

MORE PRAISE FOR  
*When Everyone Knows That Everyone Knows . . .*

“Think you know what others think about what you are thinking? It turns out you’re probably wrong. And thanks to *When Everyone Knows That Everyone Knows . . .*, now we know why. Once you read this book, you’ll never view human behavior quite the same way again.”

—Jonah Berger, author of *Contagious* and *The Catalyst*

“An expository masterpiece. Steven Pinker explains, with beautiful clarity, how common knowledge is critical to successful human interaction.”

—Eric Maskin, Nobel laureate in economics and professor at  
Harvard University

“A masterful look behind the curtain at the calculations that propel us forward. With his brilliant knack for exposing what we take for granted, celebrated cognitive psychologist Steven Pinker explores, among many phenomena, our very human tendency to reveal information strategically, only letting others see what we want them to see.”

—Annie Duke, bestselling author of *Thinking in Bets* and *How to Decide*







**ALSO BY STEVEN PINKER**

*Language Learnability and Language Development*

*Learnability and Cognition*

*The Language Instinct*

*How the Mind Works*

*Words and Rules*

*The Blank Slate*

*The Stuff of Thought*

*The Better Angels of Our Nature*

*Language, Cognition, and Human Nature: Selected Articles*

*The Sense of Style*

*Enlightenment Now*

*Rationality*

**EDITED BY STEVEN PINKER**

*Visual Cognition*

*Connections and Symbols* (with Jacques Mehler)

*Lexical and Conceptual Semantics* (with Beth Levin)

*The Best American Science and Nature Writing 2004*



# WHEN EVERYONE KNOWS THAT EVERYONE KNOWS...

COMMON KNOWLEDGE AND  
THE MYSTERIES OF MONEY,  
POWER, AND EVERYDAY LIFE

STEVEN PINKER

**Scribner**

New York Amsterdam/Antwerp London  
Toronto Sydney/Melbourne New Delhi



Scribner  
An Imprint of Simon & Schuster, LLC  
1230 Avenue of the Americas  
New York, NY 10020

For more than 100 years, Simon & Schuster has championed authors and the stories they create. By respecting the copyright of an author's intellectual property, you enable Simon & Schuster and the author to continue publishing exceptional books for years to come. We thank you for supporting the author's copyright by purchasing an authorized edition of this book.

No amount of this book may be reproduced or stored in any format, nor may it be uploaded to any website, database, language-learning model, or other repository, retrieval, or artificial intelligence system without express permission. All rights reserved. Inquiries may be directed to Simon & Schuster, 1230 Avenue of the Americas, New York, NY 10020 or [permissions@simonandschuster.com](mailto:permissions@simonandschuster.com).

Copyright © 2025 by Steven Pinker

All rights reserved, including the right to reproduce this book or portions thereof in any form whatsoever. For information, address Scribner Subsidiary Rights Department, 1230 Avenue of the Americas, New York, NY 10020.

First Scribner hardcover edition September 2025

SCRIBNER and design are trademarks of Simon & Schuster, LLC

Simon & Schuster strongly believes in freedom of expression and stands against censorship in all its forms. For more information, visit [BooksBelong.com](http://BooksBelong.com).

For information about special discounts for bulk purchases, please contact Simon & Schuster Special Sales at 1-866-506-1949 or [business@simonandschuster.com](mailto:business@simonandschuster.com).

The Simon & Schuster Speakers Bureau can bring authors to your live event. For more information, or to book an event, contact the Simon & Schuster Speakers Bureau at 1-866-248-3049 or visit our website at [www.simonspeakers.com](http://www.simonspeakers.com).

*Interior design by Karla Schweer*

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data has been applied for.

ISBN 978-1-6680-1157-7

ISBN 978-1-6682-1258-5 (Int Exp)

ISBN 978-1-6680-1160-7 (ebook)

TO MY STUDENTS



## CONTENTS

Preface	xiii
CHAPTER 1 The Emperor, the Elephant, and the Matzo Ball	1
CHAPTER 2 Common Knowledge and Common Sense	43
CHAPTER 3 Fun and Games	77
CHAPTER 4 Reading the Mind of a Mind Reader	105
CHAPTER 5 The Department of Social Relations	153
CHAPTER 6 Laughing, Crying, Blushing, Staring, Glaring	193
CHAPTER 7 Weasel Words	231
CHAPTER 8 The Canceling Instinct	263
CHAPTER 9 Radical Honesty, Radical Hypocrisy	283
Notes	301
References	324
Index	344



## PREFACE

As a cognitive scientist, I have spent my life thinking about how people think. So the ultimate subject of my fascination would have to be how people think about what other people think, and how they think about what other people think they think, and how they think about what other people think they think they think. As dizzying as this cogitation may seem, we engage in it every day, at least tacitly, and in the limit this state of awareness has a technical name, *common knowledge*.

Originating in game theory and philosophy, the theory of common knowledge can illuminate a vast range of puzzles about human social life. I first came across it through my interest in language when writing *The Stuff of Thought*. I had long wondered why people often don't say what they mean in so many words but veil their intentions in innuendo and doublespeak, counting on their listeners to read between the lines. The answer, I suggested, was that barefaced statements generate common knowledge but genteel euphemisms do not, and common knowledge is what ratifies or annuls social relationships.

In this book I'll expand on that theory and show how common knowledge also explains fundamental features of societal organization, such as political power and financial markets; some of the design specs of human nature, such as laughter and tears; and countless curiosities of private and public life, such as bubbles and crashes, road rage, anonymous donations, long goodbyes,

revolutions that come out of nowhere, social media shaming mobs, and academic cancel culture. By the time you finish the book I hope you'll be equipped to understand phenomena I never got around to explaining, such as gaslighting, Kardashian celebrity (being famous for being famous), mock outrage ("I'm shocked, shocked to find that gambling is going on in here"), "red lines" in international relations, and the psychological difference between "cc" and "bcc" in email etiquette.

This oceanic scope, I hope to convince you, is not a sign of explanatory megalomania. Common knowledge really is that powerful a concept. It is the mental feat that explains one of the hallmarks of the human condition: individual minds can coordinate their choices for mutual benefit, allowing our species to thrive in collectives ranging from couples to societies. Many of our harmonies and discords, I hope to show, fall out of our struggles to achieve, sustain, or prevent common knowledge.



This is the second of my popular books to feature my own research, and as with *Word and Rules* I dedicate it to the graduate and post-doctoral students who collaborated with me on the studies. Every professor knows that the best part of the job is learning from students, and while pursuing this research I was fortunate to have learned from Julian De Freitas, Peter DeScioli, Omar Sultan Haque, Moshe Hoffman, Yuhui Huang, James Lee, Miriam Lindner, Maxim Massenkoff, Jason Nemirow, Laura Niemi, Lawrence Ian Reed, Kyle Thomas, and Dylan Tweed. Special thanks to Peter DeScioli for his sharp and deep comments on the first draft.

Several people provided patient guidance in their areas of expertise:

Scott Aaronson, Robert John Aumann, Herbert Clark, Peter DeScioli, Rebecca Goldstein, Dacher Keltner, Eric Maskin, Dov Samet, and Jeannie Suk Gersen. Others commented on drafts of chapters: Charleen Adams, Cory Clark, Tyler Cowen, Alan Fiske, Komi Frey, Robin Hanson, Moshe Hoffman, Greg Lukianoff, Michael Macy, Jason Nemirow, Bruce Schneier, Dan Sznycer, and Jessica Tracy. Still others answered questions or offered suggestions: Paul Bloom, Yi-Chia Chen, Eve Clark, Jeffry Frieden, Bill Gates, Andrew Gelman, Joshua Goldstein, Marc Hauser, Coleman Hughes, Jillian Jordan, Peter Kinderman, Gary King, Sarah Kious, Louis Liebenberg, Lucy Matthew, Dani Passow, Dan Schacter, Fred Shapiro, Richard Shweder, Lawrence Summers, Philip Tetlock, Jeffrey Watumull, David Wolpe, and Hirschy Zarchi. Bob Woods provided skilled assistance with bibliographic and research materials. Ilavenil Subbiah designed the elegant graphics. I thank all of them.

I am grateful to my editor at Scribner, Rick Horgan, for his encouragement and guidance, and to my friend and literary agent John Brockman. Special appreciation goes to Katya Rice for copyediting the manuscript, our tenth collaboration over forty years.

Three of the dedicatees of my previous books died as I wrote this one. This is my first trade book not to have been read in draft by my mother and primary imagined reader, Roslyn Wiesenfeld Pinker. Gone too are my dear friends and intellectual inspirations John Tooby and Donald Symons. They all left a stamp on the book in their examples, ideas, and voices. Also deeply missed are two other influences, Daniel Dennett and Daniel Kahneman.

My greatest intellectual inspiration is also my life partner, and I thank Rebecca Newberger Goldstein for continually showing me what matters, intellectually and personally. It's a pleasure to acknowledge the rest of my loving family: Yael, Solly, Danielle, Kai, Susan, Martin, Eva, Carl, Eric, Rob, Kris, Jack, and David.



# The Emperor, the Elephant, and the Matzo Ball

*What common knowledge is,  
and why it matters*

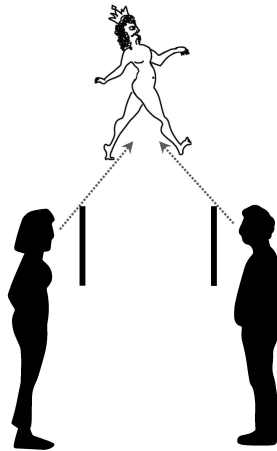
**W**hen the little boy said the emperor was naked, he wasn't telling anyone anything they didn't already know. But he added to their knowledge nonetheless. By blurting out what every onlooker could see within earshot of the others, he ensured that they now knew that everyone *else* knew what they knew, that everyone knew that everyone knew that, and so on. And that changed their relationship to the emperor, from obsequious deference to ridicule and scorn.<sup>1</sup>

Hans Christian Andersen's immortal story draws on a momentous logical distinction. With *private knowledge*, person A knows something, and person B knows it. With *common knowledge*, A knows something, and B knows it, but in addition, A *knows* that B knows it, and B knows that A knows it. On top of that, A knows that B

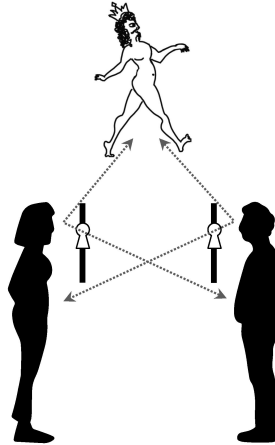
knows that A knows it, and B knows that A knows that B knows it, and so on, ad infinitum.<sup>2</sup>

“The Emperor’s New Clothes” dramatizes two features of common knowledge that make it not just a mind-blowing logical concept but a key to understanding human social life. One is that common knowledge need not be deduced from an infinite chain of musings about other people’s mental states (“I know that you know that I know that you know . . .”), which no mortal could ever think. It can be instantly imparted by a conspicuous event, like a plain sentence uttered in public. The other is that the difference between private knowledge, even when widely shared, and common knowledge is not a mere logical nicety but can unify knowers in coordinated action and sometimes explode a social status quo.

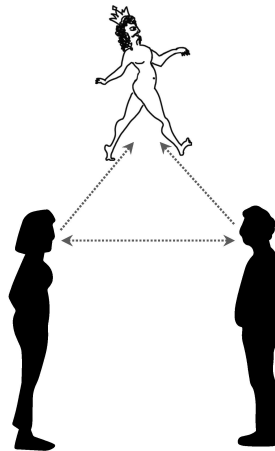
To help distinguish the different kinds of knowledge, let’s visualize them in little cartoons that depict knowing as seeing. The first is a picture of private knowledge. Each observer sees something, but neither sees the other seeing it:



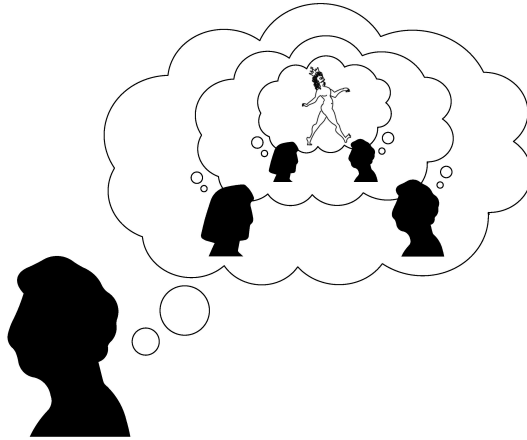
Next we have a state we will call reciprocal knowledge, in which each observer sees the event and sees the other seeing it. But because each espies the other through the anonymity of a keyhole, their awareness falls short of common knowledge; neither one knows that they have been seen seeing it:



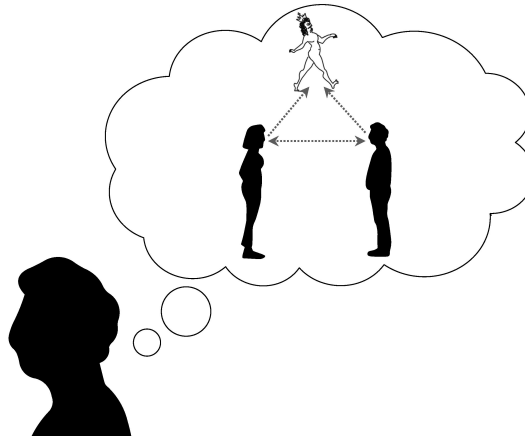
The last cartoon shows common knowledge. Each sees the other seeing the event, and sees the other seeing them seeing it, from which they can infer that each can see as many seeings of seeings as they care to ponder:



What is going on in the observers' minds when they are in a situation that provides common knowledge? It needn't be a hall of mirrors with "He knows that she knows that he knows that she knows . . .":



Our heads start to spin with just two layers of thoughts within thoughts, and common knowledge requires an infinite number of them, which can't fit inside a finite skull. Most likely it's the simple intuition that the event is "public" or "out there" or "there for all to see":



My goal in this book is to explain the obscure but momentous research on common knowledge, together with some ideas of my own, and show how the concept illuminates many enigmas of our public affairs and personal lives. The scholars I know who have worked on the concept—mathematicians, economists, philosophers, linguists, computer scientists—agree that it is a keystone in understanding the social world. But they have struggled to spread the news to a wide readership, and they often wonder how human minds handle what seems like an impossibly abstruse state of knowing. This book, written from the vantage point of psychology and cognitive science, aims to bridge these gaps.<sup>3</sup>

Here are the main ideas. First, common knowledge (in the technical sense) is logically different from private knowledge: learning about something in public, even if everyone already knows it, can change everything. Second, the main thing that common knowledge changes is the ability to *coordinate*: two or more people with common knowledge can benefit each other with complementary choices that they would have no confidence making with private knowledge. Third, because common knowledge is so potent, humans are intuitively sensitive to it, almost as if we had a sense organ for this logical concept. Fourth, this awareness is what has empowered our species to coordinate our behavior in social networks like communities, economies, and nations. For this reason, many peculiarities of public life—its mindless rituals, conventions, and norms—become intelligible as solutions to coordination problems. So do some of the pathologies of public life, including fads, mobs, panics, bubbles, and spirals of silence. Fifth, personal relationships—our bonds with family, friends, lovers, authorities, subordinates, neighbors, colleagues, and transactional partners—are also coordination games, and they, too, must be cemented by common knowledge. Sixth, because all of these coordination equilibria come with perquisites and obligations,

we often find ways to work around them by *preventing* ourselves from knowing what everyone knows, giving rise to rituals of benign hypocrisy, pretending not to know, catching someone's drift, and not going there. In other words, many of our tensions, personal and political, arise from the desire to propagate or suppress common knowledge.



Because common knowledge is the book's touchstone, I must begin with a word about the term itself, which I will be using in a different sense from its everyday meaning. In ordinary English, the expression *common knowledge* refers to something that many or most people know, especially open secrets, as in "It's common knowledge that the police around here can be bribed." This is almost the opposite of the meaning of the technical term from game theory and philosophy that we will explore in this book. Just as confusingly, *common knowledge* in the technical sense does not correspond to the literal meaning of its words, namely "knowledge that people have in common," since that could pertain to identical private knowledge, where everyone knows the same thing without necessarily knowing that the others know it.

A more transparent term would be *mutual knowledge*, with its implication of people deliberately mirroring or pooling their knowledge. Linguists sometime use that term when referring to the common knowledge between a pair of conversational partners, and that's what I called it in *The Stuff of Thought*.<sup>4</sup> But in most of the technical literature mutual knowledge has drifted into a different meaning, either widespread private knowledge (like the first diagram) or layers of reciprocal knowledge (like the second).<sup>5</sup> Other terms for common

knowledge include *open knowledge*, *conspicuous knowledge*, *public knowledge*, *interactive knowledge*, *shared reality*, *shared awareness*, *collective consciousness*, and *common ground*. But the term *common knowledge*, however misleading, has become entrenched among the experts, and that's the one I will stick with.

Happily, this concession provides an opportunity to introduce the power of common knowledge in human affairs, starting with language itself. When it comes to established vocabulary, logic and grammar are beside the point. No one cares that *awful* no longer means “filled with awe,” or that *bathroom* needn't mean “a room with a bath.” And as Voltaire quipped, the Holy Roman Empire was neither holy nor Roman nor an empire. A word or fixed phrase conveys a meaning not because people deduce it from the word's parts but simply because they expect everyone else to interpret it the way they do.<sup>6</sup>

The purpose of language is to *coordinate* our behavior—you pass me the pepper when I want pepper and the salt when I want salt. Language allows us to do this because it is a *convention*, a tacit agreement among the members of a language community to use words to refer to certain concepts, in this case that the sound [ˈpepɪ] will be used to indicate pepper and [sɒlt] to indicate salt, though it could have been the other way around if English speakers long ago had agreed on the opposite pairing. I can request the salt with confidence because I know you interpret the word as I do, and crucially, you know that I know this—if you knew that *salt* meant “salt” but thought that I thought *salt* meant “pepper,” you'd pass the pepper instead. And I know that you know that I know this, ad infinitum.

Ever since the linguist Noam Chomsky called attention to the intricacies of syntax, human language has inspired awe because of the vast number of meanings we can express by combining words into sentences.<sup>7</sup> But the power of a single word to coordinate minds is just as awesome. As the poet Craig Morgan Teicher writes, “To speak

is an incomparable act of faith. What proof do we have that when I say *mouse*, you do not think of a stop sign? The obvious response to such a question is that whoever asks it is thinking too hard about a soft thought.”<sup>8</sup> The less obvious response is that the meaning of a word is common knowledge among the speakers of the language. Children tacitly make this assumption from the start; it’s hard to see how they could master language if they had to worry that *mouse* meant “mouse” to them but “stop sign” to someone else and still other things to other people.<sup>9</sup> Experiments by the psychologists Gil Diesendruck and Lori Markson showed that indeed children don’t worry. They taught three-year-olds words for unfamiliar objects, like *mef* for a dumpling press, and found that the children immediately assumed that a stranger knew what the word meant. It’s not because children indiscriminately blur their own knowledge with everyone else’s: when the three-year-olds were taught a new *fact* about an object (such as that “my cat likes to play with it”), they didn’t assume that a stranger knew the fact.<sup>10</sup> Words are the earliest and most omnipresent exercise of common knowledge in our lives.



The idiom “to be on speaking terms” reminds us that language is the quintessential social activity, and the logic of linguistic conventions opens the door to questions about the rest of our sociality. The most basic is why we are social in the first place. Humans chat, work, play, build, and learn in ensembles, whether they are related or not, a rarity among animals, who are mainly held together by ties of blood. What are the evolutionary advantages of hanging out together so that one person can benefit another?

Evolutionary biologists think about this question by distinguishing

the two possibilities for what's in it for the helper. When one organism benefits the other at a cost to itself, that's called *altruism*. Readers of Richard Dawkins's 1976 classic *The Selfish Gene* or the dozens of books on cooperation that have appeared in its wake know that altruism is a major puzzle in biology because at first glance it seems that it could never have evolved by natural selection. Why do monkeys groom each other, each sacrificing time to pick parasites off the other, when a selfish monkey could enjoy being groomed without grooming in turn, outcompeting its generous troopmates and flooding descendant generations with its selfish genes, eventually driving grooming to extinction?

The common solution to this puzzle is reciprocal altruism, the strategy of starting out by cooperating (in this case, grooming another upon request), and thereafter doing unto others as they have done unto you: cooperate with those who cooperated, and defect (in this case, refuse to groom) against those who defected.<sup>11</sup> The problem can be modeled in game theory as a Prisoners' Dilemma, the hypothetical scenario in which two incommunicado partners in crime have no choice but to betray each other out of fear of being betrayed, leaving them both worse off than if they had cooperated. (More on this in chapter 3.) When the partners are placed in the dilemma repeatedly, strategies of reciprocal altruism (playing "Tit for Tat") can outcompete the exploitative ones, since cheaters will eventually be excluded from beneficial cooperation.<sup>12</sup>

Psychologists have pointed out that several of our mental faculties—our memory for other people and what they did to us, our sense of fairness, and our moral sentiments such as sympathy, gratitude, and anger—seem uncannily designed to implement a strategy of reciprocity, and presumably evolved as adaptations to the problem of altruistic cooperation. Not, in our case, grooming, but the myriad ways in which we trade goods, services, and favors, like bartering,

carpooling, and babysitting. Sympathy impels us to cooperate on the first move, gratitude to repay cooperation with cooperation, anger to repay defection with defection.

By now this is a familiar story; I myself have told it in five books.<sup>13</sup> It absorbs our attention because it resolves an evolutionary paradox and because it shines a light on a major theater of the human condition: our dramas of fairness, debt, obligation, exchange, guilt, appreciation, and treachery. Only recently have I come to appreciate that the story of cooperation makes up just one side of the problem of what makes humans social. The other side is *coordination*.<sup>14</sup>

When one organism benefits another, it doesn't necessarily incur a cost to itself; it may enjoy a benefit. Biologists call this second kind of helping *mutualism*, as when an oxpecker bird eats ticks off the back of a willing zebra. The oxpecker gets a meal, the zebra is tormented by fewer pests, and everyone wins (except the ticks). Reciprocity is unnecessary: the oxpecker doesn't demand that the zebra nibble ticks out of its feathers in repayment. For this reason mutualism would not seem to carry the frisson of altruism. Each party clearly gets something out of the relationship, so each has an incentive to allow the arrangement to evolve. It's not "If you scratch my back, I'll scratch yours"; it's "One hand washes the other."

But the evolution of mutualistic coordination is by no means boring. It raises another daunting evolutionary puzzle, with a different but equally fecund explanation.

Life is filled with opportunities to coordinate with other people for mutual gain. We agree on a time and place to meet, bring complementary fare to a potluck dinner, divide responsibilities on a project, dub a meeting room with a nickname, and carry opposite ends of a heavy couch up the stairs. As with the oxpecker and the zebra, there's no incentive for anyone to cheat or to fear being cheated: when

coordination works, everyone wins. This doesn't mean that it's easy to bring it about. Coordination can fail if people are not on the same page, even when they want the same thing. Schedules clash, signals get crossed, and shared goals fall through the cracks or are spoiled by too many cooks.

Consider a game that boils coordination down to its bare logic the way the Prisoners' Dilemma does for cooperation.<sup>15</sup> In *Rendezvous*, two people, James and Charlotte, enjoy each other's company and have agreed to meet for coffee, but James's cell phone goes dead before they have settled on a place. They both know that James tends to frequent the Java Joint and Charlotte usually patronizes the Coffee Connection, but neither has a real preference; they just want to end up at the same place. James predicts that Charlotte will gravitate to the Connection, so he heads there, but then realizes that she will predict that he will gravitate to the Joint, so he changes course and heads there, until he realizes that she will anticipate that he will guess that she will opt for the Connection, so he does another about-face, only for it to dawn on him that it will occur to her that he knows she is aware that he haunts the Joint, so he pirouettes once again. Meanwhile, she is whipsawed by the same futile empathy.

Note that in the game of *Rendezvous* there is no conflict of interest: the two friends want the same thing. Trust and mistrust, generosity and selfishness, honesty and deception, good deeds enjoyed and repaid, simply don't arise. James and Charlotte's problem is not motivational but cognitive. What they long for is common knowledge. It's not enough for one of them to know the other's likely intention. Each must know that the other knows what the first one knows, *ad infinitum*.

The easiest bestower of common knowledge, direct speech, is unavailable to them. But all is not lost. The next-best thing is *common salience*, also known as a *focal point*. Suppose that the Java

Joint is running a promotion and has plastered the local kiosks with ads, or the café has come up in their previous conversation, or it was recently in the news, or it's situated at the busiest intersection of the town. None of these is inherently a "good reason" to meet at the Joint, but the mere fact that a location is likely to intrude into the consciousness of each is reason enough, and they both can break their empathic impasse by heading there.

Whatever the source, when a coordination dilemma recurs in life, the parties will yearn for a focal point, any focal point, and are apt to stick with whichever solution is known to be a solution. These commonly known solutions are called *conventions*.<sup>16</sup> James and Charlotte, for instance, might adopt a personal convention to rescue them should they ever find themselves incommunicado again, such as going with the tiebreaker "Ladies first" and meeting at the Connection, or taking turns and going wherever they hadn't met the last time.

Among society-wide conventions, an obvious one we've seen is the vocabulary of a language. Other examples include closing businesses on Sunday, accepting paper currency in exchange for goods and services, using appliances that run on 110 volts, and driving on the right side of the road (or the left; it doesn't matter, as long as everyone sticks to the same side). It's irrelevant that this last convention is enforced by the police. As with many conventions, people have an incentive to conform to it as long as others do.

A sweeping narrative of the past and future of our species by the historian Yuval Noah Harari, told in his books *Sapiens*, *Homo Deus*, and *21 Lessons for the 21st Century*, is animated by a big idea, which he summarizes as follows: "Our world is built on fictions. They're all around us: nations, corporations and religion were invented in the human imagination. And if it weren't for the fictional stories we collectively tell, we might not be the dominant species on the planet."<sup>17</sup> He explains:

Ever since the stone age, self-reinforcing myths have served to unite human collectives. Indeed, *Homo sapiens* conquered this planet thanks above all to the unique human ability to create and spread fictions. We are the only mammals that can cooperate with numerous strangers because only we can invent fictional stories, spread them around, and convince millions of others to believe in them. As long as everybody believes in the same fictions, we all obey the same laws, and can thereby cooperate effectively.<sup>18</sup>

It's a worthy insight, though I would put it differently. Our world is built on *conventions* that allow us to *coordinate* effectively and are self-reinforcing because they are *common knowledge*. Conventions like the English language, Christianity, the United States of America, the euro, and Microsoft are not exactly "fictions." They are very real, even if they are not made out of physical stuff. Common knowledge creates nonphysical realities.<sup>19</sup>

While the dilemma of cooperation sets the stage for the human drama of beneficence, exploitation, and fairness, the dilemma of coordination sets the stage for its own operas, with storylines driven by privacy, publicity, precedent, fame, fads, norms, panics, rituals, piety, and outrage. The rest of this chapter will try to convince you of this ambitious claim, visiting four examples from the news that are best understood with the logic of common knowledge. I'll conclude with a peek at how the logic is experienced in our consciousness and conversation.



We have become used to social media proliferating posts about celebrity indignities, cats riding Roombas, and blue and black dresses that

some people see as gold and white.<sup>20</sup> But rarely does a viral meme originate in a math problem. That happened in 2015, when Kenneth Kong, a Singaporean television host, posted a photo of an exam question for the country's fifth-grade students on his Facebook page. The brainteaser soon took the internet by storm, attracting coverage by the *New York Times*, the *Guardian*, and the BBC, and soon getting its own Wikipedia entry.<sup>21</sup> Here it is verbatim, questionable tenses and all:

Albert and Bernard just become friends with Cheryl, and they want to know when her birthday is. Cheryl gives them a list of 10 possible dates:

May 15, May 16, May 19

June 17, June 18

July 14, July 16

August 14, August 15, August 17

Cheryl then tells Albert and Bernard separately the month and the day of her birthday respectively.

ALBERT: I don't know when Cheryl's birthday is, but I know that Bernard doesn't know too.

BERNARD: At first I don't know when Cheryl's birthday is, but I know now.

ALBERT: Then I also know when Cheryl's birthday is.

So when is Cheryl's birthday?

Try it—it's a small mental workout, but solvable without math. It helps to arrange the possibilities in a grid and work by a process of elimination.

		Day (Bernard knows)			
Month  (Albert knows)	May	15	16		19
	June			17	18
	July	14		16	
	August	14	15		17

Bernard, recall, has been told the day but not the month—that is, he knows which column Cheryl’s birthday is in but not which row. Now, if he had been told her birthday was “the 18th,” he’d know it was *June* 18, since that’s the only 18 among the possibilities (it’s the only date in its column). Similarly, if he had been told it was the 19th, he’d know it was May 19. Albert says he knows that Bernard doesn’t know the answer, so he must know that Bernard wasn’t told it was the 18th or 19th, the two sure things. But how could Albert have figured that out just by knowing the correct row, namely the birth month? He’d know only if he had been told it was July or August, which lack those easy solutions.

Now, if you and I can figure this much out, so can Bernard. And Bernard has announced that he now knows when Cheryl’s birthday is. How could he know this? Well, the day couldn’t have been the 14th, because that would not have allowed him to decide between July 14 and August 14. And if it had been the 15th or the 17th, then Albert would have no way of knowing which it was.

But Albert announces that he does know. That means it could only have been the month with a single day left among the possibilities, July, and that day is the 16th. Cheryl’s birthday is July 16.

What does the Cheryl problem tell us? For one thing, it bears the hallmark of viral social media posts: moralization about a story which turns out to be false. Commenters moaned either

about the draconian workouts inflicted on Asian schoolchildren or about the intellectual flabbiness of their Western counterparts. But in reality the problem did not come from a fifth-grade curriculum; it was drawn from an Olympiad for the nation's math-savviest teenagers.

The real reason I bring up Cheryl is to show how knowledge about knowledge is logically different from mere knowledge, and how it can be used to infer facts about the world. It also shows that people are capable—with some effort when it comes to unfamiliar situations—of thinking about other people's thoughts about still other people's thoughts. It's a talent that goes beyond the thought process that cognitive scientists call mentalizing, mind reading, intuitive psychology, or theory of mind (the "theory" here referring to an ordinary person's intuitions, not to the scientist's own theorizing). This talent can be called *recursive* mentalizing: thinking about thoughts about thoughts; reading the mind of a mind reader. The Cheryl problem calls upon us to think about Albert and Bernard's private knowledge of the birth month and day, their reciprocal knowledge of each other's private knowledge, and the common knowledge conveyed by their pronouncements.

Later in the book we'll ask how good people really are at thinking about thinking about thinking. Before that we'll exercise the ability to use states of knowledge to whittle down states of the world and will come to an astonishing conclusion: that rational, honest people cannot agree to disagree.



Those of us who have lived through the steady democratization of computers, from room-filling mainframes to refrigerator-sized

minicomputers to desktop PCs to laptops to smartphones, recall that the sharpest turning point came in 1984. That's when Apple introduced an affordable personal computer with a graphical user interface (GUI), the Macintosh. Before that time, home computers were difficult and tedious to use. Their monitors displayed twenty-four rows of eighty characters each, and their operating systems required textual commands like "rmdir c:\foobar" whose syntax had to be memorized and which could fail with an errant keystroke. Early adopters of Apple's alternative were dazzled by the windows, icons, menus, and mouse which today we take for granted. Richard Dawkins marveled at the time: "I have been an intensive programmer and user of a wide variety of digital computers for twenty-five years, and I can testify that using the Macintosh (and its imitators) is a qualitatively different experience from using any earlier type of computer. There is an effortless, natural feel to it, almost as if the virtual machine were an extension of one's own body."<sup>22</sup>

Yet the takeover by the GUI computers was not a foregone conclusion. They had been invented a decade earlier at Xerox Palo Alto Research Center—according to legend, Steve Jobs stole the idea after a visit in 1979—and Apple itself had introduced a version in 1983, the Lisa, which flopped. The problem was how to jump-start the mass acceptance of an exotic new computing platform. Enough people had to buy one that the price could come down, communities of users could share software and expertise, and a market could develop for third-party peripherals, apps, and consumables like floppy disks. Until they materialized, few consumers would take a chance at buying a computer, however "insanely great," that might leave them as oddballs and orphans. But how could Apple sell enough units to create those "network externalities," as economists call them, if no one was willing to buy one until it did?

Apple cut the knot with an ad that ran only once, during the third quarter of Super Bowl XVIII.<sup>23</sup> Directed by Ridley Scott of *Alien* and *Blade Runner* fame, it said nothing about windows, icons, menus, or mice; nothing about an effortless natural feel or an extension of one's own body; nothing about the product at all. A line of sackcloth-clad drones, prodded by riot police, trudges into a cavernous hall where a projected face intones corporate drivel about "information purification directives." The blue-gray monochrome is intercut with shots of a lithe blond woman in red shorts and a tank top sprinting into the hall with a large mallet. She hammer-throws it into the screen, which explodes in a white fireball, leaving the drones gaping. A voice-over reads the message scrolling up the screen: "On January 24th, Apple Computer will introduce Macintosh. And you'll see why 1984 won't be like '1984.'" <sup>24</sup>

The contrast between corporate conformity and youthful iconoclasm, though certainly an image Apple was happy to convey, was not the real point of the ad. The real point was that it captured people's attention during the Super Bowl, which is an American national rite, standing out in the calendar like a religious holiday. Not only were a lot of people watching, but everyone knew that a lot of people were watching, and knew that everyone knew it. A Super Bowl ad generates common knowledge. And common knowledge is necessary to entrench a convention, like the hardware and software specs of the Macintosh computer. The tens of millions of people who saw the ad knew that tens of millions of people might be intrigued by this upstart technology.

The special role of the Super Bowl ad as an instant common-knowledge creator was explained by the political scientist Michael Chwe in his 2001 book *Rational Ritual*. Chwe noted that other startups that depended on network externalities also advertised heavily on the Super Bowl, especially during the dot-com era when creating those

network effects was the key to success. Monster.com, for example, was one of the first job sites that took advantage of the vast reach of the Web, but it needed job seekers expecting employers to post ads there and employers expecting job seekers to seek jobs there. The Discover card boasted high credit limits, no annual fees, and cashback bonuses, but it was unappealing without a network of merchants who accepted it, which they would do only if there were a network of cardholders who used it.

Chwe argued that the Super Bowl also attracts a second kind of company: those selling products that depend on their brand image. American beer is American beer, and running shoes are running shoes, but it matters to consumers that they be seen as Budweiser or Miller drinkers, Nike or Adidas wearers, and certainly not as buyers of no-name generics. Also, some products are enjoyed communally, like restaurants, plays, movies, and books; it's more fun to see a movie if you can join a conversation about it with friends later. These social products are like technology standards or credit card networks: the more people adopt it, the more people want to adopt it. Chwe confirmed that goods like these which are consumed in public, like cars, beer, sodas, movies, clothing, and shoes, are more heavily advertised on the Super Bowl than goods consumed in private, like batteries, motor oil, and breakfast cereal.

Of course, the Super Bowl (and other heavily hyped happenings) have massive audiences, so maybe it's sheer numbers rather than common knowledge that appeal to advertisers. Using statistical regression techniques, Chwe did his best to control for this and other confounding factors. He confirmed that companies that sold products consumed in public were more willing to advertise, were more willing to advertise on popular shows than on niche ones, and were willing to pay more *per viewer* than companies that sold products consumed privately.

In the decades since the “1984” ad, Super Bowl ads have become a cultural spectacle, attracting almost as much attention as the game itself. This has made the ads all the more appealing for hawkers of products whose value depends on common knowledge. The climax may have been Super Bowl LVI in 2022, sometimes called the Crypto Bowl because of its plethora of high-concept ads featuring cryptocurrency exchanges, the sites or apps on which people could buy and sell cryptocurrencies like Bitcoin.

It’s not that cryptocurrencies themselves depend on common knowledge manufactured by advertising. They do depend on one kind of common knowledge, to be sure; all currencies do. I accept a green piece of paper in exchange for an old couch because I know that Stop & Shop will take that paper in exchange for some groceries, which they do because they expect that their wholesalers will accept it, and so on. At one time the common knowledge that currencies had value was reinforced by a promise that the government would exchange a dollar for a fixed amount of gold, which, according to the popular understanding, was stored in Fort Knox. Nowadays that knowledge is backstopped by government fiat, particularly the law that legal tender must be accepted to settle debts. When a government is monetarily stable, that common knowledge can become self-sustaining. People trust the American dollar because they know that everyone else does, and that the US Federal Reserve works hard to keep it that way. But when a government is unstable, the common knowledge can unravel and lead to hyperinflation, in which people furiously raise prices and wages because they see and anticipate others doing so, quickly rendering the currency worthless.

With cryptocurrency, the common knowledge is provided by a public ledger, the blockchain. In a blockchain, all transactions are indelibly recorded and protected from embezzlement or forgery by

complex cryptographic algorithms, which are transparent but impossible to hack. Everyone can see the blockchain and how it works, so there's no need for gimmicks to get people to expect others to expect that the cryptocurrency has value.

Where common knowledge has to be artificially ginned up is in cryptocurrency *speculation*. Like all currencies, the supply of crypto can grow only at a limited rate so as to prevent hyperinflation. This is accomplished by allowing people to “mine” it by solving math problems that require substantial computing time and power. If speculators anticipate that the demand will outrun this supply, fueled by the perception that crypto is the wave of the future, they might buy crypto now in hopes that they can sell it at a profit later. And that can work only if there are buyers out there who plausibly expect that *they* can sell it at a profit to still *other* buyers—“greater fools,” as investment analysts call them. What inflates a speculative bubble isn't exactly common knowledge, since there's nothing objectively to “know.” But it is a common *expectation*, which uses recursive mentalizing in a similar way: everyone expects that everyone expects something—in this case, that the asset price will rise, which does make it rise (at least for a while). Meanwhile, the exchanges skim a profit on each transaction.

A conspicuous public event like a Super Bowl ad can prop up the common expectation, and that is what the crypto exchanges were paying for in 2022. None of the exorbitant ads praised or even mentioned the virtues of cryptocurrency, such as confidentiality and protection against hyperinflation or government confiscation. Instead they paid celebrities to generate a common expectation that other people were investing in crypto, so they should too.

In one of the ads, Matt Damon, backdropped by mountain climbers, aviators, and astronauts, intoned, “In these moments of truth, these men and women—these mere mortals, just like you and me—as

they peer over the edge, they calm their minds and steel their nerves with four simple words that have been whispered by the intrepid since the time of the Romans: Fortune favors the brave.” In another, Larry David played incarnations of his cranky self at various turning points in history dismissing innovations like the wheel (“What does it do?” “It rolls.” “So does a bagel, OK? A bagel you can eat!”), the fork, the toilet, the vote, the lightbulb, and the moon program. It ended with his dismissing a representative of the FTX crypto exchange: “Ehhhh, I don’t think so. And I’m never wrong about this stuff—never!” The end tag reads, “Don’t be like Larry. Don’t miss out on crypto . . . on the next big thing.”<sup>25</sup>

Of course, it’s only so long that an asset can levitate in midair suspended by nothing but common expectation. Bubbles pop when a market starts running out of the greater fools who don’t want to miss out on the next big thing, or when a reason for doubt itself becomes common knowledge. That can upend the common expectation and send investors running for the exits, each desperate to sell the asset out of fear that others are selling it out of fear that still others are selling it.

That’s what happened a few months after the Crypto Bowl, when the value of Bitcoin sank by 75 percent and two trillion dollars of cryptocurrency value swiftly evaporated. Worse was to come. In November, a leaked balance sheet revealed that FTX had invested customer deposits in its own hedge fund, whose reserves were mostly held in a conjured-out-of-thin-air cryptocurrency, FTT. That led the CEO of Binance (a rival exchange that had also advertised on the Super Bowl) to sow doubt about the value of FTT and sell its own holdings, which FTX could not easily redeem, setting off a self-fulfilling common expectation of plunging value, that is, a bank run. Within weeks, FTX declared bankruptcy, its founder and CEO, Sam Bankman-Fried, was arrested and charged with wire fraud and