbsnews.com/news/shady-web-searches-in-missing-girl-case/.

25 Tony Pipitone, "Cops, Prosecutors Botched Casey Anthony Evidence," WKMG, November 28, 2012, https://www.clickorlando.com/news/2012/11/28/cops-prosecutors-botched-casey-anthony-evidence/.

26 K. C. Jones, "Ex-Computer Consultant Convicted in 'Google Murder' Trial," *InformationWeek*, November 30, 2005, https://www.informationweek.com/ex-computer-consultant-convicted-ingoog/174403074.

27 *People v. Zirko*, 2012 IL App (1st) 092158.

28 George Knapp and Matt Adams, "I-Team: Details the Night Attorney Susan Winters Died," 8NewsNow, February 10, 2017, https://www.8newsnow.com/news/i-team-details-the-night-attorney-susan-winters-died/.

29 "Requests for User Information," Google Transparency Report, accessed April 27, 2020, https://transparencyreport.google.com/user-data/overview.

30 "US Judge Asks Reports of Google Searches," SEL, accessed April 27, 2020, https://searchenginelaw.net/security/103-us-judge-asks-reports-of-google-searches.

31 Douglas MacMillan, "Tech's 'Dirty Secret': The App Developers Sifting Through Your Gmail," *Wall Street Journal*, July 2, 2018, https://www.wsj.com/articles/techs-dirty-secret-the-app-developers-sifting-through-your-gmail-1530544442.

32 https://support.earny.co/hc/en-us/articles/218609757-Privacy-Policy#:~:text=We%2C%20at%20Earny%20Inc.%2C,save%20money%20in%20multiple%20ways.

33 John D. McKinnon and Douglas MacMillan, "Google Says It Continues to Allow Apps to Scan Data from Gmail Accounts," *Wall Street Journal*, September 20, 2018, https://www.wsj.com/articles/google-says-it-continues-to-allow-apps-to-scan-data-from-gmail-accounts-1537459989.

34 Peter H. Lewis, "Digital Equipment Offers Web Browsers Its 'Super Spider,'" *The New York Times*, December 18, 1995, https://www.nytimes.com/1995/12/18/business/digital-equipment-offers-web-browsers-its-super-spider.html.

35 Sergey Brin and Lawrence Page, "The Anatomy of a Large-Scale Hypertextual Web Search Engine," *Computer Networks and ISDN Systems* 30, no. 1–7 (April 1998): 107–117, https://snap.stanford.edu/class/cs224w-readings/Brin98Anatomy.pdf.

36 Mark Scott, "Google Fined Record $2.7 Billion in E.U. Antitrust Ruling," *The New York Times*, June 27, 2017, https://www.nytimes.com/2017/06/27/technology/eu-google-fine.html.

37 Danah M Boyd and Nicole B. Ellison, "Social Network Sites: Definition, History, and Scholarship," *Journal of Computer-Mediated Communication* 13, no. 1 (October 1, 2007): 210–30, https://doi.org/10.1111/j.1083-6101.2007.00393.x.

38 "Friendster," June 11, 2004, https://web.archive.org/web/20040611192459/http://www.friendster.com/index.jsp.

39 Pete Cashmore, "MySpace, America's Number One," Mashable, July 11, 2006, https://mashable.com/2006/07/11/myspace-americas-number-one/.

40 "Company Info," *About Facebook*, accessed April 28, 2020, https://about.fb.com/company-info/.

41 thefacebook, "Privacy Policy," January 7, 2005, https://web.archive.org/web/20050107221705/http:/www.thefacebook.com/policy.php.

42 Bill Goodwin and Sebastian Klovig Skelton, "Facebook's Privacy Game—How Zuckerberg Backtracked on Promises to Protect Personal Data," ComputerWeekly.com, July 1, 2019, https://www.computerweekly.com/feature/Facebooks-privacy-U-turn-how-Zuckerberg-backtracked-on-promises-to-protect-personal-data.

43 Facebook, "Thoughts on Beacon," December 5, 2007, https://www.facebook.com/notes/facebook/thoughts-on-beacon/7584397130/.

44 "Facebook Asks More Than 350 Million Users Around the World to Personalize Their Privacy," *About Facebook*, December 10, 2009, https://about.fb.com/news/2009/12/facebook-asks-more-than-350-million-users-around-the-world-to-personalize-their-privacy/.

45 Kurt Opsahl, "Facebook's Eroding Privacy Policy: A Timeline," Electronic Frontier Foundation, April 28, 2010, https://www.eff.org/deeplinks/2010/04/facebook-timeline.

46 Carter Jernigan and Behram F. T. Mistree, "Gaydar: Facebook Friendships Expose Sexual Orientation," *First Monday*, September 25, 2009, https://doi.org/10.5210/fm.v14i10.2611.

47 Kashmir Hill, "Either Mark Zuckerberg Got a Whole Lot Less Private or Facebook's CEO Doesn't Understand the Company's New Privacy Settings," *Forbes*, December 10, 2009, https://www.forbes.com/sites/kashmirhill/2009/12/10/either-mark-zuckerberg-got-a-whole-lot-less-private-or-facebooks-ceo-doesnt-understand-the-companys-new-privacy-settings/.

48 Julia Angwin, "How Facebook Is Making Friending Obsolete," *The Wall Street Journal*, December 15, 2009, https://www.wsj.com/articles/SB126084637203791583.

49 Politico Staff, "Senators' Letter to Facebook," Politico, April 27, 2010, https://www.politico.com/news/stories/0410/36406.html.

50 "Facebook Redesigns Privacy," *About Facebook*, May 26, 2010, https://about.fb.com/news/2010/05/facebook-redesigns-privacy/.

51 Jennifer Valentino-DeVries and Gabriel J. X. Dance, "Facebook Encryption Eyed in Fight Against Online Child Sex Abuse," *The New York Times*, October 2, 2019, https://www.nytimes.com/2019/10/02/technology/encryption-online-child-sex-abuse.html.

52 Priti Patel et al., "Open Letter to Facebook," October 4, 2019, https://www.justice.gov/opa/press-release/file/1207081/download.

53 Cecilia Kang, "Facebook's Hands-off Approach to Political Speech Gets Impeachment Test," *The New York Times*, October 8, 2019, https://www.nytimes.com/2019/10/08/technology/facebook-trump-biden-ad.html.

CHAPTER 5

Secret Bits

How Codes Became Unbreakable

Going Dark

The section "Bits Cubed: The Snowden Files" in Chapter 2, details how Edward Snowden left his government consulting job and flew to Hong Kong (see the section), taking some thumb drives and laptop computers with him. On those devices were thousands of classified documents. Soon *The Washington Post* and other news outlets began to publish the details of a variety of secret surveillance programs run by the U.S. government and its allies.[1]

PRISM targeted the technology companies, such as Google and Yahoo!, through which vast amounts of email flowed. MUSCULAR cracked the data flows within these companies. Dishfire was specialized to text messages. XKeyscore targeted the great global fiber network that holds the Internet together. Verizon and other telecommunications companies had for years been turning over to the government information about the telephone calls of ordinary Americans—without warrants, under orders from a secret court few Americans knew existed.[2]

During a subsequent congressional inquiry, the director of the National Security Agency was unapologetic and suggested that the agency's programs needed, if anything, to be expanded. "Yes, I believe it is in the nation's best interest to put all the phone records into a lockbox that we could search when the nation needs to do it," he said.[3]

The upshot of the Snowden revelations was dramatic. Americans—and their international correspondents, diplomats, and business partners—started to want lockboxes of their own: ways to encrypt their messages so that only the sender

and the recipient could read them. The technology already existed, and in short order, email providers and cellular communications providers such as Verizon began to make it easier and more ordinary for messages to be encrypted, both in transit and at their source and destination.[4] The percentage of email that is encrypted in transit rose from barely 5% pre-Snowden[5] to at least 90% today.[6]

By 2017, encryption had become so routine and widespread that law enforcement found it increasingly difficult to decipher the communications of criminal suspects. The FBI reported that over the previous year, its lab had received 7,775 mobile devices it was unable to decrypt, in spite of having both the legal authority to do so and some of the world's best code-cracking tools.[7] "Being unable to access nearly 7,800 devices in a single year is a major public safety issue," the FBI director warned. The data was right in their hands, but it might as well have been on Pluto. There was no way to get at it. The number turned out to have been inflated, though; 2,000 phones was probably closer to the truth.[8] In any case, law enforcement was alarmed about its inability to solve crimes even when it was in possession of the critical information.

Rod Rosenstein, deputy U.S. attorney general, spoke at the U.S. Naval Academy about what was needed to keep the nation safe from criminals—and terrorists, too. Rosenstein's target was not the criminals and terrorists, however, but the technology companies that made the tools they used.[9] The companies were profiting from encryption technologies the government couldn't crack, and they had a civic obligation to help U.S. law enforcement. After all, they sometimes cooperated with foreign governments. They sometimes worked with governments to make censorship more effective, for example, and made good money doing so. Perhaps they also made good money in the United States by offering encrypted messaging, but should their profit motive be sufficient reason not to help out the U.S. government too? Rosenstein warned against "going dark...the threat to public safety that occurs when service providers, device manufacturers, and application developers deprive law enforcement and national security investigators of crucial investigative tools."

The solution, Rosenstein continued, was "responsible encryption." That is, the companies should distribute only encryption that the government could bypass or crack.

Skeptics and privacy advocates reacted with horror to Rosenstein's call for "responsible" encryption. Past studies and daily experience with data breaches had already established secure "key escrow" as an implausible idea—the notion that some third party, either the government or the technology companies themselves, could be trusted to hold keys until law enforcement asked for them. And to limit the strength of encryption technologies would only make communications less secure. In fact, it seemed, the only way to go back to the days when the government could crack every code was to repeal the laws of mathematics that made secure encryption possible. The nation had

already had these arguments only a few years before, and strong encryption was the winner. Congress had come close to making laws about encryption and backed off—for very good reasons.

Encryption in the Hands of Terrorists—and Everyone Else

September 13, 2001. Fires were still smoldering in the wreckage of the World Trade Center when Judd Gregg of New Hampshire rose to tell the Senate what had to happen. He recalled the warnings issued by the FBI years before the country had been attacked: The FBI's most serious problem was "the encryption capability of the people who have an intention to hurt America."

"It used to be," the senator went on, "that we had the capability to break most codes because of our sophistication."[10] No more. "We can give our code breakers all the money in the world, but the technology has outstripped the code breakers,"[11] he warned. Even civil libertarian cryptographer Phil Zimmermann, whose encryption software appeared on the Internet in 1991 for use by human rights workers worldwide, agreed that the terrorists were probably encoding their messages. "I just assumed," he said, "somebody planning something so diabolical would want to hide their activities using encryption."[12]

Encryption is the art of encoding messages so they can't be understood by eavesdroppers or adversaries into whose hands the messages might fall. Descrambling an encrypted message requires knowing the sequence of symbols—the "key"—that was used to encrypt it. An encrypted message may be visible to the world, but without the key, it might as well be hidden in a locked box. Without the key—exactly the right key—the contents of the box, or the message, remains secret.

Anticipating Mr. Rosenstein's appeal for the industry to act responsibly, Senator Gregg called for "the cooperation of the community that is building the software, producing the software, and building the equipment that creates the encoding technology."[13] That is, he called for cooperation enforced by legislation. The makers of encryption software would have to enable the government to bypass the locks and retrieve the decrypted messages. And what about encryption programs written abroad, which could be shared around the world in the blink of an eye, as Zimmermann's had been? The United States should use "the market of the United States as leverage" in getting foreign manufacturers to follow U.S. requirements for so-called back doors that could be used by the U.S. government.

By September 27, Gregg's legislation was beginning to take shape. The keys used to encrypt messages would be held in escrow by the government under tight security. There would be a "quasi-judicial entity," appointed by the Supreme Court, which would decide when law enforcement had made

its case for release of the keys. Civil libertarians squawked, and doubts were raised as to whether the key escrow idea could actually work. No matter, opined the senator in late September. "Nothing's ever perfect. If you don't try, you're never going to accomplish it. If you do try, you've at least got some opportunity for accomplishing it."[14]

Abruptly, three weeks later, Senator Gregg dropped his legislative plan. "We are not working on an encryption bill and have no intention to," said the Senator's spokesman on October 17.[15]

On October 24, 2001, Congress passed the USA PATRIOT Act, which gave the FBI sweeping new powers to combat terrorism. But the PATRIOT Act does not mention encryption. More than a decade passed before the Snowden revelations led the United States to make another serious attempt to legislate control over cryptographic software.

Why Not Regulate Encryption?

Throughout the 1990s, the FBI had made control of encryption its top legislative priority. Senator Gregg's proposal was a milder form of a bill, drafted by the FBI and reported out favorably by the House Select Committee on Intelligence in 1997, which would have mandated a five-year prison sentence for selling encryption products that could not be immediately decrypted by authorized officials.[16]

How could regulatory measures that law enforcement deemed critical in 1997 for fighting terrorism drop off the legislative agenda four years later, in the aftermath of the worst terrorist attack ever suffered by the United States of America?

No technological breakthrough in cryptography in the fall of 2001 had legislative significance. There also weren't any relevant diplomatic breakthroughs. No other circumstances conspired to make the use of encryption by terrorists and criminals an unimportant problem. It was just that something else about encryption had become accepted as more important: the explosion of commercial transactions over the Internet. Congress suddenly realized that it had to allow banks and their customers to use encryption tools, as well as airlines and their customers, and eBay and Amazon and their customers. Anyone using the Internet for commerce needed the protection that encryption provided. Very suddenly, there were millions of such people—so many that the entire U.S. and world economy depended on public confidence in the security of electronic transactions.

The tension between enabling secure conduct of electronic commerce and preventing secret communication among outlaws had been in the air for a decade. Senator Gregg was but the last of the voices calling for restrictions on encryption. The National Research Council had issued a report of nearly 700

pages in 1996 that weighed the alternatives. The report concluded that, on balance, efforts to control encryption would be ineffective and that their costs would exceed any imaginable benefit. The intelligence and defense establishment was not persuaded. FBI Director Louis Freeh testified before Congress in 1997 that "Law enforcement is in unanimous agreement that the widespread use of robust non-key recovery [i.e., non-escrowed] encryption ultimately will devastate our ability to fight crime and prevent terrorism."[17]

Yet only four years later, even in the face of the September 11 attack, the needs of commerce admitted no alternative to widespread dissemination of encryption software to every business in the country, as well as to every home computer from which a commercial transaction might take place. In 1997, average citizens, including elected officials, might never have bought anything online. Congress members' families might not have been regular computer users. By 2001, all that had changed: The digital explosion was happening. Computers had become consumer appliances, Internet connections were common in American homes, and awareness of electronic fraud had become widespread. Consumers did not want their credit card numbers, birth dates, and Social Security numbers exposed on the Internet.

Why is encryption so important to Internet communications that Congress was willing to risk terrorists using encryption so that American businesses and consumers could use it, too? After all, information security is not a new need. People communicating by postal mail, for example, have reasonable assurances of privacy without any use of encryption.

The answer lies in the Internet's open architecture. The data packets that move through the Internet—each about 1,500 bytes—are not like envelopes sent through postal mail, with an address on the outside and contents hidden. They are like postcards, with everything exposed for anyone to see. As the packets pass through routers, which are located at the switching points, they are stored, examined, checked, analyzed, and sent on their way. Even if all the fibers and wires could be secured, wireless networks would allow bits to be grabbed out of the air without detection.

If you send your credit card number to a store in an ordinary email, you might as well stand in Times Square and shout it at the top of your lungs. By 2001, a lot of credit card numbers were traveling as bits though glass fibers and through the air, and it was impossible to prevent snoopers from looking at them.

The way to make Internet communications secure—to make sure that no one but the intended recipient knows what is in a message—is for the sender to encrypt the information so that only the recipient can decrypt it. If that can be accomplished, then eavesdroppers along the route from sender to receiver can examine the packets all they want, but they will only find an undecipherable scramble of bits.

In a world awakening to Internet commerce, encryption could no longer be thought of as it had been from ancient times until the turn of the third millennium: as armor used by generals and diplomats to protect information critical to national security. Even in the early 1990s, the State Department demanded that an encryption researcher register as an international arms dealer. Now suddenly, encryption was less like a weapon and more like the armored cars used to transport cash on city streets, except that these armored cars were needed by everyone. Encryption was no longer a munition; it was money.

The commoditization of a critical military tool was more than a technology shift. It sparked, and continues to spark, a rethinking of fundamental notions of privacy and of the balance between security and freedom in a democratic society.

"The question," posed MIT's Ron Rivest, one of the world's leading cryptographers, during one of the many debates over encryption policy that occurred during the 1990s, "is whether people should be able to conduct private conversations, immune from government surveillance, even when that surveillance is fully authorized by a Court order."[18] In the post-2001 atmosphere that produced the PATRIOT Act, it's far from certain that Congress would have responded to Rivest's question with a resounding "Yes." But by 2001, commercial realities had overtaken the debates.

To fit the needs of electronic commerce, encryption software had to be widely available. It had to work perfectly and quickly, with no chance of anyone cracking the codes. And there was more: Although encryption had been used for more than four millennia, no method known until the late twentieth century would have worked well enough for Internet commerce. But in 1976, two young mathematicians, operating outside the intelligence community that was the center of cryptography research, published a paper that made a reality out of a seemingly absurd scenario: Two parties work out a secret key that enables them to exchange messages securely—even if they have never met and all their messages to each other are in the open, for anyone to hear. With the invention of *public-key cryptography*, it became possible for every man, woman, and child to transmit credit card numbers to Amazon more securely than any general had been able to communicate military orders on which the fate of nations depended 50 years earlier.

Historical Cryptography

Cryptography—"secret writing"—has been around almost as long as writing itself. Ciphers have been found in Egyptian hieroglyphics from as early as 2000 B.C. A *cipher* is a tool for transforming a message into an obscured form, together with a way of undoing the transformation to recover the message.

Suetonius, the biographer of the Caesars, describes Julius Caesar's use of a cipher in his letters to the orator Cicero, with whom he was planning and plotting in the dying days of the Roman Republic:

> If he [Caesar] had anything confidential to say, he wrote it in cipher, that is, by so changing the order of the letters of the alphabet, that not a word could be made out. If anyone wishes to decipher these, and get at their meaning, he must substitute the fourth letter of the alphabet, namely D, for A, and so with the others.[19]

In other words, Caesar used a letter-by-letter translation to encrypt his messages:

ABCDEFGHIJKLMNOPQRSTUVWXYZ

DEFGHIJKLMNOPQRSTUVWXYZABC

To encrypt a message with Caesar's method, replace each letter in the top row by the corresponding letter in the bottom row. For example, the opening of *Caesar's Commentaries,* "Gallia est omnis divisa in partes tres," would be encrypted as:

Plaintext: gallia est omnis divisa in partes tres

Ciphertext: jdoold hvw rpqlv glylvd lq sduwhv wuhv

The original message is called the *plaintext*, and the encoded message is called the *ciphertext*. Messages are decrypted by doing the reverse substitutions.

This method is called the *Caesar shift* or the *Caesar cipher*. The encryption/decryption rule is easy to remember: "Shift the alphabet three places." Of course, the same idea would work if the alphabet were shifted more than three places, or fewer. The Caesar cipher is really a family of ciphers, with 25 possible variations, one for each different amount of shifting.[20]

Caesar ciphers are very simple, and an enemy who knew that Caesar was simply shifting the plaintext could easily try all the 25 possible shifts of the alphabet to decrypt the message. But Caesar's method is a representative of a larger class of ciphers, called *substitution ciphers*, in which one symbol is substituted for another according to a uniform rule (the same letter is always translated the same way).

There are a great many more substitution ciphers than just shifts. For example, we could scramble the letters according to the rule

ABCDEFGHIJKLMNOPQRSTUVWXYZ

XAPZRDWIBMQEOFTYCGSHULJVKN

so that A becomes X, B becomes A, C becomes P, and so on. There is a similar substitution for every way of reordering the letters of the alphabet. The number of different reorderings is

$$26 \times 25 \times 24 \times \cdots \times 3 \times 2$$

which is about 4×10^{26} different methods—10,000 times the number of stars in the universe! It would be impossible to try them all. General substitution ciphers must be secure—or so it might seem.

Breaking Substitution Ciphers

In about 1392, an English author—once thought to be the great English poet Geoffrey Chaucer, although that is now disputed—wrote a manual for use of an astronomical instrument. Parts of this manual, titled *The Equatorie of the Planetis*,[21] were written in a substitution cipher (see Figure 5.1). This puzzle is not as hard as it looks, even though there is very little ciphertext with which to work. We know it is written in English—Middle English, actually—but let's see how far we can get thinking of it as encrypted English.

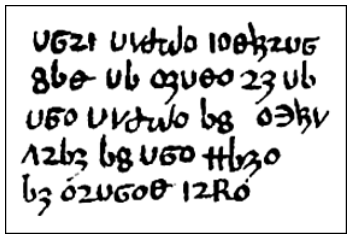

Folio 30v of Peterson MS 75.1, *The Equatorie of Planetis*, a fourteenth century manuscript held at University of Cambridge

FIGURE 5.1 Ciphertext in *The Equatorie of Planetis* (1392).

Although this looks like gibberish, it contains some patterns that may be clues. For example, certain symbols occur more frequently than others. There are 12 ◊s and 10 ∪s, and no other symbol occurs as frequently as these. In ordinary English texts, the two most frequently occurring letters are E and T, so a fair guess is that these two symbols correspond to these two letters. Figure 5.2 shows what happens if we assume that ◊ = E and ∪ = T. The pattern ∪G◊ appears twice and apparently represents a three-letter word beginning with T and ending with E. It could be TIE or TOE, but THE seems more likely, so a reasonable assumption is that G = H. If that is true, what is the four-letter word at the beginning of the text, which begins with TH? Not THAT, because it ends with a new symbol, nor THEN, because the third letter is also new.

Perhaps THIS. And there is a two-letter word beginning with T that appears twice in the second line; that must be TO. Filling in the equivalencies for H, I, S, and O yields Figure 5.3.

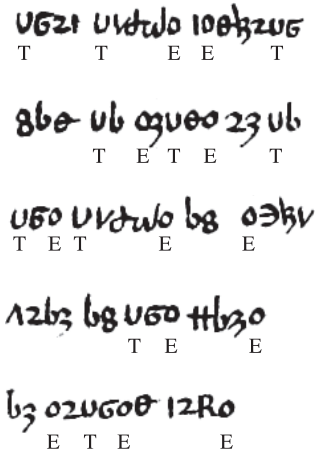

FIGURE 5.2 *Equatorie* ciphertext, with the two most common symbols assumed to stand for E and T.

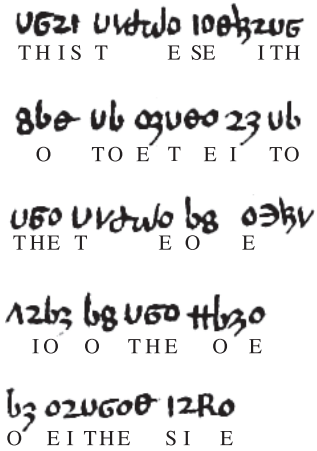

FIGURE 5.3 *Equatorie* ciphertext, with more conjectural decodings.

At this point, the guessing gets easier. Probably the last two words are EITHER SIDE—and the last few symbols can be inferred with a knowledge of Middle English and some idea of what the text is about. The complete plaintext is *This table servith for to entre in to the table of equacion of the mone on either side* (see Figure 5.4).

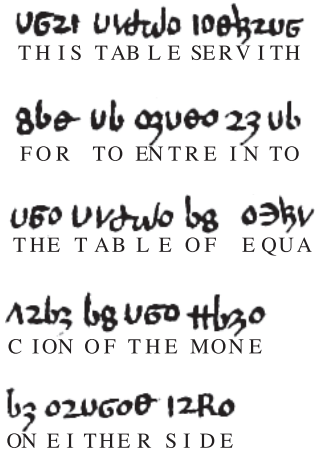

FIGURE 5.4 *Equatorie* ciphertext, fully decoded.

The technique used to crack the code is *frequency analysis*: If the cipher is a simple substitution of symbols for letters, then crucial information about which symbols represent which letters can be gathered from how often the various symbols appear in the ciphertext. This idea was first described by the Arabic philosopher and mathematician Al-Kindi, who lived in Baghdad in the ninth century.

By the Renaissance, this kind of informed guesswork had been reduced to a fine art that was well known to European governments. In a famous example of the insecurity of substitution ciphers, Mary Queen of Scots was beheaded in 1587 due to her misplaced reliance on a substitution cipher to conceal her correspondence with plotters against Queen Elizabeth I. She was not the last to have put too much confidence in an encryption scheme that looked hard to crack but wasn't. Substitution ciphers were in common use as late as the 1800s, even though they had been insecure for a millennium by that time! Edgar Allen Poe's mystery story *The Gold Bug* (1843) and A. Conan Doyle's Sherlock Holmes mystery *Adventure of the Dancing Men* (1903) both turn on the decryption of substitution ciphers.

Secret Keys and One-Time Pads

In cryptography, every advance in code-breaking yields an innovation in code-making. Seeing how easily the *Equatorie* code was broken, what could we do to make it more secure, or *stronger*, as cryptographers would say? We might use more than one symbol to represent the same plaintext letter. A method named for the sixteenth-century French diplomat Blaise de Vigenère uses multiple Caesar ciphers. For example, we can pick 12 Caesar ciphers and

use the first cipher for encrypting the 1st, 13th, and 25th letters of the plaintext; the second cipher for encrypting the 2nd, 14th, and 26th plaintext letters; and so on. Figure 5.5 shows such a Vigenère cipher. A plaintext message beginning SECURE... would be encrypted to produce the ciphertext *llqgrw*..., as indicated by the boxed characters in the figure: S is encrypted using the first row, E is encrypted using the second row, and so on. After we use the bottom row of the table, we start again at the top row and repeat the process over and over.

We can use the cipher of Figure 5.5 without having to send our correspondent the entire table. Scanning down the first column spells out *thomasbbryan*, which is the key for the message. To communicate using Vigenère encryption, the correspondents must first agree on a key. They then use the key to construct a substitution table for encrypting and decrypting messages.

When SECURE was encrypted as *llqgrw*, the two occurrences of E at the second and sixth positions in the plaintext were represented by different ciphertext letters, and the two occurrences of the ciphertext letter *l* represented different plaintext letters. This illustrates how the Vigenère cipher confounds simple frequency analysis, which was the main tool of cryptanalysts at the time. Although the idea may seem simple, the discovery of the Vigenère cipher is regarded as a fundamental advance in cryptography, and the method was considered to be unbreakable for hundreds of years.

Harvard University Archives

FIGURE 5.5 A Vigenère cipher. The key, thomasbbryan, runs down the second column. Each row represents a Caesar cipher in which the shift amount is determined by a letter of the key. (Thomas B. Bryan was an attorney who used this code for communicating with a client, Gordon McKay, in 1894.)

Cryptographers use stock figures for describing encryption scenarios: Alice wants to send a message to Bob, and Eve is an adversary who may be eavesdropping.

Cryptography and History

Cryptography (code-making) and cryptanalysis (code-breaking) have been at the heart of many momentous events in human history. The intertwined stories of diplomacy, war, and coding technology are told beautifully in two books: *The Codebreakers*[22] by David Kahn and *The Code Book*[23] by Simon Singh.

Suppose Alice wants to send Bob a message (see Figure 5.6). The lock-and-key metaphor goes this way: Alice puts the message in a box and locks the box, using a key that only she and Bob possess. (Imagine that the lock on Alice's box is the kind that needs the key to lock it as well as to open it.) If Eve intercepts the box in transit, she has no way to figure out what key to use to open it. When Bob receives the box, he uses his copy of the key to open it. As long as the key is kept secret, it doesn't matter that others can see that there is a box with something in it, and even what kind of lock is on the box. In the same way, even if an encrypted message comes with an announcement that it is encrypted using a Vigenère cipher, it will not be easy to decrypt except by someone who has the key.

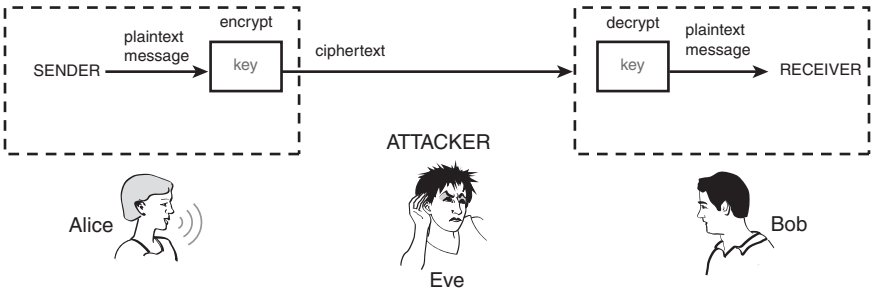

FIGURE 5.6 Standard cryptographic scenario. Alice wants to send a message to Bob. She encrypts it using a secret key. Bob decrypts it using his copy of the key. Eve is an eavesdropper. She intercepts the coded message in transit and tries to decrypt it.

Or at least that's the idea. The Vigenère cipher was actually broken in the mid-1800s by the English mathematician Charles Babbage, who is now recognized as a founding figure in the field of computing. Babbage recognized that if someone could guess or otherwise deduce the length of the key, and hence the length of the cycle on which the Vigenère cipher was repeated, the problem was reduced to breaking several simple substitutions. He then used a brilliant extension of frequency analysis to discover the length of the key. Babbage never published his technique, perhaps at the request of British

Intelligence. A Prussian Army officer, Friedrich Kasiski, independently figured out how to break the Vigenère code and published the method in 1863. The Vigenère cipher has been insecure ever since.

The sure way to beat this attack is to use a key that is as long as the plaintext so that there are no repetitions. If we wanted to encrypt a message of length 100, we might use 100 Caesar ciphers in an arrangement like that of Figure 5.5, extended to 100 rows. Every table row would be used only once. A code like this is known as a *Vernam cipher*, after its World War I–era inventor, AT&T telegraph engineer Gilbert Vernam, or, more commonly, a *one-time pad*.

The term "one-time pad" is based on a particular physical implementation of the cipher. Let's again imagine that Alice wants to get a message to Bob. Alice and Bob have identical pads of paper. Each page of the pad has a key written on it. Alice uses the top page to encrypt a message. When Bob receives it, he uses the top page of his pad to decrypt the message. Both Alice and Bob tear off and destroy the top page of the pad when they have used it. It is essential that the pages not be reused, as doing so could create patterns like those exploited in cracking the Vigenère cipher.

One-time pads were used during the Second World War and the Cold War in the form of booklets filled with digits (see Figure 5.7). Governments still use one-time pads today for sensitive communications, with large amounts of keying material carefully generated and distributed on CDs or DVDs.

National Security Agency

FIGURE 5.7 German one-time pad used for communication between Berlin and Saigon during the 1940s. Encrypted messages identified the page to be used in decryption. The cover warns, "Sheets of this encryption book that seem to be unused could contain codes for messages that are still on their way. They should be kept safe for the longest time a message might need for delivery."

A one-time pad, if used correctly, cannot be broken by cryptanalysis. There are simply no patterns to be found in the ciphertext. There is a deep relationship between information theory and cryptography, which Claude Shannon explored in 1949.[24] (In fact, it was probably his wartime research on this sensitive subject that gave birth to his brilliant discoveries about communication in general.) Shannon proved mathematically what is obvious intuitively: The one-time pad is, in principle, as good as it gets in cryptography. It is absolutely unbreakable—in theory.

But as Yogi Berra said, "In theory, there is no difference between theory and practice. In practice, there is." Good one-time pads are hard to produce. If a pad contains repetitions or other patterns, Shannon's proof that one-time pads are uncrackable no longer holds. More seriously, transmitting a pad between the parties without loss or interception is likely to be just as difficult as communicating the plaintext of the message itself without detection. Typically, the parties would share a pad ahead of time and hope to conceal it in their travels. Big pads are harder to conceal than small pads, however, so the temptation arises to reuse pages—the kiss of death for security.

The Soviet KGB fell victim to exactly this temptation, which led to the partial or complete decryption of more than 3,000 diplomatic and espionage messages by U.S. and British intelligence during the years 1942–1946.[25] The National Security Agency's VENONA project, publicly revealed only in 1995, was responsible for exposing major KGB agents such as Klaus Fuchs and Kim Philby. The Soviet messages were doubly encrypted, using a one-time pad on top of other techniques; this made the code-breaking project enormously difficult. It was successful only because, as World War II wore on and material conditions deteriorated, the Soviets reused the pads.

Because one-time pads are impractical, almost all encryption uses relatively short keys. Some methods are more secure than others, however. Computer programs that break Vigenère encryption are readily available on the Internet, and no professional would use a Vigenère cipher today. Today's sophisticated ciphers are the distant descendants of the old substitution methods. Rather than substituting message texts letter for letter, computers divide an ASCII-encoded plaintext message into blocks. They then transform the bits in the block according to some method that depends on a key. The key itself is a sequence of bits on which Alice and Bob must agree and keep secret from Eve. Unlike the Vigenère cipher, there are no known shortcuts for breaking these ciphers (or at least none known publicly). The best method to decrypt a ciphertext without knowing the secret key seems to be brute-force exhaustive search, trying all possible keys.

The amount of computation required to break a cipher by exhaustive search grows exponentially with the size of the key. Increasing the key length by one bit doubles the amount of work required to break the cipher but only slightly increases the work required to encrypt and decrypt. This is what makes these

ciphers so useful: Computers may keep getting faster—even at an exponential rate—but the work required to break the cipher can also be made to grow exponentially by picking longer and longer keys.

Lessons for the Internet Age

Let's pause for a moment to consider some of the lessons of cryptographic history—morals that were well understood by the early twentieth century. In the late twentieth century, cryptography changed drastically because of modern computer technology and new cryptographic algorithms, but these lessons are still true today. They are too often forgotten.

Breakthroughs Happen, but News Travels Slowly

Mary Stuart was beheaded when her letters plotting against Elizabeth were deciphered by using frequency analysis, which Al-Kindi had described nine centuries earlier. Older methods have also remained in use to the present day, even for high-stakes communications. Suetonius explained the Caesar cipher in the first century C.E. Yet two millennia later, the Sicilian Mafia was still using the code. Bernardo Provenzano was a notorious Mafia boss who managed to stay on the run from Italian police for 43 years. But in 2002, some *pizzini*—ciphertexts typed on small pieces of paper—were found in the possession of one of his associates. The messages included correspondence between Bernardo and his son Angelo, written in a Caesar cipher—with a shift of three, exactly as Suetonius had described it.[26] Bernardo switched to a more secure code, but the dominos started to topple. He was finally traced to a farmhouse and arrested in April 2006.

Even scientists are not immune to such follies. Although Babbage and Kasiski had broken the Vigenère cipher in the mid-nineteenth century, *Scientific American* 50 years later described the Vigenère method as "impossible of translation."[27]

Encoded messages tend to look indecipherable. The incautious, whether naïve or sophisticated, are lulled into a false sense of security when they look at apparently unintelligible jumbles of numbers and letters. Cryptography is a science, and the experts know a lot about code-breaking.

Confidence Is Good, but Certainty Would Be Better

There are no guarantees that even the best contemporary ciphers won't be broken—or haven't been broken already. Some of the ciphers have the potential to be validated by mathematical proofs, but actually providing those proofs will require deep mathematical breakthroughs. If anyone knows how to

break modern codes, it is probably someone in the National Security Agency or a comparable agency of a foreign government, and those folks don't tend to say much publicly.

In the absence of a formal proof of security, all one can do is to rely on what has been dubbed the Fundamental Tenet of Cryptography:[28] *If lots of smart people have failed to solve a problem, then it probably won't be solved (soon).*

Of course, this is not a very useful principle in practice; by definition, breakthroughs are unlikely to happen "soon." But they do happen, and when they do, indigestion among cryptographers is widespread. In 1991, the MD5 algorithm was introduced for computing crucial cryptographic operations called *message digests*, which are fundamental security elements in almost all web servers, password programs, and office products. MD5 was meant to replace the earlier MD4 algorithm after questions were raised about its security. Within a couple of years, academic researchers began to produce results suggesting that MD5 might also be vulnerable to attack. Cryptographers proposed that MD5 should itself be abandoned in favor of a stronger algorithm, SHA-1, but their warnings had limited impact since the published attacks on MD5 seemed largely theoretical and unrealistic. Then in August 2004, at an annual cryptography conference, researchers announced that they had been able to break MD5 using only an hour of computing time.[29] The push was on to switch to SHA-1, but in short order weaknesses were uncovered in SHA-1 as well.

Still, there was no reason to think that anyone except academic researchers were able to crack either MD5 or SHA-1...until 2012. That is when it was revealed that Flame, a form of espionage malware that had infected computers in Iran and other Middle Eastern countries, relied on a completely new form of MD5 attack. In other words, the creators of Flame seemed to include cryptographers as knowledgeable and creative as the ones publishing in the top academic conferences.[30]

As of this writing, SHA-1 has not been "broken," but it has been weakened: Attacks that previously were considered impossibly time-consuming have now been brought into the realm of the extremely expensive. SHA-1 has given way to new standards SHA-2 and SHA-3, which seem to be stronger—but all we really know is that they haven't been broken yet.

A provably secure encryption algorithm is one of the holy grails of computer science. Every weakness exposed in proposed algorithms yields new ideas about how to make them stronger. We aren't there yet, but progress is being made.

A provably secure encryption algorithm is one of the holy grails of computer science.

Having a Good System Doesn't Mean People Will Use It

Before we explain that unbreakable encryption may finally be possible, we need to caution that even mathematical certainty would not suffice to create perfect security if people don't change their behavior.

Vigenère published his encryption method in 1586. But foreign-office cipher secretaries commonly avoided the Vigenère cipher because it was cumbersome to use. They stayed with simple substitution ciphers—even though it was well known that these ciphers were readily broken—and they hoped for the best. By the eighteenth century, most European governments had skilled "Black Chambers" through which all mail to and from foreign embassies was routed for decryption. Finally, the embassies switched to Vigenère ciphers, which themselves continued to be used after information about how to crack them had become widely known.

And so it is today. Technological inventions, no matter how solid in theory, will not be used for everyday purposes if they are inconvenient or expensive. The risks of weak systems are often rationalized in attempts to avoid the trouble of switching to more secure alternatives.

In 1999, an encryption standard known as Wired Equivalent Privacy (WEP) was introduced for home and office wireless connections. In 2001, however, WEP was found to have serious flaws that made it easy to eavesdrop on wireless networks, a fact that became widely known in the security community.[31] Despite this, wireless equipment companies continued to sell WEP products, while industry pundits comforted that "WEP is better than nothing." A new standard, Wi-Fi Protected Access (WPA), was finally introduced in 2002, but it wasn't until September 2003 that products were required to use the new standard in order to be certified. Hackers were able to steal more than 45 million credit and debit card records from TJX, the parent company of several major retail store chains, because the company was still using WEP encryption as late as 2005.[32] That was long after WEP's insecurities were known and WPA was available as a replacement. The cost of that security breach has reached the hundreds of millions of dollars.

When encryption was a military monopoly, it was possible in principle for a commander to order everyone to start using a new code if the commander suspected that the enemy had cracked the old one. The risks of insecure encryption today arise from three forces acting in concert: the high speed at which news of insecurities travels among experts, the slow speed at which the inexpert recognize their vulnerabilities, and the massive scale at which cryptographic software is deployed. When a university researcher discovers a tiny hole in an algorithm, computers everywhere become vulnerable, and there is no central authority to give the command for software upgrades everywhere.

If the Code Is Wrong, It Doesn't Matter Whether the Algorithm Is Right

Cryptographic algorithms are subject to the same garden-variety programming bugs that plague other software. In 2014 Apple disclosed that the network security software in its Macintosh computers and iPhones had an extra "goto fail" line of code. It was just an unintentional duplicate of the line above, but it had the effect of bypassing a crucial security check and making communications vulnerable to interception by an adversary.[33]

The Enemy Knows Your System

The last lesson from history may seem counterintuitive. It is that a cryptographic method, especially one designed for widespread use, should be regarded as more reliable if it is widely known and seems not to have been broken than if the method itself has been kept secret.

The Flemish linguist Auguste Kerckhoffs articulated this principle in an 1883 essay on military cryptography.[34] As he explained it,

> The system must not require secrecy, and it could fall into the hands of the enemy without causing trouble....Here I mean by system, not the key itself, but the material part of the system: tables, dictionaries, or whatever mechanical apparatus is needed to apply it. Indeed, it's not necessary to create imaginary phantoms or to suspect the integrity of employees or subordinates, in order to understand that, if a system requiring secrecy were to find itself in the hands of too many individuals, it could be compromised upon each engagement in which any of them take part.

In other words, if a cryptographic method is put in widespread use, it is unrealistic to expect that the method can remain secret for long. Thus, it should be designed so that it will remain secure, even if everything but a small amount of information (the key) becomes exposed.

Shannon restated Kerckhoffs's Principle in his paper on systems for secret communication: "We shall assume that *the enemy knows the system being used*."[35] He went on to write:

> The assumption is actually the one ordinarily used in cryptographic studies. It is pessimistic and hence safe, but in the long run realistic, since one must expect his system to be found out eventually.

Kerckhoffs's Principle is frequently violated in modern Internet security practice. Internet startup companies routinely make bold announcements

about new breakthrough proprietary encryption methods, which they refuse to subject to public scrutiny, explaining that the method must be kept secret in order to protect its security. Cryptographers generally regard such "security through obscurity" claims with extreme skepticism.

Even well-established organizations run afoul of Kerckhoffs's Principle. The Content Scrambling System (CSS) used on DVDs (digital versatile discs) was developed by a consortium of motion picture studios and consumer electronics companies in 1996. It encrypts DVD contents in order to limit unauthorized copying. The method was kept secret to prevent the manufacture of unlicensed DVD players.[36] The encryption algorithm, which consequently was never widely analyzed by experts, turned out to be weak and was cracked within three years after it was announced.[37] CSS decryption programs, together with numerous unauthorized "ripped" DVD contents, soon circulated widely on the Internet. (See Chapter 6, "Balance Toppled," for a more detailed discussion of copy protection.)

Kerckhoffs's Principle has been institutionalized in the form of encryption standards. Data Encryption Standard (DES) was adopted as a national standard in the 1970s and is widely used in the worlds of business and finance. The design of special-purpose hardware and the inexorable progress of Moore's Law have made exhaustive search more feasible in recent years, and DES is no longer considered secure. A newer standard, Advanced Encryption Standard (AES), was adopted in 2002 after a thorough and public review.[38] It is precisely because these encryption methods are so widely known that confidence in them can be high. They have been subjected to both professional analysis and amateur experimentation, and no serious deficiencies have been discovered.

These lessons are as true today as they ever were. And yet, something else, something fundamental about cryptography, is different today. In the late twentieth century, cryptographic methods stopped being state secrets and became consumer goods.

Secrecy Changes Forever

For 4,000 years, cryptography was about making sure Eve could not read Alice's message to Bob if Eve intercepted the message en route. Nothing could be done if the key itself was somehow discovered. Keeping the key secret was therefore of inestimable importance and was a very uncertain business.

If Alice and Bob worked out the key when they met, how could Bob keep the key secret during the dangers of travel? Protecting keys was a military and diplomatic priority of supreme importance. Pilots and soldiers were instructed that, even in the face of certain death from enemy attack, their first responsibility was to destroy their codebooks. Discovery of the codes could cost thousands of lives. The secrecy of the codes was everything.

And if Alice and Bob never met, then how could they agree on a key without *already* having a secure method for transmitting the key? That seemed like a fundamental limitation: Secure communication was practical only for people who could arrange to meet beforehand or who had access to a prior method of secure communication (such as military couriers) for carrying the key between them. If Internet communications had to proceed on this assumption, electronic commerce never could have gotten off the ground. Bit packets racing through the network are completely unprotected from eavesdropping.

And then, in the 1970s, everything changed. Whitfield Diffie was a 32-year-old mathematical free spirit who had been obsessed with cryptography since his years as an MIT undergraduate. 31-year-old Martin Hellman was a hard-nosed graduate of the Bronx High School of Science and an assistant professor at Stanford. Diffie had traveled the length of the country in search of collaborators on the mathematics of secret communication. This was not an easy field to enter, since most serious work in this area was being done behind the firmly locked doors of the National Security Agency. Ralph Merkle, a 24-year-old computer science graduate student, was exploring a new approach to secure communication. In the most important discovery in the entire history of cryptography, Diffie and Hellman found a practical realization of Merkle's ideas, which they presented in a paper titled "New Directions in Cryptography."[39] This is what the paper described:

> A way for Alice and Bob, without any prior arrangement, to agree on a secret key, known only to the two of them, by using messages between them that are not secret at all.

In other words, as long as Alice and Bob can communicate with each other, they can establish a secret key. It does not matter if Eve or anyone else can hear everything they say. Alice and Bob can come to a consensus on a secret key, and there is no way for Eve to use what she overhears to figure out what that secret key is. This is true even if Alice and Bob have never met before and have never made any prior agreements.

It was revealed in 1997 that the same public-key techniques had been developed within the British secret Government Communication Headquarters (GCHQ) two years before Diffie and Hellman's work, by James Ellis, Clifford Cocks, and Malcolm Williamson.[40]

The impact of this discovery cannot be overstated. The art of secret communication was a government monopoly—and had been since the dawn of writing. Governments had the largest interests in secrets, and the smartest scientists worked for governments. But there was another reason governments had done all the serious cryptography: Only governments had the wherewithal

to assure the production, protection, and distribution of the keys on which secret communication depended. If the secret keys could be produced by public communication, everyone could use cryptography. They just had to know how; they did not need armies or brave couriers to transmit and protect the keys.

Diffie, Hellman, and Merkle dubbed their discovery "public-key cryptography." Although its significance was not recognized at the time, it is the invention that made electronic commerce possible. If Alice is you and Bob is Amazon, there is no possibility of a meeting; how could you physically go to Amazon to procure a key? Does Amazon even *have* a physical location? If Alice is to send her credit card number to Amazon securely, the encryption has to be worked out on the spot or, rather, on the two separate spots separated by the Internet. Diffie–Hellman–Merkle, and a suite of related methods that followed, made secure Internet transactions possible. If you have ever ordered anything from an online store, you have been a cryptographer without realizing it. Your computer and the store's computer played the roles of Alice and Bob.

It seems wildly counterintuitive that Alice and Bob could agree on a secret key over a public communication channel. It was not so much that the scientific community had tried and failed to do what Diffie, Hellman, and Merkle did. It had never occurred to them to try because it seemed so obvious that Alice had to give Bob the keys somehow.

Even the great Shannon missed this possibility. In his 1949 paper that brought all known cryptographic methods under a unified framework, he did not realize that there might be an alternative. "The key must be transmitted by non-interceptable means from transmitting to receiving points," he wrote.[41]

Not true. Alice and Bob can get the same secret key, even though all their messages are intercepted.

Alice and Bob can get the same secret key, even though all their messages are intercepted.

The basic picture of how Alice communicates her secret to Bob remains as shown in Figure 5.6. Alice sends Bob a coded message, and Bob uses a secret key to decrypt it. Eve may intercept the ciphertext en route.

The goal is for Alice to do the encryption in such a way that it is *impossible* for Eve to decrypt the message in any way other than a brute-force search through all possible keys. If the decryption problem is "hard" in this sense, then the phenomenon of exponential growth becomes the friend of Alice and Bob. For example, suppose they are using ordinary decimal numerals as keys, and their keys are 10 digits long. If they suspect that Eve's computers are getting powerful enough to search through all possible keys, they can switch to 20-digit keys. The amount of time Eve would require goes up by a factor of 10^{10} = 10,000,000,000. Even if Eve's computers were powerful enough to crack any

10-digit key in a second, it would then take her more than 300 years to crack a 20-digit key!

Exhaustive search is always *one* way for Eve to discover the key. But if Alice encrypts her message using a substitution or Vigenère cipher, the encrypted message will have patterns that enable Eve to find the key far more quickly. The trick is to find a means of encrypting the message so that the ciphertext reveals no patterns from which the key could be inferred.

The Key Agreement Protocol

The crucial invention was the concept of a *one-way computation*—a computation with two important properties: It can be done quickly, but it can't be undone quickly. To be more precise, the computation quickly combines two numbers x and y to produce a third number, which we'll call $x \times y$. If you know the value of $x \times y$, there is no quick way to figure out what value of y was used to produce it, even if you also know the value of x. That is, if you know the values of x and the result z, the only way to find a value of y so that $z = x \times y$ is trial-and-error search. Such an exhaustive search would take time that grows exponentially with the number of digits of z—and would be practically impossible for numbers of a few hundred digits. Diffie and Hellman's one-way computation also has an important third property: $(x \times y) \times z$ always produces the same result as $(x \times z) \times y$. (Diffie and Hellman use $x \times y$ = the remainder when x^y is divided by p, where p is a fixed industry-standard prime number.)

The key agreement protocol starts from a base of public knowledge: how to do the computation $x \times y$ and also the value of a particular large number g. All this information is available to the entire world. Knowing it, here is how Alice and Bob proceed.

Alice and Bob each choose a random number. We'll call Alice's number a and Bob's number b. We'll refer to a and b as Alice and Bob's *secret keys*. Alice and Bob keep their secret keys secret. *No one except Alice knows the value of a, and no one except Bob knows the value of b.*

1. Alice calculates $g \times a$, and Bob calculates $g \times b$ (not hard to do). The results are called their *public keys* A and B, respectively.
2. Alice sends Bob the value of A, and Bob sends Alice the value of B. It doesn't matter if Eve overhears these communications; A and B are not secret numbers.
3. When she has received Bob's public key B, Alice computes $B \times a$, using her secret key a as well as Bob's public key B. Likewise, when Bob receives A from Alice, he computes $A \times b$.

Even though Alice and Bob have done different computations, they have ended up with the same value. Bob computes $A \times b$—that is, $(g \times a) \times b$

(see step 2—A is $g \times a$). Alice computes $B \times a$—that is, $(g \times b) \times a$. Because of the third property of the one-way computation, that number is $(g \times a) \times b$ once again—the same value, arrived at in a different way!

Are We Sure No One Can Crack the Code?

No one has proved mathematically that the public-key encryption algorithms are unbreakable, in spite of determined efforts by top mathematicians and computer scientists to provide absolute proof of their security. So our confidence in them rests on the fundamental tenet: *No one has broken them so far.* If anyone knows a fast method, it's probably the National Security Agency, which operates in an environment of extreme secrecy. Maybe the NSA knows how and isn't telling. Or maybe some inventive loner has cracked the code but prefers profit to celebrity and is quietly socking away huge profits from decoding messages about financial transactions. Our bet is that no one knows how and no one will.

This shared value, call it K, is the key Alice and Bob will use for encrypting and decrypting their subsequent messages, using whatever standard method of encryption they choose.

Now here's the crucial point. Suppose Eve has been listening to Alice and Bob's communications. Can she do anything with all the information she has? She has overheard A and B, and she knows g because it is an industry standard. She knows all the algorithms and protocols that Alice and Bob are using; Eve has read Diffie and Hellman's paper, too! But to compute the key K, Eve would have to know one of the secret keys, either a or b. She doesn't; only Alice knows a, and only Bob knows b. On numbers of a few hundred digits, no one knows how to find a or b from g, A, and B without searching through impossibly many trial values.

Alice and Bob can carry out their computations with personal computers or simple special-purpose hardware. But even the most powerful computers aren't even close to fast enough to let Eve break the system—at least not by any method known.

Exploiting this difference in computational effort was Diffie, Hellman, and Merkle's breakthrough. They showed how to create shared secret keys without requiring secure channels.

Public Keys for Private Messages

Suppose Alice wants to have a way for anyone in the world to send her encrypted messages that only she can decrypt. She can do this with a small variation of the key agreement protocol. All the computations are the same as in the key agreement protocol, but they take place in a slightly different order.

Alice picks a secret key a and computes the corresponding public key A. She publishes A in a directory.

If Bob (or anyone else) now wants to send Alice an encrypted message, he gets Alice's public key from the directory. Next, he picks his own secret key b and computes B as before. He also uses Alice's public key A from the directory to compute an encryption key K, just as with the key agreement protocol: $K = A \times b$. Bob uses K as a key to encrypt a message to Alice, and he sends Alice the ciphertext, along with B. Because he uses K only once, K is like a one-time pad.

With public-key encryption, anyone *can send encrypted mail to* anyone *over an insecure, publicly exposed communication path.*

When Alice receives Bob's encrypted message, she takes the B that came with message, together with her secret key a, just as in the key agreement protocol, and computes the same $K = B \times a$. Alice now uses K as the key for decrypting the message. Eve can't decrypt it because she doesn't know the secret keys.

This might seem like just a simple variant of key agreement, but it results in a major conceptual change in how we think about secure communication. With public-key encryption, *anyone* can send encrypted mail to *anyone* over an insecure, publicly exposed communication path. The only thing on which they need to agree is to use the Diffie–Hellman–Merkle method—and knowing that is of no use to an adversary trying to decipher an intercepted message.

Digital Signatures

In addition to secret communication, a second breakthrough achievement of public-key cryptography is preventing forgeries and impersonations in electronic transactions.

Suppose Alice wants to create a public announcement. How can people who see the announcement be sure that it really comes from Alice—that it's not a forgery? What's required is a method for marking Alice's public message in such a way that anyone can easily verify that the mark is Alice's, and no one can forge it. Such a mark is called a *digital signature*.

To build on the drama we have used already, we'll continue to talk about Alice sending a message to Bob, with Eve trying to do something evil while the message is in transit. In this case, however, we are not concerned with the secrecy of Alice's message—only with assuring Bob that what he receives is really what Alice sent. In other words, the message may not be secret; perhaps it is an important public announcement. Bob needs to be confident that the signature he sees on the message is Alice's and that the message could not have been tampered with before he received it.

Digital signature protocols use public keys and secret keys, but in a different way. A digital signature protocol consists of two computations: one that Alice uses to process her message to create the signature and one that Bob uses to verify the signature. Alice uses her secret key and the message itself to create the signature. Anyone can then use Alice's public key to verify the signature. The point is that everyone can know the public key and thus verify the signature, but only the person who knows the secret key could have produced the signature. This is the reverse of the scenario of the previous section, where anyone can encrypt a message, but only the person with the secret key can decrypt it.

A digital signature scheme requires a computational method that makes signing easy if you have the secret key and verifying easy if you have the public key—and yet makes it computationally infeasible to produce a verifiable signature if you don't know the secret key. Moreover, the signature depends on the message as well as on the secret key of the person signing it. Thus, the digital signature protocol attests to the *integrity* of the message—that it was not tampered with in transit—as well as to its *authenticity*—that the person who sent it really is Alice.

In typical real systems, used to sign unencrypted email, for example, Alice doesn't encrypt the message itself. Instead, to speed up the signature computation, she first computes a compressed version of the message, called a *message digest*, which is much shorter than the message itself. It requires less computation to produce the signature for the digest than for the full message. How message digests are computed is public knowledge. When Bob receives Alice's signed message, he computes the digest of the message and verifies that it is identical to what he gets by decrypting the attached signature using Alice's public key.

The digesting process needs to produce a kind of fingerprint—something small that is nonetheless virtually unique to the original. This compression process must avoid a risk associated with using digests. If Eve could produce a different message with the same digest, then she could attach Alice's signature to Eve's message. Bob would not realize that someone had tampered with the message before he received it. When he went through the verification process, he would compute the digest of Eve's message, compare it to the result of decrypting the signature that Alice attached to Alice's message, and find them identical. This risk is the source of the insecurity of the message digest function MD5 mentioned earlier in this chapter, which is making the cryptographic community wary about the use of message digests.

RSA

Diffie and Hellman introduced the concept of digital signatures in their 1976 paper. They suggested an approach to designing signatures, but they did not present a concrete method. The problem of devising a practical digital signature scheme was left as a challenge to the computer science community.

The challenge was met in 1977 by Ron Rivest, Adi Shamir, and Len Adleman of the MIT Laboratory for Computer Science.[42] Not only was the RSA (Rivest–Shamir–Adleman) algorithm a practical digital signature scheme but it could also be used for confidential messaging. With RSA, each person generates a pair of keys: a public key and a secret key. We'll again call Alice's public key A and her secret key a. The public and private keys are inverses: If you transform a value with a, then transforming the result with A recovers the original value. If you transform a value with A, then transforming the result with a recovers the original value.

Here's how RSA key pairs are used. People publish their public keys and keep their secret keys to themselves. If Bob wants to send Alice a message, he picks a standard algorithm such as DES and a key K, and he transforms K using Alice's public key A. Alice transforms the result using her secret key a to recover K. As with all public-key encryption, only Alice knows her secret key, so only Alice can recover K and decrypt the message.

To produce a digital signature, Alice transforms the message using her secret key a and uses the result as the signature to be sent along with the message. Anyone can then check the signature by transforming it with Alice's public key A to verify that this matches the original message. Because only Alice knows her secret key, only Alice could have produced something that, when transformed with her public key, will reproduce the original message.

A breakthrough in factoring would render RSA useless and would undermine many of the current standards for Internet security.

It seems to be infeasible in the RSA cryptosystem—as in the Diffie–Hellman–Merkle system—to compute a secret key corresponding to a public key. RSA uses a different one-way computation than the one used by the Diffie–Hellman–Merkle system. RSA is secure only if it takes much longer to factor an n-digit number than to multiply two $n/2$-digit numbers. RSA's reliance on the difficulty of factoring has engendered enormous interest in finding fast ways to factor numbers. Until the 1970s, this was a mathematical pastime of theoretical interest only. One can multiply numbers in time comparable to *the number of digits*, while factoring a number requires effort comparable to *the value of the number itself*, as far as anyone knows. A breakthrough in factoring would render RSA useless and would undermine many of the current standards for Internet security.

Certificates and Certificate Authorities

There's a problem with the public-key methods we've described so far. How can Bob know that the "Alice" he's communicating with really is Alice? Anyone could be at the other end of the key agreement communication pretending

to be Alice. Or, for secure messaging, after Alice places her public key in the directory, Eve might tamper with the directory, substituting her own key in place of Alice's. Then, anyone who tries to use the key to create secret messages intended for Alice will actually be creating messages that Eve, not Alice, can read. If "Bob" is you, and "Alice" is the mayor ordering an evacuation of the city, some impostor could be trying to create a panic. If "Bob" is your computer, and "Alice" is your bank's computer, "Eve" could be trying to steal your money!

This is where digital signatures can help. Alice goes to a trusted authority, to which she presents her public key together with proof of her identity. The authority digitally signs Alice's key—producing a signed key called a *certificate*. Now, instead of just presenting her key when she wants to communicate, Alice presents the certificate. Anyone who wants to use the key to communicate with Alice first checks the authority's signature to see that the key is legitimate.[43]

People check a certificate by checking the trusted authority's signature. How do they know that the signature on a certificate really is the trusted authority's signature and not some fraud that Eve set up for the purpose of issuing fake certificates? The authority's signature is itself guaranteed by another certificate, signed by another authority, and so on, until we reach an authority whose certificate is well known. In this way, Alice's public key is vouched for not only by a certificate and a single signature but by a chain of certificates, each one with a signature guaranteed by the next certificate.

Organizations that issue certificates are called *certificate authorities*. Certificate authorities can be set up for limited use; for example, a corporation might serve as a certificate authority that issues certificates for use on its corporate network. There are also companies that make a business of selling certificates for public use. The trust you should put in a certificate depends on two things: your assessment of the reliability of the signature on the certificate and *also* your assessment of the certificate authority's policy in being willing to sign things.

Commercial Certificates

VeriSign, which is currently the major commercial certificate authority, issues three classes of personal certificates. Class 1 is for assuring that a browser is associated with a particular email address and makes no claims about anyone's real identity. Class 2 provides a modest level of identity checking. Organizations issuing Class 2 certificates should require an application with information that can be checked against employee records or credit records. Class 3 certificates require applying in person for verification of identity.

Cryptography for Everyone

In real life, none of us is aware that we are carrying out one-way computations while we are browsing the Web. But every time we order a book from Amazon, check our bank or credit card balance, or pay for a purchase using PayPal, that is exactly what happens. The tell-tale sign that an encrypted web transaction is taking place is that the URL of the website begins with https (the *s* is for *secure*) instead of http. The consumer's computer and the computer of the store or the bank negotiate the encryption, using public-key cryptography—unbeknownst to the human beings involved in the transaction. The store attests to its identity by presenting a certificate signed by a certificate authority that the consumer's computer has been preconfigured to recognize. New keys are generated for each new transaction. Keys are cheap. Secret messages are everywhere on the Internet. We are all cryptographers now.

We are all cryptographers now.

At first, public-key encryption was treated as a mathematical curiosity. Len Adleman, one of the inventors of RSA, thought that the RSA paper would be "the least interesting paper I would ever be on."[44] Even the National Security Agency, as late as 1977, was not overly concerned about the spread of these methods. It simply did not appreciate how the personal computer revolution, just a few years away, would enable anyone with a home PC to exchange encrypted messages that even NSA could not decipher.

But as the 1980s progressed, and Internet use increased, the potential of ubiquitous cryptography began to become apparent. Intelligence agencies became increasingly concerned, and law enforcement feared that encrypted communications could put an end to government wiretapping, one of its most powerful tools. On the commercial side, industry was beginning to appreciate that customers would want private communication, especially in an era of electronic commerce. In the late 1980s and early 1990s, the Bush and the Clinton administrations were floating proposals to control the spread of cryptographic systems.

In 1994, the Clinton administration unveiled a plan for an "Escrowed Encryption Standard" that would be used on telephones that provided encrypted communications. The technology, dubbed "Clipper," was an encryption chip developed by the National Security Agency (NSA) that included a *back door*—an extra key held by the government, which would let law enforcement and intelligence agencies decrypt the phone communications. According to the proposal, the government would purchase only Clipper phones for secure communication. Anyone wanting to do business with the government over a secure telephone would also have to use a Clipper

phone. Industry reception was cold, however, and the plan was dropped. But in a sequence of modified proposals beginning in 1995, the White House attempted to convince industry to create encryption products that had similar back doors. The carrot here, and the stick, was export control law. Under U.S. law, cryptographic products could not be exported without a license, and violating export controls could result in severe criminal penalties. The administration proposed that encryption software would receive export licenses only if it contained back doors.

The ensuing, often heated, negotiations, sometimes referred to as the "crypto wars," played out over the remainder of the 1990s. Law enforcement and national security argued the need for encryption controls. On the other side of the debate were the technology companies, which did not want government regulation, and civil liberties groups, which warned against the potential for growing communication surveillance. In essence, policymakers could not come to grips with the transformation of a major military technology into an everyday personal tool.

We met Phil Zimmermann at the beginning of this chapter, and his career now becomes a central part of the story. Zimmermann was a journeyman programmer and civil libertarian who had been interested in cryptography since his youth. He had read a *Scientific American* column about RSA encryption in 1977 but did not have access to the kinds of computers that would be needed to implement arithmetic on huge integers, as the RSA algorithms demanded. But computers will get powerful enough if you wait. As the 1980s progressed, it became possible to implement RSA on home computers. Zimmermann set about to produce encryption software for the people, to counter the threat of increased government surveillance. As he later testified before Congress:

> The power of computers had shifted the balance towards ease of surveillance. In the past, if the government wanted to violate the privacy of ordinary citizens, it had to expend a certain amount of effort to intercept and steam open and read paper mail, or listen to and possibly transcribe spoken telephone conversations. This is analogous to catching fish with a hook and a line, one fish at a time. Fortunately for freedom and democracy, this kind of labor-intensive monitoring is not practical on a large scale. Today, electronic mail is gradually replacing conventional paper mail, and is soon to be the norm for everyone, not the novelty it is today. Unlike paper mail, e-mail messages are just too easy to intercept and scan for interesting keywords. This can be done easily, routinely, automatically, and undetectable on a grand scale. This is analogous to driftnet fishing—making a quantitative and qualitative Orwellian difference to the health of democracy.[45]

Cryptography was the answer. If governments were to have unlimited surveillance powers over electronic communications, people everywhere needed easy-to-use, cheap, uncrackable cryptography so they could communicate without governments being able to understand them.

Zimmermann faced obstacles that would have stopped less zealous souls. RSA was a patented invention. MIT had licensed it exclusively to the RSA Data Security Company, which produced commercial encryption software for corporations, and RSA Data Security had no interest in granting Zimmermann the license he would need to distribute his RSA code freely, as he wished to do.

And there was government policy, which was, of course, exactly the problem to which Zimmermann felt his encryption software was the solution. On January 24, 1991, Senator Joseph Biden, a co-sponsor of antiterrorist legislation Senate Bill 266[46], inserted some new language into the bill:

> It is the sense of Congress that providers of electronic communications services and manufacturers of electronic communications service equipment shall ensure that communications systems permit the government to obtain the plaintext contents of voice, data, and other communications when appropriately authorized by law.

This language received a furious reaction from civil liberties groups and wound up not surviving, but Zimmermann decided to take matters into his own hands.

By June 1991, Zimmermann had completed a working version of his software. He named it PGP for "Pretty Good Privacy," after Ralph's mythical Pretty Good Groceries that sponsored Garrison Keillor's *Prairie Home Companion*. The software mysteriously appeared on several U.S. computers, available for anyone in the world to download. Soon copies were everywhere—not just in the United States but all over the world. In Zimmermann's own words: "This technology belongs to everybody." The genie was out of the bottle and was not going back in.

Zimmermann paid a price for his libertarian gesture. First, RSA Data Security was confident that this technology belonged to *it*, not to "everybody." The company was enraged that its patented technology was being given away. Second, the government was furious. It instituted a criminal investigation for violation of the export control laws, although it was not clear what laws, if any, Zimmermann had violated. Eventually MIT brokered an agreement that let Zimmermann use the RSA patent and devised a way to put PGP on the Internet for use in the United States and in conformance with export controls.

By the end of the decade, the progress of electronic commerce had overtaken the key escrow debate, and the government had ended its criminal investigation without an indictment. Zimmermann built a business around PGP, while still allowing free downloads for individuals. His website contains testimonials from human rights groups in Eastern Europe and Guatemala attesting to the liberating force of secret communication among individuals and agencies working against oppressive regimes. Zimmermann had won.

Sort of.

Cryptography Unsettled

Today, every banking and credit card transaction over the Web is encrypted. Much email and the hard drives of many laptop computers are encrypted. There is widespread concern about information security, identity theft, and degradation of personal privacy.

At the same time, cryptography is threatened by two opposing forces: indifference and fear. In contexts where individuals have to remember encryption keys—for example, to unscramble the data stored on their laptops—the inconvenience of remembering keys is enough that some users prefer not to use them at all (or to set their encryption key to KEY in the same way that they set their password to PASSWORD). And users are not the only irresponsible parties. In 2017 Apple shipped computers with a crucial password set to null, which meant that anyone could compromise the machines, even operating remotely over the Internet.

Strings of characters long enough to be hard for an adversary to guess are also hard for their rightful owners to remember. In January 2018, when a false-alarm notice went out to the people of Hawaii about an incoming ballistic missile, it took almost 40 minutes to issue a correction in part because an official couldn't remember his Twitter password. Biometrics (fingerprints, iris scans, and the like) are promoted as more convenient alternatives but raise privacy worries if such personal information is to be stored remotely. Faced with the risk of being unable to retrieve priceless data because of simple forgetfulness, some unwisely don't encrypt it at all.

At the same time, citizens hope to trust their government, feel they have nothing to hide, and know they should fear terrorists and criminals. Zimmerman's warnings about government surveillance have faded. With every report (whether justified or not) that they are in imminent danger, some are more likely to accept government surveillance and to mistrust anyone wishing to communicate in secrecy from law enforcement. "I don't care," some will

say. "Just keep me safe." Against this background, appeals like Rosenstein's for cooperation from the technology companies—or like Judd Gregg's before him—meet less resistance.

Spying On Citizens

Historically, spying on citizens required a warrant (since citizens have an expectation of privacy), but spying on foreigners did not. A series of executive orders and laws intended to combat terrorism allow the government to inspect bits that are on their way into or out of the country—perhaps even a phone call to an airline, if it is answered by a call center in India.[47] Also excluded from judicial oversight is any "surveillance directed at a person reasonably believed to be located outside of the United States," whether that person is a U.S. citizen or not. Such developments may stimulate encryption of electronic communications, and hence in the end may prove to be counterproductive.[48] That in turn might renew efforts to criminalize encryption of email and telephone communications in the United States.

The bottom-line question is this: With encryption as ordinary a tool for personal messages as it for commercial transactions, will the benefits to personal privacy, free expression, and human liberty outweigh the costs to law enforcement and national intelligence, whose capacity to eavesdrop and wiretap will be at an end?

Whatever the future of encrypted communication, encryption technology has another use. Perfect copies and instant communication have blown the legal notion of "intellectual property" into billions of bits of teenage movie and music downloads. Encryption is the tool used to lock movies so only certain people can see them and to lock songs so only certain people can hear them—to put a hard shell around this part of the digital explosion. The changed meaning of copyright is the next stop on our tour of the exploded landscape.

Endnotes

1 Barton Gellman and Ashkan Soltani, "NSA Infiltrates Links to Yahoo, Google Data Centers Worldwide, Snowden Documents Say," *Washington Post*, October 30, 2013, https://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html.

2 Barton Gellman, "NSA Broke Privacy Rules Thousands of Times per Year, Audit Finds," *Washington Post*, August 15, 2013, https://www.washingtonpost.com/world/national-security/nsa-broke-privacy-rules-thousands-of-times-per-year-audit-finds/2013/08/15/3310e554-05ca-11e3-a07f-49ddc7417125_story.html.

3 Kimberly Dozier, "Senators: Limit NSA Snooping into US Phone Records," Associated Press, September 27, 20013, https://web.archive.org/web/20131029003314/http://bigstory.ap.org/article/senators-limit-nsa-snooping-us-phone-records.

4 David Goldman, "Apps Claim They Can Keep Phone Records Secure," CNNMoney, June 6, 2013, https://money.cnn.com/2013/06/06/technology/security/verizon-call-logs/index.html.

5 Matthew Green, "The Daunting Challenge of Secure E-Mail," *The New Yorker*, accessed November 8, 2013, https://www.newyorker.com/tech/annals-of-technology/the-daunting-challenge-of-secure-email.

6 "Email Encryption in Transit," Google Transparency Report, accessed April 28, 2020, https://transparencyreport.google.com/safer-email/overview?hl=en.

7 Ellen Nakashima, "FBI Chief Calls Encryption a 'Major Public Safety Issue,'" *The Washington Post*, January 9, 2018, https://www.washingtonpost.com/world/national-security/fbi-chief-calls-encryption-a-major-public-safety-issue/2018/01/09/29a04166-f555-11e7-b34a-b85626af34ef_story.html.

8 Devlin Barrett and Ellen Nakashima, "FBI Repeatedly Overstated Encryption Threat Figures to Congress, Public," *The Washington Post*, May 22, 2018, https://www.washingtonpost.com/world/national-security/fbi-repeatedly-overstated-encryption-threat-figures-to-congress-public/2018/05/22/5b68ae90-5dce-11e8-a4a4-c070ef53f315_story.html.

9 "Deputy Attorney General Rod J. Rosenstein Delivers Remarks on Encryption at the United States Naval Academy," October 10, 2017, https://www.justice.gov/opa/speech/deputy-attorney-general-rodj-rosenstein-delivers-remarks-encryption-united-states-naval.

10 "Departments of Commerce, Justice, and State, the Judiciary, and Related Agencies Appropriations Act, 2002," *Congressional Record* 147, no. 119 (September 13, 2001), https://www.congress.gov/congressional-record/2001/9/13/senate-section/article/S9354-2.

11 John Schwartz, "Disputes on Electronic Message Encryption Take On New Urgency," *The New York Times*, September 25, 2001, https://www.nytimes.com/2001/09/25/business/disputes-on-electronic-message-encryption-take-on-new-urgency.html.

12 Schwartz.

13 "How to Address the Threat That Confronts Us Today," Federation of American Scientists, September 19, 2001, https://fas.org/irp/congress/2001_cr/s091901.html.

14 Schwartz, "Disputes on Electronic Message Encryption Take On New Urgency."

15 "Senator Backs Off Backdoors," *Wired*, October 17, 2001, https://www.wired.com/2001/10/senator-backs-off-backdoors/.

16 "Summary of H.R. 695 (105th): Security and Freedom Through Encryption (SAFE) Act," GovTrack.us, September 29, 1997, https://www.govtrack.us/congress/bills/105/hr695/summary.

17 Louis Freeh, "Encryption," Federation of American Scientists, July 8, 1997, https://fas.org/irp/congress/1997_hr/s970709f.htm.

18 Ron Rivest, "MIT Press Forum on Encryption," April 7, 1998.

19 To be precise, Caesar had only 22 shifts available because the Romans did not use the letter J, U, or W.

20 "Peterhouse MS 75.1, Folio 30" (n.d.); Geoffrey Chaucer et al., *The Equatorie of the Planetis* (Cambridge University Press, 1955); Kari Anne Rand, *The Authorship of the Equatorie of the Planetis*, Chaucer Studies 19 (DSBrewer, 1993).

21 *The Equatorie of the Planetis*, (MS Peterhouse 75.I), https://cudl.lib.cam.ac.uk/view/MS-PETERHOUSE-00075-00001.

22 David Kahn, *The Codebreakers: The Story of Secret Writing* [Rev. ed.] (Scribner, 1996).

23 Simon Singh, *The Code Book: The Evolution of Secrecy from Mary, Queen of Scots, to Quantum Cryptography* (Doubleday, 1999).

24 C. E. Shannon, "Communication Theory of Secrecy Systems," *Bell System Technical Journal* 28, no. 4 (1949): 656–715, https://ieeexplore.ieee.org/document/6769090.

25 Robert Louis Benson, *The Venona Story* (National Security Agency, Center for Cryptologic History, 2001).

26 "Language Log: The Provenzano Code," April 21, 2006, http://itre.cis.upenn.edu/~myl/languagelog/archives/003049.html.

27 "A New Cipher Code," *Scientific American* 83, no. 2143 supp (1917): 61, https://doi.org/10.1038/scientificamerican01271917-61csupp.

28 Charlie Kaufman, *Network Security: Private Communication in a Public World*, 2nd ed. (Prentice Hall PTR, 2002).

29 Xiaoyun Wang and Hongbo Yu, "How to Break MD5 and Other Hash Functions," in *Advances in Cryptology – EUROCRYPT 2005*, ed. Ronald Cramer, vol. 3494, Lecture Notes in Computer Science (Springer Berlin Heidelberg, 2005), 19–35, https://doi.org/10.1007/11426639_2.

30 Dan Goodin, "Crypto Breakthrough Shows Flame Was Designed by World-Class Scientists," *Ars Technica*, June 7, 2012, https://arstechnica.com/information-technology/2012/06/flame-cryptobreakthrough/.

31 Nikita Borisov, Ian Goldberg, and David Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11," MobiCom '01 (ACM, 2001), 180–189, https://doi.org/10.1145/381677.381695.

32 Jaikumar Vijayan, "Canadian Probe Finds TJX Breach Followed Wireless Hack," *Computerworld*, September 25, 2007, https://www.computerworld.com/article/2541162/canadian-probe-finds-tjx-breach-followed-wireless-hack.html.

33 "ImperialViolet: Apple's SSL/TLS Bug," February 22, 2014, https://www.imperialviolet.org/2014/02/22/applebug.html.

34 Aug (Auguste) Kerckhoffs, *La cryptographie militaire ou Des chiffres usités en temps de guerre avec un nouveau procédé de déchiffrement applicable aux systèmes à double clef*, Extrait du Journal des sciences militaires (Librairie militaire de LBaudoin et Cie, 1883), https://journals.openedition.org/bibnum/555.

35 Shannon, "Communication Theory of Secrecy Systems," 662.

36 J. A. Bloom et al., "Copy Protection for DVD Video," *Proceedings of the IEEE* 87, no. 7 (1999): 1267–1276, https://doi.org/10.1109/5.771077.

37 Andy Patrizio, "Why the DVD Hack Was a Cinch," *Wired*, November 2, 1999, https://www.wired.com/1999/11/why-the-dvd-hack-was-a-cinch/.

38 Morris J. Dworkin et al., "Advanced Encryption Standard (AES)," November 26, 2001, https://www.nist.gov/publications/advanced-encryption-standard-aes.

39 W. Diffie and M. Hellman, "New Directions in Cryptography," *IEEE Transactions on Information Theory* 22, no. 6 (1976): 644–654, https://doi.org/10.1109/TIT.1976.1055638.

40 J. H. Ellis, "The History of Non-Secret Encryption," *Cryptologia* 23, no. 3 (July 1, 1999): 267–273, https://www.tandfonline.com/doi/abs/10.1080/0161-119991887919.

41 Shannon, Claude. "Communication Theory of Secrecy Systems," *Bell System Technical Journal*, vol. 28(**4**), page 670, 1949.

42 R. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *Communications of the ACM* 21, no. 2 (1978): 120–126, https://doi.org/10.1145/359340.359342.

43 Originally proposed in an MIT SB thesis: Loren M Kohnfelder, "Towards a Practical Public-Key Cryptosystem" (Massachusetts Institute of Technology, 1978), https://groups.csail.mit.edu/cis/theses/kohnfelder-bs.pdf.

44 Singh, *The Code Book*, 273.

45 "Testimony of Philip Zimmermann to Subcommittee for Economic Policy, Trade, and the Environment US House of Representatives 12 Oct 1993," Federation of American Scientists, October 12, 1993, https://fas.org/irp/congress/1993_hr/931012_zimmerman.htm.

46 Joseph R. Biden, "S.266 - Comprehensive Counter-Terrorism Act of 1991," https://www.congress.gov/bill/102nd-congress/senate-bill/266.

47 Mitch McConnell, "S.1927 - Protect America Act of 2007," August 5, 2007, https://www.congress.gov/bill/110th-congress/senate-bill/1927.

48 Patricia J. Williams, "The Protect Alberto Gonzales Act of 2007," *Nation*, https://www.thenation.com/article/archive/protect-alberto-gonzales-act-2007/.

This page intentionally left blank

CHAPTER 6

Balance Toppled

Who Owns the Bits?

Stealing Music

As an inmate serving a life sentence in a Florida prison, William Demler doesn't have many creature comforts, and when the prison system took away his digital music collection, he sued.[1] Demler had purchased the only digital music player permitted in the prison and, over more than five years, filled it with more than $550 worth of music, at a cost of $1.70 a song. (Like much else in prisons, music costs more, and it can only be stored on proprietary players and in a prison-proprietary cloud.)

Prison authorities say they were just switching music-system vendors—and the new system, incompatible with the old, wouldn't let users transfer previously purchased music. Instead, those who owned older music players would be given a new tablet and a $50 credit for new music. Demler says the switch violated the promise made by the music service: "Once music is purchased, you'll always own it." The alternatives he was offered—sending the old player or a music CD to a relative outside prison—wouldn't do him any good inside, where commerce and communications are strictly limited.

Demler's lawsuit sought class action status on behalf of his fellow Florida music purchasers: All of them had purchased music they expected to be able to listen to throughout their prison terms and now found that they'd have to buy it again if they wanted to continue listening.

At the time of the suit, Florida officials were unable to explain why inmates couldn't simply transfer music from one system to another. But the answer likely has to do with a funny mix of copyright and contract controlling these particular bits.

Music, like other creative works, is protected by copyright once fixed in a "tangible medium of expression." (Digital recording counts as "fixation.") Copyright restricts the making of copies or derivative works to those with authorization from the copyright holders—and in the case of music, the copyright holders can include composers, songwriters, and performers.

With analog media—books, vinyl records, or oil paintings—when you purchase a physical copy, the copyright holder's rights over the copy end with that "first sale." You can resell the book, listen to the record as many times as you like, and charge admission to the museum in which you hang the oil painting. A library can purchase books and lend them repeatedly with no further interaction with the copyright holder. An e-book, by contrast, can't be read without making copies: from the "shelf" to a reader device and even within the random-access memory (RAM) of the device itself. If all these copies count for the purposes of copyright authorization, then they tie many more strings to an e-book than to its paper equivalent. While courts have ruled against "shrink-wrap" licenses on books, they have validated the click-through licenses routinely used on electronic media and software.

Back to Mr. Demler in Florida: While he had physical possession of the media player, most of his songs were in the provider's cloud storage, subject to the service provider's contractual terms and whims. Some terms stem from agreements with music copyright holders (usually aggregated through record labels and music publishing societies) for the "copies" the provider makes to send to each listener. When the State of Florida ended its contract with the provider, it made no provision to get that music out, leaving its inmates in silence.

You don't have to be in prison to experience the ephemerality of digital music. End a subscription to Spotify or Apple Music, and you likewise lose access to the playlists or music collections you've built there. While you're in the silo, you get access to vast amounts of recorded music, but you can't take it with you.

As we will describe, the copyright wars have escalated and morphed. What started with lawsuits against music sharers has grown to lawsuits by those unable to retain their music. In a temporary equilibrium, copyright holders have turned to subscription services for music, TV, and movies (Netflix, Hulu, Amazon Prime, Apple TV) and tacitly acknowledge "format shifting" by those determined to preserve "their" music by buying CDs and ripping them to storage. Yet this also represents a re-centralization of the entertainment business: As an artist, if you're not part of one of the big services—accepting their terms—you'll have a hard time being seen and heard. Independent media survive, especially on podcasts and SoundCloud

("if you liked that, here's my soundcloud" has become a meme) but it's too early to say if they will thrive.

Automated Crimes, Automated Justice

Tanya Andersen was home having dinner with her eight-year-old daughter in December 2005 when they were interrupted by a knock at the door.[2] It was a legal process server, armed with a lawsuit from the Recording Industry Association of America (RIAA), a trade organization representing half a dozen music publishers that together control over 90% of music distribution in the United States. The RIAA claimed that the Oregon single mother surviving on disability payments owed them close to a million dollars for illegally downloading 1,200 tracks of gangsta rap and other copyrighted music.

Andersen's run-in with the RIAA had begun nine months earlier, with a "demand letter" from a Los Angeles law firm. The letter stated that "a number of record companies" had sued her for copyright infringement and that she could settle for $4,000–$5,000 or face the consequences. She suspected the letter was a scam and protested to the RIAA that she had never downloaded any music. Andersen repeatedly offered to let the record companies verify this for themselves by inspecting her computer's hard drive, but the RIAA refused the offers. At one point, an RIAA representative admitted to her that he believed she was probably innocent. But, he warned, once the RIAA starts a lawsuit, they don't drop it, because doing so would encourage other people to defend themselves against the recording industry's claims.

Andersen found a lawyer after the December lawsuit was served, and they convinced a judge to order an inspection of the hard drive. The RIAA's own expert determined that Andersen's computer had never been used for illegal downloading. But instead of dropping the suit, the RIAA increased the pressure on Andersen to settle. They demanded that their lawyers be allowed to take a deposition from Andersen's daughter and even tried to reach the child directly by calling the apartment. An unknown woman phoned her elementary school principal, falsely claiming to be her grandmother and asking about the girl's attendance. RIAA lawyers contacted Andersen's friends and relatives and told them that Andersen was a thief who collected violent, racist music. The pressure on the 41-year-old Andersen, who suffered from a painful illness and emotional problems, forced her to abandon her hope of entering a back-to-work program. Instead, she sought additional psychiatric care. Finally, after two years, Andersen was able to file a motion for summary judgment, which required the RIAA to come to court with proof of its claims. When the RIAA could not produce proof, the case was dismissed.[3]

35,000 Lawsuits in Five Years

Between 2003 and 2008, the RIAA filed more than 35,000 lawsuits against individuals for illegal downloading.[4] It sued children; it sued seniors; it sued disabled people;[5] it sued dead people.[6] It sued people who didn't have a computer or an Internet connection at home;[7] it even sued homeless people.[8] The process began when MediaSentry, RIAA's investigative company, logged in to a file-sharing network in search of computers hosting music for download. MediaSentry connected to these computers and scanned them for music files. When it found something suspicious, it sent the computer's IP address to the RIAA's anti-piracy group, together with a list of the files it found. RIAA staff members downloaded and listened to a few of the file to verify that they were in fact copyrighted songs. Then the RIAA filed a lawsuit against "John Doe," the person who used the computer at the offending IP address. With the lawsuit as a legal basis, the RIAA subpoenaed the computer's Internet service provider, forcing disclosure of the real name of the John Doe user at that IP address. The RIAA sent the user its demand letter, naming the songs that were verified and citing the total number of songs found as the basis for damages. The letter offered an opportunity to settle; the average settlement demand was about $4,000, non-negotiable.

It was an automated sort of justice for the digital age. But these are automated sorts of crimes. File-sharing programs were commonly configured to start up and run automatically, exchanging files without human intervention. The computer's owner might not even be aware that it had been configured to upload files in the background.

It's also an error-prone form of justice. Matching names to IP addresses is unreliable; several computers on the same wireless network might share the same IP address. An Internet service provider allocating IP addresses might shift them around, so that a computer with a particular IP address today might not be the same computer that was file sharing from that IP address last week. Even if it is the same computer, there's no way to prove who was using it at the time. And maybe there was a clerical error in reporting.

The RIAA knows that the process was flawed, but given its stake in stopping downloading, the organization saw no choice. Not only was the RIAA seeing its products being distributed for free, but it might be liable to lawsuits from artists for neglecting to protect the artists' copyrights. Explains Amy Weiss, RIAA senior vice president for communications, "When you fish with a net, you sometimes are going to catch a few dolphin....But we also realize that this cybershoplifting needs to stop." Besides Andersen, other snared "dolphin" included a Georgia family that didn't own a computer, a paralyzed stroke victim in Florida sued for files downloaded in Michigan, and an 83-year-old West Virginia woman who hated computers and who, as it turned out, was deceased.

The High Stakes for Infringement

$750 a Song

The minimum damages that the court must award for infringement is $750 per infringing act. In cases where the infringement can be shown to be "willful," damages could be as high as $150,000 per infringement, or $600 million for the 4,000 songs on an iPod. For defendants who can prove that they weren't even aware of the infringement, the court still must award at least $200 per infringement—a "mere" $800,000 for 4,000 songs.[9]

Error or not, most people choose to pay when they get a demand letter. The cost of settling is less than the legal fees for contesting, and the cost of losing the lawsuit is staggering: damages of at least $750 for each song downloaded. The 4,000-song contents of a 20 GB iPod would be grounds for minimum damages of $3 million—a thousand times the cost of purchasing those songs on iTunes. (A GB, or *gigabyte*, is about a billion bytes.)

Driftnet justice, automated policing of automated crimes, and $3 million minimum damages for an iPod's worth of music are consequences of policies honed for a pre-networked world colliding with the exponentials of the digital explosion. Take the $3 million iPod. This traces to the Copyright Act of 1976, which introduced a provision letting copyright holders sue for minimum *statutory damages* of $750 per infringement.

The rationale for statutory damages is to ensure that the penalty is sufficient to deter infringement even when actual damages to the copyright holder are small. The scale of the damages has dreadful consequences in the age of digital reproduction because each song copied (uploaded or downloaded) counts as a *separate* infringement. That way of reckoning "acts of infringement" may have seemed reasonable when the standards were set in pre-Internet 1976—when people could make only a few unauthorized copies, one by one. But the damage calculations balloon into unreality when 1,000 songs can be downloaded to a home computer in a few minutes over a high-speed network connection.

Although the digital explosion may have blown the legal penalties for infringement out of realistic proportion to the offense, it has also brought a more fundamental change: that the public is now concerned with copyright at all. Before the Internet, what could an ordinary person do to infringe copyright—make 50 photocopies of a book and sell them on the street corner? That would surely be infringement. But it would also be a lot of work, and the financial loss to the copyright holder would be insignificant.

Of all the dislocations of the digital explosion, the loss of the copyright balance is the most rancorous. Ordinary people can now effortlessly copy and distribute information on a massive scale. Listeners clash with a content

Sending a Message

In October 2007, Jammie Thomas, a Minnesota single mother of two who earns $36,000 a year, was found guilty of sharing 24 songs on the Kazaa file-sharing network...and fined $222,000: $9,250 per song. This was the first of the RIAA's 16,000 lawsuits that went all the way to jury trial. In the others, people settled or, as with Tanya Andersen, the case was dismissed or dropped. Given the legal statutory damages for infringement, Thomas's fine for 24 songs could have been anywhere between $18,000 and $3.6 million.

A juror interviewed afterward reported that there were people advocating for fines at both ends of that spectrum during deliberation: "We wanted to send a message that you don't do this, that you have been warned."

Said the RIAA's lawyer after the verdict was read, "This is what can happen if you don't settle."[10]

industry whose economics relies on ordinary people not doing precisely that. As a result, millions of people are today vilified as "pirates" and "thieves," while content providers are demonized as subverters of innovation and consumer freedom trying to protect their outdated business models.

Of all the dislocations of the digital explosion, the loss of the copyright balance is the most rancorous.

The war over copyright and the Internet has been escalating for more than 25 years. It is a spiral of more and more technology that makes it ever easier for more and more people to share more and more information. This explosion is countered by a legislative response that brings more and more acts within the scope of copyright enforcement, subject to punishments that grow ever more severe. Regulation tries to keep pace by banning technology, sometimes even before the technology exists. Single mothers facing mind-numbing lawsuits are merely collateral damage in that war today. If we cannot slow the arms race, tomorrow's casualties may come to include the open Internet and the dynamic innovation that fuels the information revolution.

Sharing Becomes a Crime

Copyright infringement was not even a criminal matter in the United States until the turn of the twentieth century, although an infringer could be sued for civil damages. Infringement with a profit motive first became a crime in 1897. The maximum punishment was then a year in prison and a $1,000 fine. Things stayed that way until 1976, when Congress started enacting a series of laws that repeatedly increased the penalties, motivated largely by prompting from

the RIAA and the MPAA (Motion Picture Association of America). By 1992, an infringement conviction could result in a ten-year prison sentence and stiff fines, but only if the infringement was done "for the purpose of commercial advantage or private financial gain." Without a commercial motive, there was no crime.[11]

That changed in 1994.

During the 1980s, MIT became one of the first universities to deploy large numbers of computer workstations connected to the Internet and open to anyone on campus. Even several years later, public clusters of networked powerful computers were not very common. In December 1993, some students in one of the clusters noticed a machine that was strangely unresponsive and was strenuously exercising its disk drive. When the computer staff examined this "bug," they discovered that the machine was acting as a file-server bulletin board—a relay point where people around the Internet were uploading and downloading files. Most of the files were computer games, and there was also some word-processing software.

MIT, like most other universities, prefers to handle matters like this internally, but in this case there was a complication: The FBI had asked about this very same machine only a few days earlier. Federal agents had been investigating some crackers in Denmark who were trying to use MIT machines to break into National Weather Service computers. While measuring network traffic into and out of MIT, the bureau had noticed a lot of activity coming from this particular machine. The bulletin board had nothing to do with the Denmark operation, but MIT felt that it had to tell the FBI what was happening. An agent staked out the machine and ended up accusing an MIT undergraduate of operating the bulletin board.

The Justice Department seized on the case. The software industry was growing rapidly in 1994, and the Internet was just starting to enter the public eye—and here was the power of the Internet being turned to "piracy." The Boston U.S. Attorney issued a statement claiming that the MIT bulletin board was responsible for more than a million dollars in monetary losses, adding, "We need to respond to the culture that no one is hurt by these thefts and that there is nothing wrong with pirating software."[12]

What had occurred at MIT involved copyright infringement to be sure, but there was no commercial motive and hence no crime—no basis on which the Justice Department could act. There might have been grounds for a civil suit, but the companies whose software was involved were not interested in suing. Instead, the Boston U.S. Attorney's office, after checking with their superiors in Washington, brought a charge of wire fraud against the student, on the grounds that his acts constituted interstate transmission of stolen property.

At the trial, Federal District Judge Stearns dismissed the case, citing a Supreme Court ruling that bootleg copies do not qualify as stolen property. Stearns chastised the student, describing his behavior as "heedlessly irresponsible." The judge

suggested that Congress could modify the copyright law to permit criminal prosecutions in cases like this if it so wished. But he emphasized that changing the rules should be up to Congress, not the courts. To accept the prosecution's claim, he warned, would "serve to criminalize the conduct...of the myriad of home computer users who succumb to the temptation to copy even a single software program for private use." He cited congressional testimony from the software industry that even the industry would not consider such an outcome desirable.[13]

Two years later, Congress responded by passing the 1997 No Electronic Theft (NET) Act. Described by its supporters as "closing the loophole" demonstrated by the MIT bulletin board, NET criminalized any unauthorized copying with retail value over $1000, commercially motivated or not. This addressed Judge Stearns's suggestion, but it did not heed his caution: From now on, anyone making unauthorized copies at home, even a single copy of an expensive computer program, was risking a year in prison. After only two more years, Congress was back with the Digital Theft Deterrence and Copyright Damages Improvement Act of 1999. Its supporters argued that the NET Act had been ineffective in stopping "piracy" and that penalties needed to be increased.[14] The copyright arms race was in full swing.

The Peer-to-Peer Upheaval

The NET Act marked the first time that the Internet had triggered a significant expansion of liability for copyright infringement. It would hardly be the last.

In the summer of 1999, Sean Fanning, a student at Northeastern University, began distributing a new file-sharing program and joined his uncle in forming a company around it: Napster. Napster made it easy to share files, especially music tracks, over the Internet, and to share them on a scale never before seen.

Here is how the system worked: Suppose Napster user Mary wants to share her computer file copy of Sarah McLachlan's 1999 hit *Angel*. She tells the Napster service, which adds "Angel; Sarah McLachlan" to its directory, together with an ID for Mary's computer. Any other Napster user who would like to get a copy of *Angel*, say Beth, can query the Napster directory to learn that Mary has a copy. Beth's computer then connects directly to Mary's computer and downloads the song without any further involvement from the Napster service. The connecting and downloading are done transparently by Napster-supplied software running on Mary's and Beth's computers.

The key point is that previous file-sharing setups like the MIT bulletin board were so-called *centralized systems*. They collected files at a central computer for people to download. Napster, in contrast, maintained only a central directory showing where files on other computers could be found. The individual computers passed the files among themselves directly. This kind of system organization is called a *peer-to-peer* architecture.

Peer-to-peer architectures make vastly more efficient use of the network than centralized systems, as Figure 6.1 indicates.

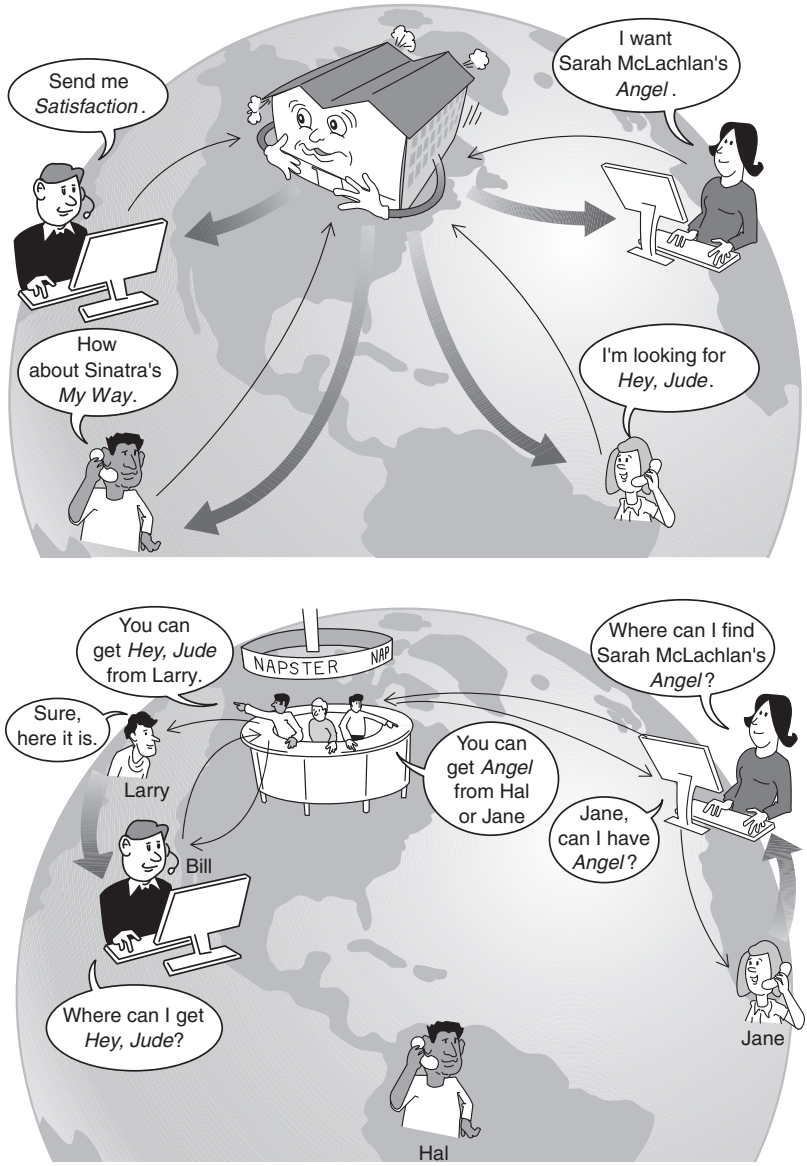

FIGURE 6.1 Underlying organization of traditional and peer-to-peer client/server network architectures. In the top figure, a traditional centralized file distribution architecture, in which files are downloaded to clients from a central server. On the bottom is a Napster-style peer-to-peer architecture in which the central server holds only directory information, and the actual files are transmitted directly between clients without passing through the server.

In a centralized system, if many users want to download files, they must all get the files from the central server, whose connection to the Internet would consequently become a bottleneck as the number of users grows. In a peer-to-peer system, the central server itself need communicate only a tiny amount of directory information, and the large network load for transmitting the files is distributed over the Internet connections of all the users. Even the slow connections common with personal computers in 1999 were enough for Napster's peer-to-peer system to let millions of users share music files...which they did. By early 2001, two years after Napster appeared, there were more than 26 million registered Napster users. At some colleges, more than 80% of the on-campus network traffic could be traced to Napster. Students held Napster parties. You hooked up a computer to some speakers and to the Internet, invited your friends over, and for any song title requested, there it was. Someone among those millions of Napster users had the song available for downloading. This was the endless cornucopia of music, the universal jukebox.

The Specter of Secondary Liability

Universal though it may have been, the Napster jukebox was collecting no quarters for the music industry. Previous escapades in file sharing, usually done on a small scale among friends, were barely annoyances from an economic perspective. Even the MIT bulletin board that engendered the No Electronic Theft Act had perhaps a few hundred users altogether. Napster was on a completely different scale, allowing anyone to readily share music files with a few hundred thousand "friends." The recording industry recognized the threat and, in December 1999, just a few months after Napster appeared, the RIAA sued it for more than $100 million in damages.

SECONDARY INFRINGEMENT

Copyright law distinguishes between two kinds of secondary infringement. The first is *contributory infringement*—that is, knowingly providing tools that enable others to infringe. The second is *vicarious infringement*—that is, profiting from the infringement of others that one is in a position to control and not preventing it. Napster was found guilty of both contributory and vicarious infringement.

Napster protested that it had no liability. After all, Napster itself wasn't copying any files. It was merely providing a directory service. How could you hold a company liable for simply publishing the locations of items on the Internet? Wasn't that publication just an example of exercising freedom of speech? Unfortunately for Napster, the California Federal District Court didn't agree, and in July 2000, the court found Napster guilty of secondary copyright infringement

(enabling others to infringe and profiting from the infringement). A year later, after an unsuccessful appeal to the Ninth Circuit, the court ordered Napster's file-sharing service to shut down.

Napster was dead, but it had captured the imagination of the technical community as a striking demonstration of the power of the Internet's fundamental architecture. No central machine controls the network; every machine in the network has equal rights to send any other machine a message. Machines connected to the Internet are, as the lingo has it, *peers*. The notion of the Internet as a network of peer machines communicating with each other directly—as opposed to a network of client machines mediated by central servers—was hardly new. Even the very first Internet technical specification, published in 1969, described the network architecture in terms of machines interacting as a network of peers. Systems incorporating peer-to-peer communication between larger computers had been in wide use since the early 1980s.[15]

Napster showed that the same principle remained valid when the peers were millions of personal computers controlled by ordinary people. Napster's use of peer-to-peer sharing was illegal, but it demonstrated the potential of the idea. Research and development in distributed computing took off. In 2000 and 2001, more than $500 million was invested in companies building peer-to-peer applications. And transcending its roots as a technical network architecture, "P2P" became enshrined in techno-pop-culture-speak as a catchword for organizations of all types—including social, corporate, and political—that harness the power of myriad cooperating individuals without reliance on central authorities. As one 2001 review gushed, "P2P is a mindset, not a particular technology or industry."[16]

Napster had also given an entire generation a taste of the Internet as universal jukebox for which people would clamor. But the recording companies, which worked together to combat illegal downloading, failed to collaborate to create a legal and profitable Internet music service to fill the vacuum left by Napster. Instead of capitalizing on file-sharing technology, they demonized it as a threat to their business. That technological rejectionism ratcheted up the rancor in the arms race, but it also did something even more shortsighted: The music companies surrendered a vast business opportunity to the profit of more imaginative entrepreneurs. Two years later, Apple would launch its iTunes music store, the first commercially successful music downloading service.

Sharing Goes Decentralized

In the meantime, new file-sharing schemes sprouted up that explored new technical architectures in attempts to tiptoe around liability for secondary infringement. Napster's legal Achilles heel had been its central directory.

As the court had ruled, control of the directory amounted to control of the file-sharing activity, and Napster was consequently liable for that activity. The new architectures got rid of central directories entirely. One of the simplest methods, called *flooding*, works like this: Each computer in the file-sharing network maintains a list of other computers in the network. When file-sharer Beth wants to find a copy of *Angel*, her computer asks all the computers in its list. Each of those computers offers to send Beth a copy of *Angel* if it has one; otherwise it relays Beth's request to all the computers on *its* list, and so on, until the request eventually reaches a computer that has the file. Figure 6.2 illustrates the process. In contrast to the Napster-style architecture in Figure 6.1, there's no central directory. Distributed architectures like this are powerful because they can be extremely robust. The network keeps working even if many individual computers fail or go offline, as long as enough computers remain to propagate the requests.

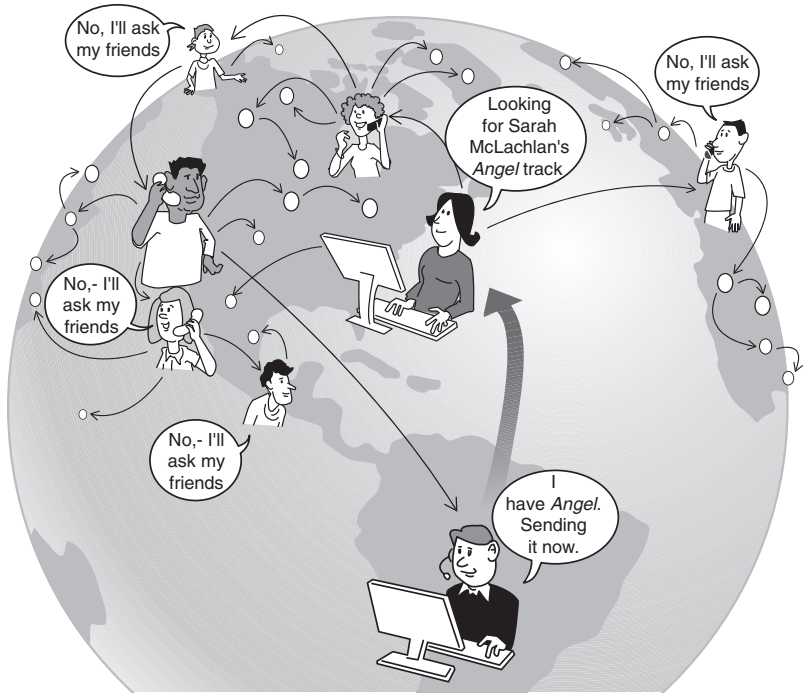

FIGURE 6.2 In contrast to Napster-style peer-to-peer systems illustrated earlier, decentralized file-sharing systems such as Grokster have no central directories.

Content-Distribution Networks

The bare-bones flooding method sketched here is too simple to support practical large networks. But the success of decentralized peer-to-peer architectures

has stimulated research into practical *content-distribution network* architectures that exploit the efficiency and robustness of peer-to-peer methods.[17]

No Safe Harbors

The companies building the new generation of file-sharing systems hoped these distributed architectures would also immunize them against liability for secondary copyright infringement. After all, once users had the software, what they did with it was beyond the companies' knowledge or control. So, how could the companies be held liable for what users did? To the recording industry, however, this was just Napster all over again: exploiting the Internet to promote copyright infringement on a massive scale. In October 2001, the RIAA sued the makers of three of the most popular systems—Grokster, Morpheus, and Kazaa—for damages of $150,000 per infringement.[18]

The three companies responded that they had no control over the users' actions. Moreover, their software was only one piece of the infrastructure that enabled file sharing, and there were many other pieces. If the three software companies were liable, wouldn't makers of the other pieces be liable as well? What about Microsoft, whose operating system lets users of one computer copy files from other computers? What about Cisco, whose routers relay the unlicensed copyrighted material? What about the computer manufacturers, whose machines run the software? Wouldn't a ruling against the file-sharing network software companies expose the entire industry to liability?

The Supreme Court had provided guidance for navigating these waters with the landmark 1984 case *Sony v. Universal Studios*.[19] In an episode that foreshadowed the *Grokster* suit 17 years later, the MPAA had sued Sony Corporation, charging Sony with secondary infringement for selling a device that was threatening to ruin the motion picture industry: the video cassette recorder. As the president of the MPAA thundered before Congress in 1982, "I say to you that the VCR is to the American film producer and the American public as the Boston strangler is to the woman home alone."[20]

In a narrow 5-to-4 decision, the Supreme Court ruled in Sony's favor, holding that even though there was widespread infringement from people using VCRs...

> ...the sale of copying equipment, like the sale of other articles of commerce, does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes. Indeed, it need merely be capable of substantial non-infringing uses.[21]

The technology industries applauded. Here was a reasonably clear criterion they could rely on in evaluating the risk in bringing new products to market. Showing that a product was capable of substantial noninfringing uses would provide a "safe harbor" against allegations of secondary infringement.

This 1984 scenario—a new technology, a threatened business model—was now being replayed in the 2001 *Grokster* suit. The file-sharing companies were quick to cite the *Sony* ruling in their defense, explaining that there were many non-infringing uses of file sharing.

In April 2003, the Central California Federal District Court agreed that the *Grokster* case was different from Napster's case and dismissed the suit, citing the *Sony* decision and commenting that the RIAA was asking the court to "expand existing copyright law beyond its well-drawn boundaries."[22] In reaction, the RIAA immediately began its campaign of suing individual users of the file-sharing software[23]—the campaign that would later snag Tanya Andersen and Jammie Thomas.

The District Court's ruling was appealed, and it was upheld by the Ninth Circuit, the same court that had ruled against Napster three years earlier:

> In short, from the evidence presented, the district court quite correctly concluded that the software was capable of substantial non-infringing uses and, therefore, that the *Sony-Betamax* doctrine applied.[24]

The RIAA naturally appealed, and when the Supreme Court agreed to review the decision, the entire networked world held its breath. Were content publishers to have no legal recourse against massive file sharing? Would the *Sony* safe harbor be overturned? In June 2005, the Court returned a unanimous verdict in favor of the RIAA:

> We hold that one who distributes a device with the object of promoting its use to infringe copyright, as shown by clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties.[25]

A Question of Intent

The content industry had won, although it ended up with less than it had hoped for. The MPAA wanted the court to be explicit in weakening the *Sony* "substantial noninfringing use" standard. Instead, the court declared that the *Sony* case was not at issue here, and it would not revisit that standard. The file-sharing companies' liability, the court said, stemmed not from the capabilities of the software but from the companies' intent in distributing it.

The technology industries (other than the three defendants, which were driven out of business) breathed an immediate sigh of relief that *Sony* had been left intact. But this was quickly followed by second thoughts. The *Grokster* decision had opened up an entirely new set of grounds on which one could be held liable for secondary infringement. As the court ruled: "Nothing in *Sony* requires courts to ignore evidence of intent to promote infringement if such evidence exists."

But what evidence? If someone accuses your company of secondary infringement, how confidently can you defend yourself against accusations of bad intent? The *Sony* safe harbor doesn't seem so safe anymore.

Take an example: The *Grokster* ruling cited "advertising an infringing use" as evidence of an active step taken to encourage infringement. Apple introduced the iTunes desktop with its CD-copying software in 2001. Early advertisements heavily promoted the product with the slogan "Rip, Mix, Burn." Was that a demonstration of Apple's bad intent? Many people certainly thought so, including the chair of Disney when he told Congress in 2002, "There are computer companies, that their ads, full-page ads, billboards up and down San Francisco and L.A., that say—what do they say?—'rip, mix, burn' to kids to buy the computer."[26]

Can your company risk introducing a product with that slogan in the post-*Grokster* era? You might expect that you would have every chance of winning an "intent" fight in court, but the risks of losing are catastrophic. In personal infringement cases like Tanya Andersen's, even the minimum statutory damage penalties of $750 per infringement could have meant a million-dollar claim over the (falsely alleged) songs on her hard drive—a staggering burden for an individual. But a technology company could conceivably be liable for damages based on *every* song illegally copied by *every* user of a device. Say you sell 14 million iPods (the number Apple sold in 2006) times 100 songs allegedly copied per iPod times $750 per song. That's more than $1 trillion in damages—more than 100 times the *total* retail revenues of the recording industry worldwide in 2006! Liability like that might seem ridiculous, but that's the law. It means that guessing wrong is a bet-the-company mistake. Better to be conservative and not introduce products with features that might prompt a lawsuit, even if you are reasonably sure that your products are legal.

No Commercial Skipping

In 2001, ReplayTV Network introduced a digital video recorder for television programs that included the ability to skip commercials automatically. It also permitted people to move recorded shows from one ReplayTV machine to another. The company, sued for secondary infringement by the major movie studios and television networks, was driven into bankruptcy before the case

was concluded. The company that bought Replay's assets settled the case, promising not to include these features in its future models.[27]

We can speculate about products and features that are unavailable today due to the uncertainties in Grokster's "intent" standard, coupled with secondary infringement penalties that could lead to nightmarish fines. Companies are naturally reluctant to give examples, but one might ask why songs shared wirelessly with Microsoft Zune players self-destruct after three plays or why some streaming services disable fast forwarding through commercials or refuse to let you move recorded movies to a PC. In 2002, the CEO of a major cable network characterized skipping commercials while watching TV as theft, although he allowed that "I guess there could be a certain amount of tolerance for going to the bathroom."[28]

But speculating about the consequences of liability alone is largely pointless because these liability risks have not been increasing in a vacuum. A second front has opened up in the copyright wars. Here, the weapons are not lawsuits but technology.

Authorized Use Only

Computers process information by copying bits—between disk and memory, between memory and networks, from one part of memory to another. Actually, most computers are able to "keep" bits in memory only by recopying them over and over, thousands of times a second. (Ordinary computers use what is called dynamic random-access memory, or DRAM. The copying is what makes it "dynamic.") The relationship of all this essential copying to the kind of copying governed by copyright law has been intellectual fodder for legal scholars—and for lawyers looking for new grounds on which to sue.

Computers cannot run programs stored on disk without copying the program code to memory. Copyright law explicitly permits this copying for the purpose of running the program. But suppose someone wants simply to *look at* the code in memory, not to run it. Does that require explicit permission from the copyright holder? In 1993, a U.S. Federal Circuit Court ruled that it does.[29]

Going further, computers cannot display images on the screen without copying them to a special part of memory called a display buffer. Does this mean that, even if you purchase a computer graphic image, you can't view the image without explicit permission from the copyright holder each time? A 1995 report from the Department of Commerce argued that it does mean exactly this. The same report went on to imply that almost any use of a digital work involves making a copy and therefore requires explicit permission.[30]

Digital Rights and Trusted Systems

Legal scholars can debate whether copyright law mandates a future of "authorized use only" for digital information. The answer may not matter much because that future is coming to pass through the technologies of digital rights management and trusted systems.

The core idea is straightforward: If computers are making it easy to copy and distribute information without permission, we need to *change computers* so that copying or distributing without permission is difficult or impossible. This is not an easy change to make; perhaps it cannot be done at all without sacrificing the computer's ability to function as a general-purpose device. But it's a change that's underway nonetheless.

Here is the issue: Suppose (fictitious) Fortress Publishers is in the business of selling content over the Web. The company would like the only people getting Fortress content to be those whose pay. Fortress can start by restricting access on its website to registered users only by requiring passwords. Much web content is sold like this today—for instance, *Wall Street Digest* or *Safari Books Online*. The method works well (or at least has worked well so far) for this type of material, but there's a problem with higher-value content. How does Fortress prevent people who've bought its material from copying and redistributing it?

One thing Fortress can do is to distribute their material in encrypted form, in such a way that it can be decrypted and processed only by programs that obey certain rules. For instance, if Fortress distributes PDF documents created with Adobe Acrobat, it can use Adobe LiveCycle Enterprise Suite to control whether people reading the PDF file with Adobe Reader are allowed to print it, modify it, or copy portions of it. Fortress can even arrange to make a document "phone home" over the Internet—that is, to notify Fortress whenever it is opened and report the IP address of the computer that is opening it. Similarly, if Fortress prepares music files for use with Windows Media Player, it can use Microsoft Windows Media Rights Manager to limit the number of times the music can be played, to control whether it can be copied to a portable player or a CD, to force it to expire after a certain period of time, or to make it phone home for permission each time it's played so that the Fortress web server can check a license and require payment if necessary.[31]

The general technique of distributing content together with control information that restricts its use is called *digital rights management* (DRM). DRM systems are widely used today, and there are industry specifications (called *rights expression languages*) that detail a wide range of restrictions that can be imposed.

DRM might appear to solve Fortress's problem, but the approach is far from airtight. How can Fortress be confident that people using their material are using it with the intended programs, the ones that obey the DRM restrictions? Encrypting the files helps, but as explained in Chapter 5, "Secret Bits," attackers break

Encryption and DRM

Chapter 5 explains public-key encryption and digital signatures—the technologies that make public distribution of encrypted material possible. The "messages" that Alice and Bob are exchanging might be not text messages but rather music, videos, illustrated documents, or anything at all. As the first koan says, "it's all just bits." Thus, the encryption technologies that Alice and Bob use for secret communication can be used by content suppliers to control the conditions under which consumers can watch movies or listen to songs.

that kind of encryption all the time—it happens regularly with PDF and Windows Media.[32] More simply, someone could modify the document reader or the media player program to save unencrypted copies of the material as they are running, and then distribute those copies all over the Internet for anyone's use.

To prevent this, Fortress could rely on the computer operating system to require that any program manipulating Fortress content be certified. Before a program is run, the operating system checks a digital signature for the program to verify that the program is approved and has not been altered. That's better, but a really clever attacker might alter the operating system so that it will run the modified program anyway. How could anyone prevent that? The answer is to build a chip into every computer that checks the operating system each time the machine is turned on. If the operating system has been modified, the computer will not boot. The chip should be tamper-proof so that any attempt to disable it will render the machine inoperable.

This basic technique was worked out during the 1980s and demonstrated in several research and advanced development projects,[33] but only since 2006 has it been ready for wide deployment in consumer-grade computers. The required chip, called a *Trusted Platform Module* (*TPM*) chip, was designed by the Trusted Computing Group, a consortium of hardware and software companies formed in 1999.[34] More than half of the computers shipped worldwide today contain TPM chips. Popular operating systems, including Microsoft Windows (beginning with Vista and continuing through Windows 10 and beyond) and several versions of GNU/Linux, can use them for security applications. One application, *trusted boot*, prevents the computer from booting if the operating system has been modified (for example, by a virus). Another application, called *sealed storage*, lets you encrypt files in such a way that they can be decrypted only on particular computers that you specify. Given today's concerns over viruses and Internet security, it's a safe bet that TPM chips will become pervasive. TPM is now used by nearly all PC and laptop manufacturers, particularly on professional products. Apple, however, does not include the technology in its products.

Asserting Control: Beyond the Bounds of Copyright

Fortress Publishers's problem could be solved in a world of digital rights management reinforced by trusted computing, but is that something we should welcome?

For one thing, it gives Fortress a level of control over use of its material that goes far beyond the bounds of copyright law. When we buy a book today, we take for granted that we have the right to read it whenever we like and as many times as we like; read it from cover to cover or skip around; lend it to a friend; resell it; copy out a paragraph for use in a book report; donate it to a school library; open it without "phoning home" to tell Fortress we are doing so. We need no permission to do any of these things. Are we willing to give up these rights when books are digital computer files? How about music? Videos? Software? Should we care?

Now leave to one side, for a moment, the dispute between music companies and listeners. DRM and trusted computing technologies, once standard in personal computers, will have other uses. The same methods that, in one country, prohibit people from playing unlicensed songs can, in another country, prevent people from listening to unapproved political speeches or reading unapproved newspapers. Developers of DRM and trusted platforms may be creating effective technologies to control the use of information, but no one has yet devised effective methods to circumscribe the limits of that control. As one security researcher warned: "Trusted computing" means that "third parties can trust that your computer will disobey your wishes."[35]

The same methods that, in one country, prohibit people from playing unlicensed songs can, in another country, prevent people from listening to unapproved political speeches or reading unapproved newspapers.

Another concern with DRM is that it increases opportunities for technology lock-in and anti-competitive mischief. It is tempting to design operating systems that run only certified applications in order to protect against viruses or bogus document readers and media players. But this can easily turn into an environment where no one can market a new media player without publishers' approval or where no one can deploy *any* application without first having it registered and approved by Microsoft, IBM, Google, or Apple. A software company that poses a competitive threat to established interests, such as publishers, operating system vendors, or computer manufacturers, might suddenly encounter "complications" in getting its products certified. One reason innovation has been so rapid in information technology is that the infrastructure is open: You don't need permission to introduce new programs and devices on the Internet. A world of trusted systems could easily jeopardize this.[36]

A third DRM difficulty is that, in the name of security and virus protection, we could easily slip into an unwinnable arms race of increasing technology lock-down that provides no real gain for content owners. As soon as attackers anywhere bypass the DRM to produce an unencrypted copy, they can distribute it—and they might be willing to go to a lot of effort to be able to do that.

Think, for example, about making unauthorized copies of movies. Very sophisticated attackers might modify the TPM hardware on their computers, putting a lot of effort into bypassing the tamper-proof chip. Here's an even easier method: Let the TPM system operate normally but hook up a video recorder in place of the computer display. That particular attack has been anticipated by the industry with a standard that requires all high-definition video to be transmitted between devices in encrypted form. There are several efforts to protect against this kind of infringement. Microsoft implemented Output Protection Management (OPM) and Intel developed HDCP (High-bandwidth Digital Content Protection) to protect video and audio content/signals. Even these protection schemes are vulnerable: You could simply point a video recorder at the screen. The result would not be high-definition quality, but once it has been digitized, it could be sent around the Internet without any further degradation.

Content owners worried about these sorts of attacks refer to them as the *analog hole*, and there seems to be no technological way to prevent them. J. K. Rowling tried to prevent unauthorized Internet copies of *Harry Potter and the Deathly Hallows* by not releasing an electronic version of the book at all. That did not stop a zealous fan from simply photographing every page and posting the entire book on the Web even before it was in bookstores.

In the words of one computer security expert, "Digital files cannot be made uncopyable, any more than water can be made not wet."[37] There is one thing for certain: The DRM approach to copyright control is difficult, frustrating, and potentially fraught with unintended consequences. Out of that frustration has emerged a third response—along with liability and DRM—to the increasing levels of copying on the Internet: outright criminalization of technology.

Forbidden Technology

The lines of text following this paragraph might be illegal to print in a book sold in the United States. We've omitted the middle four lines to protect ourselves and our publisher. Had we left them in, this would be a computer program, written in the Perl computer language, to unscramble encrypted DVDs. Informing you how to break DVD encryption so you could copy your DVDs would be a violation of 17 USC §1201, the *anti-circumvention* provision of the 1998 Digital Millennium Copyright Act (DMCA). This section of the DMCA

outlaws technology for bypassing copyright protection. Don't bother turning to the back of the book for a note telling you where to find the missing four lines. A New York U.S. District Judge ruled in 2000 that even providing so much as a web link to the code is a DMCA violation, and the Appeals Court agreed.[38]

```
s''$/=\\2048;while(<>){G=29;R=142;if((@a=unqT="C*",_)[20]&48){D=89;_unqb24.qT.@

. . . (four lines suppressed) . . .

)+=P+( F&E))for@a[128..$#a]\\}print+qT.@a}';s/[D-HO-U_]/\\$$&/g;s/q/pack+/g;eval
```

The DMCA's anti-circumvention rules do more than stop people from printing gibberish in books. They outlaw a broad class of technologies; they outlaw manufacturing them, selling them, writing about them, and even talking about them. That Congress took such a step shows the depth of the alarm and frustration at how easily DRM is bypassed. With §1201, Congress legislated not against copyright infringement but against bypassing itself, whether or not anything is copied afterward. If you find an encrypted web page that contains the raw text of the Bible and break the encryption order to read Genesis, that's not copyright infringement—but it *is* circumvention. Circumvention is its own offense, subject to many of the same penalties as copyright infringement: statutory damages and, in some cases, imprisonment. Congress intentionally chose to make the offense independent of actual infringement. Alternative proposals that would have limited the prohibition to circumvention for the purpose of copyright infringement were considered and defeated.[39]

The DMCA prohibition goes further. As §1201(a)(2) decrees:

> No person shall manufacture, import, offer to the public, provide, or otherwise traffic in any technology, product, service, device, component, or part thereof, that...is primarily designed or produced for the purpose of circumventing a technological measure that effectively controls access to a work protected under [copyright].

Here the law passes from regulating behavior (circumvention) to regulating technology itself. It's a big step, but in the words of one of the bill's supporters at the time,

> I continue to believe that we must ban devices whose major purpose is circumvention because I do not think it will work from the enforcement standpoint. That is, allowing anti-circumvention devices to proliferate freely, and outlaw only the inappropriate use of them, seems to me unlikely to deter much.[40]

In the arena of security, there is an odd asymmetry between the world of atoms and the world of bits. There are many published explanations of how to crack mechanical combination locks, and even of how to construct a physical master key for a building from a key to a single lock in the set.[41] But if the lock is digital and what is behind it is *Pirates of the Caribbean*, the rules are different. Federal law prohibits publication of any explanation of how to reverse-engineer that kind of lock.

Legislators may not have seen an effective alternative, but they crafted an awkward form of regulation that begins with a broad prohibition and then grants exemptions on a case-by-case basis. The need for exemptions became apparent even as the DMCA was being drafted. A few exemptions got written into the statute. These included permission for intelligence and law enforcement agents to break encryption during the course of investigations and permission for nonprofit libraries to break the encryption on a work—but only for the purpose of deciding whether to buy it. The law also included a complex rule that allows certain types of encryption research under certain circumstances. Recognizing that needs for new exemptions would continue to arise, Congress charged the librarian of Congress to conduct hearings to review the exemptions every three years and grant new ones if appropriate.

For instance, in November 2006, after a year-long hearing process, a new exemption[42] gave Americans the right to undo the lock-in on their mobile phones for the purpose of shifting the phone to a new cellular service provider. The ruling had a big impact nine months later, in August 2007, when Apple released its iPhone, locked to the AT&T cellular network. Users clamored to unlock their iPhones so they could be used on other networks, and several companies began selling unlocking services. But the language of the DMCA and the exemption is so murky that, while unlocking your *own* phone is legal, distributing unlocking software or even *telling* other people how to unlock their phones might still be a DMCA violation. Indeed, AT&T threatened legal action against at least one unlocking company.[43]

Copyright Protection or Competition Avoidance?

The DMCA's framework for regulation is a poor match to technology innovation because the lack of an appropriate exemption can stymie the deployment of a new device or a new application. Given the ferocity of industry competition, there's the constant temptation to exploit the broad language of the prohibition as grounds for lawsuits against competitors.

In 2002, the Chamberlain garage-door company sued a maker of universal electronic garage-door openers, claiming that the universal transmitters circumvented access controls when they sent radio signals to open and close the doors. It took two years for the case to finally die at the appeals court.[44] That

same year, Lexmark International sued a company that made replacement toner cartridges for Lexmark printers, charging that the cartridges circumvented access controls in order to function with the printer. The District Court agreed. The ruling was overturned on appeal in 2004,[45] but in the meantime, the alternative cartridges were kept off the market for a year and a half. In 2004, Storage Technology Corporation successfully convinced the Boston District Court that it was a DMCA violation for third-party vendors to service its systems. Had the appeals court not overturned the ruling, we might now be in a situation where no independent company could service computer hardware.[46] It would be as if Ford Tauruses came with their hoods sealed, and it was illegal for any mechanic not licensed by Ford to service them.

Lawsuits like these earned the DMCA the epithet "Digital Millennium Competition Avoidance." Fortunately, none of the lawsuits were ultimately successful because the courts ruled that the underlying disputes weren't sufficiently related to copyrighted material; it's unlikely that Congress intended the DMCA to apply to garage doors. But in areas where copyright enters, the anti-competitive impact of the DMCA emerges in full force.

Imagine that the 1984 Supreme Court ruling in the *Sony* case had gone the other way, and the Court had declared Sony liable for copyright infringement for selling VCRs. Would VCRs have disappeared? Almost certainly not; consumers wanted them. More likely, the electronics industry would have cut a deal with the motion picture industry, giving them control over the capabilities of VCRs. VCRs would have become highly regulated machines, regulated to meet the demands of the motion picture industry. All new VCR features would need to be approved, and any feature the MPAA didn't like would be kept off the market. The capabilities of the VCR would be under the control of the content industry.

That's the kind of world we are living in today when it comes to digital media. If a company manufactures a product that processes digital information, it needs to be concerned about copyright infringement, even without the DMCA. This is a big concern, especially after *Grokster*. But suppose the device could not be used for copyright infringement. Even then, if the digital information is restricted by DRM, the product must abide by the terms of the DRM restrictions. Otherwise, that would be circumvention, so the product couldn't be legally manufactured at all. The terms of the DRM restrictions are completely at the whim of the content provider. Once Fortress Publishers installs DRM software, it gets to dictate the behavior of any device that accesses Fortress material.

In the case of DVDs, DVD content is encrypted with an algorithm called the Content Scrambling System (CSS), developed by Matsushita and Toshiba and first introduced in 1996. As mentioned in Chapter 5, that algorithm—a textbook violation of Kerckhoffs's Principle—was quickly broken, and underground

decryption programs are today readily found on the Internet. The censored six lines of text earlier in this chapter is one such program.

Although CSS is useless for realistic copy protection, it is invaluable as an enabler of anti-competitive technology regulation. Any company marketing a product that decrypts DVDs needs a license from the DVD Copy Control Association (DVD CCA), an organization formed in 1999. The license conditions are determined by whatever the CCA decides. For example, all DVD players must obey "region coding," which limits them to playing DVDs made for one part of the world only, and an individual player's region can be changed no more than five times. Region coding has nothing to do with copyright. It is there to support a motion picture industry marketing strategy of releasing movies in different parts of the world at different times. The varied license restrictions include some that companies are not even permitted to see until after they have signed the license.

The Face of Technology Lock-in

Suppose you are a company with an idea for an innovative DVD product. Maybe it is a home entertainment system that lets people copy and store DVDs for later watching, and you have worked out a way to do this without encouraging copyright infringement. This is an actual product. Kaleidescape, the California startup that makes it, was sued by the DVD CCA in 2004 for violating a provision of the CSS license that forces DVD players to be designed to work only when there is a physical disk present. In March 2007, a California court ruled in Kaleidescape's favor, on the grounds that the license wasn't clear enough. The case was appealed and went through several reversals, with the parties finally reaching a settlement in June 2014.[47] Another startup working on a similar product at the same time folded when it failed to get venture funding, "in part due to the threat of legal action from the DVD CCA."[48]

The DVD technology lock-in has been in place since 2000. A similar lock-in is being implemented for high-definition cable TV. A campaign to extend the lock-in to all consumer media technology was promoted in Washington as the *broadcast flag initiative*. And more trial balloons keep being floated in the name of protecting copyright. A bill was introduced in Congress to ban home recording of satellite radio. NBC urged the Federal Communications Commission to force Internet service providers to filter all Internet traffic for copyright infringement (that is, to compel ISPs to check packets as they are passed around the Internet and to discard packets deemed to contain unauthorized material). In 2002, Congress considered a breathtakingly broad prohibition against *any* communications device that does not implement copyright control—a bill that had to be redrafted after it became apparent that the first draft would have banned heart pacemakers and hearing aids.[49]

So, in the United States today, a technology company is free to invent a new garage-door opener without needing its design approved by the garage-door makers. It can manufacture cheaper replacement toner cartridges without approval from the printer companies. It cannot, however, create new software applications that manipulate video from Hollywood movie DVDs without permission from the DVD CCA. It cannot in principle create *any* new product or service around DRM-restricted digital content without getting permission, often from the very people who might regard that new product as a competitive threat.

This is the regulatory posture at the present juncture in the copyright wars. People can debate the merits of this position. Some say that the DMCA is necessary. Others claim that it has been largely ineffective in curtailing infringement, as the continuing calls for ever more severe copyright penalties demonstrate. But whatever its merits, the anti-circumvention approach is poisonous to the innovation that drives the digital age. It hobbles the rapid deployment of new products and services that interoperate with existing infrastructure. The uncertain legal risks drive away the venture capital needed to bring innovations to market.

Public Knowledge (publicknowledge.org) is a Washington, DC, public-interest group that focuses on policy issues concerning digital information. See their "issues" and "policy" blogs to stay current on the latest happenings in Washington.

In essence, the DMCA has enlisted the force of criminal law in the service of the lock-in shenanigans invited by DRM. It has introduced anti-competitive regulation under the guise of copyright protection. By outlawing technology for circumventing DRM, the law has, in the words of one critic, become a tool for "circumventing competition."[50]

Copyright Koyaanisqatsi: Life Out of Balance

1982 marked the release of an astonishing film called *Koyaanisqatsi*. The title is a Hopi Indian word meaning "life out of balance." The film, which has no dialogue or narration, barrages viewers with images at once hauntingly beautiful and deeply disturbing, images that juxtapose the world of nature with the world of cities. The relentless message is that technology is destroying our ability to live harmonious, balanced lives.[51]

In the first quarter of the twenty-first century, we inhabit a world of copyright koyaanisqatsi. Virtually every salvo in the copyright war, Congressional bill introduced, lawsuit filed, court ruling issued, or advocacy piece trumpeted, pays homage to the "traditional balance of copyright" and the need to preserve it. The truth is that the balance is gone, toppled in the digital

explosion, which is likewise shattering the framework for any civil consensus over the disposition of information. The balance is gone for good reason.

Copyright (at least in the United States) is supposedly a deal the government strikes between the creator of a work and the public. The creator gets limited monopoly control over the work, for limited times, which provides the opportunity to benefit commercially. The public gets the benefit of having the work and also gets to use it without restriction after the monopoly has expired. The parameters of the deal have evolved over the years, generally in the direction of a stronger monopoly. Under the first U.S. copyright law, enacted in 1790, copyright lasted a maximum of 28 years. Today, it lasts until 70 years after the author's death. In principle, however, it's still a deal.

Digital Copyright

Digital Copyright by Jessica Litman[52] recounts the evolution of U.S. copyright law as a series of negotiated compromises. The Citizen Media Law Project (www.citmedialaw.org) offers useful information to online publishers—not just about copyright but about other legal matters as well.

It is an enormously complex deal, and it is easy see why. Today's copyright law is the outcome of 200 years of wrangling, negotiating, and compromising. The first copyright statute was printed in its entirety in two newspaper columns of the *Columbian Centinel*, shown in Figure 6.3. As the enlarged text insert shows, the law covered only maps, charts, and books, and it granted exclusive rights to "print, reprint, publish, or vend." The period of copyright was 14 years (with a 14-year renewal). Today's statute[53] runs to more than 200 pages. It's a Byzantine stew peppered with exceptions, qualifications, and arcane provisions. You can't make a public performance of a musical work unless you're an agricultural society at an agricultural fair. You can't freely copy written works, but you can if you're an association for the blind and you're making an edition of the work in Braille (but not if the work is a standardized test). A radio station can't broadcast a recording without a license from the music publisher, but it doesn't need a license from the record company—but that's only if it's an analog broadcast. For digital satellite radio, you need licenses from both (but there are exceptions).

It is a law written for specialists, not for ordinary people. Even ordinary lawyers have trouble interpreting it. But that never mattered because the copyright deal never was about ordinary people. The so-called copyright balance was largely a balancing act among competing business interests. The evolution of copyright law has been a story of the relevant players sitting down at the table and working things out, with Congress generally following suit. Ordinary people were not involved because ordinary people had no real ability to publish, and they had nothing to bring to the table.

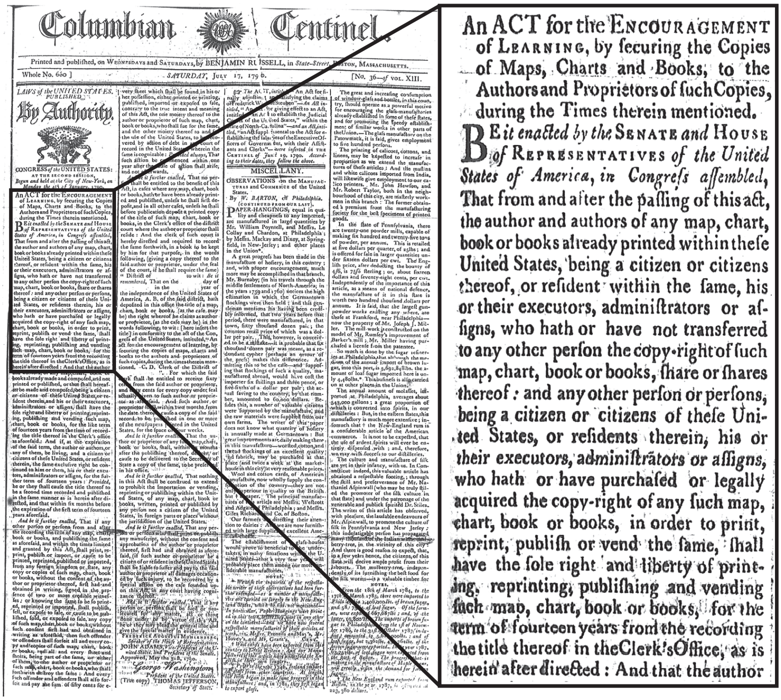

Columbian Centinel.

Printed and published, on WEDNESDAYS and SATURDAYS, by BENJAMIN RUSSELL, in State-Street, BOSTON, MASSACHUSETTS.

Whole No. 660] SATURDAY, JULY 17, 1790. [No. 36—of VOL. XIII.

An ACT for the ENCOURAGEMENT of LEARNING, by ſecuring the Copies of Maps, Charts and Books, to the Authors and Proprietors of ſuch Copies, during the Times therein mentioned.

BE it enacted by the SENATE and HOUSE of REPRESENTATIVES of the United States of America, in Congreſs aſſembled, That from and after the paſſing of this act, the author and authors of any map, chart, book or books already printed within theſe United States, being a citizen or citizens thereof, or reſident within the ſame, his or their executors, adminiſtrators or aſſigns, who hath or have not transferred to any other perſon the copy-right of ſuch map, chart, book or books, ſhare or ſhares thereof; and any other perſon or perſons, being a citizen or citizens of theſe United States, or reſidents therein, his or their executors, adminiſtrators or aſſigns, who hath or have purchaſed or legally acquired the copy-right of any ſuch map, chart, book or books, in order to print, reprint, publiſh or vend the ſame, ſhall have the ſole right and liberty of printing, reprinting, publiſhing and vending ſuch map, chart, book or books, for the term of fourteen years from the recording the title thereof in the Clerk's Office, as is herein after directed: And that the author

Harvard University Library.

FIGURE 6.3 The first U.S. copyright law—"An Act for the Encouragement of Learning." It was printed as the first two columns of the July 17, 1790, edition of the *Columbian Centinel*. Note George Washington's signature on the bill at the bottom of the second column. (Harvard University Library)

Late to the Table

The digital explosion has changed that the landscape by making it easy for anyone to copy and distribute information on a worldwide scale. We can all be publishers now. The public is now a party to the copyright deal—but the game has been going on for 200 years, and the hands were dealt long ago.

When people come to the table with their new publishing power, expecting to take full advantage of information technology, they find that there are possibilities that seem attractive, easy, and natural but for which the public's rights have already been "balanced" away. Among the lost opportunities are copying a DVD to a portable player, making the video clip equivalent of an audio mixtape, placing a favorite cartoon or a favorite song on a Facebook page, or adding your own creative input to a work of art you love and sharing that with the world.

People resent it when acts like these are denounced as theft and piracy. As a contributor to a computer bulletin board quipped, "My first-grade teacher told me I should share, and now they're telling me it's illegal."

Can You Copy Music CDs to Your Computer?

Of course, you *can* easily copy CDs to your computer hard drive: There are dozens of software packages designed to do just that, and millions of people do it regularly. Yet the legal issues in CD copying are both murky and confusing—a striking example of the mismatch of copyright law and public understanding.

In testimony at the Jammie Thomas trial in October 2007 (see the sidebar earlier in this chapter), Jennifer Pariser, the head of litigation for Sony BMG, suggested that ripping your own legally purchased CD, even for personal use, is illegal, asserting that making a copy of a purchased song is just "a nice way of saying 'steals just one copy.'" The RIAA website specifically states that there is no legal right to copy music CDs, although it allows that copying music "usually won't raise concerns" so long as the copy is for personal use, and it warns that it's illegal to give your copy away or lend it to others to copy.

The growth of streaming services such as Spotify, Pandora, Apple Music, and Amazon music might have resulted in a reduction in music copying by making content readily available. While these services often explicitly prohibit downloading, the Internet offers plenty of instruction on how to do it.

That resentment can easily grow to a sense of moral outrage. In the words of Electronic Frontier Foundation founder John Gilmore:

> What is wrong is that we have invented the technology to eliminate scarcity, but we are deliberately throwing it away to benefit those who profit from scarcity. We now have the means to duplicate any kind of information that can be compactly represented in digital media....We should be rejoicing in mutually creating a heaven on earth! Instead, those crabbed souls who make their living from perpetuating scarcity are sneaking around, convincing co-conspirators to chain our cheap duplication technology so that it *won't* make copies—at least not of the kind of goods *they* want to sell us. This is the worst sort of economic protectionism—beggaring your own society for the benefit of an inefficient local industry.[54]

But one person's sharing can be another person's theft, and the other side in the copyright war has no shortage of its own moral outrage. The motion picture industry estimates that the retail value of unauthorized movie copies floating around the Internet is more than $7 billion. As Dan Glickman, then president of the MPAA, put it:

> We will not welcome...theft masquerading as technology. No business, including the movies, can keep its doors open, its employees paid, and

> its customers satisfied if pirates and thieves are allowed to run ramshackle over this country's basic protection of the right of individuals to the ownership of their creative expressions, and to benefit from those expressions and that ownership.[55]

This is not "balance." It's a nasty firefight filled with indignation, recriminations, and a path of escalating punishments and anti-competitive regulation in the name of copyright law. As collateral damage of the battle, innovation is being held hostage.

Toward De-escalation

Getting off that path requires freeing ourselves of old ideas and perspectives. Difficult as that seems, there are grounds for optimism. During 2007, the recording industry made a major shift away from reliance on digital rights management. In addition to restraints it imposes on technology, DRM is an inconvenience both for consumers and publishers. There has been an increasing public acknowledgment of the downsides of DRM not only by consumer groups but by the industry itself.

One of the first visible moves was an announcement in February 2007 by Apple's Steve Jobs, in the form of an open letter to recording industry executives, asking them to relax the licensing restrictions that required Apple to implement DRM on iTunes music. In Jobs's view, a world of online stores selling DRM-free music that could play on any player would be "clearly the best alternative for consumers, and Apple would embrace it in a heartbeat."[56] The industry reacted coldly, but other groups chimed in to agree with Jobs. In March, MusicLoad, one of Europe's largest online music retailers, came out against DRM, noting that 75% of its customer service calls were due to DRM. MusicLoad asserted that DRM makes using music difficult for consumers and hinders the development of a mass market for legal downloads. In November, the British Entertainment Retailers Association also came out against DRM. Its director general claimed that copy protection mechanisms were "stifling growth and working against the consumer interest."[57]

By the summer of 2007, Apple iTunes and (separately) Universal Music Group began releasing music tracks that could be freely copied.[58] The iTunes tracks contained information ("watermarks") identifying the original purchaser from iTunes. That way, if large numbers of unauthorized copies would appear on the Internet, the original purchaser could be traced and held accountable.

A few months later, even that level of restriction was vanishing.[60] By the beginning of 2008, all four major music labels—Universal, EMI, Warner, and Sony/BMG—were releasing music for sale through Amazon without watermarks that identified individual buyers. It was a remarkable about-face over

the course of a year. When Jobs made his February 2007 proposal, Warner Music CEO Edgar Bronfman flat-out rejected the idea as "completely without logic or merit."[61] Before the end of the year, Warner was announcing that it would sell DRM-free music on Amazon,[62] with Bronfman explaining in a note to employees:

> By removing a barrier to the sale and enjoyment of audio downloads, we bring an energy-sapping debate to a close and allow ourselves to refocus on opportunities and products that will benefit not only WMG, but our artists and our consumers as well.[63]

Using Watermarking

Using watermarking rather than copy restrictions and access control is an example of a general approach to regulation through *accountability* rather than *restriction*. This idea is to not try to prohibit violations in advance but make it possible to identify violations when they occur and deal with them then. The same perspective can apply in privacy, as mentioned in Chapter 3, "Who Owns Your Privacy?" where one can focus on the appropriate use of personal information rather than restricting access to it.[59]

The increasing recognition that the DRM approach was failing sparked experiments with other models for distributing music on the Internet. Universal talked to Sony and other labels about a subscription service, where users would pay a fixed fee and then get as much music as they wanted. One plan linked the service to a new hardware device, where the price of the service would be folded into the price of the hardware.[64]

The general failure of DRM to provide an acceptable user experience was one, of many, factors that led to a seismic shift in how music is purchased and consumed. Record collections are relegated to museums. CDs and DVDs are sufficiently rare that laptops no longer include readers for them. Apple experienced an economic renaissance when it began selling songs for $.99 on iTunes. After 18 years, iTunes was finally laid to rest, replaced by Apple Music.

In 2019, according to the Nielsen mid-year report, 78% of music was consumed through streaming services, and only 5% through digital track sales.[65] The most popular song of the year, "Old Town Road," sold 958,000, but was streamed more than 1.3 billion times.

A complementary approach to streaming services promotes sharing of music and other creative works in a way that enriches the common culture, by making it easy for creators to distribute their own work and to build on each other's work. *Creative Commons* is an organization that provides technical and legal tools to encourage sharing. This organization distributes a

family of copyright licenses that creators can use for publishing their works on the Internet, including licenses that permit open sharing. The licenses are expressed both as legal documents and as computer code that can support new applications. If a work appears on the Web with the appropriate Creative Commons code, for example, search engines might return references to it when asked to find material that can be used under specified licensing conditions. Stimulating open sharing on the Internet is an example of moving toward a *commons*—that is, a system of sharing that minimizes the need for fine-grained property restrictions. (Chapter 8, "Bits in the Air," includes more on the notion of a commons.)[66]

Experience with these and other approaches will show whether there are economically viable models for distributing music that do not rely upon DRM. Success could pave the way for the motion picture industry and other publishers to get off the anticircumvention path—a dead end that has been more effective at harming innovation than at stopping infringement, and which even some of the original architects of the policy are now acknowledging as a failed approach.

CREATIVE COMMONS LICENSES

If you've created works that you want to publish on the Internet, you can use the Creative Commons license chooser at creativecommons.org to obtain a license fit for your needs. With the license, you can retain specified rights of your choice while granting blanket permission for other uses.

Even then, however, the larger problems created by the DMCA would not fade away because policies locked into law are not easily unlocked. If the content industry moves to better business models and the DRM battles subside, the DMCA's anticircumvention provisions may continue to be anti-consumer, anti-competitive blots on the digital landscape. Unless repealed from the legal code, they would remain as battlefield relics of a war that was settled by peaceful means—unexploded ordnance that a litigious business could still use in ways unrelated to the law's original intent.

The Limits of Property

For decades, the fights over digital music and digital video have been the front line of the copyright wars. Perhaps innovations and experiments that are already under way will help defuse those battles. The enormous potential of the Internet for good—and for profit—need not be sacrificed to combat its abuse. If you do not like what others are doing with the Internet, the Internet does not have to become your enemy—unless you make it your enemy.

Free Culture

Lawrence Lessig's *Free Culture: How Big Media Uses Technology and the Law to Lock Down Culture and Control Creativity*[67] compellingly traces the story of how overbroad copyright restrictions are jeopardizing the future of a robust and vibrant public culture.

The indignation over copyright is intense. The interest in new approaches, such as accountability and commons, suggests the deeper source of the discomfort with the metaphors of property and theft when applied to words and music. The copyright balance that is being toppled by digitization is not just the traditional tension between creator and the public. It is the balance between the individual and society that underlies our notions of property itself. Accountability and commons are attempts to find substitutes for the ever-expanding property restrictions imposed in the name of digital copyright law.

When we characterize movies, songs, and books as "property," we evoke visceral metaphors of freedom and independence: "my parcel of land versus your parcel of land." But the digital explosion is fracturing these property metaphors. "My parcel of land" might be different from "your parcel of land," but when both parcels are blown to clouds of bits, the clouds swirl together. The property lines that would separate them vanish in a fog of network packets.

And perhaps the fences have started to come down. January 1, 2019, was the first day in 20 years that a tranche of new works entered the public domain. After the 1998 Copyright Term Extension Act added 20 years to existing terms of copyright, works stopped becoming freely available regularly. The Sonny Bono Act (named after the deceased entertainer whose wife and congressional successor said "Sonny wanted the term of copyright protection to last forever") extended copyright to a period 70 years after the death of the author, or 95 years for works of corporate authorship. By that standard Windows 95 wouldn't come out of copyright until 2090, long after its bits had become irrelevant.

Some in the copyright trenches had feared that the copyright lobby would try to extend the terms further, like Disney, seeking to extend the life of Mickey Mouse. In 2019, however, the "limited times" of the constitutional clause finally kicked back in, and works published in 1923 reached the public free and clear.[68] Americans can now reprint Robert Frost's Stopping by *Woods on a Snowy Evening* and Kahlil Gibran's *The Prophet*, improvise to George Gershwin's *American in Paris*, and even remix Cecil B. DeMille's *The Ten Commandments*.

Learning to Fly Through the Digital Clouds

In 2004, Google embarked on a project to index the book collections of several large libraries for Google's search engine. The idea was that when you

searched on the Web, you could find books relevant to your search query, together with a snippet of text from the book. As Google described it, the company wanted to create "an enhanced card catalog of the world's books," and this should be no more controversial than any other card catalog.

The Association of American Publishers (AAP) and the Authors' Guild objected to the Google book project and sued Google for copyright infringement. In the words of AAP President Patricia Schroeder, "Google is seeking to make millions of dollars by freeloading on the talent and property of authors and publishers." The president of the Authors' Guild equates including a book in the project with stealing the work. At issue is the fact that Google is scanning the books and making copies in order to create the search index, and the case would turn on legal technicalities about whether this scanning constituted copyright infringement.

After a seven-year legal battle, the parties reached a settlement agreement with the American Association of Publishers. The detailed terms are confidential, though in their press releases, the parties acknowledged the complexity of allowing digital access while maintaining copyright protections.[69]

That wasn't the end of the story. The AAP and Author's Guild parted ways in the middle of the lawsuit. The Author's Guild continued their litigation through 2016 when the Supreme Court, by refusing to hear the appeal, allowed the rule in favor of Google to stand.

The library project would benefit Google by making its search engine more valuable, and Google was indeed scanning the books without permission from the copyright holders. Is the company "appropriating property" and extracting value from it without compensating the owners and without even asking for permission? Should Google be permitted to do that? In 2020, as the COVID-19 pandemic shuttered physical lending libraries, the non-profit Internet Archive stepped into the gap with a "National Emergency Library" offering short-term controlled digital lending—and publishers sued.[70] If you write a book, and that's "property" that you "own," how far should the limits of your ownership extend?

As a society, we have faced this kind of question before. If a stream runs through your land, do you own the water in the stream? Are there limits to your ownership? Can you pump out that water and sell it—even if that would cause water shortages downstream? What about the obligations of landowners upstream from you? These were major controversial issues in the western United States in the nineteenth century, which eventually resulted in codifying a system of limited property rights that landowners have to the water running through their land.

Suppose an airplane flies over your land. Is that trespassing? Suppose the plane is flying very low. How far upward do your property rights extend? From ancient times, property rights were held to reach upward indefinitely.

Copyright and Web Searching

If you believe that the Google library project violates copyright, you might wonder whether search engines themselves infringe copyright by caching and indexing websites and providing links. This claim has been the source of lawsuits, but the courts have been rejecting it. In *Field v. Google* (2006)[71], a Nevada District Court ruled that Google's caching and indexing of websites is permissible. One of the factors in the ruling was that Google stores web pages in its cache only temporarily. In *Perfect 10 v. Google* (2007),[72] the Ninth Circuit Court denied an adult magazine's request for a preliminary injunction to prevent Google from linking to its site and posting thumbnail images from it.

Perhaps airlines should be required to seek permission from every landowner whose property their planes traverse. Imagine being faced with that regulatory question at the dawn of the aviation age. Should we require airlines to obtain that permission out of respect for property and ownership? That might have seemed reasonable at a time when planes flew at only a few thousand feet. But had society done that, what would have been the implications for innovation in air travel? Would we ever have seen the emergence of transcontinental flight, or would the path to that technology have been blocked by thickets of regulation? Congress forestalled the growth of those thickets by nationalizing the navigable airspace in 1926.

Similarly, should we require Google to get permission from every book's copyright holder before including it in the index? It seems somewhat reasonable—and in fact other book indexing projects are underway that do seek that permission. Yet perhaps book search is the fledging digital equivalent of the low-flying aircraft. Can we envision the future transcontinental flights, where books, music, images, and videos are automatically extracted, sampled, mixed, and remixed; fed into massive automated reasoning engines; assimilated into the core software of every personal computer and every cell phone—and thousands of other things for which the words don't even exist yet?

What's the proper balance? How far "upward" into the bursting information space should property rights extend? What should ownership even mean when we're talking about bits? We don't know, and finding answers won't be easy. But somehow, we must learn to fly.

The digital explosion casts information every which way, breaching established boundaries of property. Technologies have confounded copyright—the rules that would regulate and restrain bits in their flight. Technological solutions have been brought to bear on the problems technology created. Those solutions created *de facto* policies of their own, bypassing the considerations of public interest on which copyright was balanced.

Property lines are not the only boundaries the explosion is breaching, and copyright is not the only arena in which information regulation is challenged. Bits fly across national borders. They fly into private homes and public places carrying content that is unwanted, even harmful—content that has historically been restricted, not by copyright, but by regulations against defamation and pornography. Yet the bits fly anyway, and that is the conundrum to which we now turn in the next chapter.

Endnotes

1 Ben Conarck, "Florida Prisoners Could Form Class Action to Demand Refund on Confiscated Media Players and Files," *The Florida Times-Union*, accessed February 19, 2019, https://www.jacksonville.com/news/20190219/florida-prisoners-could-form-class-action-to-demand-refund-on-confiscated-media-players-and-files.

2 Ashbel S. Green, "Music Goliath Unloads on Wrong David," *ORian* (Portland, OR), 2007.

3 Nationwide Class Action Allegation, US District Court, Portland OR. https://www.wired.com/images_blogs/threatlevel/files/andersenclassaction.pdf

4 David Silverman, "Why the Recording Industry Really Stopped Suing Its Customers," *Harvard Business Review*, December 22, 2008.

5 "Recording Industry vs. The People" (blog), "RIAA Sues Stroke Victim in Michigan," March 13, 2008. http://recordingindustryvspeople.blogspot.com/2007/03/riaa-sues-stroke-victim-in-michigan.html.

6 Eric Bangeman, "'I Sue Dead People...,'" *Ars Technica*, February 4, 2005, https://arstechnica.com/uncategorized/2005/02/4587-2/.

7 AfterDawn (blog), "RIAA Lawsuit Hits Family with No Computer or Internet Access," https://www.afterdawn.com/news/article.cfm/2006/04/25/riaa_lawsuit_hits_family_with_no_computer_or_internet_access.

8 Boing, Boing (blog), "RIAA's Lawsuit Against Homeless Man Not Going Entirely Smoothly," April 18, 2008. https://boingboing.net/2008/04/18/riaas-lawsuit-agains.html.

9 "17 U.S. Code § 504. Remedies for Infringement: Damages and Profits," Legal Information Institute, accessed April 30, 2020, https://www.law.cornell.edu/uscode/text/17/504.

10 David Kravets, "RIAA Jury Finds Minnesota Woman Liable for Piracy, Awards $222,000," *Wired*, October 4, 2007, https://www.wired.com/2007/10/riaa-jury-finds/.

11 Lydia Pallas Loren, "Digitization, Commodification, Criminalization: The Evolution of Criminal Copyright Infringement and the Importance of the Willfulness Requirement," *Washington University Law Review* 77, no. 3 (January 1, 1999): 835–899.

12 Josh Hartmann, "Student Indicted on Piracy Charges," *The Tech* 114, no. 19 (April 8, 1994), http://tech.mit.edu/V114/N19/piracy.19n.html.

13 *United States v. LaMacchia*, 871 F. Supp. 535 (D. Mass. 1994), https://law.justia.com/cases/federal/district-courts/FSupp/871/535/1685837/.

14 "Departments of Commerce, Justice, and State, the Judiciary, and Related Agencies Appropriations Act, 2002," Congressional Record 114, no. 119 (September 13, 2001), https://www.congress.gov/congressional-record/2001/9/13/senate-section/article/S9354-2.

15 S. Crocker, "Host Software," Internet Engineering Task Force, April 7, 1961, https://tools.ietf.org/html/rfc1.

16 Dale Dougherty et al., *2001 P2P Networking Overview: The Emergent P2P Platform of Presence, Identity, and Edge Resources* (O'Reilly Media, 2001).

17 I. Balakrishnan et al., "Looking Up Data in P2P Systems," *Communications of the ACM* 46, no. 2 (2003): 43–48.

18 *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster, Ltd.*, 259 F. Supp. 2d 1029 (C.D. Cal. 2003), https://law.justia.com/cases/federal/district-courts/FSupp2/259/1029/2362925/.

19 *Sony Corporation of America v. Universal City Studios, Inc.*, 464 U.S. 417 (1984). https://www.oyez.org/cases/1982/81-1687.

20 "Jack Valenti Testimony at 1982 House Hearing on Home Recording of Copyrighted Works," May 30, 2002, https://cryptome.org/hrcw-hear.htm.

21 *Sony Corp. v. Universal City Studios*, 464 U.S. 417 (1984), https://supreme.justia.com/cases/federal/us/464/417/.

22 *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster, Ltd.*, 259 F. Supp. 2d 1029 (C.D. Cal. 2003).

23 "Statement from the RIAA on File-Sharing," *Wall Street Journal*, June 25, 2003, https://www.wsj.com/articles/SB105656559944720700.

24 *Metro-Goldwyn-Mayer Studios, Inc.; et al., Plaintiffs-Appellants, v. Consumer Empowerment Bv, Aka Fasttrack; et al., Defendants-Appellees*, 380 F.3d 1154 (9th Cir. 2004), https://law.justia.com/cases/federal/appellate-courts/F3/380/1154/533557/.

25 *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.* (Syllabus), 545 U.S. 913 (U.S. Supreme Court 2005). https://www.oyez.org/cases/2004/04-480.

26 Michael Eisner, "Protecting Content in a Digital Age," § U.S. Senate Committee on Commerve, Science, and Transportation (2002).

27 F. Von Lohmann and W. Seltzer, "Death by DMCA," *IEEE Spectrum* 43, no. 6 (2006): 24–30, https://doi.org/10.1109/MSPEC.2006.1638041.

28 Staci D Kramer, "Content's King," *Cableworld*, April 29, 2002, https://www.2600.com/news/050102-files/jamie-kellner.txt.

29 *Mai Systems Corp. v. Peak Computer, Inc.*, 991 F.2d 511 (9th Cir. 1993), https://h2o.law.harvard.edu/collages/34140.

30 Bruce A. Lehman, "Intellectual Property and the National Information Infrastructure, The Report of the Working Group on Intellectual Property Rights," 1995, https://eric.ed.gov/?id=ED387135.

31 Microsoft, "PlayReady DRM," February 8, 2017, https://docs.microsoft.com/en-us/windows/uwp/audio-video-camera/playready-client-sdk.

32 "Microsoft Windows Media Copy Protection Broken," *Informitv*, September 12, 2006, https://informitv.com/2006/09/12/microsoft-windows-media-copy-protection-broken/.

33 Stephen Thomas Kent, "Protecting Externally Supplied Software in Small Computers," doctoral thesis (Massachusetts Institute of Technology, 1980), http://www.dtic.mil/docs/citations/ADA104678.

34 "Welcome to Trusted Computing Group," Trusted Computing Group, accessed April 30, 2020, https://trustedcomputinggroup.org.

35 Bryan Alexander, "The State of Digital Rights Management," Mindjack, March 21, 2003, http://mindjack.com/feature/drm.html.

36 Jonathan Zittrain, *The Future of the Internet—and How to Stop It* (Yale University Press, 2008).

37 "Crypto-Gram," Schneier on Security, May 15, 2001, https://www.schneier.com/cryptogram/archives/2001/0515.html.

38 "Law School Case Brief: Universal City Studios v. Corley - 273 F.3d 429 (2d Cir. 2001)," accessed April 30, 2020, https://www.lexisnexis.com/community/casebrief/p/casebrief-universal-city-studios-v-corley.

39 John Ashcroft, "S.1146 - Digital Copyright Clarification and Technology Education Act of 1997," 105th Congress (1997–1998), September 3, 1997, https://www.congress.gov/bill/105th-congress/senate-bill/1146.

40 U.S. Rep. Barney Frank, "Letter to Hal Abelson," July 6, 1998.

41 Matt Blaze, "Cryptology and Physical Security: Rights Amplification in Master-Keyed Mechanical Locks," October 22, 2002, http://eprint.iacr.org/2002/160.

42 "Exemption to Prohibition on Circumvention of Copyright Protection Systems for Access Control Technologies," *Federal Register* 71, no. 227 (November 27, 2006) https://www.copyright.gov/fedreg/2006/71fr68472.html.

43 David Kravets, "Unlocking Your IPhone Is Legal; Distributing the Hack, Maybe Not," *Wired*, August 27, 2007, https://www.wired.com/2007/08/to-unlock-the-i/.

44 "Chamberlain Group Inc. v. Skylink Technologies Inc.," Electronic Frontier Foundation, July 1, 2011, https://www.eff.org/cases/chamberlain-group-inc-v-skylink-technologies-inc.

45 *Lexmark Int'l, Inc. v. Static Control Components, Inc.*, 572 U.S. 118 (2014), https://supreme.justia.com/cases/federal/us/572/118/.

46 William Curtis Bryson, *Storage Technology Corp. v. Custom Hardware Engineering & Consulting, Inc.*, Wikisource, August 24, 2005.

47 Franklin D. Elia, (2009-10-17), "DVD Copy Control Association v. Kaleidescape, Inc.," Case H031631. Court of Appeal of the State of California, Sixth Appellate District, https://caselaw.findlaw.com/ca-court-of-appeal/1385010.html.

48 Rick Merritt, "Judge Rules Against DVD Consortium," *EETimes*, March 29, 2007, https://www.eetimes.com/judge-rules-against-dvd-consortium/.

49 Ernest F. Hollings, "S.2048 - Consumer Broadband and Digital Television Promotion Act," 107th Congress (2001–2002):, March 21, 2002, https://www.congress.gov/bill/107th-congress/senate-bill/2048.

50 "Circumventing Competition: The Perverse Consequences of the Digital Millennium Copyright Act," Cato Institute, March 21, 2006, https://www.cato.org/publications/policy-analysis/circumventing-competition-perverse-consequences-digital-millennium-copyright-act.

51 "Koyaanisqatsi: Life out of Balance," http://www.koyaanisqatsi.org/films/koyaanisqatsi.php. Ironically, the film was unavailable through much of the 1980s due to a copyright dispute.

52 Jessica Litman, *Digital Copyright* (Prometheus Books, 2001).

53 "Copyright Law of the United States," U.S. Copyright Office," accessed May 1, 2020, https://www.copyright.gov/title17/.

54 John Gilmore, "What's Wrong with Copy Protection," The Ethical Spectacle, 2001, http://www.spectacle.org/0501/gilmore.html.

55 Simon Avery, "Cuban Backs Grokster," *The Globe and Mail*, March 29, 2005, https://www.theglobeandmail.com/report-on-business/cuban-backs-grokster/article1116498/.

56 "Steve Jobs's Statement on DRM," *The Wall Street Journal*, February 6, 2007, https://www.wsj.com/articles/SB117079254216799934.

57 Andrew Edgecliffe-Johnson, "Anti-Piracy Moves 'Hurt Sales,'" *FT.Com*, 2007.

58 Ken Fisher, "Universal to Track DRM-Free Music Online via Watermarking," *Ars Technica*," August 15, 2007, https://arstechnica.com/uncategorized/2007/08/universal-to-track-drm-free-music-online-via-watermarking/.

59 Daniel J. Weitzner et al., "Information Accountability," *Communications of the ACM* 51, no. 6 (June 1, 2008): 82–87, https://doi.org/10.1145/1349026.1349043.

60 Eliot Van Buskirk, "Some of Amazon's MP3 Tracks Contain Watermarks," *Wired*, September 25, 2007, https://www.wired.com/2007/09/some-of-amazons/.

61 Gregg Keizer, "Warner Chief Calls Jobs' DRM Fight 'Without Logic,'" *Computerworld*, February 9 2007, https://www.computerworld.com/article/2543284/warner-chief-calls-jobs--drm-fight--without-logic-.html.

62 Peter Sayer, "Warner to Offer Music via Amazon without DRM," *Computerworld*, December 28, 2007, https://www.computerworld.com/article/2538422/warner-to-offer-music-via-amazon-without-drm.html.

63 Jessica Mintz, "Warner Music Group in Deal with Amazon.Com to Sell Songs Online Free of Copy Protection," *The Seattle Times*, December 27, 2007, https://www.seattletimes.com/business/warner-music-group-in-deal-with-amazoncom-to-sell-songs-online-free-of-copy-protection/.

64 Peter Burrows, "Universal Music Takes on iTunes," *Bloomberg Businessweek*, 2007.

65 https://www.nielsen.com/wp-content/uploads/sites/3/2019/06/nielsen-us-music-mid-year-report-2019.pdf.

66 For a thorough discussion of commons and especially their relationship to the digital environment, see Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (Yale University Press, 2006).

67 Lawrence Lessig, *Free Culture: How Big Media Uses Technology and the Law to Lock Down Culture and Control Creativity* (Penguin, 2004).

68 In 1923, all copyrights were issued for a term of years, since extended to 95. Life of the author became a factor in the 1976 Copyright Act. See https://www.copyright.gov/circs/circ15a.pdf.

69 https://www.prnewswire.com/news-releases/publishers-and-google-reach-agreement-172650721.html.

70 Matt Enis, Publishers' Lawsuit Against Internet Archive Continues Despite Early Closure of Emergency Library, *Library Journal*, August 17, 2020, https://www.libraryjournal.com/?detailStory=publishers-lawsuit-against-internet-archive-continues-despite-early-closure-of-emergency-library.

71 *Field v. Google, Inc.*, 412 F.Supp. 2d 1106 (D. Nev. 2006).

72 *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007).

This page intentionally left blank

CHAPTER 7

You Can't Say That on the Internet

Guarding the Frontiers of Digital Expression

Child Sex Trafficking Goes Digital

M. A. was 13 years old when she snuck off to a back-to-school party with some friends. Then she disappeared for 270 days.[1] After months of desperate searching, her mother clicked on an ad for "escort" services on the website backpage.com and found a photo of her daughter being offered for sex, along with other girls. The ads used heart emojis to pitch the girls' youth and innocence—and an umbrella emoji if the ground rules required the use of condoms.

After being abducted, M. A. had been sold for forcible sex multiple times a day. To gain control of her, the pimps had beaten and stabbed her. They addicted her to drugs, so that even after her mother rescued her, she ran back to her captors to feed her dependency. According to the National Center for Missing and Exploited Children, there are thousands—likely tens of thousands—of similarly abused victims of child sex trafficking.[2]

Until it was shut down in 2017, backpage.com was the favored site for online child sex trafficking. Most of the ads in its "Adult" section were for sex of one kind or another, and a significant number of those suggested that the sex being offered was with a minor child.

Child sex trafficking did not begin with the Internet, but the Internet has made it easy, efficient, and highly profitable. An ad lists a phone number to call; arrangements are made between the client and a pimp or the child herself or himself; the pimp can schedule the child in half-hour blocks and move the child from place to place. As observed in the documentary *I Am Jane Doe*, the

horrifying efficiency of digital sex trafficking—some girls are sold as many as 20 times a day—makes it more profitable than the drug trade.

M. A.'s mother and the parents of other trafficked children sued backpage.com. There was no question that sex with minor children is illegal everywhere. And yet backpage.com won every legal case brought against it for facilitating child sex trafficking. It even won a case *it* brought against Illinois sheriff Thomas Dart, who had prevailed upon credit card companies to stop accepting payments for services offered in backpage.com.[3] The site got the court to restore its credit card privileges. In every case, the basic logic was the same for the findings in favor of backpage.com: The site was a publisher, not a sex trafficker. In the United States, publishers have a lot of latitude in terms of what they can publish—and digital publishers enjoy some extra protections.

M. A.'s mother was understandably bewildered that the justice system considered the website blameless for the trauma her daughter had experienced. Some of the nation's top lawyers took similar cases against backpage.com and lost them all. It took three years before federal agents finally shut down backpage.com—for money-laundering, among other things—and a judge dismissed backpage's suit against Sheriff Dart.[4]

What law could possibly immunize Backpage's flagrant invitations to buy children for sex? Ironically, something called the Communications Decency Act (CDA). The CDA says, in essence, that a website isn't responsible for content that it posts that was written by others. This part of the CDA was designed to protect newspapers and blogs with comment sections, so the editors wouldn't have to check all the wild claims and misstatements offered by members of the public.

This chapter is the story of the dilemmas created when pre-Internet metaphors for speaking, writing, and publishing are applied to a medium in which everyone can participate. The explosion in digital communications has confounded long-held assumptions about human relationships—how people meet, how they come to know each other, and how they decide if they can trust each other. At the same time, the explosion in digital information has put at the fingertips of millions material that only a few years ago no one could have found without great effort and expense. Anyone can tell a story on a web page, in photographs and videos, and via posts to social networks. Political dissidents in Chinese Internet cafés can (if they dare) read pro-democracy blogs. People all around the world who are ashamed about their illnesses, starved for information about their sexual identity, or eager to connect with others of their minority faith can find facts, opinion, advice, and companionship. And children too small to leave home by themselves can see lurid pornography on their families' home computers. Can societies control what their members see and to whom they talk?

Metaphors for Something Unlike Anything Else

The latest battle in a long war between conflicting values is the SESTA–FOSTA package—the Stop Enabling Sex Traffickers Act, coupled with the Allow States and Victims to Fight Online Sex Trafficking Act. Passed into law in April 2018,[5] it amended the Communications Decency Act to assign some liability to sites that knowingly facilitate sex trafficking. It did so at the cost of limiting certain forms of speech on the Internet in ways arguably inconsistent with First Amendment protections of free speech. Since SESTA–FOSTA says nothing about the traffickers themselves—their activities were already illegal—critics (including some in law enforcement) argued that it might simply drive the trafficking business underground or, to be precise, back into dark alleys. The moral victory over sleazy sites like backpage.com, and the satisfaction at seeing their owners pay a price, may come at a price not only to free expression but to the children being trafficked, if the commerce moves off the Internet, where at least it can be monitored. The politicians may proclaim that they have done something about child sex trafficking, but perhaps all they have done is to hide it from view. Indeed, within months of the passage of SESTA–FOSTA, reports began to appear that sex workers, unable to communicate with their clients electronically, were returning to the age-old and far more dangerous practice of soliciting them on the streets, under the direction of pimps.[6] At the same time, SESTA–FOSTA creates liability for the proprietors of online networks who can't monitor their every user and establishes a precedent for restrictions on other forms of disfavored speech.

SESTA-FOSTA was resisted because of concerns that it might affect much more of the Internet than the disreputable sites to which it was meant to respond. This controversy is the latest in a series of conflicts set off by the Internet's unprecedented communications capabilities. On the one hand, society has an interest in protecting children, and the Internet brings torrents of digital information of every kind directly into our households. On the other hand, society has an interest in maximizing open communication. The U.S. Constitution largely protects the freedom to speak and the right to hear. Over and over, society has struggled to find a metaphor for electronic communication that captures the ways in which it is the same as the media of the past and the ways in which it is different. Laws and regulations are built on traditions; only by understanding the analogies can the speech principles of the past be extended to the changed circumstances of the present—or be consciously transcended.

What laws should apply? The Internet is not exactly like anything else. If you put up a website, that is something like publishing a book, so perhaps the laws about books should apply. But that was Web 1.0—a way for "publishers" to publish and viewers to view. In today's digital universe, sites and services

such as Facebook change constantly in response to user postings. If you send a text message, or contribute to a blog, that is something like placing a telephone call, or maybe a conference call, so maybe laws about telephones should be the starting point. Neither metaphor is perfect. Maybe television is a better analogy, since browsing the Web is like channel surfing—except that the Internet is two-way, and there is no limit to the number of "channels."

Underneath the web and application software is the Internet itself. The Internet just delivers packets of bits, not knowing or caring whether they are parts of books, movies, text messages, or voices, nor whether the bits will wind up in a web browser, a telephone, or a movie projector. John Perry Barlow, former lyricist for the Grateful Dead and co-founder of the Electronic Frontier Foundation, used a striking metaphor to describe the Internet as it burst into public consciousness in the mid-1990s. The world's regulation of the flow of information, he said, had long controlled the transport of wine bottles. In "meatspace," the physical world, different rules applied to books, postal mail, radio broadcasts, and telephone calls—different kinds of bottles. Now the wine itself flowed freely through the network, nothing but bits freed from their packaging. Anything could be put in, and the same kind of thing would come out. But in between, it was all the same stuff—just bits. What are the rules of cyberspace? What are the rules for the bits themselves?[7]

When information is transmitted between two parties, whether the information is spoken words, written words, pictures, or movies, there is a source and a destination. There may also be some intermediaries. In a lecture hall, the listeners hear the speaker directly, although whoever provided the hall also played an important role in making the communication possible. Books have authors and readers, but also publishers and booksellers in between. It is natural to ascribe similar roles to the various parties in an Internet communication, and, when things go wrong, to hold any and all of the parties responsible.

The Internet has a complex structure. The source and destination may be friends texting each other, they may be a commercial website and a customer sitting at home, or they may be one office of a company sending a mockup of an advertising brochure to another office halfway around the world. The source and destination each has an ISP. Connecting the ISPs are routing switches, fiber-optic cables, satellite links, and so on. A packet that flows through the Internet may pass through devices and communication links owned by dozens of different parties. In the style of Jonathan Zittrain, we'll depict the collection of devices that connect the ISPs to each other as a cloud. As shown in Figure 7.1, speech on the Internet goes from the source to an ISP, into the cloud, out of the cloud to another ISP, and to its destination.

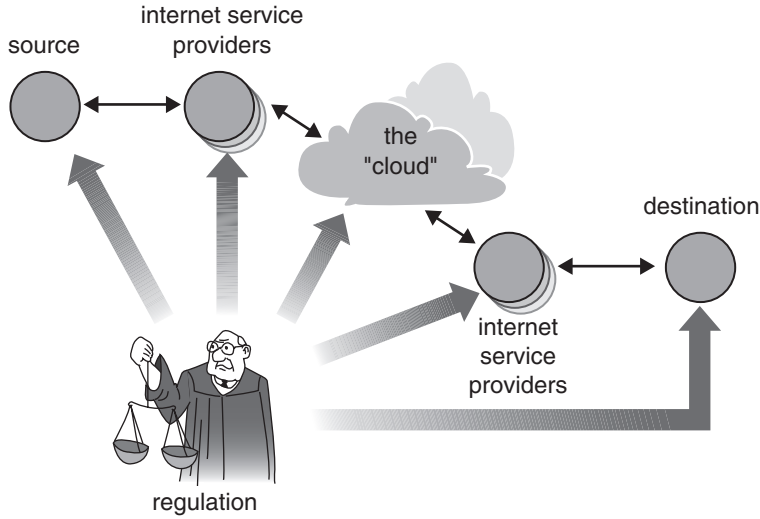

FIGURE 7.1 Where to regulate the Internet? (Based on figure by Jonathan Zittrain)

If a government seeks to control speech, it can attack at several different points. It can try to control the speaker or the speaker's ISP, by criminalizing certain kinds of speech. But that won't work if the speaker isn't in the same country as the listener. It can try to control the listener, by prohibiting possession of certain kinds of materials. In the United States, using copyrighted software without an appropriate license is unlawful because software can't be used without being copied. Also unlawful is distribution of other copyrighted material with the intent to profit. If citizens have reasonable privacy rights, however, it is hard for the government to know what its citizens have. In a society such as the United States, where citizens have reasonable rights of due process, one-at-a-time prosecutions for possession are unwieldy. As a final alternative, the government can try to control the intermediaries.

Very early, defamation laws had to adapt to the Internet. In the United States, speech is defamatory if it is false, communicated to third parties, and damages one's reputation.

In the physical world, when the speaker defames someone, the intermediaries between the speaker and the listener sometimes share responsibility with the speaker—and sometimes not. If we defame someone in this book, we could be sued, but so could the book's publisher, which might have known that what we were writing was false. On the other hand, the truckers who transported the book to the bookstore probably aren't liable, even though they too helped get our words from us to our readers. Are the various electronic intermediaries more like publishers or truckers? The cases against backpage.com rest on the answer to this question.

Defaming Public Figures

Damaging statements about public figures, even if false, are not defamatory unless they were made with malicious intent. This extra clause protects news media against libel claims by celebrities who are offended by the way the press depicts them. What President Donald Trump described as America's "very weak" libel laws[8] are barely 50 years old, however. The pivotal case was *New York Times Co. v. Sullivan*,[9] in which the newspaper was sued by officials in Alabama on the basis of a pro-civil-rights advertisement it published. The story is detailed, along with a readable history of the First Amendment, in *Make No Law* by Anthony Lewis.[10] For a later account of First Amendment struggles, see Anthony Lewis's *Freedom for the Thought That We Hate*.[11]

Society has struggled to identify the right metaphors to describe the parties to an electronic communication. To understand this part of the story of electronic information, we have to go back to pre-Internet electronic communication.

Publisher or Distributor?

CompuServe was an early provider of computer services, including bulletin boards and other electronic communities users could join for a fee. One of these fora, Rumorville USA, provided a daily newsletter of reports about broadcast journalism and journalists. CompuServe didn't screen or even collect the rumors posted on Rumorville. It contracted with a third party, Don Fitzpatrick Associates (DFA), to provide the content. CompuServe simply posted whatever DFA provided, without reviewing it. And for a long time, no one complained.

In 1990, a company called Cubby, Inc. started a competing service, Skuttlebut, which also reported gossip about TV and radio broadcasting. Items appeared on Rumorville describing Skuttlebut as a "new start-up scam" and alleging that its material was being stolen from Rumorville. Cubby cried foul and went after CompuServe, claiming defamation. CompuServe acknowledged that the postings were defamatory but claimed it was not acting as a publisher of the information—just a distributor. It simply was sending on to subscribers what other people gave it. It wasn't responsible for the contents, any more than a trucker is responsible for libel that might appear in the magazines he handles.

What was the right analogy? Was CompuServe more like a newspaper or more like the trucker who transports the newspaper to its readers?

More like the trucker, ruled the court. A long legal tradition held distributors blameless for the content of the publications they delivered. Distributors can't be expected to have read all the books on their trucks. Grasping for a better analogy, the court described CompuServe as "an electronic for-profit library." Distributor or library, CompuServe was independent of DFA and couldn't be held responsible for libelous statements in what DFA provided. The case of *Cubby v. CompuServe*[12] was settled decisively in CompuServe's favor. Cubby might go after the source, but that wasn't CompuServe. CompuServe was a blameless intermediary.

When *Cubby v. CompuServe* was decided, providers of computer services everywhere exhaled. If the decision had gone the other way, electronic distribution of information might have become a risky business that few dared to enter. Computer networks created an information infrastructure unprecedented in its low overhead. A few people could connect tens of thousands, even millions, to each other at very low cost. If everything disseminated had to be reviewed by human readers before it was posted, to ensure that any damaging statements were truthful, its potential use for participatory democracy would be severely limited. For a time, a spirit of freedom ruled.

Neither Liberty nor Security

"The law often demands that we sacrifice some liberty for greater security. Sometimes, though, it takes away our liberty to provide us less security."[13] So wrote law professor Eugene Volokh in the fall of 1995, commenting on a court case that looked similar to *Cubby v. CompuServe*, but in one crucial respect wasn't the same.

Prodigy was a provider of computer services, much like CompuServe. But in the early 1990s, as worries began to rise about the sexual content of materials available online, Prodigy sought to distinguish itself as a family-oriented service. It pledged to exercise editorial control over the postings on its bulletin boards. "We make no apology," Prodigy stated, "for pursuing a value system that reflects the culture of the millions of American families we aspire to serve. Certainly no responsible newspaper does less." Prodigy's success in the market was due in no small measure to the security families felt in accessing its fora, rather than the anything-goes sites offered by other services.

One of Prodigy's bulletin boards, called "Money Talk," was devoted to financial services. In October 1994, someone anonymously posted comments on Money Talk about the securities investment firm Stratton Oakmont. The firm, said the unidentified poster, was involved in "major criminal fraud." Its president was "soon to be proven criminal." The whole company was a "cult of brokers who either lie for a living or get fired."[14]

The saga of Stratton Oakmont was dramatized in the 2013 film *The Wolf of Wall Street,* produced by Martin Scorsese and starring Leonardo DiCaprio. The film's portrayal of Stratton Oakmont aligns with the claims made by Money Talk's anonymous poster.

Stratton Oakmont sued Prodigy for libel, claiming that Prodigy should be regarded as the publisher of these defamatory comments. It asked for $200 million in damages. Prodigy countered that it had zero responsibility for what its posters said. The matter had been settled several years earlier by the *Cubby v. CompuServe* decision. Prodigy wasn't the publisher of the comments, just the distributor.

In a decision that stunned the Internet community, a New York court ruled otherwise. By exercising editorial control in support of its family-friendly image, said the court, Prodigy became a publisher, with the attendant responsibilities and risks. Indeed, Prodigy had likened itself to a newspaper publisher and could not at trial claim to be something less.

It was all quite logical, as long as the choice was between two metaphors: distributor or publisher. In reality, though, a service provider wasn't exactly like either. Monitoring for bad language was a pretty minor form of editorial work. That was a far cry from checking everything for truthfulness.

Be that as it may, the court's finding undercut efforts to create safe districts in cyberspace. After the decision, the obvious advice went out to bulletin board operators: Don't even consider editing or censoring. If you do, *Stratton Oakmont v. Prodigy* means you may be legally liable for any malicious falsehood that slips by your review. If you don't even try, *Cubby v. CompuServe* means you are completely immune from liability.

This was fine for the safety of the site operators, but what about the public interest? Freedom of expression was threatened, since fewer families would be willing to roam freely through the smut that would be posted. At the same time, security would not be improved, since defamers could always post their lies on the remaining services with their all-welcome policies.

The Nastiest Place on Earth

Every communication technology has been used to control, as well as to facilitate, the flow of ideas. Barely a century after the publication of the Gutenberg Bible, Pope Paul IV issued a list of 500 banned authors. In the United States, the First Amendment protects authors and speakers from government interference: "Congress shall make no law...abridging the freedom of speech, or of the press." But First Amendment protections are not absolute. No one has the right to publish obscene materials. The government can destroy materials

it judges to be obscene, as postal authorities did in 1918 when they burned magazines containing excerpts of James Joyce's *Ulysses*.

What exactly counts as obscene has been a matter of much legal wrangling over the course of U.S. history. The prevailing standard today is the one the Supreme Court used in 1973 in deciding the case of *Miller v. California*, which is therefore called the *Miller test*.[15] To determine whether material is obscene, a court must consider the following:

1. Whether the average person, applying contemporary community standards, would find that the work, taken as a whole, appeals to the prurient interest
2. Whether the work depicts or describes, in a patently offensive way, sexual conduct specifically defined by the applicable state law
3. Whether the work, taken as a whole, lacks serious literary, artistic, political, or scientific value

Only if the answer to each part is "yes" does the work qualify as obscene. The Miller decision was a landmark because it established that there are no national standards for obscenity. There are only "community" standards, which could be different in Mississippi than in New York City. But there were no computer networks in 1973. What is a "community" in cyberspace?

In 1992, the infant World Wide Web was hardly worldwide, but many Americans were using dial-up connections to access information on centralized electronic bulletin boards. Some bulletin boards were free and united communities of interest—lovers of baseball or birds, for example. Others distributed free software. Bob and Carleen Thomas of Milpitas, California, ran a different kind of bulletin board, called Amateur Action. In their advertising, they described it as "The Nastiest Place on Earth."

For a fee, anyone could download images from Amateur Action. The pictures were of a kind not usually shown in polite company but readily available in magazines sold in the nearby cities of San Francisco and San Jose. The Thomases were raided by the San Jose police, who thought they might have been distributing obscene materials. After looking at their pictures, the police decided that the images were not obscene by local standards.

Bob and Carleen were not indicted, and they added this notice to their bulletin board: "The San Jose Police Department as well as the Santa Clara County District Attorney's Office and the State of California agree that Amateur Action BBS is operating in a legal manner."[16]

Two years later, in February 1994, the Thomases were raided again, and their computer was seized. This time, the complaint came from Agent David

Dirmeyer, a postal inspector—in *western Tennessee*. Using an assumed name, Dirmeyer had paid $55 and had downloaded images to his computer in Memphis. Nasty stuff indeed, particularly for Memphis: bestiality, incest, and sadomasochism. The Thomases were arrested. They stood trial in Memphis on federal charges of transporting obscene material via common carrier and via interstate commerce. They were convicted by a Tennessee jury, which concluded that their Milpitas bulletin board violated the community standards of Memphis. Bob was sentenced to 37 months incarceration and Carleen to 30.

The Thomases appealed their conviction, on the grounds that they could not have known where the bits were going and that the relevant community, if not San Jose, was a community of cyberspace. The appeals court did not agree. Dirmeyer had supplied a Tennessee postal address when he applied for membership in Amateur Action. The Thomases had called him at his Memphis telephone number to give him the password; they had known where he was. The Thomases, concluded the court, should have been more careful where they sent their bits, once they started selling them out of state. Shipping the bits was just like shipping a videotape by UPS (a charge of which the Thomases were also convicted).[17] The laws of meatspace applied to Cyberspace—and one city's legal standards sometimes applied thousands of miles away.

The Most Participatory Form of Mass Speech

Pornography was part of the electronic world from the moment it was possible to store and transmit words and images. The Thomases learned that bits were like books, and the same obscenity standards applied.

In the mid-1990s, something else happened. The spread of computers and networks vastly increased the number of digital images available and the number of people viewing them. Digital pornography became not just the same old thing in a new form; it seemed to be a brand-new thing because there was a lot of it, and it was easy to obtain in the privacy of the home. Nebraska Senator James Exon attached an anti-Internet-pornography amendment to a telecommunications bill, but it seemed destined for defeat on civil liberties grounds. And then all hell broke loose.

On July 3, 1995, *Time Magazine* blasted "CYBERPORN" across its cover. The accompanying story, based largely on a single university report, stated:

> What the Carnegie Mellon researchers discovered was: THERE'S AN AWFUL LOT OF PORN ONLINE. In an 18-month study, the team surveyed 917,410 sexually explicit pictures, descriptions, short stories, and film clips. On those Usenet newsgroups where digitized images are stored, 83.5% of the pictures were pornographic.[18]

The article later noted that this statistic referred to only a small fraction of all data traffic but failed to explain that the offending images were mostly on limited-access bulletin boards, not openly available to children or anyone else. It mentioned the issue of government censorship, and it quoted John Perry Barlow on the critical role of parents. Nonetheless, when Senator Chuck Grassley of Iowa read the *Time Magazine* story into the Congressional Record, attributing its conclusions to a study by the well-respected Georgetown University Law School, he called on Congress to "help parents who are under assault in this day and age" and to "help stem this growing tide."

Grassley's speech, and the circulation in the Capitol building of dirty pictures downloaded by a friend of Senator Exon, galvanized the Congress to save the children of America. In February 1996, the Communications Decency Act (CDA) passed almost unanimously and was signed into law by President Clinton.

The CDA made it a crime to use "any interactive computer service to display in a manner available to a person under 18 years of age, any comment, request, suggestion, proposal, image, or other communication that, in context, depicts or describes, in terms patently offensive as measured by contemporary community standards, sexual or excretory activities or organs." Criminal penalties would also fall on anyone who "knowingly permits any telecommunications facility under such person's control to be used" for such prohibited activities. And finally, it criminalized the transmission of materials that were "obscene or indecent" to persons known to be under 18.

These "display provisions" of the CDA vastly extended existing anti-obscenity laws, which already applied to the Internet. The dual prohibitions against *making offensive images available to a person under 18* and against transmitting *indecent materials to persons known to be under 18* were unlike anything that applied to print publications. "Indecency," whatever it meant, was something short of obscenity, and only obscene materials had been illegal prior to the CDA. A newsstand could tell the difference between a 12-year-old customer and a 20-year-old, but how could anyone check ages in cyberspace?

When the CDA was enacted, John Perry Barlow saw the potential of the Internet for the free flow of information challenged. He issued a now-classic manifesto against the government's effort to regulate speech:[19]

> Governments of the Industrial World, you weary giants of flesh and steel, I come from Cyberspace, the new home of Mind. On behalf of the future, I ask you of the past to leave us alone. You have no sovereignty where we gather....We are creating a world that all may enter without privilege or prejudice accorded by race, economic power, military force, or station of birth. We are creating a world where anyone, anywhere may express his or her beliefs, no matter how singular,

> without fear of being coerced into silence or conformity....In our world, all the sentiments and expressions of humanity, from the debasing to the angelic, are parts of a seamless whole, the global conversation of bits....[Y]ou are trying to ward off the virus of liberty by erecting guard posts at the frontiers of Cyberspace.

Brave and stirring words, even if the notion of cyberspace as a "seamless whole" had already been rendered doubtful. At a minimum, bits had to meet different obscenity standards in Memphis than in Milpitas, as the Thomases had learned. In fact, the entire metaphor of the Internet as a "space" with "frontiers" was fatally flawed, and misuse of that metaphor continues to plague laws and policies to this day.

Civil libertarians joined the chorus challenging the Communications Decency Act. In short order, a federal court and the U.S. Supreme Court ruled in the momentous case of *ACLU v. Reno*. The display provisions of the CDA were unconstitutional. "The Government may only regulate free speech for a compelling reason," wrote Judge Dalzell in the district court decision, "and in the least restrictive manner." It would chill discourse unacceptably to demand age verification over the Internet from every person who might see material that any adult has a legal right to see.

The government had argued that the authority of the Federal Communications Commission (FCC) to regulate the content of TV and radio broadcasts, which are required not to be "indecent," provided an analogy for government oversight of Internet communications.

The courts disagreed. The FCC analogy was wrong, they ruled, because the Internet was far more open than broadcast media. Different media required different kinds of laws, and the TV and radio laws were more restrictive than laws were for print media or should be for the Internet. Judge Dalzell wrote:

> I have no doubt that a Newspaper Decency Act, passed because Congress discovered that young girls had read a front page article in the *New York Times* on female genital mutilation in Africa, would be unconstitutional....The Internet may fairly be regarded as a never-ending worldwide conversation. The Government may not, through the CDA, interrupt that conversation. As the most participatory form of mass speech yet developed, the Internet deserves the highest protection from governmental intrusion.[20]

The CDA's display provisions were dead. In essence, the court was unwilling to risk the entire Internet's promise as a vigorous marketplace of ideas to serve the narrow purpose of protecting children from indecency. Instead, it transferred the burden of blocking unwanted communications from source

ISPs to the destination. Legally, there seemed to be nowhere else to control speech except at the point where it came out of the cloud and was delivered to the listener.

Defending Electronic Freedoms

The Electronic Frontier Foundation, www.eff.org, is the leading public advocacy group defending First Amendment and other personal rights in cyberspace. Ironically, it often finds itself in opposition with media and telecommunications companies. In principle, communications companies should have the greatest interest in unfettered exchange of information. In actual practice, they often benefit financially from policies that limit consumer choice or expand surveillance and data gathering about private citizens. The EFF was among the plaintiffs bringing suit in the case that overturned the CDA.

Lost in the 1995–1996 Internet indecency hysteria was the fact that the "Carnegie Mellon report" that started the legislative ball rolling had been discredited almost as soon as the *Time Magazine* story appeared. The report's author, Martin Rimm, was an electrical engineering undergraduate. His study's methodology was flawed—and perhaps fraudulent. For example, he told adult bulletin board operators that he was studying how best to market pornography online and that he would repay them for their cooperation by sharing his tips.[21] His conclusions were unreliable. Why hadn't that been caught when his article was published? Because the article was not a product of Georgetown University, as Senator Grassley had said. Rather, it appeared in the *Georgetown Law Review*, a student publication that used neither peer nor professional reviewers. Three weeks after publishing the "Cyberporn" article, *Time* acknowledged that Rimm's study was untrustworthy. In spite of this repudiation, Rimm salvaged something from his efforts: He published a book called *The Pornographer's Handbook: How to Exploit Women, Dupe Men, & Make Lots of Money*.

Protecting Good Samaritans—and a Few Bad Ones

The *Stratton Oakmont v. Prodigy* decision, which discouraged ISPs from exercising any editorial judgment, had been handed down in 1995, just as Congress was preparing to enact the Communications Decency Act to protect children from Internet porn. Congress recognized that the consequences of *Stratton Oakmont* would be fewer voluntary efforts by ISPs to screen their

sites for offensive content. So, the bill's sponsors added a "Good Samaritan" provision to the CDA.

The intent was to allow ISPs to act as editors without running the risk that they would be held responsible for the edited content, thus putting themselves in the jam in which Prodigy had found itself. So the CDA included a provision absolving ISPs of liability on account of anything they did, in good faith, to filter out "obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable" material. For good measure, the CDA pushed the *Cubby* court's "distributor" metaphor to the limit, and beyond. ISPs should *not* be thought of as publishers or as sources either. "No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider."[22] This was the bottom line of Section 230 of the CDA, and it meant that there would be no more *Stratton Oakmont v. Prodigy* catch-22s.

When the U.S. Supreme Court struck down the CDA in 1996, it negated only the display provisions, the clauses that threatened the providers of "indecent" content. The Good Samaritan clause was allowed to stand and remains the law today. ISPs can do as much as they want to filter or censor their content, without any risk that they will assume publishers' liabilities in the process—or as little as they choose, as Ken Zeran learned to his sorrow a few years later.

The CDA and Discrimination

The "Good Samaritan" clause envisioned a sharp line between "service providers" (which got immunity) and "content providers" (which did not). But as the technology world evolved, the distinction became fuzzy. A roommate-matching service was sued in California, on the basis that it invited users to discriminate by categorizing their roommate preferences (women only, for example). A court ruled that the operators of the website were immune as service providers. An appeals court reversed the decision, on the basis that the website became a content provider by filtering the information applicants provided: People seeking female roommates would not learn about men looking for roommates. There was nothing wrong with *that*, but the principle that the roommate service had *blanket* protection, under the CDA, to filter as it wished would mean that with equal impunity, it could ask about racial preferences and honor them. That form of discrimination would be illegal in newspaper ads. "We doubt," wrote the appeals court judge, "this is what Congress had in mind when it passed the CDA."[23]

The worst terrorist attack in history on U.S. soil prior to the 2001 destruction of New York's World Trade Center was the bombing of the Alfred P. Murrah Federal Building in Oklahoma City on April 19, 1995. The bombing killed 168 people, some of them children in a day care center. Hundreds more were injured when the building collapsed around them and glass and rubble rained down on the neighborhood.

Less than a week later, someone with screen name "Ken ZZ03" posted an advertisement on an America Online (AOL) bulletin board. Ken had "Naughty Oklahoma T-Shirts" for sale. Among the available slogans were "Visit Oklahoma—it's a Blast" and "Rack'em, Stack'em, and Pack'em—Oklahoma 1995." Others were even cruder and more tasteless. To get your T-shirt, said the ads, you should call Ken. The posting gave Ken's phone number.

The number belonged to Ken Zeran, an artist and filmmaker in Seattle, Washington. Zeran had nothing to do with the posting on AOL. It was a hoax.

Ken Zeran started to receive calls. Angry, insulting calls. Then death threats.

Zeran called AOL and asked them to take down the posting and issue a retraction. An AOL employee promised to take down the original posting but said retractions were against company policy.

The next day, an anonymous poster with a slightly different screen name offered more T-shirts for sale, with even more offensive slogans.

Call Ken. And by the way—there's high demand. So if the phone is busy, call back.

Zeran kept calling AOL to ask that the postings be removed and that further postings be prevented. AOL kept promising to close down the accounts and remove the postings, but it didn't. By April 30, Ken was receiving a phone call every two minutes. Ken's art business depended on that phone number; he couldn't change it or fail to answer it without losing his livelihood.

About this time, Shannon Fullerton, the host of a morning drive-time radio talk show on KRXO in Seattle, received by email a copy of one of the postings. Usually his show was full of light-hearted foolishness, but after the bombing, Fullerton and his radio partner had devoted several shows to sharing community grief about the Oklahoma City tragedy. Fullerton read Ken's T-shirt slogans over the air. And he read Ken's telephone number and told his listeners to call Ken and tell him what they thought of him.

Zeran got even more calls and more death threats. Fearing for his safety, he obtained police surveillance of his home. Most callers were not interested in hearing what Ken had to say when he answered the phone, but he managed to keep one on the line long enough to learn about the KRXO broadcast. Zeran contacted the radio station. KRXO issued a retraction, after which the number of calls Ken received dropped to 15 per day. Eventually, a newspaper exposed the hoax. AOL finally removed the postings, after having left them visible for a week. Ken's life began to return to normal.

Zeran sued AOL, claiming defamation, among other things. By putting up the postings and leaving them up long after it had been informed that they were false, AOL had damaged him severely, he claimed.

The decision went against Zeran, and the lower court's decision held up on appeal. AOL certainly had behaved like a publisher, by communicating the postings in the first place and by choosing not to remove them when informed that they were fraudulent. Unlike the defendant in the *Cubby v. CompuServe* case, AOL knew exactly what it was publishing. But the Good Samaritan provision of the CDA specifically stated that AOL should not legally be *treated* as a publisher. AOL had no liability for Zeran's woes.[24]

Zeran's only recourse was to identify the actual speaker, the pseudonymous Ken ZZ03 who made the postings. And AOL would not help him do that. Everyone felt sorry for Ken, but the system gave him no help.

Was the Radio Station Liable?

Zeran sued the radio station separately but failed in that effort as well. Much as he may have suffered, reasoned the court, it wasn't defamation because none of the people who called him even knew who Ken Zeran was—so his reputation couldn't possibly have been damaged when the radio station spoke ill of "Ken"![25]

The posters could evade responsibility as long as they remained anonymous, as they easily could on the Internet. And Congress had given the ISPs a complete waiver of responsibility for the consequences of false and damaging statements, even when the ISPs knew they were false. Had anyone in Congress thought through the implications of the Good Samaritan clause?

Laws of Unintended Consequences

The Good Samaritan provision of the CDA has been the friend of free speech and a great relief to Internet service providers. Yet its application has defied logical connection to the spirit that created it. In ruling for backpage.com, a federal district court echoed the frustration voiced by M. A.'s mother and others:

> Let me make it clear that the court is not unsympathetic to the tragic plight described by Jane Doe No. 1, Jane Doe No. 2, and Jane Doe No. 3. Nor does it regard the sexual trafficking of children as anything other than an abhorrent evil. Finally, the court is not naïve—I am fully aware that sex traffickers and other purveyors of illegal wares ranging from drugs to pornography exploit the vulnerabilities of the Internet as a marketing tool. Whether one agrees with its stated policy or not..., Congress has made the determination that the balance between

> suppression of trafficking and freedom of expression should be struck in favor of the latter in so far as the Internet is concerned.[26]

The CDA also created a safe space for libel. Sidney Blumenthal was an aide of Bill Clinton while he was president. Blumenthal's job was to dish dirt on the president's enemies. On August 11, 1997, conservative online columnist Matt Drudge reported, "Sidney Blumenthal has a spousal abuse past that has been effectively covered up." The White House denied it, and the next day Drudge withdrew the claim. The Blumenthals sued AOL, which had a deal with Drudge—and had deeper pockets: The Blumenthals asked for $630,000,021. AOL was as responsible for the libel as Drudge, claimed the Blumenthals, because AOL could edit what Drudge supplied. AOL could even insist that Drudge delete items AOL did not want posted. The court sided with AOL and cited the Good Samaritan clause of the CDA. AOL couldn't be treated like a publisher, so it couldn't be held liable for Drudge's falsehoods. Case closed.[27]

It was against this context that the Good Samaritan clause of the Communications Decency Act strangely came to protect the use of the Internet for sex crimes.

In 1998, Jane and John Doe, a mother and her minor son, sued AOL for harm inflicted on the son. The Does alleged that AOL chat rooms were used to sell pornographic images of the boy made when he was 11 years old. They claimed that in 1997, Richard Lee Russell had lured John and two other boys to engage in sexual activities with each other and with Russell. Russell then used AOL chat rooms to market photographs and videotapes of these sexual encounters.

Jane Doe complained to AOL. Under the terms of its agreement with its users, AOL specifically reserved the right to terminate the service of anyone engaged in such improper activities. And yet AOL did not suspend Russell's service or even warn him to stop what he was doing. The Does wanted compensation from AOL for its role in John Doe's sexual abuse.

The Does lost. Citing the Good Samaritan clause and the precedent of the *Zeran* decision, the Florida courts held AOL blameless. Online service providers that knowingly allow child pornography to be marketed on their bulletin boards could not be treated as though they had published ads for kiddie porn.

The Does appealed and lost again. The decision in AOL's favor was 4–3 at the Florida Supreme Court. Judge J. Lewis fairly exploded in his dissenting opinion. The Good Samaritan clause was an attempt to remove disincentives from the development of filtering and blocking technologies, which would assist parents in their efforts to protect children. "It is inconceivable," wrote Lewis, that Congress intended the CDA to shield from potential liability an ISP alleged to have taken absolutely no actions to curtail illicit activities... while profiting from its customer's continued use of the service." The law had

been transformed into one "which both condones and exonerates a flagrant and reprehensible failure to act by an ISP in the face of...material unquestionably harmful to children." This made no sense, opined Lewis. The sequence of decisions, he wrote, "thrusts Congress into the unlikely position of having enacted legislation that encourages and protects the involvement of ISPs as silent partners in criminal enterprises for profit."[28]

The problem, as Judge Lewis saw it, was that it wasn't enough to say that ISPs were not like publishers. They really were more like distributors—as Ken Zeran had tried to argue—and distributors are not *entirely* without responsibility for what they distribute. A trucker who knows he is carrying child pornography and is getting a cut of the profits has *some* legal liability for his complicity in illegal commerce. His role is not that of a publisher, but it is not nothing either. The *Zeran* court had created a muddle by using the wrong analogy. Congress had made the muddle possible by saying nothing about the right analogy after saying that publishing was the wrong one.

All this led, some 20 years later, to the harrowing frustration of M. A.'s mother. In denying the appeal in the case against backpage.com, the First Circuit Court of Appeals made clear that the CDA was intentionally broad, and if it needed to be narrowed, it was up to Congress to alter it, not the courts to reinterpret it. "Congress did not sound an uncertain trumpet when it enacted the CDA, and it chose to grant broad protections to Internet publishers. Showing that a website operates through a meretricious business model is not enough to strip away those protections. If the evils that the appellants have identified are deemed to outweigh the First Amendment values that drive the CDA, the remedy is through legislation, not through litigation."[29]

Can the Internet Be Like a Magazine Store?

After the display provision of the CDA was ruled unconstitutional in 1997, Congress went back to work to protect America's children. The Child Online Protection Act (COPA), passed into law in 1998, contained many of the key elements of the CDA but sought to avoid the CDA's constitutional problems by narrowing it. It applied only to "commercial" speech and criminalized knowingly making available to minors "material harmful to minors." For the purposes of this law, a "minor" was anyone under 17. The statute extended the Miller test for obscenity to create a definition of material that was not obscene but was "harmful to minors":

> The term "material that is harmful to minors" means any communication...that (A) the average person, applying contemporary community standards, would find, taking the material as a whole and with respect

> to minors, is designed to appeal to...the prurient interest; (B) depicts, describes, or represents, in a manner patently offensive with respect to minors,...[a] sexual act, or a lewd exhibition of the genitals or post-pubescent female breast; and (C) taken as a whole, lacks serious literary, artistic, political, or scientific value for minors.[30]

COPA was challenged immediately and never took effect. A federal judge enjoined the government from enforcing it, ruling that it was likely to be unconstitutional. The matter bounced between courts through two presidencies. The case started out as *ACLU v. Reno*, for a time was known as *ACLU v. Ashcroft*, and was decided as *ACLU v. Gonzalez*. The judges were uniformly sympathetic to the intent of Congress to protect children from material they should not see. But in March 2007, the ax finally fell on COPA. Judge Lowell A. Reed, Jr., of U.S. District Court for the Eastern District of Pennsylvania, confirmed that the law went too far in restricting speech.

Part of the problem was with the vague definition of material "harmful to minors." The prurient interests of a 16-year-old were not the same as those of an 8-year-old; and what had literary value for a teenager might be valueless for a younger child. How would a website designer know which standard to use to avoid the risk of imprisonment?

But there was an even more basic problem. COPA was all about keeping away from minors material that would be perfectly legal for adults to have. It put a burden on information distributors to ensure that recipients of such information were of age. COPA provided a "safe harbor" against prosecution for those who in good faith checked the ages of their customers. Congress imagined a magazine store where the clerks wouldn't sell dirty magazines to children who could not reach the countertop and might ask for identification of any who appeared to be of borderline age. The law envisioned that something similar would happen in cyberspace:

> It is an affirmative defense to prosecution under this section that the defendant, in good faith, has restricted access by minors to material that is harmful to minors (A) by requiring use of a credit card, debit account, adult access code, or adult personal identification number; (B) by accepting a digital certificate that verifies age; or (C) by any other reasonable measures that are feasible under available technology.

The big problem was that these methods either didn't work or didn't even exist. Not every adult has a credit card, and credit card companies don't want their databases used to check customers' ages. And if you don't know what is meant by an "adult personal identification number" or a "digital certificate that verifies age," don't feel badly—neither do we. Clauses (B) and (C) were

basically a plea from Congress for the industry to come up with some technical magic for determining age at a distance.

In the state of the art, however, computers can't reliably tell if the party on the other end of a communications link is human or is another computer. For a computer to tell whether a human is over or under the age of 17, even imperfectly, would be very hard indeed. Mischievous 15-year-olds could get around any simple screening system that could be used in the home. The Internet just isn't like a magazine store.

Even if credit card numbers or personal identification systems could distinguish children from adults, Judge Reed reasoned, such methods would intimidate computer users. Fearful of identity theft or government surveillance, many computer users would refuse interrogation and would not reveal personal identifying information as the price for visiting websites deemed "harmful to minors." The vast electronic library would, in practice, fall into disuse and start to close down, just as an ordinary library would become useless if everyone venturing beyond the children's section had to endure a background check.

Congress's safe harbor recommendations, concluded Judge Reed, if they worked at all, would limit Internet speech drastically. Information adults had a right to see would, realistically, become unavailable to them. The filtering technologies noted when the CDA was struck down had improved, so the government could not credibly claim that limiting speech was the only possible approach to protecting children. And even if the free expression concerns were calmed or ignored, and even if everything COPA suggested worked perfectly, plenty of smut would still be available to children. The Internet was borderless, and COPA's reach ended at the U.S. frontier. COPA couldn't stop the flood of harmful bits from abroad.

Summing up, Reed quoted the thoughts of Supreme Court Justice Kennedy about a flag-burning case: "The hard fact is that sometimes we must make decisions we do not like. We make them because they are right, right in the sense that the law and the Constitution, as we see them, compel the result." Much as he was sympathetic to the end of protecting children from harmful communications, Judge Reed concluded, "perhaps we do the minors of this country harm if First Amendment protections, which they will with age inherit fully, are chipped away in the name of their protection."[31]

Let Your Fingers Do the Stalking

Newsgroups for sharing sexual information and experiences started in the early 1980s. By the mid-1990s, there were specialty sites for every orientation and inclination. So when a 28-year-old woman entered an Internet chat room in 1998 to share her sexual fantasies, she was doing nothing out of the

ordinary. She longed to be assaulted, she said, and invited men reading her email to make her fantasy a reality. "I want you to break down my door and rape me," she wrote.

What *was* unusual was that she gave her name and address—and instructions about how to get past her building's security system. Over a period of several weeks, nine men took up her invitation and showed up at her door, often in the middle of the night. When she sent them away, she followed up with a further email to the chat room, explaining that her rejections were just part of the fantasy.[32]

In fact, the "woman" sending the emails was Gary Dellapenta, a 50-year-old security guard whose attentions the actual woman had rebuffed.[33] The victim of this terrifying hoax did not even own a computer. Dellapenta was caught because he responded directly to emails sent to entrap him. He was convicted and imprisoned under a recently enacted California anti-"cyberstalking" statute. The case was notable not because the events were unusual but because it resulted in a prosecution and conviction. Most victims are not so successful in seeking redress. Most states lacked appropriate laws, and most victims could not identify their stalkers. Sometimes the stalker did not even know the victim—but simply found her contact information somewhere in cyberspace.

Speeches and publications with frightening messages have long received First Amendment protections in the United States, especially when their subject is political. Only when a message is likely to incite "imminent lawless action" (in the words of a 1969 Supreme Court decision) does speech become illegal—a test rarely met by printed words.[34] This high threshold for government intervention builds on a "clear and present danger" standard explained most eloquently by Justice Louis Brandeis in a 1927 opinion:

> Fear of serious injury cannot alone justify suppression of free speech.... No danger flowing from speech can be deemed clear and present, unless the incidence of the evil apprehended is so imminent that it may befall before there is opportunity for full discussion.[35]

Courts apply the same standard to websites. An anti-abortion group listed the names, addresses, and license plate numbers of doctors performing abortions on a website called the "Nuremberg Files." It suggested stalking the doctors and updated the site by graying out the names of those who had been wounded and crossing off those who had been murdered. The website's creators acknowledged that abortion was legal and claimed not to be threatening anyone—only collecting dossiers in the hope that the doctors could at some point in the future be held accountable for "crimes against humanity." The anti-abortion group was taken to court in a civil action. After a long legal

process, the group was found liable for damages because "true threats of violence were made with the intent to intimidate."

The courts had a very difficult time with the question of whether the Nuremberg Files website was threatening or not, but there was nothing intrinsic to the mode of publication that complicated that decision. In fact, the same group had issued paper "WANTED" posters, which were equally part of the materials at issue. Reasonable jurists could, and did, come to different conclusions about whether the text on the Nuremberg Files website met the judicial threshold.[36]

But the situation of Dellapenta's victim and other women in similar situations seemed to be different. The scores being settled at the expense of these women had no political dimensions. There were already laws against stalking and telephone harassment; the Internet was being used to recruit proxy stalkers and harassers. Following the lead of California and other states, Congress passed a federal anti-cyberstalking law.

Like an Annoying Telephone Call?

The 2005 Violence Against Women and Department of Justice Reauthorization Act[37] (signed into law in early 2006) assigned criminal penalties to anyone who:

> ...utilizes any device or software that can be used to originate telecommunications or other types of communications that are transmitted, in whole or in part, by the Internet...without disclosing his identity and with intent to annoy, abuse, threaten, or harass any person.

The clause was little noticed when the act was passed in the House on a voice vote and in the Senate unanimously.

Civil libertarians again howled, this time about a single word in the legislation. It was fine to outlaw abuse, threats, and harassment by Internet. Those terms had some legal history. Although it was not always easy to tell whether the facts fit the definitions, at least the courts had standards for judging what these words meant.

But "annoy"? People put lots of annoying things on websites and say lots of annoying things in chat rooms. There is even a website, annoy.com, devoted to posting annoying political messages anonymously. Could Congress really have intended to ban the use of the Internet to annoy people?

Congress had extended telephone law to the Internet, on the principle that harassing VoIP calls should not receive more protection than harassing landline telephone calls. In using broad language for electronic communications, however, it created another in the series of legal muddles about the aptness of a metaphor.

The Telecommunications Act of 1934[38] made it a criminal offense for anyone to make "a telephone call, whether or not conversation ensues, without disclosing his identity and with intent to annoy, abuse, threaten, or harass any person at the called number." In the world of telephones, the ban posed no threat to free speech because a telephone call is one-to-one communication. If the person you are talking to doesn't want to listen, your free speech rights are not infringed. The First Amendment gives you no right to be sure anyone in particular hears you. If your phone call is unwelcome, you can easily find another forum in which to be annoying. The CDA, in a clause that was not struck down along with the display provisions, extended the prohibition to faxes and emails—still, basically, person-to-person communications. But harassing VoIP calls were not criminal under the Telecommunications Act. In an effort to capture all telephone-like technologies under the same regulation, the same clause was extended to all forms of electronic communication, including the vast "electronic library" and "most participatory form of mass speech" that is the Internet.

Defenders of the law assured alarmed bloggers that "annoying" sites would not be prosecuted unless they also were personally threatening, abusive, or harassing. This was an anti-cyberstalking provision, they argued, not a censorship law. Speech protected by the First Amendment would certainly be safe. Online publishers, on the other hand, were reluctant to trust prosecutors' judgment about where the broadly written statute would be applied. And based on the bizarre and unexpected uses to which the CDA's Good Samaritan provisions had been put, there was little reason for confidence that the legislative context for the law would restrict its application to one corner of cyberspace.

The law was challenged by The Suggestion Box, which described itself as helping people send anonymous emails for reasons such as to "report sensitive information to the media" and to "send crime tips to law enforcement agencies anonymously."[39] The law, as the complaint argued, might criminalize the sort of employee whistleblowing that Congress encouraged in the aftermath of scandals about corporate accounting practices. The Suggestion Box dropped its challenge when the government stated that mere annoyances would not be prosecuted—only communications meant "to instill fear in the victim." So the law remained in force, with many left wishing that Congress would be more precise with its language!

Digital Protection, Digital Censorship, and Self-Censorship

The First Amendment's ban on government censorship complicates government efforts to protect the safety and security of U.S. citizens. Given a choice between protection from personal harm and some fool's need to spout

profanities, most of us would opt for safety. Security is immediate, and freedom is long term, and most people are short-range thinkers. Most people think of security as a personal thing and gladly leave it to the government to worry about the survival of the nation.

Given a choice between protection from personal harm and some fool's need to spout profanities, most of us would opt for safety.

But in the words of one scholar, the bottom line on the First Amendment is that "in a society pledged to self-government, it is never true that, in the long run, the security of the nation is endangered by the freedom of the people."[40] The Internet censorship bills have passed Congress by wide margins because members of Congress dare not be on record as voting against the safety of their constituents—and especially against the safety of children. Relatively isolated from political pressure, the courts have repeatedly undone speech-restricting legislation passed by elected officials.

Free speech precedes the other freedoms enumerated in the Bill of Rights, but not just numerically. In a sense, it precedes them logically as well. In the words of Supreme Court Justice Benjamin Cardozo, it is "the matrix, the indispensable condition, of nearly every other form of freedom."[41]

INTERNET FREEDOM

A great many organizations devote significant effort to maintaining the Internet's potential as a free marketplace of ideas. In addition to EFF, mentioned earlier in this chapter, some others include the Electronic Privacy Information Center (www.epic.org); The Free Expression Network (freeexpression.org), which is actually a coalition; the American Civil Liberties Union (www.aclu.org); Lumen (lumendatabase.org), which catalogs material that has been removed from the Web due to legal threats; and the Freedom House reports (https://freedomhouse.org/report-types/freedom-net), which catalog Internet censorship globally.

For most governments, the misgivings about censoring electronic information are less profound.

In mainland China, you can't get to Gmail, YouTube, Pinterest, or Facebook. In Saudi Arabia, you can't get to www.sex.com. Many religiously conservative nations censor sites that promote religions deviating from the national norm. Denying the holocaust is unlawful in Germany, and sites promoting such denials are inaccessible there.[42]

The disparity of information freedom standards between the United States and other countries creates conflicts when electronic transactions involve two nations. As discussed in Chapter 3, "Who Owns Your Privacy?," China insists that Google not help its citizens get information the government

does not want them to have. If you try to get to certain websites from your hotel room in Shanghai, you suddenly lose your Internet connection, with no explanation. You might think there was a glitch in the network somewhere, except that you can reconnect and visit other sites with no problem. Visitors to China routinely avoid surveillance by using VPN connections, but connecting to the Internet in China is risky for other reasons. It is so likely that users there will be infected with malware that it is now standard practice for businesspeople to use laptops, tablets, and phones that have on them no data they aren't willing to disclose and to have these devices scrubbed to their factory settings after they are brought back from abroad.

Self-censorship by Internet companies is also increasing—and it is the price they pay for doing business in certain countries. A Google official described censorship as the company's "No. 1 barrier to trade."[43] Stirred by the potential costs in lost business and legal battles, Internet companies have become outspoken information libertarians, even as they do what must be done to meet the requirements of foreign governments. It is easy for Americans to shrug their isolationist shoulders over such problems. As long as all the information is available in the United States, one might reason, who cares what version of Google or YouTube runs in totalitarian regimes abroad? That is for those countries to sort out.

But the free flow of information into the United States is threatened by the laws of other nations about the operation of the press. Consider the case of Joseph Gutnick and *Barron's* magazine.

On October 30, 2000, the financial weekly *Barron's* published an article suggesting that Australian businessman Joseph Gutnick was involved in money laundering and tax evasion. Gutnick sued Dow Jones Co., the publisher of *Barron's*, for defamation.[44] The suit was filed in an Australian court. Gutnick maintained that the online edition of the magazine, available in Australia for a fee, was in effect published in Australia. Dow Jones countered that the place of "publication" of the online magazine was New Jersey, where its web servers were located. The suit, it argued, should have been brought in a U.S. court and judged by the standards of U.S. libel law, which are far more favorable to the free speech rights of the press. The Australian court agreed with Gutnick, and the suit went forward. Gutnick ultimately won an apology from Dow Jones and $580,000 in fines and legal costs.[45]

The implications seem staggering. Americans on American soil expect to be able to speak very freely, but the Australian court claimed that the global Internet made Australia's laws applicable, regardless of where the bits reaching Australian soil may have originated. The Amateur Action conundrum about what community standards apply to the borderless Internet had been translated to the world of global journalism. Will the freedom of the Internet press henceforth be the minimum applying to any of the nations of the earth?

Is it possible that a rogue nation could cripple the global Internet press by extorting large sums of money from alleged defamers or by imposing death sentences on reporters it claimed had insulted their leaders?[46]

The American press tends to fight hard for its right to publish the truth, but the censorship problems reach into Western democracies more insidiously for global corporations not in the news business. It is sometimes easier for American companies to meet the minimum "world" standards of information freedom than to keep different information available in the United States. There may even be reasons in international law and trade agreements that make such accommodations to censorship more likely. Consider the trials of Yahoo! France.

In May 2000, the League Against Racism and Anti-Semitism (LICRA, in its French acronym) and the Union of French Jewish Students (UEJF) demanded to a French court that Yahoo! stop making Nazi paraphernalia available for online auction, stop showing pictures of Nazi memorabilia, and prohibit the dissemination of anti-Semitic hate speech on discussion groups available in France. Pursuant to the laws of France, where the sale and display of Nazi items is illegal, the court concluded that what Yahoo! was doing was an offense to the "collective memory" of the country and a violation of Article R654 of the Penal Code. It told Yahoo! that the company was a threat to "internal public order" and that it had to make sure no one in France could view such items.

Yahoo! removed the items from the yahoo.fr site ordinarily available in France. LICRA and UEJF then discovered that from within France, they could also get to the American site, yahoo.com, by slightly indirect means. Reaching across the ocean in a manner reminiscent of the Australian court's defamation action, the French court demanded that the offending items, images, and words be removed from the American website as well.

Yahoo! resisted for a time, claiming it couldn't tell where the bits were going—an assertion somewhat lacking in credibility since the company tended to attach French-language advertising to web pages if they were dispatched to locations in France. Eventually, Yahoo! made a drastic revision of its standards for the U.S. site. Hate speech was prohibited under Yahoo!'s revised service terms with its users, and most of the Nazi memorabilia disappeared. But Nazi stamps and coins were still available for auction on the U.S. site, as were copies of *Mein Kampf*. In November 2000, the French court affirmed and extended its order: *Mein Kampf* could not be offered for sale in France. The fines were adding up.

Yahoo! sought help in U.S. courts. It had committed no crime in the United States, it stated. French law could not leap the Atlantic and trump U.S. First Amendment protections. Enforcement of the French order would have a chilling effect on speech in the United States. A U.S. district court agreed, and the

decision was upheld on appeal by a three-judge panel of the Court of Appeals for the Ninth Circuit (northern California).

But in 2006, the full 11-member court of appeals reversed the decision and found against Yahoo!. The company had not suffered enough, according to the majority opinion, nor tried long enough to have the French change their minds, for appeal to First Amendment protections to be appropriate. A dissenting opinion spoke plainly about what the court seemed to be doing. "We should not allow a foreign court order," wrote Judge William Fletcher, "to be used as leverage to quash constitutionally protected speech."[47]

Such conflicts will be more common in the future, as more bits flow across national borders. The laws, trade agreements, and court decisions of the next few years will shape the world of the future. For example, The European Union has implemented a "right to be forgotten," which the European Court of Justice has interpreted to require search engines and other third parties to remove links to personal information that a European individual asserts is "inaccurate, inadequate, irrelevant or excessive"—including information that may have been truthful at the time of its posting. The Court of Justice of the European Union ruled that Mario Costeja González, a Spanish citizen, was entitled to have Google Spain remove a link to an article about an old foreclosure auction on a debt he had subsequently paid.[48] In the United States, by contrast, the First Amendment would protect search engines and other publishers of truthful—even if outdated—information.

Since then, search engines have received more and more requests to remove links on the grounds of right to be forgotten, and many are now complying with those demands in Europe—but preserving the relevant links in non-European searches. Some European data protection commissioners have argued that such local removal is insufficient, and the "forgotten" links should be removed globally. This dispute is ongoing.

It would be a sad irony if information liberty, so stoutly defended for centuries in the United States, would fall in the twenty-first century to a combination of domestic child protection laws and international money-making opportunities. But as one British commentator said when a photo-hosting site removed photos to conform with orders from Singapore, Germany, Hong Kong, and Korea, "Libertarianism is all very well when you're a hacker. But business is business."[49]

What About Social Media?

Can you say whatever you want on a social network like Facebook? Are social platforms more like publishers or distributors? On the one hand, Facebook could not function at all if, like a publisher, it were held responsible for every

word posted by every one of its users. Facebook users generate 4 petabytes of data every day.[50] That is a massive amount of data. So the protections provided by Section 230 of the CDA mean that if Alice says something about Bob on Facebook that Bob doesn't like, Facebook can't be held liable.

On the other hand, Facebook doesn't have to observe First Amendment guarantees of free speech. It can create rules of its own for what can be said in words or shown in imagery. And it *does* create these rules so that people will want to use Facebook. Many users would drop away if they found the experience unpleasant. So there are rules, and exceptions to the rules, and exceptions to the exceptions.

The rules get complicated because Facebook's core mission is to connect people, lots of people, and different people are interested in talking about, showing, and seeing lots of different things. As Facebook explains,

> To ensure that everyone's voice is valued, we take great care to craft policies that are inclusive of different views and beliefs, in particular those of people and communities that might otherwise be overlooked or marginalized.[51]

Facebook allows no nudity, although baby photos are okay—but not if they show child abuse. And no bared breasts—unless the women are breastfeeding or have uncovered their breasts in an act of protest. Bared buttocks are generally out of bounds, but they are okay if they are satirically photoshopped onto an image of a public figure. There are rules about when acts of violence may be shown, what counts as impermissible hate speech, and a variety of other topics. And every rule is subject to interpretation and to judgment about intent. Who's a "public figure"? What makes violence "gratuitous" and therefore impermissible? In the public sphere, these questions would be settled by courts and lawmakers, with the First Amendment severely limiting how restrictive the standards could be. If a government promulgated and enforced rules like Facebook's, Americans would shout "censorship," though in China they might be regarded as necessary measures to create a "harmonious society."

As a private corporation, Facebook can do more or less what it likes. And yet it is not so simple. Almost half of the United States gets news via Facebook. After Facebook was used during the 2016 presidential election to spread politically damaging "fake news"—fabricated stories that look like news—Facebook took upon itself the role of sometimes policing the truth. It says it won't generally remove material it knows to be false unless it violates some other one of its standards. But its algorithms have been tweaked so that users are less likely to see information it has identified as false.

So denials of the Holocaust stay up but are given low display priority. Even viciously false anti-Semitic diatribes are not banned, but bare buttocks are. The complexities of First Amendment law are being replayed inside the walls of social media companies but without any bedrock principles. For example, fake news will be taken down if it is likely to result in bodily harm. What then of Pizzagate, a conspiracy theory that members of the Hillary Clinton presidential campaign were involved in a child sex trafficking ring run out of Comet Ping Pong, a Washington, DC, pizza parlor? The outlandish rumor was refuted, but the owner of the pizza parlor received death threats, and a man fired a gun in the restaurant as part of his personal "investigation." Does that mean that Pizzagate should not be discussed on Facebook?

There is something severely discomfiting about all this, beyond the obvious fact that no set of decisions about what to allow will make everyone happy. The pioneers of the Internet imagined a re-creation, in cyberspace, of the public square of the early American democracy. In the words of John Perry Barlow, "We are creating a world where anyone, anywhere may express his or her beliefs, no matter how singular, without fear of being coerced into silence or conformity."[52] In practice, such expression may have stayed out of government control in the United States, but much of it under the control of a few people in Menlo Park, California—hardly a modern analog of an eighteenth-century New England town meeting. Private fora facilitate remarkable conversations but are moderated by private parties who set the rules of the conversation. Unmoderated fora are so rife with such bullying and threats that participants don't feel safe enough to exercise their nominal freedom of speech. There may be no happy medium.

Takedowns

The SESTA–FOSTA amendments spelled the end of backpage.com by making operators of Internet sites liable if their users were blatantly using the fora for sex trafficking. This limited step back from the broad guarantees of CDA Section 230 do not significantly change the basic picture: Operators of Internet discussion fora are largely immune to liability for material posted by their users. But there is one important exception, in addition to the SESTA–FOSTA carve-outs: copyrighted material. Under 17 U.S. Code Section 512, a website operator has to take down material when notified that a user has posted it in violation of the owner's copyright. To be precise, the operator of the site must take it down or be prepared to defend the posting as if the operator were the one who had posted the material. This seems fair, but with little cost to the complainant for making a complaint and potentially ruinous legal costs to the

site owner to defend the legality of material posted by a third party, Section 512 provides an effective domestic censorship tool. If you don't like what someone is saying about something you said, make a copyright infringement complaint about their use of your words (see Chapter 6, "Balance Toppled").

Information freedom on the Internet is a tricky business. Technological changes happen faster than legal changes. When a technology shift alarms the populace, legislators respond with overly broad laws. By the time challenges have worked their way through the courts, another cycle of technology changes has happened, and the slow heartbeat of lawmaking pumps out another poorly drafted statute.

The technology of radio and television has also challenged the legislative process, but in a different way. In the broadcast world, strong commercial forces are arrayed in support of speech-restricting laws that have long since outgrown the technology that gave birth to them. We now turn to those changes in the radio world.

Endnotes

1 This narrative is based on the documentary film *I Am Jane Doe*, https://www.iamjanedoefilm.com.

2 National Center for Missing and Exploited Children, "NCMEC Data," https://www.missingkids.org/ourwork/ncmecdata, accessed May 1, 2020.

3 Jack Bouboushian, "Backpage.com Wins Injunction Against Sheriff," Courthouse News Service, December 1, 2015, https://www.courthousenews.com/backpage-com-wins-injunction-against-sheriff/.

4 Tessa Weinberg, "Backpage.Com Lawsuit Against Cook County Sheriff Dismissed," *Chicago Tribune*, June 1, 2018, https://www.chicagotribune.com/news/breaking/ct-met-backpage-lawsuit-against-sheriff-dismissed-20180601-story.html.

5 Ann Wagner, "H.R.1865 - Allow States and Victims to Fight Online Sex Trafficking Act of 2017," 115th Congress (2017–2018), April 11, 2018, https://www.congress.gov/bill/115th-congress/house-bill/1865/text.

6 Lux Alptraum, "Congress Is Forcing Sex Workers to Revert Back to a More Dangerous, Pre-Internet Era," *The Verge*, May 1, 2018, https://www.theverge.com/2018/5/1/17306486/sex-work-online-fosta-backpage-communications-decency-act; Samantha Cole, "Pimps Are Preying on Sex Workers Pushed Off the Web Because of FOSTA-SESTA," *Vice*, April 30, 2018, https://www.vice.com/en_us/article/bjpqvz/fosta-sesta-sex-work-and-trafficking.

7 John Perry Barlow, "The Economy of Ideas," *Wired*, March 1, 1994, https://www.wired.com/1994/03/economy-ideas/.

8 Tristan Lejeune, "Trump Says He'll Take a 'Strong Look' at Libel Laws in Response to Book," TheHill, January 10, 2018, https://thehill.com/homenews/

administration/368309-trump-says-hell-take-a-strong-look-at-libel-laws-in-response-to-book.

9 *New York Times Co. v. Sullivan*, 376 U.S. 254 (1964), https://supreme.justia.com/cases/federal/us/376/254/.

10 Anthony Lewis, *Make No Law: The Sullivan Case and the First Amendment* (Vintage Books, 1992).

11 Anthony Lewis, *Freedom for the Thought That We Hate: A Biography of the First Amendment* (Basic Books, 2007).

12 *Cubby, Inc. v. CompuServe Inc.*, 776 F. Supp. 135 (S.D.N.Y. 1991), https://law.justia.com/cases/federal/district-courts/FSupp/776/135/2340509/.

13 "The Law: Chilled Prodigy," *Reason.com*, August 1, 1995, https://reason.com/1995/08/01/chilled-prodigy/.

14 *Stratton Oakmont, Inc. v. Prodigy Services Co.*, accessed May 1, 2020, https://h2o.law.harvard.edu/cases/4540.

15 *Miller v. California*, 413 U.S. 15 (1973), https://supreme.justia.com/cases/federal/us/413/15/.

16 David Loundy, "Whose Standards? Whose Community?" Chicago Daily Law Bulletin," August 1, 1994, http://www.loundy.com/CDLB/AABBS.html.

17 *United States v. Robert Alan Thomas (94-6648) and Carleen Thomas (94-6649)*, 74 F.3d 701, https://www.courtlistener.com/opinion/711150/united-states-v-robert-alan-thomas-94-6648-and-carleen-thomas-94-6649/.

18 Philip Elmer-Dewitt, "Online Erotica: On a Screen Near You," *Time*, July 3, 1995, http://content.time.com/time/magazine/article/0,9171,983116,00.html.

19 John Perry Barlow, "A Declaration of the Independence of Cyberspace," Electronic Frontier Foundation, February 8, 1996, https://www.eff.org/cyberspace-independence.

20 "*ACLU v. Reno*," February 15, 1996, https://www.epic.org/free_speech/censorship/lawsuit/TRO.html.

21 "Dirty Business at CMU," *Internet World*, October 1995, https://www.tnl.net/who/bibliography/dirty-business-cmu/.

22 "47 U.S. Code § 230 - Protection for Private Blocking and Screening of Offensive Material," Legal Information Institute, accessed May 1, 2020, https://www.law.cornell.edu/uscode/text/47/230.

23 *Fair Housing Council of San Fernando Valley v. Roommate.com*, 489 F.3d 921 (9th Cir. 2007), Nos. 04-56916, 04-57173. https://caselaw.findlaw.com/us-9th-circuit/1466388.html.

24 *Zeran v. America Online, Inc.*, 958 F. Supp. 1124 (E.D. Va. 1997), https://law.justia.com/cases/federal/district-courts/FSupp/958/1124/1881560/.

25 *Zeran v. Diamond Broadcasting, Inc.*, 203 F.3d 714 (2000), https://www.quimbee.com/cases/zeran-v-diamond-broadcasting-inc

26 *Doe No. 1 v. Backpage.com, LLC*, No. 15-1724 (1st Cir. 2016), https://law.justia.com/cases/federal/appellate-courts/ca1/15-1724/15-1724-2016-03-14.html

27 *Blumenthal v. Drudge*, 992 F. Supp. 44 (D.D.C. 1998), https://www.lexisnexis.com/community/casebrief/p/casebrief-blumenthal-v-drudge.

28 *Jane Doe v. America Online, Inc.*, 783 So.2d 1010 (Fla. S.C. 2001), https://www.eff.org/issues/cda230/cases/jane-doe-v-america-online-inc.

29 *Jane Doe No. 1 v. Backpage.Com, LLC*, 817 F.3d 12 (1st Cir. 2016), https://casetext.com/case/doe-v-backpagecom-llc-1.

30 Michael G. Oxley, "H.R.3783 - Child Online Protection Act," 105th Congress (1997–1998), October 8, 1998, https://www.congress.gov/bill/105th-congress/house-bill/3783.

31 *American Civil Liberties Union v. Gonzales*, 237 F.R.D. 120 (2006), https://cite.case.law/frd/237/120/.

32 Greg Miller and Davan Maharaj, "N. Hollywood Man Charged in 1st Cyber-Stalking Case," *Los Angeles Times*, January 22, 1999, https://www.latimes.com/archives/la-xpm-1999-jan-22-mn-523-story.html.

33 The Associated Press, "Computer Stalking Case a First for California," *The New York Times*, January 25, 1999, https://www.nytimes.com/1999/01/25/us/national-news-briefs-computer-stalking-case-a-first-for-california.html; Valerie Alvord, "Cyberstalkers Must Beware of the E-Law," *USA Today*, November 8, 1999.

34 *Brandenburg v. Ohio*, 395 U.S. 444 (1969), https://supreme.justia.com/cases/federal/us/395/444/.

35 *Whitney v. California*, 274 U.S. 357 (1927), https://supreme.justia.com/cases/federal/us/274/357/.

36 Planned Parenthood of the Columbia/Willamette, Inc. v. American Coalition of Life Activists, 290 F.3d 1058 (9th Cir. 2002), https://casetext.com/case/planned-parenthood-v-amer-coalition-of-life; Rene Sanchez, "Abortion Foes' Internet Site on Trial," *The Washington Post*, January 15, 1999.

37 "H.R.3402 - Violence Against Women and Department of Justice Reauthorization Act of 2005," 109th Congress (2005–2006), January 5, 2006, https://www.congress.gov/bill/109th-congress/house-bill/3402.

38 Communications Act of 1934 (1934), https://transition.fcc.gov/Reports/1934new.pdf.

39 "Free Anonymous Email and Private Email, Private Label Email," TheAnonymous Email.Com, October 5, 2007, https://web.archive.org/web/20071005080638/http://www.theanonymousemail.com/.

40 Alexander Meiklejohn, "Meaning of the First Amendment," § Senate Judiciary Committee on Constitutional Rights (1955).

41 Palko v. Connecticut, 302 U.S. 319 (1937). https://supreme.justia.com/cases/federal/us/302/319/.

42 "OpenNet Initiative," accessed May 3, 2020, https://opennet.net.

43 Associated Press, "Google Joins Lobby Against Censors," *Daily News*, June 24, 2007, https://www.dailynews.com/2007/06/24/google-joins-lobby-against-censors/.

44 *Dow Jones & Company Inc. v Gutnick*, December 10, 2002, https://en.wikisource.org/wiki/Dow_Jones_%26_Company_Inc._v_Gutnick.

45 "Dow Jones Settles Defamation Suit," *The Wall Street Journal*, November 12, 2004, https://www.wsj.com/articles/SB110029740775772893.

46 Felicity Barringer, "THE MEDIA BUSINESS; Internet Makes Dow Jones Open to Suit in Australia," *The New York Times*, December 11, 2002, https://www.nytimes.com/2002/12/11/business/the-media-business-internet-makes-dow-jones-open-to-suit-in-australia.html.

47 *Yahoo! Inc., a Delaware Corporation, Plaintiff-appellee, v. La Ligue Contre Le Racisme et L'antisemitisme, a French Association; L'union Des Etudiants Juifs De France, a French Association, Defendants-appellants*, 433 F.3d 1199 (9th Cir. 2006), https://law.justia.com/cases/federal/appellate-courts/F3/433/1199/546158/.

48 BBC News, "EU court backs 'right to be forgotten' in Google case," 13 May 2014, https://www.bbc.com/news/world-europe-27388289.

49 John Naughton, "The Germans Get Their Flickrs in a Twist over 'Censorship,'" *The Observer*, June 17, 2007, https://www.theguardian.com/media/2007/jun/17/newmedia.business.

50 https://research.fb.com/blog/2014/10/facebook-s-top-open-data-problems/.

51 "Community Standards," accessed May 3, 2020, https://www.facebook.com/communitystandards/introduction.

52 John Perry Barlow, "A Declaration of the Independence of Cyberspace," Electronic Frontier Foundation, February 8, 1996, https://www.eff.org/cyberspace-independence.

This page intentionally left blank

CHAPTER 8

Bits in the Air

Old Metaphors, New Technologies, and Free Speech

Censoring the Candidate

On October 8, 2016, barely six weeks before the U.S. presidential election, *The Washington Post* dropped a bombshell.[1] The *Post* had received a video recording in which then candidate Donald Trump used coarse language in boasting of sexual aggression toward women. As the news story reported, Mr. Trump said,

> "I did try and f--- her. She was married."..."And when you're a star, they let you do it....You can do anything....Grab them by the p---y."

That is the way *The Post* printed it, with hyphens replacing certain letters, but with enough context for the words to be reconstructed unambiguously. The story linked to the full, unexpurgated video, so anyone with an Internet connection could hear what Mr. Trump later characterized as his "locker room talk."

The New York Times chose to print the offensive words in full.[2] The cable television network CNN showed the full video, but its reporters used "the F word" and "the P word" where the *Post* had used hyphens. These decisions were matters of editorial judgment and could have gone either way. *Post* editor Marty Baron said, "We make our best judgments in weighing taste against clarity about what was said," while *Times* editor Carolyn Ryan explained, "we decided the vulgar words themselves were newsworthy, and omitting or describing them would have been less than forthright."

These media outlets could come to different conclusions about what to report, but no law or government regulation was relevant. The First

Amendment protected the right of newspapers and cable TV stations to report on the tape as they wished.

But not broadcast television stations. Although most Americans now view their ABC, CBS, and NBC shows via cable boxes or over the Internet, the terms "broadcast" and "over the air" still differentiate the stations that can be received by antennas picking up radio signals emanating from giant transmitting antennas. And those stations are subject to regulation by the Federal Communications Commission (FCC): No indecent or profane language during hours when children might hear it—and Mr. Trump's language probably breached that threshold. According to the FCC, indecent content portrays sexual or excretory organs or activities in a way that does not meet the Miller test for obscenity, and profane content includes "grossly offensive" language that is considered a public nuisance.

The broadcast stations bleeped out the offensive words when they showed the video. They had no choice: They might have been subject to large FCC fines or even the loss of their broadcast licenses if they had not.

Under the First Amendment, the government is generally not in the speech-restricting business. It can't force its editorial judgments on newspapers, even to increase the range of information available to readers. The Supreme Court struck down as unconstitutional a Florida law assuring political candidates a simple "right to reply" to newspaper attacks on them. Cable television stations need not have worried about FCC complaints, though some were filed; that medium is largely beyond the reach of government censors.

Nonetheless, in 2016, an agency of the federal government was keeping words off broadcast television, using rules that covered even the remarks of a presidential candidate in the middle of a political campaign. We are in an era of heightened sensitivity about programming that children may see, but Americans remain generally opposed to having the government nanny their television shows. Why does the FCC get to regulate what can be said over the airwaves?

How Broadcasting Became Regulated

The FCC gained its authority over what is said on radio and TV broadcasts when there were fewer ways to distribute information. The public airways were scarce, went the theory, and the government had to make sure they were used in the public interest. As radio and television became universally accessible, a second rationale emerged for government regulation of broadcast speech. Because the broadcast media have "a uniquely pervasive presence in the lives of all Americans," as the Supreme Court put it in 1978, the government had a special interest in protecting a defenseless public from objectionable radio and television content.

The explosion in communications technologies has confused both rationales. In the digital era, there are far more ways for bits to reach the consumer, so broadcast radio and television are hardly unique in their pervasiveness. With minimal technology, anyone can sit at home or in Starbucks and choose from among billions of web pages and tens of millions of blogs. Shock jock Howard Stern left broadcast radio for satellite radio, where the FCC has no authority to regulate what he says. Almost 90% of American television viewers get their TV signal through similarly unregulated cable or satellite rather than through broadcasts from rooftop antennas.[3] RSS feeds supply up-to-date information to millions of on-the-go cell phone users. Radio stations and television channels are today neither scarce nor uniquely pervasive.

In the digital era, there are far more ways for bits to reach the consumer, so broadcast radio and television are hardly unique in their pervasiveness.

For the government to protect children from all offensive information arriving through any communication medium, its authority would have to be expanded greatly and updated continuously. Though some proposals have been made, Congress has passed no law extending FCC indecency regulations for broadcast media to satellite and cable television. What is shown on cable and satellite TV is limited by what viewers and advertisers will accept, but not by what any government authority might dictate.

The explosion in communications raises another possibility, however. If almost anyone can now send information that many people can receive, perhaps the government's interest in restricting transmissions should be less than what it once was, not greater. In the absence of scarcity, perhaps the government should have no more authority over what gets said on radio and TV than it does over what gets printed in newspapers. In that case, rather than expand the FCC's censorship authority, Congress should eliminate it entirely, just as the Supreme Court ended Florida's regulation of newspaper content.

Parties who already have spots on the radio dial and the TV channel lineup respond that the spectrum—the public airwaves—remains a limited resource, requiring government protection. No one is making any more radio spectrum, goes the theory, and it needs to be used in the public interest.

But look around you. There are still only a few stations on the AM and FM radio dials. But thousands, maybe tens of thousands, of radio communications are passing through the air around you. Most Americans walk around with two-way radios in their pockets—devices we call cell phones. To be precise, we walk around with cell phones in our *hands*, since many of us would rather risk walking into lampposts than delay reading our text messages by even a few seconds. If you are listening to music on a Bluetooth headset and browsing

the Web over a Wi-Fi connection, that's two more radio connections you are using. Radios and television sets could be much smarter than they now are and could make better use of the airwaves, just as cell phones do.

Engineering developments have vitiated the government's override of the First Amendment on radio and television. The Constitution demands, under these changed circumstances, that the government stop its verbal policing. Indeed, when the U.S. Supreme Court threw out the fines the FCC was imposing when celebrities threw a "fleeting expletive" in their live-broadcast remarks, it restricted the scope of its decision but hinted that it might soon be time for a look at the whole broadcast censorship question.

As a scientific argument, the claim that the spectrum is necessarily scarce is now very weak. Yet that view is still forcefully advanced by the very industry that is being regulated. The incumbent license holders—existing broadcast stations and networks—have an incentive to protect their "turf" in the spectrum against any risk, real or imagined, that their signals might be corrupted. By deterring technological innovation, incumbents can limit competition and avoid capital investments. These oddly intertwined strands—the government's interest in artificial scarcity to justify speech regulation and the incumbents' interest in artificial scarcity to limit competition and costs—today impair both cultural and technological creativity, to the detriment of society.

To understand the confluent forces that have created the world of today's radio and television censorship, we have to go back to the inventors of the technology.

From Wireless Telegraph to Wireless Chaos

Red, orange, yellow, green, blue—the colors of the rainbow—are all different and yet are all the same. Any child with a crayon box knows that they are all different. They are the same because they are all the result of electromagnetic radiation striking our eyes. The radiation travels in waves that oscillate very quickly. The only physical difference between red and blue is that red waves oscillate around 450,000,000,000,000 times per second and blue waves about 50% faster.

Because the spectrum of visible light is continuous, an infinity of colors exists between red and blue. Mixing light of different frequencies creates other colors—for example, half blue waves and half red creates a shade of pink known as magenta, which does not appear in the rainbow.

In the 1860s, British physicist James Clerk Maxwell realized that light consists of electromagnetic waves. His equations predicted that there might be waves of other frequencies—waves that people couldn't sense. Indeed, such waves have been passing right through us from the beginning of time. They

shower down invisibly from the sun and the stars, and they radiate when lightning strikes. No one suspected they existed until Maxwell's equations said they should. Indeed, there should be a whole spectrum of invisible waves of different frequencies, all traveling at the same great speed as visible light.

In 1887, the radio era began with a demonstration by Henrich Hertz. He bent a wire into a circle, leaving a small gap between the two ends. When he set off a big electric spark a few feet away, a tiny spark jumped the gap of the almost-completely-circular wire. The big spark had set off a shower of unseen electromagnetic waves, which had traveled through space and caused electric current to flow in the other wire. The tiny spark was the current completing the circuit. Hertz had created the first antenna and had revealed the radio waves that struck it. The unit of frequency is named in his honor: One cycle per second is 1 hertz, or Hz for short. A kHz (kilohertz) is a thousand cycles per second, and a MHz (megahertz) is a million cycles per second. These are the units on the AM and FM radio dials.

Guglielmo Marconi was neither a mathematician nor a scientist. He was an inventive tinkerer. Only 13 years old at the time of Hertz's experiment, Marconi spent the next decade developing, by trial and error, better ways of creating bursts of radio waves and antennas for detecting them over greater distances.

In 1901, Marconi stood in Newfoundland and received a single Morse code letter transmitted from England. On the strength of this success, the Marconi Wireless Telegraph Company was soon enabling ships to communicate with each other and with the shore. When the *Titanic* left on its fateful voyage in 1912, it was equipped with Marconi equipment. The main job of the ship's radio operators was to relay personal messages to and from passengers, but they also received at least 20 warnings from other ships about the icebergs that lay ahead.[4]

The words "Wireless Telegraph" in the name of Marconi's company suggest the greatest limitation of early radio. The technology was conceived as a device for point-to-point communication. Radio solved the worst problem of telegraphy. No calamity, sabotage, or war could stop wireless transmissions by severing cables. But there was a compensating disadvantage: Anyone could listen in. The enormous power of broadcasting to reach thousands of people at once was at first seen as a liability. Who would pay to send a message to another person when anyone could hear it?

As wireless telegraphy became popular, another problem emerged—one that has shaped the development of radio and television ever since. If several people were transmitting simultaneously in the same geographic area, their signals couldn't be kept apart. The *Titanic* disaster demonstrated the confusion that could result. The morning after the ship hit the iceberg, American newspapers reported excitedly that all passengers had been saved and the

ship was being towed to shore. The mistake resulted from a radio operator's garbled merger of two unrelated segments of Morse code. One ship inquired if "all Titanic passengers safe?" A completely different ship reported that it was "300 miles west of the *Titanic* and towing an oil tank to Halifax."[5] All the ships had radios and radio operators. But there were no rules or conventions about whether, how, or when to use them.

Listeners to Marconi's early transmitters were easily confused because they had no way to "tune in" a particular communication. For all of Marconi's genius in extending the range of transmission, he was using essentially Hertz's method for generating radio waves: big sparks. The sparks splattered electromagnetic energy across the radio spectrum. The energy could be stopped and started to turn it into dots and dashes, but there was nothing else to control. One radio operator's noise was like any other's. When several transmitted simultaneously, chaos resulted.

The many colors of visible light look white if all blended together. A color filter lets through some frequencies of visible light but not others. If you look at the world through a red filter, everything is a lighter or darker shade of red because only the red light comes through. What radio needed was something similar for the radio spectrum: a way to produce radio waves of a single frequency, or at least a narrow range of frequencies, and a receiver that could let through those frequencies and screen out the rest. Indeed, that technology already existed.

In 1907, Lee De Forest patented a key technology for the De Forest Radio Telephone Company—dedicated to sending voice and even music over the radio waves. When he broadcast Enrico Caruso from the Metropolitan Opera House in New York singing *Pagliacci* on January 13, 1910, the singing reached ships at sea. Amateurs huddled over receivers in New York and New Jersey. The effect was sensational. Hundreds of amateur broadcasters sprang into action over the next few years, eagerly saying whatever they wanted, and playing whatever music they could, to anyone who happened to be listening.

With no clear understanding about what frequencies to use, radio communication was a hit-or-miss affair.

But with no clear understanding about what frequencies to use, radio communication was a hit-or-miss affair. Even what *The New York Times* described as the "homeless song waves" of the Caruso broadcast clashed with another station that, "despite all entreaties," insisted on broadcasting at the identical 350 kHz frequency. Some people could "catch the ecstasy" of Caruso's voice, but others got only some annoying Morse code from the other broadcaster: "I took a beer just now."[6]

Radio Waves in Their Channels

The emerging radio industry could not grow under such conditions. Commercial interests complemented the concerns of the U.S. Navy about amateur interference with its ship communications. The *Titanic* disaster, although it owed little to the failures of radio, catalyzed government action. On May 12, 1912, William Alden Smith called for radio regulation on the floor of the U.S. Senate. "When the world weeps together over a common loss...," proclaimed the senator, "why should not the nations clear the sea of its conflicting idioms and wisely regulate this new servant of humanity?"[7]

The Radio Act of 1912[8] limited broadcasting to license holders. Radio licenses were to be "granted by the Secretary of Commerce and Labor upon application therefor." In granting a license, the secretary would stipulate the frequencies "authorized for use by the station for the prevention of interference and the hours for which the station is licensed for work." The act reserved for government use the choice frequencies between about 200 and 500 kHz, which permitted the clearest communications over long distances. Amateurs were pushed off to "short wave" frequencies above 1500 kHz, considered useless for technological reasons. The frequency 1000 kHz was reserved for distress calls, and licensed stations were required to listen to it every 15 minutes (the one provision that might have helped the *Titanic*, since the radio operators of a nearby ship had gone off-duty and missed the *Titanic's* rescue pleas). The rest of the spectrum the secretary could assign to commercial radio stations and private businesses. Emphasizing the nature of radio as "wireless telegraphy," the act made it a crime for anyone hearing a radio message to divulge it to anyone except its intended recipient.

High Frequencies

Over the years, technological improvements have made it possible to use higher and higher frequencies. Early TV was broadcast at what were then considered "very high frequencies" (VHF) because they were higher than AM radio. Technology improved again, and more stations appeared at "ultra high frequencies" (UHF). The highest frequency in commercial use today is 77 GHz—that is, 77 gigahertz, which is 77,000 MHz. In general, high-frequency signals fade with distance more than low-frequency signals, and they are therefore mainly useful for localized or urban environments. Short waves correspond to high frequencies because all radio waves travel at the same speed, which is the speed of light.

Much has changed since 1912. The uses of radio waves have become more varied, the allocation of spectrum blocks has changed, and the range of usable frequencies has grown. The current spectrum allocation picture has grown

into a dense, disorganized quilt, the product of decades of Solomonic FCC judgments (see Figure 8.1). But still, the U.S. government stipulates what parts of the spectrum can be used for what purposes. It prevents users from interfering with each other and with government communications by demanding that they broadcast at limited power and only at their assigned frequencies. As long as there weren't many radio stations, the implied promise in the act of 1912 that licenses would be granted "upon application therefor" caused no problems. With the gossip of the pesky amateurs pushed into remote radio territory, there was plenty of spectrum for commercial, military, and safety use.

FIGURE 8.1 Frequency allocation of the U.S. radio spectrum. The spectrum from 3 kHz to 300 GHz is laid out from left to right and top to bottom, with the scale 10 times denser in each successive row. For example, the large block in the second row is the AM radio dial, about 1 MHz wide. The same amount of spectrum would be about .00002 inch wide in the bottom row.[9]

Within a decade, that picture had changed dramatically. On November 2, 1920, a Detroit station broadcast the election of Warren Harding as president of the United States, relaying to its tiny radio audience the returns it was receiving by telegraph. Radio was no longer just point-to-point communication. A year later, a New York station broadcast the World Series between the Giants and the Yankees, pitch by pitch. Sports broadcasting was born with

a broadcaster drearily repeating the ball and strike information telephoned by a newspaper reporter at the ballpark.[10]

Public understanding of the possibilities grew rapidly. The first five radio stations were licensed for broadcasting in 1921. Within a year, there were 670.[11] The number of radio receivers jumped in a year from fewer than 50,000 to more than 600,000, perhaps a million.[12] Stations using the same frequency in the same city divided up the hours of the day. Radio broadcasting became a profitable business, but the growth could not go on forever.

On November 12, 1921, the New York City broadcast license of Intercity Radio Co. expired. Herbert Hoover, then the secretary of Commerce, refused to renew it, on the grounds that there was no frequency on which Intercity could broadcast in the city's airspace without interfering with government or other private stations. Intercity sued Hoover to have its license restored—and won.[13] Hoover, said the court, could choose the frequency, but he had no discretion to deny the license. As the congressional committee proposing the 1912 Radio Act had put it, the licensing system was "substantially the same as that in use for the documenting upward of 25,000 merchant vessels." The implied metaphor was that Hoover should keep track of the stations like ships in the ocean. He could tell them what shipping lanes to use, but he couldn't keep them out of the water.

The radio industry begged for order. Hoover convened a National Radio Conference in 1922 in an attempt to achieve consensus on new regulations before chaos set in. The spectrum was "a great national asset," he said, and "it becomes of primary public interest to say who is to do the broadcasting, under what circumstances, and with what type of material."[14] "The large mass of subscribers need protection as to the noises which fill their instruments," and the airwaves need "a policeman" to detect "hogs that are endangering the traffic."[15]

Hoover divided the spectrum from 550 kHz to 1350 kHz in 10 kHz bands—called "channels," consistent with the nautical metaphor—to squeeze in more stations. Empty "guard bands" were left on each side of allocated bands because broadcast signals inevitably spread out, reducing the amount of usable spectrum. Persuasion and voluntary compliance helped Hoover limit interference. As stations became established, they found it advantageous to comply with Hoover's prescriptions. Startups had a harder time breaking in. Hoover convinced representatives of a religious group that to warn of the coming apocalypse, they should buy time on existing stations rather than build a station of their own. After all, their money would go farther that way: In six months, after the world had ended, they would have no further use for a transmitter.[16] Hoover's effectiveness made Congress complacent; the system was working well enough without laws.

But as the slicing got finer, the troubles got worse. WLW and WMH in Cincinnati broadcast on the same frequency in 1924 until Hoover brokered a deal for three stations to share two frequencies in rotating time slots. Finally, the system broke down.[17] In 1925, Zenith Radio Corporation was granted a license to use 930 kHz in Chicago, but only on Thursday nights, only from 10 p.m. to midnight, and only if a Denver station didn't wish to broadcast then. Without permission, Zenith started broadcasting at 910 kHz, a frequency that was more open because it had been ceded by treaty to Canada. Hoover fined Zenith; Zenith challenged Hoover's authority to regulate frequencies and won in court.[18] The secretary then got even worse news from the U.S. attorney general: The 1912 Act, drafted before broadcasting was even a concept, was so ambiguous that it probably gave Hoover no authority to regulate anything about broadcast radio—frequency, power, or time of day.

Hoover threw up his hands. Anyone could start a station and choose a frequency—there were 600 applications pending—but in doing so, they were "proceeding entirely at their own risk."[19] The result was the "chaos in the air" that Hoover had predicted. It was worse than before the 1912 act because many more transmitters existed, and they were much more powerful. Stations popped up, jumped all over the frequency spectrum in search of open air, and turned up their transmission power to the maximum to drown out competing signals. Radio became virtually useless, especially in cities. Congress finally was forced to act.

The Spectrum Nationalized

The premises of the Radio Act of 1927[20] are still in force. The spectrum has been treated as a scarce national resource ever since, managed by the government. The purpose of the act was to

> maintain the control of the United States over all the channels of... radio transmission; and to provide for the use of such channels, but not the ownership thereof, by individuals, firms, or corporations, for limited periods of time, under licenses granted by Federal authority.

The public could use the spectrum, under conditions stipulated by the government, but could not own it. A new authority, the Federal Radio Commission (FRC), made licensing decisions. The public had a qualified expectation that license requests would be granted:

> The licensing authority, if public convenience, interest, or necessity will be served thereby,...shall grant to any applicant therefor a station license.

The act recognized that demand for licenses could exceed the supply of spectrum. In case of competition among applicants,

> the licensing authority shall make such a distribution of licenses, bands of frequency..., periods of time for operation, and of power among the different States and communities as to give fair, efficient, and equitable radio service to each.[21]

The language about "public convenience, interest, or necessity" echoes Hoover's 1922 speech about a "national asset" and the "public interest." It is also no accident that this law was drafted as the Teapot Dome Scandal was cresting. Oil reserves on federal land in Wyoming had been leased to Sinclair Oil in 1923, with the assistance of bribes paid to the secretary of the Interior. It took several years for congressional investigations and federal court cases to expose the wrongdoing; the secretary was eventually imprisoned. By early 1927, the fair use of national resources in the public interest was a major concern in the United States.

THE "RADIO COMMISSION" GROWS

In 1934, the FRC's name was changed to the Federal Communications Commission—the FCC—when telephone and telegraph regulation came under the commission's oversight. When a separate chunk of radio spectrum was allocated for television, the FCC assumed authority over video broadcasts as well.

With the passage of the act of 1927, the radio spectrum became federal land. International treaties followed, to limit interference near national borders. But within the United States, just as Hoover had asked it to five years earlier, the federal government took control over who would be allowed to broadcast, which radio waves they could use—and even what they could say.

Goat Glands and the First Amendment

The Radio Act of 1927 stipulated that the FRC could not abridge free speech over the radio:

> Nothing in this Act shall be understood or construed to give the licensing authority the power of censorship..., and no regulation or condition...shall interfere with the right of free speech by means of radio communications.[22]

Inevitably, a case would arise exposing the implicit conflict: On the one hand, the commission had to use a public interest standard when granting

and renewing licenses. On the other, it had to avoid censorship. The pivotal case was over the license for KFKB radio, the station of the Kansas goat-gland doctor John Romulus Brinkley (see Figure 8.2). The wrath brought down on CBS in 2004 for showing a flash of Janet Jackson's breast[23]—and the $325,000 fine[24] assessed on TV station WBDJ of Roanoke, Virginia, in 2015 for accidentally showing a graphic image for three seconds—descend from the FRC's action against this classic American charlatan.

PREACHES FUNDAMENTALISM PRACTICES GOAT-GLAND SCIENCE

How a Famous Surgeon Combines Old-Time Religion and New-Fangled Operations on a Strange Medico-Gospel Farm

SCIENTIST AND PROTEGE. Dr. John R. Brinkley, Eminent Surgeon and Ardent Fundamentalist, with "Billy," Now 5 Years Old, the First "Goat-Gland Baby."

WIFE AND SWISS KIDS. Mrs. Minnie Brinkley Holding Two Fine Young Specimens of Toggenberg Goats, Used in Her Husband's Delicate Organ-Graftings.

MEET the most unusual scientist - fundamentalist in the whole world, Dr. John R. Brinkley, of Milford, Kan., who saves souls with the word of God and repairs human bodies with glands from lively goats.

A surgeon of distinction, whose services have been recognized here and abroad, the doctor is today a devout and literal believer in the teachings of the Bible. This strain of faith strongly permeates Brinkley Institute, a community centre ("teaching and promoting Bible, health, citizenship, recreation and music") at Milford. But it does not detract from the fact that Dr. Brinkley began his career, in 1911, as the first goat-gland transplanting expert in the United States, some say in the world.

Today, chief surgeon and owner of the Kansas General Research Hospital, at Milford, ten miles from the stone marking the geographical centre of this country, Dr. Brinkley gained scientific distinction during the World War. At the onset of hostilities between America and the central empires, as an A. E. F. veterinary surgeon, he continued to make his daring gland transplantation experiments to renew functional powers that in men and women had been weakened or worn out.

He discovered, among other things, that certain quadrupeds' glands wouldn't do for human beings—those of the bull were "too strenuous," those of apes were too short-lived and too prone to physical decay. Finally, he came to the conclusion that goats, preferably the Toggenberg species of Switzerland, approached the ideal, since they were practically diseaseless and their glands strongly resembled those of *homo sapiens*.

With the war at an end, he found it possible to devote the major part of his time to the importation of the frisky Toggenbergs and the conversion of their glands into human agencies. More than 4,000 individuals of both sexes, many of them prominent people, were treated by him—and successfully.

Brinkley's signal triumphs in this newest and most delicate form of medical science were not received with a unanimous shout of approbation. Old-time surgeons got quite hot under the collar, and he was even charged with practicing under a license obtained by fraud. Dropped from the list of practitioners in Missouri, Connecticut and California, he appealed to the courts, receiving a complete vindication.

Rightfully flushed with his victory on native soil, Dr. Brinkley went abroad for study and research. The Royal University of Pavia, Italy, which had awarded degrees to Napoleon, Michael Angelo and various Italian rulers, licensed him to practice. He was the only American ever accorded such an honor. In London, later, the British Medical Association gave him similar laurels after he had passed its rigid examination.

Back in the States, Brinkley found a storm brewing in his home state, Tennessee. The great Scopes trial was in process of incubation. The returning scientist, himself a Tennessean, saw and sympathized with the mountaineers' efforts to keep their beloved religion untouched by scientific fingers. Perhaps some hidden mental echo of his boyhood in that region awoke within him the latest sparks of devotional fervor. Perhaps the white-hot oratory of the late William Jennings Bryan struck a response in his heart.

Whatever the cause, Brinkley opened his hands to fundamentalism. He converted his spacious grounds at Milford into a glorified camp-meeting ground. At the same time he reserved the hospital for continued gland transplantations. Thus soul-saving and bodily repairing went hand in hand under the keen eye of one of the best scientists and firmest "believers" on the globe.

A first-hand pen picture of this extraordinary man is supplied by a friend of long standing. "Probably," says this individual, "there never was so odd a mixture of talents and tendencies as John. He is a thirty-second degree Mason. In 1923 he erected a broadcasting station at a cost of $50,000 to put his medical lectures on the air and also to establish the Manhattan Agricultural College as 'the first college of the air.' When he first went to Italy, he sold this station at 'a total loss.' Back home, having gotten religion, he repurchased the station to broadcast fundamentalist doctrines and denounce evolutionists and 'other workers of the devil'."

Mrs. Minnie Brinkley, wife of the religious scientist and press representative of his institute, has embodied in letters to friends in the East the working pattern of the doctor's current philosophy. "You and your evolution theory are all wrong," she wrote to one New York acquaintance of the family. "You can't deny the Supreme Ruler of the universe and get away with it—with us. Remember, we are Tennessee people, not of the mountains, but the same heart, and we are a success in this world. . . . We are afraid to know anyone who blasphemes as you did in your letter. You can't refuse our God and get away with it. We are not the kind of Christians who just pray; we work."

Further insight into the workings of Dr. Brinkley's system of religious training, as exemplified in his choice of an assistant, is afforded in another of his wife's communications. "Dr. Charles Draper, of Kansas City, Fort Worth, St. Louis and Chicago, where he had done big work in representing his Almighty God, in a positive defense of the truth as seen by the miracles of Jesus Christ, was told of this project at Milford.

"Dr. Brinkley hired Dr. Draper at Dr. Draper's own price and paid for a picture machine and reels and built a picture house and bandstand for the local band and built benches in the park and a platform, and put a piano on it. Every Sunday and Wednesday nights moral and religious pictures are shown to remind people of the Greatest Man ever known."

Tennis and basketball courts, children's playgrounds, a merry-go-round and other features are a part of Dr. Draper's vision of what should be at Milford. To build up this project subscriptions are being volunteered by neighbors. Some farmers donate about $5 apiece. Dressmakers give the price they charge for one gown.

Such is the remarkable labor which Dr. Brinkley has set himself, as an international scientist, and one who, later, "following the gleam," has determined to stick by his forefathers' creeds—glands or no glands.

SUNDAY SCHOOL. Children on the Playground at Milford, Kansas, Where Rigid Fundamentalism Is Taught, Despite the Adjacency of Dr. Brinkley's Gland Hospital Activities.

PREACHES GOSPEL. Dr. Charles Draper, Pastor of the Brinkley Institute, Who Doesn't Believe in Scientific Evolution.

FIGURE 8.2 A planted newspaper article about "Dr." Brinkley's goat-gland clinic. The doctor himself is shown at the left, holding the first baby—named "Billy," of course—conceived after a goat-gland transplant. (New York Evening Journal, September 11, 1926. Microfilm courtesy of the Library of Congress.)

Brinkley, born in 1885, became a "doctor" licensed to practice in Kansas by buying a degree from the Eclectic Medical University in Kansas City. He worked briefly as a medic for Swift & Co., the meatpackers. In 1917, he set up

his medical practice in Milford, a tiny town about 70 miles west of Topeka. One day, a man came for advice about his failing virility, describing himself as a "flat tire." Drawing on his memory of goat behavior from his days at the slaughterhouse, Brinkley said, "You wouldn't have any trouble if you had a pair of those buck glands in you." "Well, why don't you put 'em in?" the patient asked. Brinkley did the transplant in a back room, and a business was born. Soon he was performing 50 transplants a month, at $750 per surgery. In time, he discovered that promising sexual performance was even more lucrative than promising fertility.[25]

As a young man, Brinkley had worked at a telegraph office, so he knew the promise of communication technology. In 1923, he opened Kansas's first radio station, KFKB—"Kansas First, Kansas Best" radio, or sometimes "Kansas Folks Know Best." The station broadcast a mixture of country music, fundamentalist preaching, and medical advice from Dr. Brinkley himself. Listeners sent in their complaints, and the advice was almost always to buy some of Dr. Brinkley's mail-order patent medicines. "Here's one from Tillie," went a typical segment:

> She says she had an operation, had some trouble 10 years ago. I think the operation was unnecessary, and it isn't very good sense to have an ovary removed with the expectation of motherhood resulting therefrom. My advice to you is to use Women's Tonic No. 50, 67, and 61. This combination will do for you what you desire if any combination will, after three months persistent use.[26]

KFKB had a massively powerful transmitter, heard halfway across the Atlantic. In a national poll, it was the most popular station in America—with four times as many votes as the runner-up.[27] Brinkley was receiving 3,000 letters a day and was a sensation throughout the Plains States. On a good day, 500 people might show up in Milford. But the American Medical Association—prompted by a competing local radio station—objected to his quackery. The FRC concluded that "public interest, convenience, or necessity" would not be served by renewing the license. Brinkley objected that the cancellation was nothing less than censorship.

An appeals court sided with the FRC in a landmark decision. Censorship, the court explained, was prior restraint, which was not at issue in Brinkley's case. The FRC had "merely exercised its undoubted right to take note of appellant's past conduct." An arguable point—as Albert Gallatin said more than 200 years ago about prior restraint of the press, it was "preposterous to say, that to punish a certain act was not an abridgment of the liberty of doing that act."[28]

The court used the public land metaphor in justifying the FRC's action:

> Because the number of available broadcasting frequencies is limited, the commission is necessarily called upon to consider the character and quality of the service to be rendered.... Obviously, there is no room in the broadcast band for every business or school of thought.

"Necessarily" and "obviously." It is always wise to scrutinize arguments that proclaim loudly how self-evident they are.

Judge Felix Frankfurter, in an opinion on a different case in 1943, restated the principle in a form that has often been quoted:

> The plight into which radio fell prior to 1927 was attributable to certain basic facts about radio as a means of communication—its facilities are limited; they are not available to all who may wish to use them; the radio spectrum simply is not large enough to accommodate everybody. There is a fixed natural limitation upon the number of stations that can operate without interfering with one another.[29]

These were facts of the technology of the time. They were true, but they were contingent truths of engineering. They were never universal laws of physics and are no longer limitations of technology. Thanks to engineering innovations, there is practically no significant "natural limitation" on the number of broadcast stations. Arguments about inevitable scarcity can no longer justify U.S. government denials of the use of the airwaves.

The vast regulatory infrastructure, built to rationalize use of the spectrum by much more limited radio technology, has adjusted slowly—as it almost inevitably must: Bureaucracies don't move as quickly as technological innovators. The FCC tries to anticipate resource needs centrally and far in advance. But technology can cause abrupt changes in supply, and market forces can cause abrupt changes in demand. Central planning works no better for the FCC than it did for the Soviet Union.

Moreover, plenty of stakeholders in old technology are happy to see the rules remain unchanged. Like tenants enjoying leases on public land, incumbent radio license holders have no reason to encourage competing uses of the assets they control. The more money that is at stake, the greater the leverage of the profitable ventures. Radio licenses had value almost from the beginning, and as scarcity increased, so did price. By 1925, a Chicago license was sold for $50,000. As advertising expanded and stations bonded into networks, transactions reached seven figures. After the 1927 act, disputes between stations had to be settled by litigation, trips to Washington, and pressure by

friendly congressional representatives—all more feasible for stations with deep pockets than for the rest. At first, there were many university stations, but the FRC squeezed them as the value of the airwaves went up. As nonprofits, university stations could not hold their ground. Eventually, most educational stations sold out to commercial broadcasters. *De facto*, as one historian put it, "while talking in terms of the public interest,...the commission actually chose to further the ends of the commercial broadcasters."[30]

The Path to Spectrum Deregulation

Today, we are all radio broadcasters and radio receivers. The smartphone in your pocket is using radio waves to post your Instagram photos, to send your search query to Google, to text message your brother, and of course to transmit your voice, wirelessly, should you make an old-fashioned phone call. But you use radio waves for countless other devices that don't get far from your body. Your Bluetooth headphones use very short-range radio signals for encoded musical tunes coming from your smartphone or iPod. The keyless lock fob in your pocket uses another signal to unlock your car rather than the one parked next to it. One of us wears an insulin pump that chats wirelessly with his blood glucose sensor to simulate, very roughly, a human biochemical regulatory cycle that has failed him.

Each of these signals uses a bit of the spectrum. They obey the same basic physical laws as WBZ's radio broadcasts in Boston, which have been going on continuously since WBZ became the first eastern U.S. commercial station in 1921. But the new radio broadcasts are different in two critical respects. There are billions of them going on every day. And whereas WBZ's broadcast power is 50,000 watts, a car keyfob's power is less than .0002 watt.

If the government still had to license every radio transmitter—as Congress authorized in the aftermath of the radio chaos of the 1920s—neither radio keys nor any of hundreds of other innovative uses of low-power radio could have come about. The law and the bureaucracy would have snuffed this part of the digital explosion.

Another development also lay behind the wireless explosion. Technology had to change so that the available spectrum could be used more efficiently. Digitalization and miniaturization changed the communications world. The story of cell phones and wireless Internet and many conveniences as yet unimagined is a knot of politics, technology, and law. You can't understand the knot without understanding the strands, but in the future, the strands need not remain tied up in the same way as they are today.

From a Few Bullhorns to Billions of Whispers

Forty years ago, there were no cell phones. A handful of business executives had mobile phones, but the devices were bulky and costly. Miniaturization helped change the mobile phone from the perk of a few corporate bigwigs into the birthright of every American teenager. But the main advance was in *spectrum allocation*—in rethinking the way the radio spectrum was used.

In the era of big, clunky mobile phones, the radio phone company had a big antenna and secured from the FCC the right to use a few frequencies in an urban area. The executive's phone was a little radio station, which broadcast its call. The mobile phone had to be powerful enough to reach the company's antenna, wherever in the city the phone might be located. The number of simultaneous calls was limited to the number of frequencies allocated to the company. The technology was the same as broadcast radio stations used, except that the mobile phone radios were two-way. The scarcity of spectrum, still cited today in limiting the number of broadcast channels, then limited the number of mobile phones. Hoover understood this way back in 1922. "Obviously," he said, "if 10,000,000 telephone subscribers are crying through the air for their mates...the ether will be filled with frantic chaos, with no communication of any kind possible."[31]

Cell phones, Wi-Fi, Bluetooth—these technologies exploit Moore's Law. Radio transmitters and receivers have become faster, cheaper, and smaller. Take cell phones, for example. Because cell phone towers are only a mile or so apart, cell phones need only be powerful enough to send their signals less than a mile. Once received by an antenna, the signal is sent on to the cell phone company by "wireline"—that is, by copper or fiber-optic cables on poles or underground. There need be only enough radio spectrum to handle the calls within the "cell" surrounding a tower, since the same frequencies can be used simultaneously to handle calls in other cells. A lot of fancy dancing has to be done to prevent a call from being dropped as an active phone is carried from cell to cell, but computers, including the little computers inside cell phones, are smart and fast enough to keep up with such rearrangements.

Cell phone technology illustrates an important change in the use of radio spectrum. Most radio communications are now over short distances. They are transmissions between cell phone towers and cell phones. Between wireless routers and the computers of office workers and coffee drinkers. Between cordless telephone handsets and their bases. Between highway toll booths and the transponders mounted on commuters' windshields. Between electronic keys and the cars they unlock. Between video gamers and their games. Between cell phones and Bluetooth earbuds.

Even "satellite radio" transmissions sometimes go from a nearby antenna to a customer's receiver, not directly from a satellite orbiting in outer space. In

urban areas, so many buildings lie between the receiver and the satellite that the radio companies have installed "repeaters"—antennas connected to each other by wireline. When you listen to SiriusXM in your car driving around a city, the signal is probably coming to you from an antenna a few blocks away.[32]

5G cellular technology is another development of the same kind. 5G achieves higher data rates by using a different, higher-frequency part of the spectrum than its predecessor 4G technology. A price must be paid for using this high-data-rate, high-frequency radio: The signals weaken rapidly with distance from the transmitter. So 5G cells are smaller than 4G cells, and many more cellular transmitters must be installed. That is why 5G is most widely deployed in urban areas.

Details aside, the radio spectrum is no longer mainly for long-range signaling. Spectrum policies were set when the major use of radio was for ship-to-shore transmissions, SOS signaling from great distances, and broadcasting over huge geographic areas. As the nation has become wired, most radio signals travel only a few feet or a few hundred feet. Under these changed conditions, the old rules for spectrum management don't make sense.

Can We Just Divide the Property Differently?

Even parts of the spectrum that are "allocated" to licensees may be drastically underused in practice. A Federal Communications Committee Report puts it this way: "The shortage of spectrum is often a *spectrum access problem*. That is, the spectrum resource is available, but its use is compartmented by traditional policies based on traditional technologies."[33] The committee came to this conclusion in part by listening to the air waves in various frequency blocks to test how often nothing at all was being transmitted. Most of the time, even in the dense urban settings of San Diego, Atlanta, and Chicago, important spectrum bands were nearly 100% idle. The public would be better served if others could use the otherwise idle spectrum.

For about 20 years, the FCC has utilized "secondary spectrum marketing." Someone wanting some spectrum for temporary use may be able to lease it from a party that has a right to use it but is willing to give it up in exchange for a payment. A university radio station, for example, might need the capacity to broadcast at high power only on a few Saturday afternoons to cover major football games—a time when the stock markets are closed and some spectrum is not heavily used by financial businesses. Or perhaps instead of reserving a band exclusively for emergency broadcasts, it could be made available to others for entertainment, with the understanding—enforced by codes wired into the transmitters—that the frequency would be yielded on demand for public safety broadcasts.

Computerized auctions can result in very efficient distribution of goods, whether used stuff on eBay or tiny spectrum bands. The use of particular pieces of the spectrum—at particular times, and in particular geographic areas—creates efficiencies if licensees of underutilized spectrum bands have an incentive to sell some of their time to other parties.

But secondary markets don't change the basic model: A frequency band belongs to one party at a time. Such auction ideas change the allocation scheme. Rather than having a government agency license spectrum statically to a single party with exclusive rights, several parties can divide it up and make trades. But these schemes retain the fundamental notion that spectrum is like land to be split up among those who want to use it.

Sharing the Spectrum

In his 1943 opinion, Justice Frankfurter used an analogy that unintentionally pointed toward another way of thinking. Spectrum was inevitably scarce, he opined: "Regulation of radio was therefore as vital to its development as traffic control was to the development of the automobile."

Just as the spectrum is said to be, the roadways are a national asset. They are controlled by federal, state, and local governments, which set rules for their use. You can't drive too fast. Your vehicle can't exceed height and weight limits, which may depend on the road.

But everyone shares the roads. There aren't any special highways reserved for government vehicles. Trucking companies can't get licenses to use particular roads and keep out their competitors. Everybody shares the capacity of the roads to carry traffic.[34]

The roads are what is known in law as a "commons" (a notion introduced in Chapter 6, "Balance Toppled"). The ocean is also a commons, a shared resource subject to international fishing agreements. In theory at least, the ocean need not be a commons. Fishing boats could have exclusive fishing rights in separate sectors of the ocean's surface. If the regions were large enough, it might be possible to earn a good living by fishing under these conditions. But such an allocation of the resources of the ocean would be dreadfully inefficient for society as a whole. The oceans better satisfy human needs if they are treated as a commons and fishing boats move with the fish—under agreed limits about the intensity of fishing.

Yochai Benkler's site, www.benkler.org, has several important and readable papers for free download, including the classic "Overcoming Agoraphobia."[35] His book *The Wealth of Networks*[36] details these and other concepts.

The spectrum can be shared rather than split up into pieces. There is a precedent in electronic communications. The Internet is a digital commons. Everyone's packets get mixed with everyone else's on the fiber optics and satellite links of the Internet backbone. The packets are coded. Which packet belongs to whom is sorted out at the ends. Anything confidential can be encrypted.

In spectrum policymaking, there is a choice between allocating spectrum—perhaps with some capacity for deal-making to improve utilization—and a more open, commons-based, Internet-like approach. As wireless technology surged in the first two decades of the twenty-first century, Congress moved to free up some underutilized sections of the spectrum, previously assigned to broadcast television, and adopted a combination of the "licensed" and "open" approaches to its use.[37]

To make any kind of spectrum sharing work, two ideas are key: First, using lots of bandwidth need not cause interference and can greatly increase transmission capacity; and second, putting computers into radio receivers can greatly improve the utilization of the spectrum.[38]

The Most Beautiful Inventor in the World

Spread spectrum was discovered and forgotten several times and in several countries.[39] Corporations (ITT, Sylvania, and Magnavox), universities (especially MIT), and government laboratories doing classified research all shared in giving birth to this key component of modern telecommunications—and were often unaware of each other's activities.

By far the most remarkable precedent for spread spectrum was a patented invention by Hollywood actress Hedy Lamarr—"the most beautiful woman in the world," in the words of movie mogul Louis Mayer—and George Antheil, an avant-garde composer known as "the bad boy of music."

Lamarr made a scandalous name for herself in Europe by appearing nude in 1933, at the age of 19, in the Czech movie *Ecstasy*. She became the trophy wife of Fritz Mandl, an Austrian munitions maker whose clients included both Hitler and Mussolini. In 1937, she disguised herself as a maid and escaped Mandl's house, fleeing first to Paris and then to London. There she met Mayer, who brought her to Hollywood. She became a star—and the iconic beauty of her screen generation (see Figure 8.3).

In 1940, Lamarr arranged to meet Antheil. Her upper torso could use some enhancement, she thought, and she hoped Antheil could give her some advice. Antheil was a self-styled expert on female endocrinology and had written a series of articles for *Esquire* magazine with titles such as "The Glandbook for the Questing Male."[40] Antheil suggested glandular extracts.[41] Their conversation then turned to other matters—specifically, to torpedo warfare.

A torpedo—just a bomb with a propeller—could sink a massive ship. Radio-controlled torpedoes had been developed by the end of World War I but were far from foolproof. An effective countermeasure was to jam the signal controlling the torpedo by broadcasting loud radio noise at the frequency of the control signal. The torpedo would go haywire and likely miss its target. From observing Mandl's business, Lamarr had learned about torpedoes and why it was hard to control them.

Lamarr had become fiercely pro-American and wished to help the Allied war effort. She conceived the idea of transmitting the torpedo control signal in short bursts at different frequencies. The code for the sequence of frequencies would be held identically within the torpedo and the controlling ship. Because the sequence would be unknown to the enemy, the transmission could not be jammed by flooding the airwaves with noise in any limited frequency band. Too much power would be required to jam all possible frequencies simultaneously.

(Source: © Bettmann/CORBIS)

FIGURE 8.3 Hedy Lamarr, at about the age when she and George Antheil made their spread spectrum discovery.

Antheil's contribution was to control the frequency-hopping sequence by means of a player piano mechanism—with which he was familiar because he had scored his masterpiece, *Ballet Mécanique*, for synchronized player pianos. As he and Lamarr conceived the device (it was never built), the signal would therefore hop among 88 frequencies, like the 88 keys on a piano keyboard. The ship and the torpedo would have identical piano rolls—in effect, encrypting the broadcast signal.

In 1941, Lamarr and Antheil assigned their patent (see Figure 8.4) to the Navy. A small item on the "Amusements" page of the *New York Times* quoted an army engineer as describing their invention as so "red hot" that he could not say what it did, except that it was "related to the remote control of apparatus employed in warfare."[42] Nonetheless, the Navy seems to have done nothing with the invention at the time. Instead, Lamarr went to work selling war bonds. Calling herself "just a plain gold-digger for Uncle Sam," she sold kisses and once raised $4.5 million at a single lunch.[43] The patent was ignored for more than a decade. Romuald Ireneus Scibor-Marchocki, who was an engineer for a Naval contractor in the mid-1950s, recalls being given a copy when he was put to work on a device for locating enemy submarines. He didn't recognize the patentee because she had not used her stage name.

The story of Antheil and Lamarr, and the place of their invention in the history of spread spectrum, is told in *Spread Spectrum* by Rob Walters.[44]

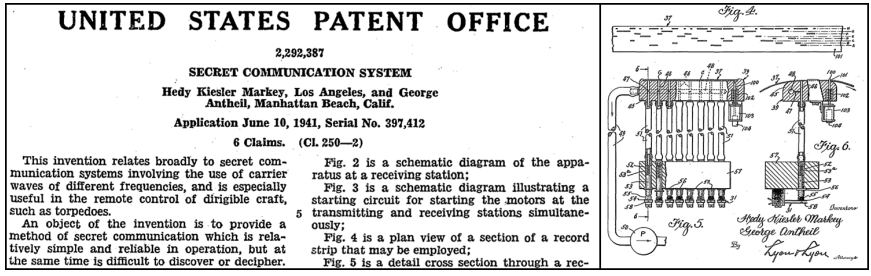

UNITED STATES PATENT OFFICE

2,292,387

SECRET COMMUNICATION SYSTEM

Hedy Kiesler Markey, Los Angeles, and George Antheil, Manhattan Beach, Calif.

Application June 10, 1941, Serial No. 397,412

6 Claims. (Cl. 250—2)

This invention relates broadly to secret communication systems involving the use of carrier waves of different frequencies, and is especially useful in the remote control of dirigible craft, such as torpedoes.

An object of the invention is to provide a method of secret communication which is relatively simple and reliable in operation, but at the same time is difficult to discover or decipher.

Fig. 2 is a schematic diagram of the apparatus at a receiving station;

Fig. 3 is a schematic diagram illustrating a starting circuit for starting the motors at the transmitting and receiving stations simultaneously;

Fig. 4 is a plan view of a section of a record strip that may be employed;

Fig. 5 is a detail cross section through a rec-

Fig. 4.

Fig. 6.

Fig. 5.

Hedy Kiesler Markey
George Antheil

FIGURE 8.4 Original spread spectrum patent by Hedy Lamarr (née Kiesler—Gene Markey was her second husband, of six) and George Antheil. On the left, the beginning of the patent itself. On the right, a diagram of the player-piano mechanism included as an illustration in the patent. (U.S. Patent Office)

And that, in a nutshell, is the strange story of serendipity, teamwork, vanity, and patriotism that led to the Lamarr–Antheil discovery of spread spectrum. The connection of these two to the discovery of spread spectrum was made only in the 1990s. By that time, the influence of their work had become

entangled with various lines of classified military research. Whether Hedy Lamarr was more a Leif Erikson than a Christopher Columbus of this new conceptual territory, she was surely the most unlikely of its discoverers. In 1997, the Electronic Frontier Foundation honored her for her discovery; she welcomed the award by saying, "It's about time." When asked about her dual achievements, she commented, "Films have a certain place in a certain time period. Technology is forever."

Channel Capacity

Lamarr and Antheil had stumbled on a particular way of exploiting a broad frequency range—"spreading" signals across the spectrum. The theoretical foundation for spread spectrum was one of the remarkable mathematical results of Claude Shannon in the late 1940s. Although no digital telephones or radios existed at the time, Shannon derived many of the basic laws by which they would have to operate. The Shannon–Hartley Theorem predicted spread spectrum in the same way that Maxwell's equations predicted radio waves.

Shannon's result (building on work by Ralph Hartley two decades earlier) implies that "interference" is not the right concept for thinking about how much information can be carried in the radio spectrum. Signals can overlap in frequency and yet be pulled apart perfectly by sufficiently sophisticated radio receivers.

Early engineers assumed that communication errors were inevitable. Send bits down a wire, or through space using radio waves, and some of them would probably arrive incorrectly because of noise. You could make the channel more reliable by slowing the transmission, the engineers supposed, in the same way that people talk more slowly when they want to be sure that others understand them—but you could never guarantee that a communication is errorless.

Shannon showed that communication channels actually behave quite differently. Any communication channel has a certain *channel capacity*—a number of bits per second that it can handle. If your Internet connection is advertised as having a bit rate of 3 Mbit/sec (or 3 Mpbs, 3 million bits per second), that number is the channel capacity of the particular connection between you and your Internet service provider (or should be—not all advertisements tell the truth). If the connection is over telephone wiring and you switch to a service that runs over fiber-optic cables, the channel capacity should increase.

However large it is, the channel capacity has a remarkable property, which Shannon proved: Bits can be transmitted through the channel, from the source to the destination, *with negligible probability of error* as long as the transmission rate does not exceed the channel capacity. Any attempt to push bits down the channel at a rate higher than the channel capacity will inevitably result

in data loss. With sufficient cleverness about the way data from the source is encoded before it is put in the channel, the error rate can be essentially zero, as long as the channel capacity is not exceeded. Only if the data rate exceeds the channel capacity do transmission errors become inevitable.

Errors and Delays

Although transmission errors can be made unlikely, they are never impossible. However, errors can be made far less probable than, for example, the death of the intended recipient in an earthquake that just happens to occur while the bits are on their way. Guaranteeing correctness requires adding redundant bits to the message—in the same way that fragile postal shipments are protected by the addition of Styrofoam or air-filled packing material. Attaining data rates close to the "Shannon limit" involves pre-processing the bits. That may increase *latency*—the time delay between the start of the "packing" process and the insertion of bits into the channel. Latency can be a problem in applications such as voice communication, where delays annoy the communicants. Happily, phone calls don't require error-free transmission; we are all used to putting up with a little bit of static.

Power, Signal, Noise, and Bandwidth

The capacity of a radio channel depends on the frequencies at which messages are transmitted and the amount of power used to transmit them. It's helpful to think about these two factors separately.

A radio broadcast is never "at" a single frequency. It always uses a range, or *band*, of frequencies to convey the actual sounds. The only sound that could be carried at a single, pure frequency would be an unvarying tone. The *bandwidth* of a broadcast is the size of the frequency band—that is, the difference between the top frequency and the bottom frequency of the band. Hoover, to use this language, allotted 10 kHz of bandwidth for each AM station.

Bandwidth

Because channel capacity depends on frequency bandwidth, the term *bandwidth* is used informally to mean "amount of information communicated per second." But technically, bandwidth involves electromagnetic communication, and even then it is only one of the factors affecting the capacity to carry bits.

If you can transmit so many bits per second with a certain amount of bandwidth, you can transmit twice that many bits per second if you have twice as much bandwidth. The two

transmissions could simply go on side by side, not interacting with each other in any way. So, *channel capacity is proportional to bandwidth.*

The relationship to signal power is more surprising. To use simple numbers for clarity, suppose you can transmit 1 bit, either 0 or 1, in 1 second. If you could use *more power* but *no more time or bandwidth*, how many bits could you transmit?

One way a radio transmission might distinguish between 0 and 1 is for the signals representing these two values to have different signal powers. To continue to oversimplify, assume that zero power represents 0, and a little more power, say 1 watt, represents 1. Then, to distinguish a 1 from a 0, the radio receiver has to be sensitive enough to tell the difference between 1 watt and 0 watts. The uncontrollable noise—radio waves arriving from sunspots, for example—also must be weak enough that it does not distort a signal representing 0 so that it is mistaken for a signal representing 1.

Under these conditions, four times as much power would enable transmission of 2 bits at once, still in 1 second. Power level 0 could represent 00; 1 watt, 01; 2 watts, 10; and 3 watts could represent 11. Successive power levels have to be separated by at least a watt to be sure that one signal is not confused with another. If the power levels were closer together, the unchanged noise might make them impossible to distinguish reliably. To transmit 3 bits at a time, you'd need eight times as much power, using levels 0 through 7 watts—that is, the amount of power needed increases exponentially with the number of bits to be transmitted at once (see Figure 8.5).

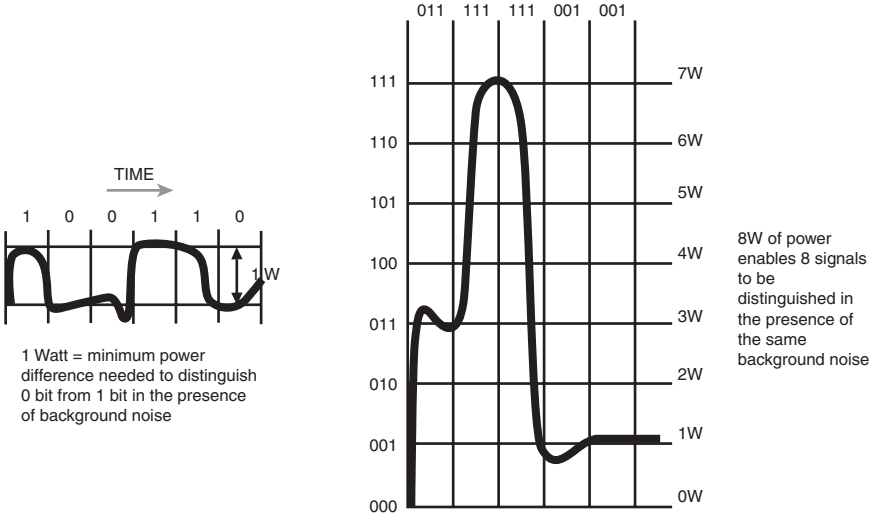

FIGURE 8.5 Shannon–Hartley. Signal levels must be far enough apart to be distinguishable in spite of the distortion caused by noise. Tripling the bit rate requires eight times as much power.

So the Shannon–Hartley result says that channel capacity depends on *both bandwidth and signal power, but more bandwidth is exponentially more valuable than more signal power*. You'd have to get more than *1,000 times* more signal power to get the same increase in channel capacity as you could get from having just 10 times more bandwidth (because $1{,}024 = 2^{10}$). Bandwidth is precious indeed.

My Signal Is Your Noise

The consequences of the Shannon–Hartley result about the value of bandwidth are quite astonishing. If WBZ were transmitting digitally with its 50,000-watt transmitter, it could transmit the same amount of information (over shorter distances) using less power than a flat screen TV if it could get 100 kHz of bandwidth rather than the 10 kHz the FCC has allowed it.

Of course, no station could get exclusive use of 100 kHz. Even giving each station 10 kHz uses up the spectrum too quickly. The spectrum-spreading idea works only if the spectrum is regarded as a commons. And to see the consequences of many signals broadcasting in the same spectrum, one more crucial insight is needed.

The power level that affects the capacity of a radio channel is not actually the signal power but the ratio of the signal power to the noise power—the *signal-to-noise ratio*. In other words, you could transmit at the same bit rate with 1 watt of power as with 10 watts—*if* you could also reduce the noise by a factor of 10. And *"noise" includes other people's signals*. It really doesn't matter whether the interference is coming from other human broadcasts or from distant stars. All the interfering broadcasts can share the same spectrum band to the extent that they can coexist with the equivalent amount of noise.

"Spread-Spectrum Radio"

A readable account of spread spectrum radio appeared in 1998: "Spread-Spectrum Radio" by David R. Hughes and DeWayne Hendricks.[45]

A surprising consequence of Shannon–Hartley is that *there is some channel capacity even if the noise (including other people's signals) is stronger than the signal*. Think of a noisy party: You can pick out a conversation from the background noise if you focus on a single voice, even if it is fainter than the rest of the noise. But the Shannon–Hartley result predicts even more: *The channel can transmit bits flawlessly, if slowly, even if the noise is many times more powerful than the signal. And if you could get a lot of bandwidth, you could drastically reduce the signal power without lowering the bit rate at all* (see Figure 8.6). What would seem to be just noise to anyone listening casually on a particular frequency would actually have a useful signal embedded within it.

The Shannon–Hartley Theorem is a mathematician's delight—a tease that hints at what is possible in theory. It is like Einstein's $E = mc^2$—which at once says nothing and everything about nuclear reactors and atomic bombs. Hedy Lamarr's frequency hopping was one of the spread spectrum techniques that would eventually be practical, but other ingenious inventions, named by odd acronyms, would emerge in the late twentieth century.

Two major obstacles stood between the Shannon–Hartley result and usable spread spectrum devices. The first was engineering: Computers had to become fast, powerful, and cheap enough to process bits for transmission of high-quality audio and video to consumers. That wouldn't happen until the 1980s. The other problem was regulatory. Here the problem was not mathematical or scientific. Bureaucracies change more slowly than the technologies they regulate.

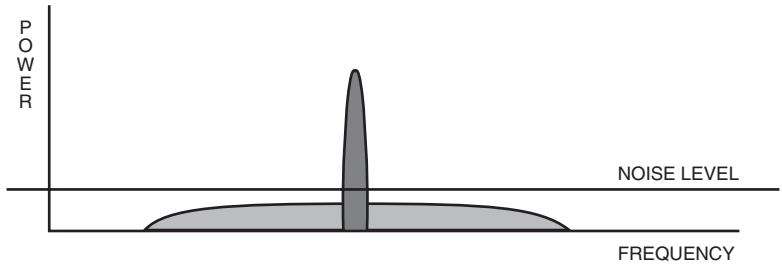

FIGURE 8.6 The spread spectrum principle. The same bit rate can be achieved at much lower power by using more bandwidth, and the signal power can even be less than the noise.

Spectrum Deregulated

Today, three-quarters of American households[46] have Wi-Fi Internet access. Hotel rooms and office suites have wireless Internet access. Even in buildings built less than 20 years ago, thousands of miles of cables are "dark"; they were installed to carry bits when Internet usage was exploding but are no longer needed because computers are connected wirelessly.

Wi-Fi happened because a tiny piece of the spectrum, a slice less than a millimeter wide in Figure 8.1, was deregulated and released for experimental use by creative engineers. It is an example of how deregulation can stimulate industrial innovations and about how existing spectrum owners prefer a regulatory climate that maintains their privileged position.[47]

Michael Marcus is an improbable revolutionary. An MIT-trained electrical engineer, he spent three years as an Air Force officer during the Vietnam War,

designing communications systems for underground nuclear test detection at a time when ARPANET—the original, military-sponsored version of the Internet—was first in use. After finishing active duty, he went to work at a Pentagon think tank, where he explored potential military uses of emerging communications technologies.

In the summer of 1979, Marcus attended an Army workshop on electronic warfare. As was typical at Army events, attendees were seated alphabetically. Marcus's neighbor was Steve Lukasik, the FCC's chief scientist. Lukasik had been director of ARPA during the development of ARPANET and then an ARPANET visionary at Xerox. He came to the FCC, not generally considered a technologically adventurous agency, because Carter administration officials were toying with the idea that existing federal regulations might be stifling innovation. Lukasik asked Marcus what he thought could stimulate growth in radio communications. Marcus answered, among other things, "spread spectrum." His engineering was sound, but not his politics. People would not like this idea.

The military's uses of spread spectrum were little known to civilians, since the Army likes to keep its affairs secret. The FCC prohibited all civil use of spread spectrum, since it would require, in the model the commission had used for decades, trespassing on spectrum bands of which incumbents had been guaranteed exclusive use. Using lots of bandwidth, even at low power levels, was simply not possible within FCC regulations. Lukasik invited Marcus to join the FCC, to champion the development of spread spectrum and other innovative technologies. That required changing the way the FCC had worked for years.

Shortly after the birth of the Federal Radio Commission, the U.S. plummeted into the worst depression it had ever experienced. In the 1970s, the FCC was still living with the culture of the 1930s, when national economic policies benevolently reined in free-market capitalism. As a general rule, innovators hate regulation, and incumbent stakeholders love it—when it protects their established interests. In the radio world, where spectrum is a limited, indispensable, government-controlled raw material, this dynamic can be powerfully stifling.[48]

Incumbents, such as existing radio and TV stations and cell phone companies, have spectrum rights granted by the FCC in the past, perhaps decades ago, and renewed almost automatically. Incumbents have no incentive to allow use of "their" spectrum for innovations that could threaten their business. Innovators can't get started without a guarantee from regulators that they will be granted use of spectrum, since investors won't fund businesses that are reliant on resources the government controls and may decide not to provide.

Regulators test proposals to relax their rules by inviting public comment, and the parties they hear from most are the incumbents—who have the resources to send teams to lobby against change. Their complaints predict disaster if the rules are relaxed. In fact, their doomsday scenarios are often exaggerated in the hope that the regulators will exclude competition. Eventually, the regulators lose sight of their ultimate responsibility, which is to the public good and not to the good of the incumbents. It is just easier to leave things alone. They can legitimately claim to be responding to what they are being told, however biased by the huge costs of travel and lobbying. Regulatory powers meant to prevent electromagnetic interference wind up preventing competition instead.

And then there is the revolving door. Most communications jobs are in the private sector. FCC employees know that their future lies in the commercial use of the spectrum. Hundreds of FCC staff and officials, including all recent FCC chairs, have gone to work for or represented the businesses they regulated.[49] These movements from government to private employment violate no government ethics rules. But FCC officials can be faced with a choice between angering a large incumbent that is a potential employer and disappointing a marginal startup or a public interest nonprofit. It is not surprising that they remember that they will have to earn a living after leaving the FCC.

In 1981, Marcus and his colleagues invited comment on a proposal to allow low-power transmission in broad frequency bands.[50] The incumbents who were using those bands almost universally howled. The FCC beat a retreat and attempted, in order to break the regulatory logjam, to find frequency bands where there could be few complaints about possible interference with other uses. They hit on the idea of deregulating three "garbage bands," so called because they were used only for "industrial, scientific, and medical" (ISM) purposes. Microwave ovens, for example, cook food by pummeling it with 2.450 GHz electromagnetic radiation (in the ISM bands). There should have been no complaints about deregulating the bands: Microwave ovens were unaffected by "interference" from radio signals, and the telecommunications industry did not use these bands.[51]

RCA and GE complained anyway about possible low-power interference, but their objections were determined to be exaggerated.[52] This spectrum band was opened to experimentation in 1985, with the proviso that frequency hopping or a similar technique be used to limit interference.

Marcus did not know what might develop, but engineers were waiting to take advantage of the opportunity. Irwin Jacobs founded QUALCOMM a few months later, and by 1990, the company's cell phone technology was in widespread use, using a spread spectrum technique called CDMA.

Over the next few years, several groups worked to develop protocols for wireless local area networking—ways for computers and other devices located

only a few feet apart to communicate with each other using newly deregulated spectrum bands. Low power consumption is desirable in local area applications, such as home use for connecting wireless keyboards to computers: Desktop and consumer devices, perhaps powered by batteries, should use little power, and the radius of their signals should be limited. Spectrum spreading makes it possible to achieve high bit rate with low power consumption.

One local area networking group branded its protocol "Wi-Fi"—though the name doesn't actually mean anything. NCR began manufacturing Wi-Fi devices that were inexpensive enough for consumer use. Steve Jobs recognized the potential, and Apple incorporated the NCR technology into the trademarked Airport wireless routers in 1997. When the FCC approved the 802.11 standard and the spectrum bands were finally available for public use, the press barely noticed, though the auction of other spectrum bands for cell phone use was widely reported. Within three years, wireless networking was everywhere, and virtually all personal computers now come ready for Wi-Fi.

Bluetooth is another low-power wireless technology that relies on spread spectrum to connect devices a few feet apart; this is referred to as "personal area networking," as opposed to local area networking. A great variety of headsets, keyboards, trackpads, and medical devices now connect wirelessly to computers and cell phones because a tiny segment of the radio spectrum (around 2.4 GHz) was deregulated and opened up for anyone to use, no license required.

Michael Marcus's website, www.marcus-spectrum.com, has interesting materials, and opinions, about spectrum deregulation and spread spectrum history.

For his efforts to open up the radio spectrum to competition, Marcus was sent into internal exile within the FCC for seven years. He emerged in the Clinton era and returned to spectrum policy work. He is now retired and working as a consultant in the private sector.

What Does the Future Hold for Radio?

In the world of radio communications, as everywhere else in the digital explosion, time has not stopped. In fact, digital communications have advanced less far than computer movie-making or voice recognition or weather prediction, because only in radio does the weight of federal regulation retard the explosive increase in computational power. The deregulation that is possible has only begun to happen.

What If Radios Were Smart?

Spread spectrum is a way of making better use of the spectrum. Another dramatic possibility comes with the recognition that ordinary radios are extremely stupid in comparison with what is computationally possible today. If taken back in time, today's radios could receive the broadcasts of 80 years ago, and the AM radios of 80 years ago would work as receivers of today's broadcasts. To achieve such total "backward compatibility," a great deal of efficiency must be sacrificed. The reason for such backward compatibility is not that many 80-year-old radios are still in service. It's that *at any moment in time*, the incumbents have a strong interest in retaining their market share and, therefore, in lobbying against efforts to make radios "smarter" so more stations can be accommodated.

If radios were intelligent and active, rather than dumb and passive, vastly more information could be made available through the airwaves.

If radios were intelligent and active, rather than dumb and passive, vastly more information could be made available through the airwaves. Rather than broadcasting at high power so that signals could travel great distances to reach passive receivers, low-power radios could pass signals on to each other. A request for a particular piece of information could be transmitted from radio to radio, and the information could be passed back. The radios could cooperate with each other to increase the information flux received by all of them. Or multiple weak transmitters could occasionally synchronize to produce a single powerful beam for long-range communication.

What Does "Smart" Mean?

"Intelligent" or "smart" radio goes by various technical names. The two most commonly used terms are *software-defined radio* (*SDR*) and *cognitive radio*. Software-defined radio refers to radios capable of being reprogrammed to change characteristics usually implemented in hardware today (such as recognizing AM, FM, or some other form of modulation). Cognitive radio refers to radios that use artificial intelligence to increase the efficiency of their spectrum utilization. "Smart radio" is also used as a marketing term to describe receivers that connect to the Internet as well as to AM and FM broadcast stations.

Such "cooperation gains" are already being exploited in *wireless sensor networking*. Small, low-power, radio-equipped computers are equipped with

sensors for temperature or seismic activity, for example. These devices can be scattered in remote areas with hostile environments, such as the rim of a smoldering volcano, or the Antarctic nesting grounds of endangered penguins. At far lower cost and greater safety than human observers could achieve, the devices can exchange information with their neighbors and eventually pass on a summary to a single high-power transmitter.

There are vast opportunities to use "smart" radios to increase the number of broadcast information options—if the regulatory stranglehold on the industry can be loosened and the incentives for innovation increased.

Radios can become "smarter" in another respect. Even under the "narrowband" model for spectrum allocation, where one signal occupies only a small range of frequencies, cheap computation can make a difference. The very notion that it is the government's job to prevent "interference," enshrined in legislation since the 1912 Radio Act, is now anachronistic.

Radio waves don't really "interfere," the way people in a crowd interfere with each other's movements. The waves don't bounce off each other; they pass right through each other. If two different waves pass through the antenna of a dumb old radio, neither signal can be heard clearly.

To see what might be possible in the future, ask a man and a woman to stand behind you, reading from different books at about the same voice level. If you don't focus, you will hear an incoherent jumble. But if you concentrate on one of the voices, you can understand it and block out the other. If you shift your focus to the other voice, you can pick that one out. This is possible because your brain performs sophisticated signal processing. It knows something about male and female voices. It knows the English language and tries to match the sounds it is hearing to a lexicon of word-sounds it expects English speakers to say. Radios could do the same thing—if not today, then soon, when computers become a bit more powerful.

But there is a chicken-and-egg cycle. No one will buy a "smart" radio unless there is something to listen to. No one can undertake a new form of broadcasting without raising some capital. No investor will put up money for a project that is dependent on uncertain deregulation decisions by the FCC. Dumb radios and inefficient spectrum use protect the incumbents from competition, so the incumbents lobby against deregulation.

Moreover, the incumbent telecommunications and entertainment industries are among the leading contributors to congressional election campaigns. Members of Congress often pressure the FCC to go against the public interest and in favor of the interests of the existing stakeholders. This problem was apparent even in the 1930s, when an early history of radio regulation stated that "no quasi-judicial body was ever subject to so much congressional pressure as the Federal Radio Commission."[53] The pattern has not changed.

In other technologies, such as the personal computer industry, there is no such cycle. Anyone who wants to innovate needs to raise money. Investors are inhibited by the quality of the technology and the market's expected reaction to it—but not by the reactions of federal regulators. Overextended copyright protections have chilled creativity, as discussed in Chapter 6, but lawmakers are to blame for that problem, not unelected commissioners.

But Do We Want the Digital Explosion?

Technologies converge. In 1971, Anthony Oettinger foresaw the line blurring between computing and communications. He called the emerging single technology "compunication."[54] Today's computer users don't even think about the fact that their data is stored thousands of miles away—until their Internet connection fails. Telephones were first connected using copper wires, and television stations first broadcast using electromagnetic waves, but today most telephone calls go through the air, and most television signals go through wires.

Laws, regulations, and bureaucracies change much more slowly than the technologies they govern. The FCC still has separate "Wireless" and "Wireline" bureaus. Special speech codes apply to broadcast radio and television, although "broadcasting" is an engineering anachronism. In a decision he signed in 2009, Justice Clarence Thomas signaled that he would be open to reconsidering the special speech codes that apply to broadcast radio and television:

> Dramatic technological advances have eviscerated the factual assumptions underlying those decisions. Broadcast spectrum is significantly less scarce than it was 40 years ago.... The extant facts that drove this Court to subject broadcasters to unique disfavor under the First Amendment do not exist today.[55]

The silo organization of the legal structures inhibits innovation in today's layered technologies. Regulation of the content layer should not be driven by an outdated understanding of the engineering limits of the physical layer. Investments made in developing the physical layer should not enable the same companies to control the content layer. The public interest is in innovation and efficiency; it is not in the preservation of old technologies and revolving doors between regulators and the incumbents of the regulated industry.

But if the spectrum is freed up—used vastly more efficiently than it now is and made available for innovative wireless inventions and far more "broadcast" channels—will we like the result?

There are general economic and social benefits from innovations in wireless technology. Electronic car keys, Xboxes, and highway toll transponders do not save lives, but wireless fire detectors and Global Positioning System devices do. The story of Wi-Fi illustrates how rapidly an unforeseen technology can become an essential piece of both business and personal infrastructure.

But what about television and radio? Would we really be better off with a million channels than we were in the 1950s with 13 or are today with a few hundred on satellite and cable? Won't this profusion of sources cause a general lowering of content quality and a societal splintering as *de facto* authoritative information channels wither? And won't it become impossible to keep out the smut, which most people don't want to see, whatever the rights of a few? Do we really want the airwaves to look like the undisciplined mess the Internet has turned into?

But there is another way to look at it. As a society, we simply have to confront the reality that our mindset about radio and television is wrong. It has been shaped by decades of the scarcity argument. That argument is now brain-dead, kept breathing on artificial life support by institutions that gain from the speech control it rationalizes. Without the scarcity argument, TV and radio stations become less like private leases on public land, or even shipping lanes, and more like...books.

There will be a period of social readjustment as television becomes more like a library. But the staggering—even frightening—diversity of published literature is not a reason not to have libraries. To be sure, there should be determined efforts to minimize the social cost of getting the huge national investment in old TV sets retired in favor of million-channel TV sets. But we know how to do that sort of thing. There is always a chicken-and-egg problem when a new technology comes along, such as FM radios or personal computers.

When market forces govern what gets aired, we may *not* be happy with the results, however plentiful. But if what people want is assurance about what they *won't* see, then the market will develop channels without dirty words and technologies to lock out the others. The present system stays in place because of the enormous financial and political influence of the incumbents—and because the government likes speech control.

How Much Government Regulation Is Needed?

Certainly, where words end and actions begin, people need government protection. Dr. Brinkley lost his medical license, which was right then and would be right today.

In the new wireless world, government needs to enforce the rules for spectrum sharing—technologies that can work only if everyone respects power and bandwidth restraints. The government has to ensure that manufactured devices obey the rules and that rogues don't violate them. The government also has to help develop and endorse standards for "smart" radios. There is a clear government interest in radio communications between, for example, autonomous vehicles and each other and the various traffic control devices they sense and signal to.

Government also has the ultimate responsibility for deciding if the dire warnings of incumbents about the risks imposed by new technologies are scientifically valid and, if valid, of sufficiently great social importance to block the advancement of engineering. A typical caution was the one issued in the fall of 2007 by the National Association of Broadcasters as it rolled out a national advertising campaign to block a new technology to locate unused parts of the TV spectrum for Internet service: "While our friends at Intel, Google, and Microsoft may find system errors, computer glitches, and dropped calls tolerable, broadcasters do not."[56] Scientific questions about interference should be settled by science, not by advertisements or congressional meddling. We will always need an independent body, like the FCC, to make these judgments rationally and in the public interest.

If we let science run the show, the scarcity problem will disappear. At that point, government authority over content should—and constitutionally *must*—drop back to its level for other non-scarce media, such as newspapers and books. Obscenity and libel laws would remain in place for wireless communication as for other media. So would any other lawful restrictions Congress might adopt, perhaps for reasons of national security.

Other regulation of broadcast words and images should end. Its legal foundation survives no longer in the newly engineered world of information. There are too many ways for the information to reach us. We need to take responsibility for what we see and what our children are allowed to see. And they must be educated to live in a world of information plenty.

If there were more channels, the government would not have any need, or authority, to second-guess the editorial judgment of broadcasters. Artificial spectrum scarcity has, in the words of Justice William O. Douglas, enabled "administration after administration to toy with TV or radio in order to serve its sordid or its benevolent ends."[57] That view—expressed in 1973, before cell phones and the Internet dominated our media lives—continues to be true today. Justice Frankfurter's claim that "there is no room in the broadcast band for every business or school of thought" is now false.

Bits are bits, whether they represent movies, payrolls, expletives, or poems. Bits are bits, whether they are moved as electrons in copper wire, light pulses in glass fiber, or modulations in radio waves. Bits are bits, whether they are

stored in gigantic data warehouses, in cell phones with 256 GB of memory, or on flash drives on keychains. The regulation of free speech on broadcast radio and television is but one example of the lingering social effects of historical accidents of technology. There are many others—in telephony, for example. Laws and policies regulating information developed around the technologies in which that information was embodied.

The digital explosion has reduced all information to its lowest common denominator: sequences of 0s and 1s. There are now adapters at all the junctions in the worldwide networks of information. A telephone call, a personal letter, and a television show all reach you through the same mixture of media. The bits are shunted between radio antennas, fiber-optic switching stations, and telephone wiring many times before they reach you.

Bits are bits, whether they represent movies, payrolls, expletives, or poems.

The universality of bits gives humankind a rare opportunity. We are in a position to decide on an overarching view of information. We can be bound in the future by first principles, not historical contingencies. In the United States, the digital explosion has blown away much of the technological wrapping obscuring the First Amendment. Knowing that information is just bits, all societies will be faced with stark questions about where information should be open, where it should be controlled, and where it should be banned.

Endnotes

1 David A. Fahrenthold, "Trump Recorded Having Extremely Lewd Conversation About Women in 2005," *The Washington Post*, October 8, 2016, https://www.washingtonpost.com/politics/trump-recorded-having-extremely-lewd-conversation-about-women-in-2005/2016/10/07/3b9ce776-8cb4-11e6-bf8a-3d26847eeed4_story.html.

2 Alexander Burns et al., "Donald Trump Apology Caps Day of Outrage over Lewd Tape," *The New York Times*, October 7, 2016, https://www.nytimes.com/2016/10/08/us/politics/donald-trump-women.html; Al Tompkins, "As Profanity-Laced Video Leaks, Outlets Grapple with Trump's Language," Poynter, October 7, 2016, https://www.poynter.org/reporting-editing/2016/as-profanity-laced-video-leaks-outlets-grapple-with-trumps-language/.

3 "Rise in Broadband Only Television Homes," Informitv, July 12, 2017, https://informitv.com/2017/07/12/rise-in-broadband-only-television-homes/.

4 knjazmilos, "The Ice Warnings Received By Titanic," *Titanic-Titanic.Com*, June 17, 2019, http://www.titanic-titanic.com/tag/ice/.

5 Karl Baarslag, *S O S to the Rescue* (Oxford University Press, 1935), p. 72.

6 "Wireless Melody Jarred," *New York Times*, January 14, 1910.

7 William Alden Smith, "United States Senate Inquiry Report," Titanic Inquiry Project, May 28, 1912, https://www.titanicinquiry.org/USInq/USReport/AmInqRepSmith01.php.

8 "An Act to Regulate Radio Communication," August 13, 1912, http://earlyradiohistory.us/1912act.htm.

9 "United States Frequency Allocation Chart," National Telecommunications and Information Administration, 2003, https://www.ntia.doc.gov/files/ntia/publications/2003-allochrt.pdf.

10 Erik Barnouw, *A Tower in Babel* (Oxford University Press, 1966), 69.

11 Barnouw, 91.

12 "Asks Radio Experts to Chart the Ether," *New York Times*, February 28, 1922.

13 *Hoover, Secretary of Commerce, v. Intercity Radio Co, Inc.*, Decision No. 3766 (52 App.D.C.339, 286 Fed. 1003).

14 Herbert Hoover, *Memoirs*, vol. 2 (The Macmillan Company, 1952), p. 140, https://heinonline-org.ezpprod1.hul.harvard.edu/HOL/Index?index=presidents%2Fmherbhv&tcollection=presidents.

15 "Asks Radio Experts to Chart the Ether."

16 *The Reminiscences of Herbert Clark Hoover*, vol. Radio Unit, Oral History Research Project, 1951, 12.

17 "End Cincinnati Radio Row," *The New York Times*, February 15, 1925.

18 *United States v. Zenith Radio Corporation*, 12 F.2d 614 (N.D. Ill. 1926), https://law.justia.com/cases/federal/district-courts/F2/12/614/1490149/.

19 "Hoover Asks Help to Avoid Air Chaos," *The New York Times*, July 19, 1926.

20 Radio Act of 1927, Pub. L. No. HR 9971 (1927), https://www.fcc.gov/document/radio-act-1927-established-federal-radio-commission.

21 Radio Act of 1927, 1.

22 Radio Act of 1927, 29.

23 Cecilia Kang, "Court Knocks Down FCC's Fine for Janet Jackson's 'Wardrobe Malfunction,'" *The Washington Post*, November 2, 2011, https://www.washingtonpost.com/business/economy/court-knocks-down-fccs-fine-for-janet-jacksons-wardrobe-malfunction/2011/11/02/gIQA98BpgM_story.html.

24 Ralph Berrier, "FCC Hits WDBJ with Proposed $325,000 Indecency Fine," *Roanoke Times*, March 23, 2015, https://www.roanoke.com/news/local/fcc-hits-wdbj-with-proposed-325-000-indecency-fine/article_f9c2a1b6-0f9a-50a9-8f9b-d02c0f2079ac.html.

25 Barnouw, *A Tower in Babel*, p. 169; Gerald Carson, *The Roguish World of Doctor Brinkley* (Rinehart, 1960), p. 33; Pope Brock, *Charlatan: America's Most Dangerous Huckster, the Man Who Pursued Him, and the Age of Flimflam* (Crown Publishers, 2008).

26 *KFKB Broadcasting Association v. Federal Radio Commission* (D.C. Cir. 1930), http://archive.org/details/dc_circ_1930_5240_kfkb_broad_assn_v_fed_radio_commn.

27 Carson, *The Roguish World of Doctor Brinkley*, p. 143.

28 Anthony Lewis, *Make No Law: The Sullivan Case and the First Amendment* (Vintage Books, 1992), p. 60.

29 *National Broadcasting Co., Inc. v. United States*, 319 U.S. 190 (1943), https://supreme.justia.com/cases/federal/us/319/190/.

30 E. Herring, "Politics and Radio Regulation," *Harvard Business Review* 13 (1935): 167–178.

31 "Asks Radio Experts to Chart the Ether."

32 Mark Lloyd, "The Strange Case of Satellite Radio," Center for American Progress, February 8, 2006, https://www.americanprogress.org/issues/democracy/news/2006/02/08/1829/the-strange-case-of-satellite-radio/.

33 "Report of the Spectrum Efficiency Working Group," Federal Communications Commission, Spectrum Policy Task Force, November 15, 2002, https://transition.fcc.gov/sptf/files/SEWGFinalReport_1.pdf.

34 Eli Noam, "Taking the Next Step Beyond Spectrum Auctions: Open Spectrum Access," 1995, http://www.columbia.edu/dlc/wp/citi/citinoam21.html.

35 Yochai Benkler, "Overcoming Agoraphobia: Building the Commons of the Digitally Networked Environment," *Harvard Journal of Law and Technology* 11 (1998), http://www.benkler.org/agoraphobia.pdf.

36 Yochai Benkler, *The Wealth of Networks* (Yale University Press, 2007).

37 Yochai Benkler, "Open Wireless vs. Licensed Spectrum: Evidence from Market Adoption," SSRN Scholarly Paper, November 7, 2012, https://papers.ssrn.com/abstract=2211680; Middle Class Tax Relief and Job Creation Act of 2012, Pub. L. No. 112–96 (n.d.), https://www.govinfo.gov/content/pkg/PLAW-112publ96/pdf/PLAW-112publ96.pdf.

38 Kevin Werbach, "Open Spectrum: The New Wireless Paradigm," October 2002, http://werbach.com/docs/new_wireless_paradigm.htm.

39 R. A. Scholtz, "The Origins of Spread-Spectrum Communications," *IEEE Transactions on Communications* 30, no. 5 (1982): 822–854, https://doi.org/10.1109/TCOM.1982.1095547; R. Scholtz, "Notes on Spread-Spectrum History," *IEEE Transactions on Communications* 31, no. 1 (1983): 82–84, https://doi.org/10.1109/TCOM.1983.1095718; R. Price, "Further Notes and Anecdotes on Spread-Spectrum Origins," *IEEE Transactions on Communications* 31, no. 1 (1983): 85–97, https://doi.org/10.1109/TCOM.1983.1095725; Rob Walters, *Spread Spectrum* (Book surge LLC, 2005).

40 George Antheil, "Glands on a Hobby Horse," *Esquire*, April 1936; George Antheil, "Glandbook for the Questing Male," *Esquire*, May 1936; George Antheil, "The Glandbook in Practical Use," *Esquire*, June 1936.

41 George Antheil, *Bad Boy of Music* (Doubleday, Doran & Co., 1945), 327.

42 "Hedy Lamarr Inventor," *The New York Times*, October 1, 1941.

43 "$4,547,000 Bonds," *The New York Times*, September 2, 1942; "Hollywood Puts on a Show," *Time*, October 12, 1942.

44 Walters, *Spread Spectrum*.

45 David R. Hughes and DeWayne Hendricks, "Spread-Spectrum Radio," *Scientific American*, April 1998.

46 Parks Associates, ">75% of U.S. Households Use WiFi for In Home Connectivity," IEEE Communications Society, June 12, 2017, http://techblog.comsoc.org/2017/06/02/parks-associates-75-of-u-s-households-use-wifi-for-in-home-connectivity/

47 "Early Civil Spread Spectrum History," Marcus Spectrum Solutions LLC, accessed May 7, 2020, http://www.marcus-spectrum.com/page4/SSHist.html.

48 Debora L. Spar, *Ruling the Waves: Cycles of Discovery, Chaos, and Wealth from the Compass to the Internet* (Harcourt, 2001).

49 John Dunbar et al., "Networks of Influence," Center for Public Integrity, February 2, 2006, https://publicintegrity.org/inequality-poverty-opportunity/networks-of-influence/.

50 "Early Civil Spread Spectrum History."

51 "A Brief History of Wi-Fi," *The Economist*, June 12, 2004, https://www.economist.com/technology-quarterly/2004/06/12/a-brief-history-of-wi-fi.

52 "Amendment of the Rules to Authorize Spread Spectrum and Other Wideband Emissions; Development of Appropriate Test Procedures to Determine Extent of Harmful Interference," Federal Communications Commission, December 3, 2018, https://www.fcc.gov/document/amendment-rules-authorize-spread-spectrum-and-other-wideband-1.

53 Laurence F. Schmeckebier, *The Federal Radio Commission* (The Brookings Institution, 1932), p. 55, https://babel.hathitrust.org/cgi/pt?id=uc1.$b99479&view=1up&seq=7.

54 Martin Greenberger et al., *Computers, Communications, and the Public Interest* (Johns Hopkins Press, 1971).

55 *FCC v. Fox TV Stations, Inc.*, 556 U.S. 502, 129 S. Ct. 1800 (2009), https://www.law.cornell.edu/supct/html/07-582.ZC.html.

56 Adario Strange, "NAB Launches Campaign Against 'White Space Devices,'" *Wired*, September 11, 2007, https://www.wired.com/2007/09/nab-launches-ca/.

57 *CBS v. Democratic Nat'l Committee*, 412 U.S. 94 (1973), https://supreme.justia.com/cases/federal/us/412/94/.

CHAPTER 9

The Next Frontier

AI and the Bits World of the Future

In Chapter 1, "Digital Explosion," we met Nicolette, the woman who did not get a second job interview after her first interview was conducted by a computer program. Nicolette's frustration with not knowing the basis for the decision against her is a canonical manifestation of misgivings about the very nature and use of artificial intelligence. In the past, many thinkers have speculated on the potential of intelligent devices and have formulated abstract principles that could be used to guide their performance. In 1950, science fiction writer Isaac Asimov posited three "Laws of Robotics":

> **First Law**
>
> A robot may not injure a human being or, through inaction, allow a human being to come to harm.
>
> **Second Law**
>
> A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
>
> **Third Law**
>
> A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.[1]

These simple rules interact in surprising ways, especially in the hands of an expert writer of fiction. But the world has moved on from imaging what might happen to experiencing what does.

We close our tour of the bits world by probing some of the dilemmas and opportunities posed by our breathtaking and yet limited technological success.

Thrown Under a Jaywalking Bus

In Ningbo, in eastern China, executive Dong Mingzhu, the "iron lady[2]" chair of the country's largest air conditioner manufacturer, was featured on a billboard with the caption "lawbreaker." Intersections in China's big cities often feature glittering LED screens. Plenty of them feature advertising, but they are also used for the latest in traffic policing: Live-action photographs combine with facial recognition software to name and shame jaywalkers caught crossing against the light. Dong Mingzhu wasn't a scofflaw. She wasn't even there at the time. The photo displayed there had indeed captured her face crossing the intersection—in an advertisement on the side of a bus.[3]

Ms. Dong escaped with a laugh, but those who don't have her standing or public support can find errors in "artificially intelligent" systems much more damaging. This small event tells us quite a bit about the state of artificial intelligence (AI), about its risks and potential. It portends a future in which machines of all sorts are capable not only of thinking and deciding but of acting without our knowledge or control. Just when we had come to the point in our story where we thought we understood the implications of the digital explosion, we realize that we are entering a whole new world.

How did this happen? At a basic level, the system that caught Ms. Dong worked exactly as it was designed to do: A camera trained on the intersection caught a picture of a person in the crosswalk when the light was against pedestrians crossing, matched the face with a massive database, and—in real time—displayed the infraction for all the world to see. There is quite a lot to explore here. This "system" interacted directly with its environment, watched

what was happening in the intersection, knew when the light was red or green, detected the presence of a person—or at least what it mistakenly thought was a person—matched up the photo against images of everyone it knew, decided what to do next, and took action.

This brief moment, seconds at most, captures many aspects of artificial intelligence, algorithmic decision making, privacy, ethics, fallibility, inherent bias, transparency, and accountability. We will explore all this and more. But first, let's consider some basics: artificial intelligence, machine learning and its cousin deep learning, and algorithmic decision making—where the pieces all come together.

What's Intelligent About Artificial Intelligence?

It has long been a dream of those working with technology to build machines that can learn and ultimately outgrow their programmers. Even Homer, some 28 centuries ago, imagined the divine blacksmith Hephaestus assisted by robots:

> In support of their master moved his attendants. These are golden, and in appearance like living young women. There is intelligence in their hearts, and there is speech in them and strength, and from the immortal gods they have learned how to do things. These stirred nimbly in support of their master.[4]

The concept of AI is often attributed to Alan Turing, the great British mathematician whose 1948 paper on intelligent machinery established a vision for machines that could mimic human behavior in areas as varied as games, language learning, cryptography, and mathematics. In 1950, in a paper titled "Computing Machinery and Intelligence,"[5] Turing introduced the concept of the imitation game (which we now refer to as the Turing test) as a way to measure whether machines could be considered "intelligent."

Seventy years after Turing's prescient vision, this form of AI behavior is commonplace. Six-year-old kids say "Alexa, show me a movie of Darth Vader" and expect the device to recognize who they are, understand the request, impose the suitable parental controls, and show the video. We take language comprehension—something completely inconceivable not all that long ago—entirely for granted. Our phones, cars, and TV remotes are capable of understanding continuous speech, placing it in context, and taking appropriate actions.

Machines that mimic human behavior—recognizing speech, understanding language, even passing the Turing test—are AI 1.0. Turing's earliest musings posited the basis of AI 2.0—machine learning. His 1948 report on intelligent machinery explored the notion that machines could learn from experience and modify themselves, much as humans do.

The Turing Test

British computer scientist Alan Turing proposed this test for intelligent machines: A human judge asks questions of two contestants, one computer and one human, communicating only through written messages. If the judge is unable to distinguish the computer from the human, the computer "passes" the test. While the Turing test was written as a thought experiment, it has sparked continued discussion and actual competitions. So far, at least, no computer has fooled a majority of its human judges.

Turing's notion spawned a whole range of chatbots—programs that interact with people through text. The first chat-bot, Joseph Weizenbaum's 1965 Eliza,[6] played a psychotherapist. Eliza convinced some users it was human despite using a simple algorithm that repurposed the question text and threw in some generalities when it couldn't do so. For example:

Men are all alike.

IN WHAT WAY?

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE?

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed much of the time.

I AM SORRY TO HEAR YOU ARE DEPRESSED

Machine Learning: I'll Figure It Out

Programmers write code to tell computers what to do. At least that's the way it has always been. The history of computer science has largely been the history of finding new algorithms to get computers to do increasingly complex and useful things, along with making the machines themselves faster and cheaper. Machine learning (ML) changes all that.

With ML, rather than programmers writing code to do something, the programmers write code that enables the computer to learn. The implications are profound. Before ML, you could reasonably ask how a computer program arrived at a particular result. Not so in this brave new world. The ML code the programmers wrote provided the framework. That piece of software was then exposed to data from which it "learned" how to do the task at hand: translate from Russian to English, decide who should get four years in prison and who

should get eight, buy or sell some stocks, or stomp on the brakes and avoid a car wreck. Most computer software arrives at a definitive result. Give it all the info about what you earned last year, and it will compute what you owe in taxes. ML programs make their best guesses about something they haven't seen before, based on what they have seen in the past.

At heart, ML systems are programs that observe and predict. Is this message spam? Does this photograph depict a cat or a dog? Is the person asking for a loan a good credit risk? Who is the person crossing the street while the light is red? Programmers generate mathematical models, train them on data for which they know the answers, and then apply them to unknown inputs. The trick of machine learning is that you don't start from scratch every time but build a general model, an artificial neural network, that you can give more specific input filters and train to solve new kinds of problems.

For each new task (classify an email as spam or a news story as fake), the developers of an ML solution decide what information should be considered. In the case of email, it might include the sender's email address, subject, a library of key phrases ("make money fast," "grow more hair"), list of known spammers, and more. Next, the system is trained by being allowed to process a set of previously categorized inputs—emails that are known to be good or known to be spam. The software adjusts the weight, a measure of importance, that it gives to each of the characteristics it is considering. In use, that learning process often continues every time the user tells the software that it made the wrong decision.

That's the simplest form of machine learning. Programmers code up the basic rules, and then the software adjusts its discrimination weights after looking at an adequately large set of sample data. While this is conceptually simple, in practice it requires quite powerful machines and very large data sets for training. These issues of scale were largely responsible for the long gap in implementation between Turing's original idea of learning machines and their practical realization.

ML systems operate on structured data. These systems assume that it's possible to get structured input, both for the training data and the actual operation. That's fine for something like classifying email as spam, but it wouldn't work for a self-driving car. Deep learning takes this model one very important step further.

In 1943, five years before Turing's intelligent machines report, Warren McCulloch and Walter Pitts published a paper[7] that described the potential to structure logical decision making in a manner based on a network of neurons, each of which takes inputs, has a threshold weight, and produces an output. This seminal work led, some 60 years later, to neural networks for classification and decision making. Neural networks apply feedback between layers of

nodes to "learn" at several different levels of generality. Importantly, unlike simpler ML systems, neural networks do not need structured data.

Figure 9.1 shows an example from the paper that introduced the model in 1958.[8] Each circle represents a "node"—that is, a set of interconnected neurons that process information before sending it on to nodes to its right or, in the case of rightmost nodes, back to nodes in the previous layer. Neural networks apply feedback between layers of nodes to "learn" at several different levels of generality. Each artificial neuron gets inputs, processes them based on an activation condition, and sends possibly modified output to the next layer or back to a previous layer.

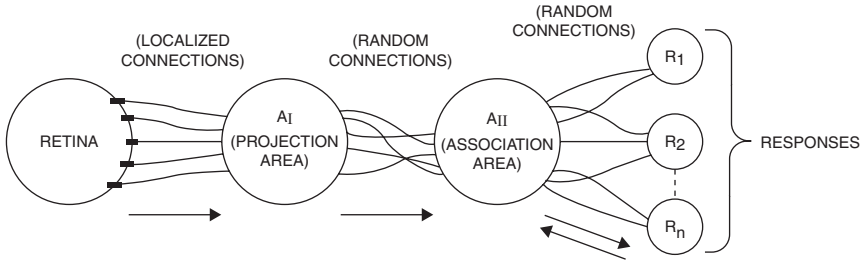

FIGURE 9.1 An artificial neural net modeled on the processing of signals from the retina to the brain.

Deep learning systems, built on artificial neural net platforms, are even more computationally intensive than traditional ML and frequently require special-purpose processors to achieve results in real time. Tesla claims that the AI processor in one of its cars (in 2019) is capable of 144 trillion operations per second—roughly equivalent to the power of 1,000 PCs.

Language translation is a good example of the impact of deep learning. Early attempts at machine translation—from Russian to English, for example—relied on attempts to create a logical model of grammar and vocabulary. As interesting as these linguistic models were, they ultimately failed as a means of doing language translation. Even early on, it was assumed that some form of machine learning would be better. Why not? After all, that's how we all learned our native language. You don't teach a two-year-old to conjugate verbs. You talk to her and correct her mistakes. And when, at six, the same child learns a second language, she doesn't start by learning sentence structure. She repeats the original process.

It took advances in algorithm development and computing power, and it also took a large resource of documents in multiple languages to make neural translation possible. The growth of the Internet made that a reality, and now

we have Google Translate, which currently supports more than 100 languages, from Afrikaans to Zulu.

ML isn't magical. It depends on powerful computers and lots of training data. It can seem like magic, though, because it can cut through tasks that are time-consuming or challenging for humans, And when a result comes out the other end, it often comes without any explanation—just a number, a rating, a score—yes, you get an interview, no you don't; your risk of recidivism is high or low.

Machine Learning and Training Data

Modeling a complex environment for a machine requires lots of labeled training data: lots of pictures of street signs and traffic signals with their meanings spelled out, lots of medical symptoms labeled to distinguish benign conditions from dangerous ones, lots of conversations transcribed and annotated. A great deal of human effort is therefore required to bootstrap the machines, and the data gathering raises numerous questions related to privacy for the data subjects included in training sets, working conditions for the people who do the labeling, and competitive implications of the reliance on big data.

Privacy

Jillian York, a technology and civil liberties activist, was on vacation when a friend contacted her to ask whether she knew her face appeared in an online database of "celebrity images."[9] She didn't, and upon investigation, was surprised to find included in the data set several years' worth of photographs captured by friends and still images extracted from videos. The images were included in the IARPA Janus Benchmark–C data set, a collection of images the U.S. Government's National Institute of Science and Technology (NIST) had made available as part of a public challenge to improve the state of the art in face recognition for "unconstrained in-the-wild face images." The challenge was "intended to drive research and development into face detection, verification, identification, and identity clustering."[10]

York was disturbed that her casual photos had been collected and indexed, but for NIST and IARPA, that was the point. "In-the-wild" meant training and testing recognition on faces with a variety of poses, backgrounds, and settings, so they gathered images that had been posted online, under Creative Commons licenses permitting their reproduction. That copyright permission didn't mean the photos' subjects expected to find themselves in a face ID database, however. Were it not for deep learning's voracious appetite for data, this privacy issue may never have arisen.

Labor

Training an AI system requires human labor to label its data sets: this image has a stop sign; that one shows a red light; that's a normal epithelial cell, this is a malignant one. Some companies distribute this new piece work through Amazon's Mechanical Turk. Others contract with workers who spend their days in call center–like cubicles, circling suspected polyps on colonoscopy videos. Without signing up for any of these jobs, you may have contributed to a training data set by answering the question "which of these photo squares contains a crosswalk?" in the captcha presented on sign-in to a new website. Sometimes this hidden labor comes to light, as when a number of automated voice "assistant" services updated their privacy policies to indicate that sometimes a human might listen to your conversation to help improve system performance.

And then sometimes, unbeknownst to us, systems monitor what we say and do and learn from our interactions. Is it right for corporations to use customers as unpaid labor?

Competition

The need for a great deal of training data means that data processing has economies of scale: Those who can gather more data can learn more from it, and as they improve their services, those services enable them to gather more data from the service's users. Every click of a Google search result has an effect—taking its searcher to a page Google believes will be useful—and a side effect—teaching Google that the page was responsive to the user's query. In the world of bits, scale really matters and provides the potential for insurmountable competitive advantage.

Nearly every Tesla car on the road sends back data that can be used to improve self-driving capabilities. Every time a Tesla driver intervenes when the car is steering to make a correction, data is captured and sent back for analysis. By mid-2020, Tesla had logged 3 billion miles driven on Autopilot.[11] Every mile driven, every automated lane change, every driver intervention provides data to Tesla that are used to improve the software. No other automaker comes close to having this sort of data source.

"Every time the customers drive the car, they're training the systems to be better. I'm just not sure how anyone competes with that."—Elon Musk

Algorithmic Decisions: I Thought Only People Could Do That

For the most part, computers take inputs, do some processing, and produce outputs. Increasingly, however, computers are taking over tasks that involve making decisions. Some are benign, but others can be life altering: which patients should enter hospice care,[12] who should receive a kidney transplant,[13] who gets to have a second-round interview for a job, which convict is likely to commit another crime.[14] Not all algorithmic decision systems rely on some form of machine learning, but many do. The very nature of these systems raises a host of concerns.

Just because a computer made a decision does not mean that it was right. Sometimes the results are funny—and sometimes tragic.

I'll Take Drosophila Biology for $23M

A postdoctoral researcher looking to buy a classic text on fruit fly biology found *The Making of a Fly* listed on Amazon for more than $1 million—$1,730,045.91 (+ $3.99 shipping), to be precise. He found this hard to imagine for a 1992 book (list price $70), when, 19 years later, a used copy could be ordered for as little as $35.54. Over several days, though, the price spiraled ever higher, until one seller was offering the book for $23,698,655.93.[15]

What was going on? Michael Eisen, a research biologist at University of California, Berkeley, found that only two sellers were offering new copies of the book, and as he tracked the book's pricing over a week, he discovered that the two had engaged in an algorithmic cat-and-mouse game. Both were using the same strategy. When one saw the other increasing its price, it increased its price, too. But the numbers were slightly different because the two were aiming at different markets. "bordeebook" aimed to top the market, while "profnath" aimed for a price just barely below the top price. As bordeebook's bot saw the market price shift, it raised its demand to 1.27059 times profnath's price; profnath's bot saw the market jump and, in turn, repriced to offer the book for sale at 99.83% of bordeebook's price, still a significant increase over its own previous price. bordeebook noticed profnath increasing its price and increased its own, and so on. With no human in the loop to check the algorithmic prices, the two bots ran amok, ratcheting up the price beyond what any person would pay for even a very good book (it averages 4.1 stars on Amazon) on Drosophila development.

Amazon isn't the only marketplace where bots can run wild. Consultants offer algorithmic pricing models for eBay auctions and Craigslist sales.

Ride-share services Uber and Lyft set "surge pricing" based on computerized observations that demand exceeds supply—and then they need human intervention or new limit conditions to stop price-gouging during natural disasters or emergencies.[16] Even the stock markets are filled with algorithmic traders. To keep up with other algorithms and gain a competitive advantage, traders program machines to make stock trades at high frequency without any human oversight. But automated trades inspire fears of a "flash crash"—fears that a rapid drop in the markets could be triggered by an input error or a connectivity glitch, setting off a cascade of automated sales and accidentally precipitating widespread economic chaos.

As we deal with buggy automation, we trade speed and efficiency against the risk of compound failures. Our risk tolerance for these algorithmic errors depends on the stakes. While we may find it acceptable to let rare book prices careen skyward until a person notices and pulls them down, we're less likely to tolerate that behavior for staple goods or stock markets, instead demanding regulation that slows the pace and provides oversight.

Algorithmic Justice League

Joy Buolamwini is a researcher at MIT's Media Lab, studying artificial intelligence—in particular, facial recognition. When she began her research, however, she found that her own face wasn't reliably detected as a face.[17] She had to wear a mask to be seen by the systems she was studying. Joy is black, and the programs she was testing had been trained on sets of white faces. The camera didn't see her because it hadn't been programmed with a wide enough range of human appearances.

Instead of just preparing a mask, Joy founded the Algorithmic Justice League to call attention to the problems of biased data sets. Her research has shown algorithms that mis-identify former First Lady Michelle Obama and women members of Congress at a much higher rate than their white male colleagues. Even if an algorithm can be designed without any racist intent or gender bias, if it is shown predominantly white male faces, those biases will color its facial recognition.

Unfortunately, these data biases not only persist but affect real-world deployments. The city of Detroit, which is majority black, has installed thousands of video cameras with facial recognition as part of "Project Green Light," a crime deterrence program.[18] However, one commercially available algorithm has a false positive rate of 1 in 1,000 for black faces, as compared with 1 in 10,000 for white faces. In a city with five times as many black faces as white, if a random set of the city population walks past the cameras, 50 black people will be wrongly identified for every 1 misidentified white person. Anyone who

appears in the state's driver's license database or other city records is liable to be caught up in the dragnet. Is being wrongly tagged for suspicion yet another hazard of "driving while black"?

In summer 2018, the ACLU used Amazon's Rekognition facial recognition application to run photos of sitting members of the U.S. Congress against a mugshot database. Rekognition nabbed 28 possible criminal matches.[19] (Whatever you think of their voting records, none of the congresspeople was pictured among the 25,000 public mugshots in the ACLU database.) Again, the errors were not uniformly distributed: "The false matches were disproportionately of people of color, including six members of the Congressional Black Caucus, among them civil rights legend Rep. John Lewis (D–Ga.)." Amazon responded that setting a higher threshold (rather than using the default) eliminated these errors.[20] That is to say, Amazon suggested that its software was fine but had been used incorrectly. If false positives (identifying a match that really wasn't a match) were a big problem, the ACLU should have set a higher confidence requirement for reporting a match (say, 99% confidence rather than 75% confidence). Of course, that explanation says nothing about the software's apparent propensity to misidentify black people.

Systematic errors like these prompted the city of San Francisco to enact a ban on facial recognition technologies—perhaps better described as face surveillance technologies, since no one objects to the use of face recognition for unlocking one's own cell phone. Civil liberties advocates brought their concerns to the city's board of supervisors, who voted 8–1 for the ban, persuaded that the dangers of discrimination and error outweighed the technology's potential utility.

Black-Box Justice

Algorithms in the justice system help determine whether an accused defendant will be released on bail or held in jail until trial, and they help set the length of a sentence someone will get after pleading guilty.

Investigative journalist Julia Angwin, writing in ProPublica, compared the use of AI in the cases of two people, each accused of $80 thefts.[21] Brisha Borden, age 18, was walking with a friend when they passed a child's bicycle and scooter and took the conveyances for brief rides before dropping them again. Vernon Prater, age 41, was caught after shoplifting $86.53 of tools from Home Depot. As each was booked into jail, a computer program gave a rating of their likelihood of re-offending. Borden, a black woman who had one juvenile misdemeanor charge, was deemed high risk, and Prater, a white man with an armed robbery conviction and another pending charge, was deemed low risk.

The computer program providing these re-offense risks asked a set of background questions about the defendants and delivered a number. But later events didn't match the computed predictions. Prater, the white man fingered by the computer as low risk, was re-arrested for a major electronics theft and is serving eight years in prison. Borden, identified as high risk, has avoided further legal trouble. However, the friend arrested along with her reports that the earlier arrest has made it difficult for her to find employment.

Programs such as Northpointe's COMPAS are used around the country as inputs to judges' sentencing decisions or bail hearings. The programs themselves are opaque: Fed a series of inputs, they produce a number but no explanation of their reasoning or opportunities to give feedback on their results. When defense attorneys ask for more detail to help their clients argue against mistakes, they are told that the programs are commercial trade secrets that can't be revealed.

Large-scale analyses show inconsistencies in the programs' application. In Angwin's analysis at Pro Publica, COMPAS was systematically biased against African Americans, marking them as more likely to re-offend than white defendants in otherwise similar circumstances. It will take further investigation to determine whether that bias was introduced deliberately in the programming, whether it was embedded in data used to train the system, or whether it emerged from data patterns that give an incomplete prediction for the future. The program itself is a black box: We see only its inputs and outputs, not what happens inside. But it becomes embedded in the justice system. It may save time, and people may truly believe its provider's claims of improved decision making, but it also can deflect responsibility for making hard decisions: "It wasn't my choice; the computer told me so."

Opacity is a fundamental characteristic of deep learning-based decision systems. These are systems that learned from experience, that formulated their own decision rules, and, most often, that have no mechanism for explaining how they arrived at their conclusions.

AI systems, by their very nature, are black boxes. The conclusions they draw, the classifications they make, the judgments they render result not from an algorithm defined and implemented by humans but rather from an accumulation of knowledge gained through observation. An opaque process cannot be seen to be fair both to the accused and to the observing public. Neither the accused nor the public can take guidance on how to stay on the right side of the law, as they can from reasoned judicial decision. When not even its creators know why an algorithm reached its conclusions or can explain what factors would change its outcome, and yet its product is given weight in a judicial proceeding, those seeking to challenge algorithmic results are denied due process.

What's Next

AI 2.0 solutions that utilize deep learning have the potential to be transformative beyond what we have already seen from the digital explosion. Their capacity to synthesize complex information and render judgments in real time presents the double-edged sword of opportunity and risk. A host of questions remain to be explored in the next years.

Responsibility

Elaine Herzberg was walking her bicycle across the road in Tempe, Arizona, on a dark Sunday night in 2018 when she was struck and killed by an autonomous Uber vehicle, becoming the first reported American death caused by a self-driving car.[22] In the days following, Uber engineers and law enforcement investigators pored over the data to try to determine what had gone wrong. The self-driving car was designed to avoid collisions but had failed to do so here, with fatal result. Was the problem in the software or hardware? In the sensors, computer vision, processing, or actuated response? And who was responsible—the software developers, the car company, the car owner?

Many of the car's components recorded data, logging inputs and outputs in a manner similar to an aircraft black-box flight recorder. Those logs could show whether sensors detected an object and whether a command to apply the brakes had issued. At some of the intermediate stages, however, questions of interpretation become more difficult: If the car "saw" an object but failed to identify the object as a person for whom it should stop (even if it had to brake suddenly), where might the recognition have erred?[23] Much of this software's operation is discontinuous: Two very similar scenes can present widely varied appearance to AI, and a small error or difference between test and real-world conditions can have dramatic consequences. Enumerating all the possibilities in testing, or even recording enough information for an after-the-fact audit, can be daunting.

In Uber's collision with Ms. Herzberg, a National Transportation Safety Board investigation found that an emergency braking system had been disabled, and the human "backup driver" didn't react fast enough to stop the car.

According to data obtained from the self-driving system, the system first registered radar and LIDAR observations of the pedestrian about 6 seconds before impact, when the vehicle was traveling at 43 mph. As the vehicle and pedestrian paths converged, the self-driving system software classified the pedestrian as an unknown object, as a vehicle, and then as a bicycle with varying expectations of future travel path. At 1.3 seconds before impact, the self-driving system determined that an emergency braking maneuver was

needed to mitigate a collision. According to Uber, emergency braking maneuvers are not enabled while the vehicle is under computer control, to reduce the potential for erratic vehicle behavior. The vehicle operator is relied on to intervene and act. The system is not designed to alert the operator.[24]

The car's computer vision systems saw a moving object, but, uncertain what kind of object, couldn't confidently predict whether it—she—would be in the car's path. By the time the system identified the pedestrian walking a bicycle and realized the car would hit her, its only corrective move was one that had been disabled. Designers of the system apparently made earlier choices that put the system in this bind. For example, although they could have instructed the car to slow at any hint of danger, that would slow the trip; unexpected braking could introduce a greater risk of being hit from behind.

Habituation

As we design artificially intelligent systems, we face choices of what parameters to set and what risks to take. Paradoxically, as systems get better, the choices may become more explicit. A human driver reacts instinctually (or fails to react) to an object darting into the road; an autonomous vehicle must be programmed to anticipate such obstacles and also to choose among a series of imperfect alternatives that might include braking hard (at some risk to its occupants and cars behind), swerving into oncoming traffic (at risk to occupants and the opposite lane), or veering into a crowded sidewalk. Calculating the risks and their potential harms seems both necessary and unfair: How can we compare human lives?

Trolley Problem

Philosopher Philippa Foot first described the "trolley problem" in a 1967 article.[25] She asks us to consider the choice facing "the driver of a runaway tram which he can only steer from one narrow track on to another; five men are working on one track and one man on the other; anyone on the track he enters is bound to be killed." Suppose his trolley is running straight toward the five, with no room to apply the brakes, but he could throw a switch to divert to a spur track on which only one is working. Should he throw the switch? Must he?

This thought experiment, and a variety of similarly gruesome choices (should you push the fat man off the overpass to stop the train? kill a healthy person if her organs can help five sick people to live?), can help illuminate the factors that go into our thinking about "fairness" or, as our hair-splitting

get ever more tangled, remind us that intuitive judgments may be clouded by elements that should be irrelevant to a principled fairness.

The trolley problem has obvious appeal to those thinking about autonomous vehicles, where the vehicle might be in the pre-switch position, caught between alternatives in which any choice causes harm: Should the car protect its passenger or another car with two passengers? Its passenger or a pedestrian? Its owner's wallet against damage to the car? Some observers note that each of these decision situations comes only as the last step in a long chain of events and suggest that we should break out of the trolley problem framing altogether by changing the parameters of driving: Build interlock switches to keep the trolley off the track entirely when repairs are underway. Instead of putting autonomous vehicles on the road with human-driven cars, give them separate lanes and software that coordinates the actions of the fleet. Any of these lenses refocuses questions of systemic design and justice: Whose interests are considered and given priority, and who might be overlooked?

Transparency: Why Did You Do That?

And so we come back to Nicolette, who didn't get the job after the computer interviewed her. There are many similar examples. Christian Sandvig and Karrie Karahalios want to help us understand algorithmic systems by auditing them. They have tried, for example, preparing Facebook profiles with slight variations to compare the advertisements shown to each.[26] Will someone who describes herself as a single childless woman see the same housing ads as someone who gushes about her young children? As a man?

When presented with black-box AI, or even with a fully disclosed algorithm whose workings are opaque, we can try to understand its operation through audits. An after-the-fact audit examines the inputs and outputs, pulling out individual data points and aggregate results to look for anomalies or unexpected behavior.

Because their tests involve creating possibly false or adversarial information to interact with computer systems, however, Sandvig and Karahalios face challenges under the Computer Fraud and Abuse Act (CFAA) if their "fake profiles" go against the terms of service of one of the platforms they seek to test. The black box has rules against testing how it works. Landlords argued that auditors were trespassing when they came to test an apartment they weren't planning to rent (though they lost their suit). As Karahalios and Sandvig fight to enforce the right to conduct live-audit tests for possible machine discrimination, they are seeking to enshrine a similar protection for assessment of the digital environment. When does the right to examine a system's bias outweigh the system operator's right to limit access?

Is Better Good Enough?

Elaine Herzberg's death was a tragedy. Even though the Tempe police concluded that the accident was unavoidable,[27] Uber terminated its testing of self-driving cars as a result of this accident.[28] That is entirely understandable. We naturally respond to individual deaths in this way, and there is broad concern about the overall safety of autonomous vehicles even when accidents don't involve fatalities. How, though, should we think about statistical advantages? Consider the reported results for Tesla cars with Autopilot:

> In the 1st quarter, we registered one accident for every 4.68 million miles driven in which drivers had Autopilot engaged. For those driving without Autopilot but with our active safety features, we registered one accident for every 1.99 million miles driven. For those driving without Autopilot and without our active safety features, we registered one accident for every 1.42 million miles driven. By comparison, NHTSA's most recent data shows that in the United States there is an automobile crash every 479,000 miles.[29]

The makers of the HireVue automated screening system similarly argue that their system works better than the humans it replaces. Granting that the system offers no explanation for its decision, the company claims that:

> Decades of research have shown that traditional interviews are full of implicit and explicit bias, and tremendous inconsistency. The HireVue approach has been proven to be measurably more accurate at predicting performance than human evaluators and is audited, tested, retrained, and audited again to ensure that there is no adverse impact.[30]

So, which is it—biased or not? Are we fundamentally better off with a system whose bias can be evaluated, even if it cannot be interrogated?

The Future of Work

We have already seen the tremendous impact on productivity that has resulted from the digital explosion. We no longer have typing pools, and travel agents are all but extinct. The accounts payable department is a tiny fraction of the size it once was, as technology has taken over a wide range of tasks.

Until the advent of AI 2.0—of machine learning systems capable of perceiving, understanding, and interacting with the physical world—the majority of the impact on labor has been limited to information-intensive tasks. That is about to change, as we see computer systems that are capable of learning,

capable of self-improvement, and capable of making seemingly complex decisions in real time.

One of the most common jobs in the United States is "driver"—trucks, buses, taxis, tractors, fork lifts, Uber, and Lyft. How long before autonomous vehicles of one sort or another take over these jobs? And what will happen in all sorts of other professions—in tax preparation, reading X-rays, customer service, and more?

The Role of Regulation

AI, in general, and deep learning systems that render opaque judgments in particular, present a critical need for intelligent regulation.

Asimov's laws of robotics have given birth to 1,000 alternatives and descendants. Even standards just for algorithmic fairness and transparency are numerous. The U.S. Association for Computing Machinery proposed the following principles as a starting point:[31]

1. Awareness: Owners, designers, builders, users, and other stakeholders of analytic systems should be aware of the possible biases involved in their design, implementation, and use and the potential harm that biases can cause to individuals and society.

2. Access and redress: Regulators should encourage the adoption of mechanisms that enable questioning and redress for individuals and groups that are adversely affected by algorithmically informed decisions.

3. Accountability: Institutions should be held responsible for decisions made by the algorithms that they use, even if it is not feasible to explain in detail how the algorithms produce their results.

4. Explanation: Systems and institutions that use algorithmic decision-making are encouraged to produce explanations regarding both the procedures followed by the algorithm and the specific decisions that are made. This is particularly important in public policy contexts.

5. Data Provenance: A description of the way in which the training data was collected should be maintained by the builders of the algorithms, accompanied by an exploration of the potential biases induced by the human or algorithmic data-gathering process. Public scrutiny of the data provides maximum opportunity for corrections. However, concerns over

privacy, protecting trade secrets, or revelation of analytics that might allow malicious actors to game the system can justify restricting access to qualified and authorized individuals.

6. Auditability: Models, algorithms, data, and decisions should be recorded so that they can be audited in cases where harm is suspected.

7. Validation and Testing: Institutions should use rigorous methods to validate their models and document those methods and results. In particular, they should routinely perform tests to assess and determine whether the model generates discriminatory harm. Institutions are encouraged to make the results of such tests public.

But how do we get policymakers, designers, engineers, and even consumers to take such principles seriously when motives for profit, market share, innovation, and simple excitement are pressing on everyone to get the product built and out the door without, in the words of Norbert Wiener, the father of Cybernetics, "exert[ing] the full strength of our imagination to examine where the full use of our new modalities may lead us"[32]?

Optimism for the Future

Despite all the inherent risks and complexities, we should recognize that AI and ML have vast potential for good. Just a few examples of the solutions in development now are accelerated pharmaceutical development, increased crop yields, reduced automotive injuries and deaths, lower-cost health care, better fraud and crime detection, and improved manufacturing efficiencies.

AI in general and machine learning in particular create the potential for exponential acceleration in the capabilities of digital systems as they learn to improve themselves, much as we humans have done for millennia.

Bits Lighting Up the World

So what happens after the digital explosion? To be sure, we are today nowhere near "after." More bits than ever before are now being produced, analyzed, stored, and used as training data for systems that consume and generate more bits. We are still near the beginning of the explosion. But it is not too soon to try to view it as a whole.

In Greek mythology, Prometheus stole Zeus's fire and brought it from Olympus to Earth, along with the useful arts of civilization. Zeus retaliated for

Prometheus's trickery by visiting upon humanity the ills and evils that beset us. We have been trying to make the best of things ever since.

The Prometheus myth is about technology. Technology, like fire, is neither good nor bad; its value depends on how we use it. And once we start using a technology, society itself changes. As William Yeats wrote, "All changed, changed utterly: A terrible beauty is born."[33]

Information technologies spark a special kind of fire. Bits are the atomic particles of the information flames. With our information tools, we can do things, both good and bad, that we could not have done unassisted. For better or worse, these technologies enable us to think, reason, create, express, debate, compromise, learn, and teach in ways never before possible. They connect people across physical space, both in pairs and in groups. They extend the reach of our voices and the range of our hearing. They also amplify our capacity to frighten, harass, and hate other people and to misrepresent ourselves to others. They enable us to earn and to spend money without going anywhere and also to steal money from the comfort of our homes.

So central was Prometheus, the fire-bringer, to the Greek conception of humanity that in later retellings of the myth, he is credited with creating the human species itself. What changes to society will information technologies yield, in a decade or two, when the ongoing digital explosion has unimaginable power?

We don't know, of course. But if things go on changing as they are changing today, there are likely to be dramatic changes to three distinctive aspects of human culture: our sense of personal identity and privacy, our capacity for free speech, and the creativity that drives human progress.

Privacy and Personhood

As the digital explosion was beginning, the struggle over privacy seemed to be a war. Individuals wanted to protect themselves from invasive forces. Institutions, both corporations and government, wanted the benefit of information that individuals would rather not reveal. In actual practice, as us versus them, good versus bad, or individuals versus institutions.

In the digital explosion, as technologies improved, data gathering became easier and less annoying. Modest incentives induced individuals to sacrifice their personal privacy, often before they understood what they were giving up. Relatively few people today worry about stores keeping track of their purchases. Even without loyalty cards, a credit card swipe together with bar code scans at the cash register link a customer's name to his preferences in candy and condoms. You have to give up many conveniences to protect your privacy, and most people are not willing to do it.

The next generation may not even see the loss of privacy as a sacrifice. Socrates said that the unexamined life is not worth living, but people who have grown up with social networks may find life fully exposed to public view simply normal. As Sun Microsystem CEO Scott McNealy quipped, "You have zero privacy anyway. Get over it."

But getting over it is not so simple when social interactions happen through the computer screen. When most personal interactions were face-to-face or over the telephone, we mistrusted people claiming to represent our bank and trusted people we felt we had gotten to know. In the electronic world, we do the opposite: We trust our bank's website with large sums of our money, but we have to be reminded that close electronic friends may be impostors. Where is the border for children between the personal and the public? Will we need laws about fraudulent friendships?

As electronic privacy becomes lost in the cloud of bits and as caution gives way to social networking, what societal structures will break down? What will evolve to replace them? Society as we know it functions because of a web of trusting relationships between parties who are independently responsible for their own actions. What will replace that if the concept of personal identity becomes meaningless? Will the very notions of privacy and identity be destroyed in the explosion?

What Can We Say, and Who Will Be Listening?

The digital explosion revolutionizes human communication. Earlier technologies for disseminating text, spoken words, and images also changed the world—but all included choke points. A million eyes might have read your book, but only if you could get it published. You might have discovered a scandal that would bring down a government, but only if you could get a newspaper to expose it to public view. A million ears might have heard your speeches, but only if you could control a radio station.

No longer are speakers bound by the whims of those who control the loudspeakers and printing presses. In the United States, anyone can say anything, without permission from church or state, and be audible to millions. No one has to listen, but it is easy to put the message where millions can hear it.

An Earlier Information Revolution

Victor Hugo said of printing in *The Hunchback of Notre Dame:* "It is the mother of revolution. It is the mode of expression of humanity which is totally renewed; it is human thought stripping off one form and donning another; it is the complete and definitive change of skin of that symbolical serpent which since the days of Adam has represented intelligence."

And yet there is a cost. It is not a financial cost; after all, it costs nothing to spread the word via email or Twitter or Facebook or YouTube. The cost is that the speaker relies on many intermediaries to handle the messages, and so there are many opportunities for snooping, eavesdropping, filtering, and censoring, not to mention the dilemma of knowing the source and reliability of the information you receive. The choke points have multiplied and become more diffuse, but they have not disappeared.

The very technological miracles that have created the communication revolution have also created a Big Brother revolution. Massive surveillance in China has been automated based on face recognition coupled with monitoring of cell phone signals.[34] U.S. law enforcement has embraced Clearview AI's face surveillance capabilities with equal enthusiasm.[35] With the success of speech recognition and language understanding, we have to expect that it is already feasible to monitor every voice communication passing over telephone wires or the Internet by means of the technological equivalent of a human listener. Machines will be waiting attentively for someone to say the "wrong" thing—whatever that is deemed to be.

Governments eavesdrop to protect national security, political opposition, and public morality. Communication companies want to listen to what their networks are being used for so that they can tailor their service to the content in the most profitable way—a soft form of corporate censorship, in which unwanted communications are slowed down or made costly. Service providers want to listen in so they can add advertising to the content they deliver.

In spite of the unimaginable expansion of communications over the past quarter century, the jury is still out on whether speech will be freer or less free in the future than it was in the past, even in the United States, with its uncompromising First Amendment. And like the tree falling in the forest, of what use will free speech be if no one is listening? The dramatic pluralism of our information sources threatens to create a society where no one learns anything from people with whom they disagree. It is simply too easy for people to decide whom they want to hear and to ignore everyone else. Will the digital explosion in fact make information more limited?

In spite of the unimaginable expansion of communications over the past quarter century, the jury is still out on whether speech will be freer or less free in the future than it was in the past.

A Creative Explosion or a Legal Explosion?

In the same letter quoted in Chapter 1, Thomas Jefferson wrote, "He who receives an idea from me, receives instruction himself without lessening mine;

as he who lights his taper at mine, receives light without darkening me." Will the digital explosion be used to enlighten the world or to create illusions and to blind us to the truth?

What would Jefferson have to say about the viral spread of misinformation through online social networks? Four years after the 2016 U.S. elections, social media companies and governments everywhere have found no simple means to balance the right to express political views with the reality that interesting lies spread quickly and unexciting truths can take a long time and immense labor to reveal. And not all misinformation is political or commercial. When disease spreads rapidly, it can be hard for reliable information to penetrate where unfounded conspiracy theories have taken root. During the novel coronavirus epidemic that began in 2019, anti-vaccination misinformation spread before the first vaccine trials had even begun. Only a year earlier, 83 people, mostly young children, died on the tiny island of Samoa from the entirely preventable disease measles because an anti-vaccination campaign had reduced vaccination rates on the island to 34%.[36]

Manipulation of speech is not the only form of technologically enabled control of information. Patent and copyright laws in the United States were designed to promote individual creativity in the interest of the progress of society. The law struck a balance between providing financial incentive to the creator and high social benefit to the population at large. The term for which artists and inventors maintained exclusive control over their creations was designed to be long enough to provide a financial return and short enough to provide an incentive for continued creativity. And there was a high threshold on what could be protected at all so that the system did not encourage lawyerly inventiveness rather than artistic and engineering creativity.

As mechanical tools have been supplanted by information-processing tools, and all manner of writing, music, and art have gone digital, the rules of the game have changed. The parties that receive the strongest protections are now major corporations rather than the original creators or the ultimate consumers. At a time when information technology promises disintermediation—getting rid of intermediaries—those intermediaries are becoming more powerful, not less.

Unexpected—and unintended—consequences attend any speech-limiting regulation. In the midst of the Covid-19 crisis, copyright takedown notices were issued to Google to remove informative articles from search results. Google purged a news story about two visitors to Vietnam who had become sick—information that could have helped others learn that they had been exposed to the virus. The story appeared on a Vietnamese government–affiliated news website, naming the hotel, bars, and restaurants the tourists visited and urging readers who patronized the establishments to take precautions. Someone who wanted this information removed backdated a blog

post with the identical information and then complained that the news story infringed his copyright. Google removed search links to the original story even though the fake was dated more than four months before the tourists visited. The blog with the fake post included only seven others, all cited in copyright complaints filed with Google.[37]

The legal power of strong intermediaries protecting their economic interests has increased at the same time as new technologies have empowered the creators to reach their consumers directly.

Similar tensions are visible in the world of invention. The power of the incumbent radio and television broadcast industries to exclude newcomers from the airwaves restrains both speech and invention, limiting radio communications and keeping useful devices off the market. And there is good reason to ask if the same pattern is repeating itself in the domains of information search—where Google is dominant—and social networking, with Facebook's overwhelming position in large parts of the world.

Will the United States move toward being an information democracy or an information oligarchy? Whose hands will be on the controls that regulate the way we produce and use bits in the future?

A Few Bits in Conclusion

The worldwide bits explosion is lighting up the world (see Figure 9.2).[38] Most of the illumination today is in Europe and North America, but it is growing brighter almost everywhere. There is no physical reason it can't continue to grow. Bits are not like oil or coal. Their production requires almost no raw materials and only tiny amounts of electricity. They flow through glass fibers in astonishing numbers, and they radiate through space, over short distances and long. With our cameras and computers, we produce them at will, in unintelligibly larger numbers every year. Existing dark spots—North Korea, for example—may remain black for a time, but eventually even these regions may glow brighter. And all that data and thought-stuff, all those atoms of light, can be captured and stored electronically for eternity.

The explosion happened through technological inventions supported by political and economic freedoms. Gutenberg laid the foundation when he invented the printing press, and Morse's telegraph, Bell's telephone, and Edison's phonograph were all precursors. Claude Shannon was the bits Prometheus. After the Second World War, his mathematical insights lit the flame of communication and computing technologies, which have now illuminated the earth with bits.

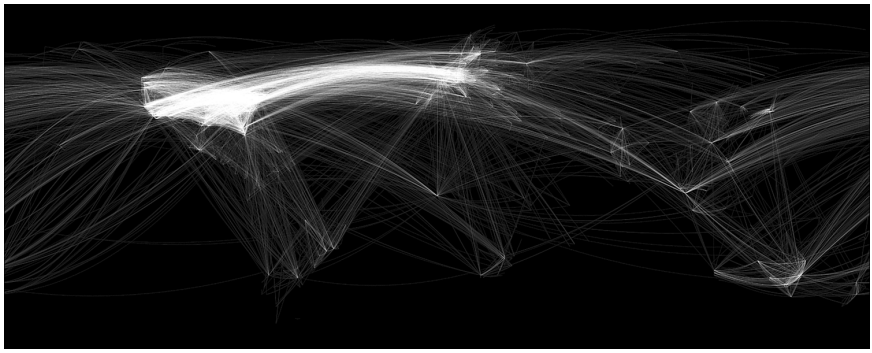

FIGURE 9.2 A map of the world, showing the number of Internet connections between routers. At present, the United States and Europe are heavily interconnected. If the volume of data transmissions were depicted instead (giving more prominence, for example, to areas with heavily used Internet cafés), Africa, Asia, and South America might show more prominently.

The bits explosion is not over. We are in the middle of it. But we don't know whether it will be destructive or enlightening. The time for deciding who will control the explosion may soon be past. Bits are still a new phenomenon—a new natural resource whose regulatory structures and corporate ownership are still up for grabs. The legal and economic decisions being made today—not just about bits but about everything that depends on bits—will determine how our descendants will lead their lives. The way the bits illuminate or distort the world will shape the future of humanity.

Endnotes

1 Isaac Asimov, *I, Robot* (Gnome Press, 1950).

2 Russell Flannery, "2017 Forbes China 100 Top Businesswomen List," *Forbes*, February 6, 2017.

3 Xinmei Shen, "Facial Recognition Camera Catches Top Businesswoman 'Jaywalking' Because Her Face Was on a Bus," *Abacus*, November 22, 2018, https://www.abacusnews.com/digital-life/facial-recognition-camera-catches-top-businesswoman-jaywalking-because-her-face-was-bus/article/2174508.

4 Homer, *The Iliad of Homer* (University of Chicago Press, 2011).

5 A. M. Turing, "Computing Machinery and Intelligence," *Mind* 49, no. 236 (October 1950): 433–460.

6 Weizenbaum, Joseph. "ELIZA-a Computer Program for the Study of Natural Language Communication between Man and Machine." *Communications of the ACM* 9, no. 1 (1966): 36–45. https://doi.org/10.1145/365153.365168.

7 Warren McCulloch and Walter Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bulletin of Mathematical Biophysics* 5, no. 4 (1943): 115–133, https://link.springer.com/article/10.1007/BF02478259

8 Frank Rosenblatt, "The Perceptron," *Psychological Review* 65, no. 6 (1958).

9 Max Harlow and Madhumita Murgia, "Who's Using Your Face? The Ugly Truth About Facial Recognition," April 19, 2019, https://www.ft.com/content/cf19b956-60a2-11e9-b285-3acd5d43599e.

10 "Face Challenges," National Institute of Science and Technology, June 10, 2015, https://www.nist.gov/programs-projects/face-challenges.

11 Karpathy, Andrej, Presentation at Scaled Machine Learning Conference, April 2020, https://youtu.be/hx7BXih7zx8

12 Kris Newby, "Compassionate Intelligence: Can Machine Learning Bring More Humanity to Health Care?" Stanford Medicine, Summer 2018, https://stanmed.stanford.edu/2018summer/artificial-intelligence-puts-humanity-health-care.html.

13 Loren Larsen, "HireVue Assessments and Preventing Algorithm Bias," HireVue, June 22, 2018, https://www.hirevue.com/blog/hirevue-assessments-and-preventing-algorithmic-bias.

14 Noel L. Hillman, "The Use of Artificial Intelligence in Gauging Risk of Recidivism," *The Judges' Journal*, January 1, 2019, https://www.americanbar.org/groups/judicial/publications/judges_journal/2019/winter/the-use-artificial-intelligence-gauging-risk-recidivism/.

15 Michael Eisen, "Amazon's $23,698,655.93 Book About Flies," It Is Not Junk, April 22, 2011, http://www.michaeleisen.org/blog/?p=358; John D. Sutter, "Amazon Seller Lists Book at $23,698,655.93 -- plus Shipping," CNN, April 25, 2011, http://www.cnn.com/2011/TECH/web/04/25/amazon.price.algorithm/index.html.

16 Victor Lyckerson, "Uber Agrees to Limit Surge Pricing During Emergencies, Disasters," *Time*, July 8, 2014, https://time.com/2967490/uber-agrees-to-limit-surge-pricing-during-emergencies-disasters/; Mike Isaac, "Uber Reaches Deal with New York on Surge Pricing in Emergencies," *Bits Blog*, July 8, 2014, https://bits.blogs.nytimes.com/2014/07/08/uber-reaches-agreement-with-n-y-on-surge-pricing-during-emergencies/.

17 Steve Lohr, "Facial Recognition Is Accurate, If You're a White Guy," *The New York Times*, February 9, 2018, https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.

18 Amy Harmon, "As Cameras Track Detroit's Residents, a Debate Ensues Over Racial Bias," *The New York Times*, July 8, 2019, https://www.nytimes.com/2019/07/08/us/detroit-facial-recognition-cameras.html.

19 Jacob Snow, "Amazon's Face Recognition Falsely Matched 28 Members of Congress with Mugshots," American Civil Liberties Union, July 26, 2018, https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28.

20 Ry Crist, "Amazon's Rekognition Software Lets Cops Track Faces: Here's What You Need to Know," CNET, March 19, 2019, https://www.cnet.com/news/what-is-amazon-rekognition-facial-recognition-software/.

21 Julia Angwin et al., "Machine Bias," ProPublica, May 23, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

22 Daisuke Wakabayashi, "Self-Driving Uber Car Kills Pedestrian in Arizona, Where Robots Roam," *The New York Times*, March 19, 2018, https://www.nytimes.com/2018/03/19/technology/uber-driver-lessfatality.html.

23 Aarian Marshall and Alex Davies, "Uber's Self-Driving Car Didn't Know Pedestrians Could Jaywalk," *Wired*, November 5, 2019, https://www.wired.com/story/ubers-self-driving-car-didnt-know-pedestrians-could-jaywalk/.

24 "Preliminary Report, Highway HWY18MH010," National Traffic Safety Board, accessed August 25, 2020, https://www.ntsb.gov/investigations/AccidentReports/Reports/HWY18MH010-prelim.pdf.

25 Philippa Foot, "The Problem of Abortion and the Doctrine of the Double Effect," *Oxford Review*, no. 5 (1967): 5–15.

26 Christian Sandvig et al., "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms," paper presented to "Data and Discrimination: Converting Critical Concerns into Productive Inquiry," a preconference at the 64th Annual Meeting of the International Communication Association, May 22, 2014, Seattle, WA, https://www.semanticscholar.org/paper/Auditing-Algorithms-%3A-Research-Methods-for-on-Sandvig-Hamilton/b7227cbd34766655dea10d0437ab10df3a127396?p2df.

27 Uriel J. Garcia and Katrina Bland, "Tempe Police Chief: Fatal Uber Crash Likely 'Unavoidable' for Any Kind Of Driver," *The Arizona Republic*, March 20, 2018.

28 Carolyn Said, "Uber Puts the Brakes on Testing Robot Cars in California After Arizona Fatality," *San Francisco Chronicle*, March 27, 2018. https://www.sfchronicle.com/business/article/Uber-pulls-out-of-all-self-driving-car-testing-in-12785490.php#:~:text=Uber%20plans%20to%20end%20all,Motor%20Vehicles%20to%20the%20company.&text=Uber's%20permit%20will%20expire%20Saturday%2C%20the%20letter%20said.

29 Tesla Motors, "Tesla Vehicle Safety Report," https://www.tesla.com/VehicleSafetyReport.

30 Loren Larsen, "HireVue Assessments and Preventing Algorithm Bias,"

31 "Statement on Algorithmic Transparency and Accountability," Association for Computing Machinery, U.S. Public Policy Council, January 12, 2017, https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf.

32 *Science* 06 May 1960:Vol. 131, Issue 3410, pp. 1355-1358. https://science.sciencemag.org/content/131/3410/1355.

33 William Butler Yeats, "Easter 1916," https://www.poetryfoundation.org/poems/43289/easter-1916.

34 Paul Mozur and Aaron Krolik, "A Surveillance Net Blankets China's Cities, Giving Police Vast Powers," - *The New York Times*," accessed May 14, 2020, https://www.nytimes.com/2019/12/17/technology/china-surveillance.html.

35 Kashmir Hill, "The Secretive Company That Might End Privacy as We Know It," *The New York Times*, December 17, 2019, https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facialrecognition.html.

36 Renee DiResta, "Health Experts Don't Understand How Information Moves," *The Atlantic*, May 6, 2020, https://www.theatlantic.com/ideas/archive/2020/05/health-experts-dont-understand-how-information-moves/611218/.

37 Andrea Fuller et al., "Google Hides News, Tricked by Fake Claims," *Wall Street Journal*, May 15, 2020, https://www.wsj.com/articles/google-dmca-copyright-claims-takedown-online-reputation-11589557001.

38 Chris Harrison, Human–Computer Interaction Institute, Carnegie Mellon University, www.chrisharrison.net/projects/InternetMap/high/worldBlack.png.

This page intentionally left blank

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Y

Z

This page intentionally left blank

VIDEO TRAINING FOR THE **IT PROFESSIONAL**

LEARN QUICKLY
Learn a new technology in just hours. Video training can teach more in less time, and material is generally easier to absorb and remember.

WATCH AND LEARN
Instructors demonstrate concepts so you see technology in action.

TEST YOURSELF
Our Complete Video Courses offer self-assessment quizzes throughout.

CONVENIENT
Most videos are streaming with an option to download lessons for offline viewing.

Learn more, browse our store, and watch free, sample lessons at
informit.com/video

Save 50%* off the list price of video courses with discount code **VIDBOB**

Pearson

informIT®
the trusted technology learning source

*Discount code VIDBOB confers a 50% discount off the list price of eligible titles purchased on informit.com. Eligible titles include most full-course video titles. Book + eBook bundles, book/eBook + video bundles, individual video lessons, Rough Cuts, Safari Books Online, non-discountable titles, titles on promotion with our retail partners, and any title featured as eBook Deal of the Day or Video Deal of the Week is not eligible for discount. Discount may not be combined with any other offer and is not redeemable for cash. Offer subject to change.

Business & Marketing Books, eBooks & Video

Pearson and FT Press, an imprint of Pearson, publish content from the world's best minds on the most important business and management topics. InformIT is the online presence of Pearson imprints focused on professional workforce topics including IT, business applications, and leadership.

Visit **informit.com/business** to shop and read sample chapters from featured products including these bestsellers.

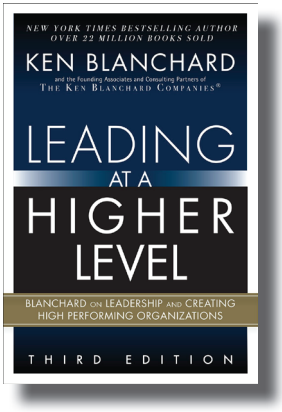

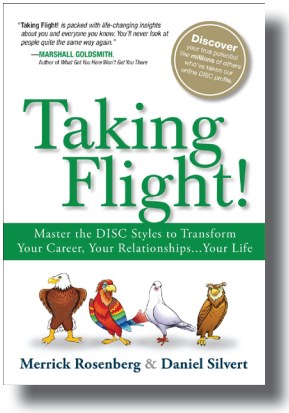

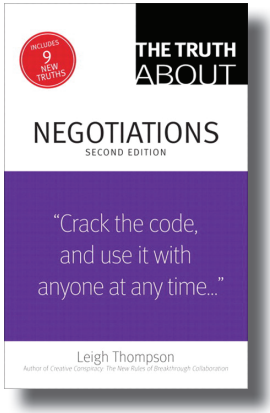

Pearson

informIT.com
the trusted technology learning source

Register Your Product at informit.com/register

Access additional benefits and **save 35%** on your next purchase

- Automatically receive a coupon for 35% off your next purchase, valid for 30 days. Look for your code in your InformIT cart or the Manage Codes section of your account page.
- Download available product updates.
- Access bonus material if available.*
- Check the box to hear from us and receive exclusive offers on new editions and related products.

**Registration benefits vary by product. Benefits will be listed on your account page under Registered Products.*

InformIT.com—The Trusted Technology Learning Source

InformIT is the online home of information technology brands at Pearson, the world's foremost education company. At InformIT.com, you can:

- Shop our books, eBooks, software, and video training
- Take advantage of our special offers and promotions
- Sign up to receive special offers and monthly newsletter
- Access thousands of free chapters and video lessons

Connect with InformIT—Visit informit.com/community

informIT®
the trusted technology learning source

Addison-Wesley · Adobe Press · Cisco Press · Microsoft Press · Pearson IT Certification · Prentice Hall · Que · Sams · Peachpit Press

Pearson