



SCOTT E. PAGE

# THE MODEL THINKER

what you need to know  
to make data work for you

## Copyright

Copyright © 2018 by Scott E. Page  
Cover design by Chin-Yee Lai  
Cover © 2018 Hachette Book Group, Inc.

Hachette Book Group supports the right to free expression and the value of copyright. The purpose of copyright is to encourage writers and artists to produce the creative works that enrich our culture.

The scanning, uploading, and distribution of this book without permission is a theft of the author's intellectual property. If you would like permission to use material from the book (other than for review purposes), please contact [permissions@hbgusa.com](mailto:permissions@hbgusa.com). Thank you for your support of the author's rights.

Basic Books  
Hachette Book Group  
290 Avenue of the Americas, New York, NY 10104  
[www.basicbooks.com](http://www.basicbooks.com)

First Edition: November 2018

Published by Basic Books, an imprint of Perseus Books, LLC, a subsidiary of Hachette Book Group, Inc. The Basic Books name and logo is a trademark of the Hachette Book Group.

The Hachette Speakers Bureau provides a wide range of authors for speaking events. To find out more, go to [www.hachettespeakersbureau.com](http://www.hachettespeakersbureau.com) or call (866) 376-6591.

The publisher is not responsible for websites (or their

# CONTENTS

*Cover*

*Title Page*

*Copyright*

*Dedication*

*Epigraph*

*Prologue*

1 The Many-Model Thinker

2 Why Model?

3 The Science of Many Models

4 Modeling Human Actors

5 Normal Distributions: The Bell Curve

6 Power-Law Distributions: Long Tails

7 Linear Models

8 Concavity and Convexity

9 Models of Value and Power

10 Network Models

11 Broadcast, Diffusion, and Contagion

12 Entropy: Modeling Uncertainty

- [13 Random Walks](#)
- [14 Path Dependence](#)
- [15 Local Interaction Models](#)
- [16 Lyapunov Functions and Equilibria](#)
- [17 Markov Models](#)
- [18 Systems Dynamics Models](#)
- [19 Threshold Models with Feedbacks](#)
- [20 Spatial and Hedonic Choice](#)
- [21 Game Theory Models Times Three](#)
- [22 Models of Cooperation](#)
- [23 Collective Action Problems](#)
- [24 Mechanism Design](#)
- [25 Signaling Models](#)
- [26 Learning Models](#)
- [27 Multi-Armed Bandit Problems](#)
- [28 Rugged-Landscape Models](#)
- [29 Opioids, Inequality, and Humility](#)

[\*About the Author\*](#)

[\*Notes\*](#)

[\*Bibliography\*](#)

[\*Index\*](#)

*To Michael D. Cohen*

*(1945–2013)*

It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience.

—ALBERT EINSTEIN

## Prologue

*To me success means effectiveness in the world, that I am able to carry my ideas and values into the world—that I am able to change it in positive ways.*

—Maxine Hong Kingston

This book began as the result of a chance meeting with Michael Cohen in 2005 near the flower garden in the mall adjacent to the University of Michigan’s West Hall. Michael, a scholar known for his generosity, made a comment that altered my teaching career. With a twinkle in his eyes, Michael said, “Scottie, I once taught a course called Introduction to Modeling for Social Scientists, based on a book written by Charles Lave and James March. You should resurrect the course. It needs you.”

It needed me? I returned to my office a little confused, so I chased down an old course syllabus. I discovered that Michael had misled me. The course did not need me. I needed it. I had been wanting to develop a course that would introduce students to the core ideas of complex systems—networks, diversity, learning, large events, path dependence, tipping points—that would be relevant to their daily lives and future careers. By teaching modeling, I could make students better thinkers while introducing them to complexity. I could teach them tools that would improve their abilities to reason, explain, predict, design, communicate, act, and explore

The course’s motivating idea would be that we must confront the complexity of the modern world with multiple models. At semester’s end, rather than see the world from a

particular angle, students would see the world through many lenses. They would be standing in houses with many windows, able to look in multiple directions. My students would be better prepared for the complex challenges before them—improving education, reducing poverty, creating sustainable growth, finding meaningful work in an age of artificial intelligence, managing resources, and designing robust financial, economic, and political systems.

The next fall, I resurrected the course. I contemplated rebranding it as *Thirty-Two Models That Will Turn You into a Genius*, but the culture at Michigan frowns on hyperbole, so I stuck with Michael's title: *An Introduction to Modeling*. Lave and March's book proved to be a brilliant introduction. However, modeling had made huge advances in the intervening decades. I needed an updated version that included models of long-tailed distributions, networks, rugged landscapes, and random walks. I needed a book that discussed complexity.

So I began to write. For two years, the ground proved rocky. My plow moved at a slow place. One spring day, I again ran into Michael, this time in the arch-way underneath West Hall. I had been questioning the course, which was now drawing twenty students. Were models too abstract for undergraduates? Should I teach a different course on a specific issue or policy domain? Michael offered up a smile, noting that any endeavor worth pursuing merited questioning. As we parted, Michael commented on the importance and value of helping people think clearly. He told me not to give up, that he took joy in my challenges.

In the fall of 2012, the ground under the course shifted. Vice Provost Martha Pollack asked me to teach an online version—what is now called a MOOC. With a tablet computer, a \$29 camera, and a \$90 microphone, *Model Thinking* was born. With assistance from too many people at Michigan, Coursera, and Stanford University to thank properly (a quick shout-out to Tom Hickey, who did yeoman's work), I reorganized my lectures into a form suitable for an online



course, dividing each subject into modules and removing all copyrighted material. With my dog Bounder as an audience, I taped and retaped lectures.

The first offering of Model Thinking drew 60,000 students. That number now approaches a million. The popularity of the online course led me to abandon the book. I thought the project unnecessary, but, over the next two years, my email inbox began to fill with requests for a book to complement the online lectures. Then Michael Cohen lost his battle with cancer, and I felt that I needed to finish the book. I reopened the manuscript folder.

Writing a book requires large blocks of time and spaces that allow for clear thought. The poet Wallace Stevens wrote, “Perhaps the truth depends on a walk around the lake.” I relied on a close analog: mind-clearing swims across Winans Lake, where my family spends our summers. Throughout the writing process, the continuous life I share with the love of my life, Jenna Bednar, our sons, Orrie and Cooper, and our enormous dogs, Bounder, Oda, and Hildy, has brought laughter, comfort, and opportunities—among them Orrie having one week to correct the penultimate draft’s mathematical errors and Jenna having two weeks to identify instances of unclear writing, logical flaws, and muddled thinking. As has been true of most of my written work, this manuscript might be best described as an original draft by Scott Page with substantial revision by Jenna Bednar.

During the seven-year period of writing this book, my children have transitioned from pre-teens to young adults. Orrie is now off to college. Cooper follows next year. In the interval between sketching the initial outline and submitting the final version, my family has consumed copious amounts of bibimbap, pasta carbonara, and oatmeal chocolate chip cookies, taken the saws and loppers to scores of fallen branches and limbs, repaired dozens of breaks in the backyard fence, embarked on numerous failed initiatives to reduce the entropy in the basement and garage, and wished and hoped for the ice on the lake to be suitable for skating.

We have also had to accept loss. Midway through the project, my mother, Marilyn Tamboer Page, died from a sudden heart attack while enjoying the bliss of her routine daily walk with her dog. Not a day goes by when I do not reflect on the love she showered on her family and the support she gave to others.

The book before you is as complete as it can be at this moment in time. Doubtless, new models will be created, and old models will find new uses creating gaps in this current offering. As I humbly send the manuscript out into the world, I feel that my efforts will have been repaid if you, the reader, find the models and ideas within to be useful and generative, and that you are able to carry them out into the world and change it in positive ways.

If one day, when sitting in some professor's or graduate student's office, preferably at a college or university in my beloved Midwest, I scan the bookshelves and find this book leaning, as it has during its writing, on a well-worn copy of Lave and March, then my efforts will have been all the sweeter.

# 1. *The Many-Model Thinker*

*To become wise you've got to have models in your head. And you've got to array your experience—both vicarious and direct—on this latticework of models.*

—Charlie Munger

This is a book about models. It describes dozens of models in straightforward language and explains how to apply them. Models are formal structures represented in mathematics and diagrams that help us to understand the world. Mastery of models improves your ability to reason, explain, design, communicate, act, predict, and explore.

This book promotes a *many-model thinking* approach: the application of ensembles of models to make sense of complex phenomena. The core idea is that many-model thinking produces wisdom through a diverse ensemble of logical frames. The various models accentuate different causal forces. Their insights and implications overlap and interweave. By engaging many models as frames, we develop nuanced, deep understandings. The book includes formal arguments to make the case for multiple models along with myriad real-world examples.

The book has a pragmatic focus. Many-model thinking has tremendous practical value. Practice it, and you will better understand complex phenomena. You will reason better. You exhibit fewer gaps in your reasoning and make more robust decisions in your career, community activities, and personal life. You may even become wise.

Twenty-five years ago, a book of models would have been

intended for professors and graduate students studying business, policy, and the social sciences along with financial analysts, actuaries, and members of the intelligence community. These were the people who applied models and, not coincidentally, they were also the people most engaged with large data sets. Today, a book of models has a much larger audience: the vast universe of knowledge workers, who, owing to the rise of big data, now find working with models a part of their daily lives.

Organizing and interpreting data with models has become a core competency for business strategists, urban planners, economists, medical professionals, engineers, actuaries, and environmental scientists among others. Anyone who analyzes data, formulates business strategies, allocates resources, designs products and protocols, or makes hiring decisions encounters models. It follows that mastering the material in this book—particularly the models covering innovation, forecasting, data binning, learning, and market entry timing—will be of practical value to many.

Thinking with models will do more than improve your performance at work. It will make you a better citizen and a more thoughtful contributor to civic life. It will make you more adept at evaluating economic and political events. You will be able to identify flaws in your logic and in that of others. You will learn to identify when you are allowing ideology to supplant reason and have richer, more layered insights into the implications of policy initiatives, whether they be proposed greenbelts or mandatory drug tests.

These benefits will accrue from an engagement with a variety of models—not hundreds, but a few dozen. The models in this book offer a good starting collection. They come from multiple disciplines and include the Prisoners' Dilemma, the Race to the Bottom, and the SIR model of disease transmission. All of these models share a common form: they assume a set of entities—often people or organizations—and describe how they interact.

The models we cover fall into three classes:

simplifications of the world, mathematical analogies, and exploratory, artificial constructs. In whatever form, a model must be tractable. It must be simple enough that within it we can apply logic. For example, we cover a model of communicable diseases that consists of infected, susceptible, and recovered people that assumes a rate of contagion. Using the model we can derive a contagion threshold, a tipping point, above which the disease spreads. We can also determine the proportion of people we must vaccinate to stop the disease from spreading.

As powerful as single models can be, a collection of models accomplishes even more. With many models, we avoid the narrowness inherent in each individual model. A many-models approach illuminates each component model's blind spots. Policy choices made based on single models may ignore important features of the world such as income disparity, identity diversity, and interdependencies with other systems.<sup>1</sup> With many models, we build logical understandings of multiple processes. We see how causal processes overlap and interact. We create the possibility of making sense of the complexity that characterizes our economic, political, and social worlds. And, we do so without abandoning rigor—model thinking ensures logical coherence. That logic can be then grounded in evidence by taking models to data to test, refine, and improve them. In sum, when our thinking is informed by diverse logically consistent, empirically validated frames, we are more likely to make wise choices.

## **Models in the Age of Data**

The appearance of a book on models may seem out of place in the era of big data. Today, data exists in unprecedented dimensionality and granularity. Customer purchase data, which used to arrive in monthly aggregates on printed paper, now streams instantaneously with geospatial, temporal, and consumer tags. Student academic performance data now

includes scores on every homework, paper, quiz, and exam, as opposed to semester-end summary grades. In the past, a farmer might mention dry ground at a monthly Grange meeting. Now, tractors transmit instantaneous data on soil conditions and moisture levels in square-foot increments. Investment firms track dozens of ratios and trends for thousands of stocks and use natural-language processing tools to parse documents. Doctors can pull up page upon page of individual patient records that can include relevant genetic markers.

A mere twenty-five years ago, most of us had access to little more than a few bookshelves' worth of knowledge. Perhaps your place of work had a small reference library, or at home you had a collection of encyclopedias and a few dozen reference books. Academics and government and private-sector researchers had access to large library collections, but even they had to physically visit the material. As late as the turn of the millennium, academics could be found shuttling back and forth between card catalog rooms, microfiche collections, library stacks, and special collections in search of information.

That has all changed. Content that had been paper-bound for centuries now flows in tiny packets through the air. So too does the information about the here and now. News that arrived on our doorsteps on newsprint once a day now flows in a continuous digital stream into our personal devices. Stock prices, sports scores, and news of political events and cultural happenings can all be accessed with a swipe or query.

As impressive as the data may be, it is no panacea. We now know what has happened and is happening, but, owing to the complexity of the modern world, we may be less capable of understanding why it happened. Empirical findings may be misleading. Data on piece-rate work often shows that the more people are paid per unit of output, the less they produce. A model in which pay depends on work conditions can explain those data. If conditions are poor so that producing output is difficult, per unit pay may be high. If

conditions are good, per unit pay may be low. Thus, higher pay does not lead to less productivity. Instead, more difficult work conditions require higher per unit pay.<sup>2</sup>

In addition, most of our social data—that is, data about our economic, social, and political phenomena—documents only moments or intervals in time. It rarely tells us universal truths. Our economic, social, and political worlds are not stationary. Boys may outscore girls on standardized tests in one decade and girls may outscore boys the next. The reasons people vote today may differ from the reasons they vote in coming decades.

We need models to make sense of the fire-hose-like streams of data that cross our computer screens. Thus, it is because we have so much data that this might also be called the age of many models. Look across the academy, government, the business world, and the nonprofit sector, and you struggle to find a domain of inquiry or decision not informed by models. Consulting giants McKinsey and Deloitte build models to formulate business strategies. Financial firms such as BlackRock and JPMorgan Chase apply models to select investments. Actuaries at State Farm and Allstate use models to calibrate risk when pricing insurance policies. The people team at Google builds predictive analytic models to evaluate its more than three million job applicants. College and university admissions officers construct predictive models to select from among tens of thousands of applicants.

The Office of Management and Budget constructs economic models to predict the effects of tax policies. Warner Brothers applies data analytics to create models of audience responses. Amazon develops machine learning models to make product recommendations. Researchers funded by the National Institutes of Health build mathematical models of human genomics to search for and evaluate potential cures for cancer. The Gates Foundation uses epidemiological models to design vaccination strategies. Even sports teams use models to evaluate draft prospects and trade opportunities and to formulate within-game strategies. By relying on

models to select players and strategies, the Chicago Cubs won a World Series championship after more than a century of failures.

To people who use models, the rise of model thinking has an even simpler explanation: *models make us smarter*. Without models, people suffer from a laundry list of cognitive shortcomings: we overweight recent events, we assign probabilities based on reasonableness, and we ignore base rates. Without models, we have limited capacity to include data. With models, we clarify assumptions and think logically. And, we can leverage big data to fit, calibrate, and test causal and correlative claims. With models, we think better. In head-to-head competitions between models and people, models win.<sup>3</sup>

## Why We Need Many Models

In this book we advocate using not just one model in a given situation but many models. The logic behind the many-model approach builds on the age-old idea that we achieve wisdom through a multiplicity of lenses. This idea traces back to Aristotle, who wrote of the value of combining the excellences of many. A diversity of perspectives was also a motivation for the great-books movement, which collected 102 important transferable ideas in *The Great Ideas: A Syntopicon of Great Books of the Western World*. The approach finds a modern voice in the work of Maxine Hong Kingston, who wrote in *The Woman Warrior*, “I learned to make my mind large, as the universe is large, so that there is room for paradoxes.” It is also the basis for pragmatic actions in the world of business and policy. Recent books argue that if we want to understand of international relations, we should not model the world exclusively as a group of self-interested nations with well-defined objectives, or only as an evolving nexus of multinational corporations and intergovernmental organizations. We should do both.<sup>4</sup>



As commonsensical as the many-model approach may seem, keep in mind that it runs counter to how we teach models and the practice of modeling. The traditional approach—the one taught in high school—relies on a one-to-one logic: one problem requires one model. For example: now we apply Newton's first law; now we apply the second; now the third. Or: here we use the replicator equation to show the size of the rabbit population in the next period. In this traditional approach, the objective is to (a) identify the one proper model and (b) apply it correctly. Many-model thinking challenges that approach. It advocates trying many models. Had you used many-model thinking in ninth grade, you might have been held back. Use it now, and you will move forward.

Academic papers, for the most part, follow the one-to-one approach as well, even though they use those single models to explain complex phenomena: Trump voters in the 2016 election were those who had been left behind economically. Or: the quality of a child's second-grade teacher determines how economically successful that child will be as an adult.<sup>5</sup> A stream of best-selling nonfiction titles present cures for our ills based on single-model thinking: Educational success depends on grit. Inequality results from concentrations of capital. Our nation's poor health is due to sugar consumption. Each of these models may be true, but none is comprehensive. To confront the complexity of these challenges, to create a world of broader educational achievement, will require lattices of models.

By learning the models in this book, you can begin to build your own lattice. The models originate from a broad spectrum of disciplines, addressing phenomena as varied as the causes of income inequality, the distribution of power, the spread of diseases and fads, the conditions that precede social uprisings, the evolution of cooperation, the emergence of order in cities, and the structure of the internet. The models vary in their assumptions and their structure. Some describe small numbers of rational, self-interested actors. Others describe large populations of rule-following altruists.

Some describe equilibrium processes. Others produce path dependence and complexity. The models also differ in their uses. Some help predict and explain. Others guide actions, inform designs, or facilitate communication. Still others create artificial worlds for our minds to explore.

The models share three common characteristics: First, they simplify, stripping away unnecessary details, abstracting from reality, or creating anew from whole cloth. Second, they formalize, making precise definitions. Models use mathematics, not words. A model might represent beliefs as probability distributions over states of the world or preferences as rankings of alternatives. By simplifying and making precise, they create tractable spaces within which we can work through logic, generate hypotheses, design solutions, and fit data. Models create structures within which we can think logically. As Wittgenstein wrote in his *Tractatus Logico-Philosophicus*, “Logic takes care of itself; all we have to do is to look and see how it does it.” The logic will help to explain, predict, communicate, and design. But the logic comes at a cost, which leads to their third characteristic: *all models are wrong*, as George Box noted.<sup>6</sup> That is true of all models; even the sublime creations of Newton that we refer to as laws hold only at certain scales. Models are wrong because they simplify. They omit details. By considering many models, we can overcome the narrowing of rigor by crisscrossing the landscape of the possible.

To rely on a single model is hubris. It invites disaster. To believe that a single equation can explain or predict complex real-world phenomena is to fall prey to the charisma of clean, spare mathematical forms. We should not expect any one model to produce exact numerical predictions of sea levels in 10,000 years or of unemployment rates in 10 months. We need many models to make sense of complex systems. Complex systems like politics, the economy, international relations, or the brain exhibit ever-changing emergent structures and patterns that lie between ordered and random. By definition, complex phenomena are difficult to

explain, evolve, or predict.<sup>7</sup>

Thus, we confront a disconnect. On the one hand, we need models to think coherently. On the other hand, any single model with a few moving parts cannot make sense of high-dimensional, complex phenomena such as patterns in international trade policy, trends in the consumer products industry, or adaptive responses within the brain. No Newton can write a three-variable equation that explains monthly employment, election outcomes, or reductions in crime. If we hope to understand the spread of diseases, variation in educational performance, the variety of flora and fauna, the effect of artificial intelligence on job markets, the impact of humans on the earth's climate, or the likelihood of social uprisings, we must come at them with machine learning models, systems dynamics models, game theory models, and agent-based models.

## The Wisdom Hierarchy

To sketch the argument for many-model thinking, we begin with a query from poet and dramatist T. S. Eliot: "Where is the wisdom we have lost in knowledge? Where is the knowledge we have lost in information?" To that we might add, where is the information we have lost in all this data?

Eliot's questioning can be formalized as the *wisdom hierarchy*. At the bottom of the hierarchy lie *data*: raw, uncoded events, experiences, and phenomena. Births, deaths, market transactions, votes, music downloads, rainfall, soccer matches, and speciation events. Data can be long strings of zeros and ones, time stamps, and linkages between pages. Data lack meaning, organization, or structure.

*Information* names and partitions data into categories. Examples clarify the distinction between data and information. Rain falling on your head is data. Total rainfall for the month of July in Burlington, Vermont, and Lake Ontario's water level are information. The bright red peppers

and yellow corn on farmers' stands surrounding the capitol in Madison, Wisconsin, on market Saturdays are data. The farmers' total sales are information.




Figure 1.1: How Models Transform Data into Wisdom

We live in an age of abundant information. A century and a half ago, knowing information brought great economic and social status. Jane Austen's *Emma* asks if Frank Churchill is "a young man of information." Today she would not care. Churchill, like everyone else, would have a smartphone. The question is whether he could put that information to use. As Fyodor Dostoyevsky writes in *Crime and Punishment*, "We've got facts, they say. But facts aren't everything; at least half the battle consists in how one makes use of them!"

Plato defined *knowledge* as justified true belief. More modern definitions refer to it as understandings of correlative, causal, and logical relationships. Knowledge organizes information. Knowledge often takes model form. Economic models of market competition, sociological models of networks, geological models of earthquakes, ecological models of niche formation, and psychological models of learning all embed knowledge. Those models explain and predict. Models of chemical bonds explain why metallic bonds prevent us from putting our hands through steel doors while hydrogen bonds yield to our weight when we dive into a lake.<sup>8</sup>

Atop the hierarchy lies *wisdom*, the ability to identify and apply relevant knowledge. Wisdom requires many-model thinking. Sometimes, wisdom consists of selecting the best model, as if drawing from a quiver of arrows. Other times, wisdom can be achieved by averaging models; this is common when making predictions. (We discuss the value of model

averaging in the next section.) When taking actions, wise people apply multiple models like a doctor's set of diagnostic tests. They use models to rule out some actions and privilege others. Wise people and teams construct a dialogue across models, exploring their overlaps and differences.

Wisdom can consist of selecting the correct knowledge or model; consider the following physics problem: A small stuffed cheetah falls from an airplane's hold at 20,000 feet. How much damage will it do upon landing? A student might know a gravity model and a terminal velocity model. The two models give different insights. The gravity model predicts that the stuffed animal would tear through a car's roof. The terminal velocity model predicts that the toy cheetah's speed tops out at around 10 mph.<sup>9</sup> Wisdom consists of knowing to apply the terminal velocity model. A person could stand on the ground and catch the soft cheetah in her hands. To quote the evolutionary biologist J. B. S. Haldane, "You can drop a mouse down a thousand-yard mine shaft; and, on arriving at the bottom, it gets a slight shock and walks away, provided that the ground is fairly soft. A rat is killed, a man is broken, a horse splashes."

In the stuffed-cheetah problem, arriving at the correct solution requires information (the weight of the toy), knowledge (the terminal velocity model), and wisdom (selecting the correct model). Business and policy leaders also rely on information and knowledge to make wise choices. On October 9, 2008, the value of Iceland's currency, the króna, began a free fall. Eric Ball, then treasurer of software giant Oracle, was faced with a decision. A few weeks prior he had dealt with the domestic repercussions of the home mortgage crisis. Iceland's situations posed an international concern. Oracle held billions of dollars in overseas assets. Ball considered network contagion models of financial collapse. He also thought of economic models of supply and demand in which the magnitude of a price change correlates with the size of the market shock. In 2008, Iceland had a GDP of \$12 billion, or less than six months' revenues for McDonald's

Corporation. Ball recollected thinking, “Iceland is smaller than Fresno. Go back to work.”<sup>10</sup> The key to understanding this event, and many-model thinking generally, lies in recognizing that Ball did not search among many models to find one that supported an action that he had already decided to take. He did not use many models to find one that justified his action. Instead, he evaluated two models as possibly useful and then chose the better one. Ball had the right information (Iceland is small), chose the right model (supply and demand), and made a wise choice.

We next show how to create a dialogue among multiple models by reconsidering two historical events: the 2008 global financial market collapse, which reduced total wealth (or what had been thought to be wealth) by trillions of dollars, resulting in a four-year global recession, and the 1961 Cuban missile crisis, which nearly resulted in nuclear war.

The 2008 financial collapse has multiple explanations: too much foreign investment, over-leveraged investment banks, lack of oversight in the mortgage approval process, blissful optimism among home-flipping consumers, the complexity of financial instruments, a misunderstanding of risk, and greedy bankers who knew the bubble existed and expected a bailout. Superficial evidence aligns with each of these accounts: money flowed in from China, loan originators wrote toxic mortgages, investment banks had high leverage ratios, financial instruments were too complex for most to understand, and some banks expected a bailout. With models we can adjudicate between these accounts and check the internal consistency of these accounts: Do they make logical sense? We can also calibrate the models and test the magnitude of the effects.

The economist Andrew Lo, exercising many-model thinking, evaluates twenty-one accounts of the crisis. He finds each to be lacking. It does not make sense that investors would contribute to a bubble that they knew would lead to a global crisis. Hence, the extent of the bubble must have been a surprise to many. Financial firms may well have assumed

the other firms had done due diligence when in fact they had not. Second, what were, in retrospect, clearly toxic (low-quality) bundles of mortgages found buyers. Had global collapse been a foregone conclusion, the buyers would not have existed. And while leverage ratios had increased since 2002, they were not much higher than they had been in 1998. And as for the notion that the government would bail out the banks, Lehman Brothers collapsed on September 15, 2008; with over \$600 billion in holdings, it was the largest bankruptcy in US history. The government did not intervene.

Lo finds that each account contains a logical gap. The data, such as it is, privileges no single explanation. As Lo summarizes: “We should strive at the outset to entertain as many interpretations of the same set of objective facts as we can, and hope that a more nuanced and internally consistent understanding of the crisis emerges in the fullness of time.” He goes on to say, “Only by collecting a diverse and often mutually contradictory set of narratives can we eventually develop a more complete understanding of the crisis.” No single model suffices.<sup>11</sup>

In *Essence of Decision*, Graham Allison undertakes a many-model approach to explain the Cuban missile crisis. On April 17, 1961, a CIA-trained paramilitary group landed on the shores of Cuba in a failed attempt to overthrow Fidel Castro’s communist regime, increasing tensions between the United States and the Soviet Union, Cuba’s ally. In response, Soviet premier Nikita Khrushchev moved short-range nuclear missiles to Cuba. President John F. Kennedy responded by blockading Cuba. The Soviet Union backed down, and the crisis ended.

Allison interprets events with three models. He applies a rational-actor model to show that Kennedy had three possible actions: start a nuclear war, invade Cuba, or impose a blockade. He chose the blockade. The rational-actor model assumes that Kennedy draws a game tree with each action followed by the possible responses by the Soviets. Kennedy then thinks through the Soviets’ optimal response. If, for

example, Kennedy launched a nuclear attack, the Soviets would strike back, resulting in millions dead. If Kennedy imposed a blockade, he would starve the Cubans. The Soviet Union could either back down or launch missiles. Given that choice, the Soviet Union should back down. The model reveals the central strategic logic at play and provides a rationale for Kennedy's bold choice to blockade Cuba.

Like all models, though, it is wrong. It ignores relevant details, allowing it to initially appear a better explanation than it really is. The model neglects to add a stage in which the Soviets put the missiles in Cuba. If the Soviets had been rational, they should have drawn the same tree as Kennedy and realized that they would have to remove the missiles. The rational-actor model also fails to explain why the Soviets did not hide the missiles.

Allison applies an organizational process model to explain these inconsistencies. A lack of organizational capacity explains the Soviets' failure to hide the missiles. The same model can explain Kennedy's choice to blockade. At the time, the United States Air Force lacked the capacity to wipe out the missiles in a single strike. If even a single missile remained, it could kill millions of Americans. Allison deftly combines the two models. An insight from the organizational model changes the payoffs in the rational-choice model.

Allison adds a governmental process model. The other two models reduce countries to their leaders: Kennedy acts for the United States and Khrushchev for the Soviet Union. The government process model recognizes that Kennedy had to contend with Congress and that Khrushchev needed to maintain a political base of support. Thus, Khrushchev's placing of the missiles in Cuba signaled strength.

Allison's book shows the power of models alone and in dialogue. Each model clarifies our thinking. The rational-actor model identifies possible actions once the missiles have arrived and allows us to see the implications of those actions. The organizational model draws our attention to the fact that organizations, not individuals, carry out those actions. The



governmental process model highlights the political cost of invasion. By evaluating events through all three lenses, we gain a broader and deeper understanding. All models are wrong; many are useful.

In both examples, the different models explicate distinct causal forces. Multiple models can also focus on different scales. In an oft-repeated tale, a child claims that the Earth rests on the back of a giant elephant. A scientist asks the child what the elephant stands on, to which the child replies, “A giant turtle.” Anticipating what’s about to come next, the child quickly adds, “Don’t even ask. It’s turtles all the way down.”<sup>12</sup> If the world were turtles all the way—if the world were self-similar—then a model of the top level would apply at every level. But the economy, the political world, and society are not turtles all the way down, nor is the brain. At the sub-micron level, the brain is made up of molecules that form synapses, which in turn form neurons. The neurons combine in networks. The networks overlap in elaborate ways that can be studied with brain imaging. These neuronal networks exist on a scale below that of functional systems such as the cerebellum. Given that the brain differs at each level, we need multiple models, and those models differ. The models that characterize the robustness of neuronal networks bear little resemblance to the molecular biology models used to explain brain cell function, which in turn differ from the psychological models used to explain cognitive biases.

The success of many-model thinking depends on a degree of separability. In analyzing the 2008 financial crisis, we rely on separate models of foreign purchases of assets, of the bundling of assets, and of increased leverage ratios. Allison drew implications from the game theoretic model without considering the organizational model. In studying the human body, doctors separate the skeletal, muscular, limbic, and nervous systems. That said, many-model thinking does not require that these distinct models divide the system into independent parts. Confronted with a complex system, we

cannot, to paraphrase Plato, carve the world at its joints. We can partially isolate the major causal threads and then explore how they are interwoven. In doing so, we will find that the data produced by our economic, political, and social systems exhibits coherence. Social data is more than sequences of incomprehensible hairballs that might have been spit up by the family cat.

## **Summary and Outline of the Book**

To summarize, we live in a time awash in information and data. The same technological advances generating those data shrink time and distance. They make economic, political, and social actors more agile, capable of responding to economic and political events in an instant. They also increase connectedness, and therefore complexity. We face a technologically induced paradox: we know more about the world, but that world is more complex. In light of that complexity, any single model will be more likely to fail. We should not though abandon models. To the contrary, we should privilege logical coherence over intuition and double, triple, and even quadruple down on models and become many-model thinkers.

Becoming a many-model thinker requires learning multiple models of which we gain a working knowledge; we need to understand the formal descriptions of the models and know how to apply them. We need not be experts. Hence, this book balances accessibility and depth. It can function both as a resource and as a guide. The formal descriptions are isolated in stand-alone boxes. It avoids line after line of equations, which overwhelm even the most dedicated readers. The formalism that remains should be engaged and absorbed. Modeling is a craft, mastered through engagement; it is not a spectator sport. It requires deliberate practice. In modeling, mathematics and logic play the role of an expert coach. They correct our flaws.

The remainder of the book is organized as follows: Chapters 2 and 3 motivate the many-model approach. Chapter 4 discusses the challenges of modeling people. The next twenty or so chapters cover individual models or classes of models. By considering one type of model at a time, we can better wrap our heads around its assumptions, implications, and applications. This structure also means that we can pull the book from our bookshelves or open it in our browsers and find self-contained analyses of linear models, prediction models, network models, contagion models, and models of long-tailed distributions, learning, spatial competition, consumer preferences, path dependence, innovation, and economic growth. Interspersed throughout the chapters are applications of many-model thinking to a variety of problems and issues. The book concludes with two deeper dives into the opioid epidemic and income inequality.

## 2. Why Model?

*Knowing reality means constructing systems of transformations that correspond, more or less adequately, to reality.*

—Jean Piaget

In this chapter, we define types of models. Models are often described as simplifications of the world. They can be, but models can also take the form of analogies or be fictional worlds mined for ideas and insights. We also describe the uses of models. In school, we apply models to explain data. In practice, we can also use models to predict, design, and take actions. We can use models to explore ideas and possibilities. And we can use models to communicate ideas and understandings.

The value of models also resides in their ability to reveal conditions under which results hold. Most of what we know holds only in some cases: the square of the longest side of a triangle equals the sum of the squares of the other sides only if the longest side is opposite a right angle. Models reveal similar conditions for our intuitions. With models we can parse out when diseases spread, when markets work, when voting leads to good outcomes, and when crowds make accurate predictions. None of those is a sure thing.

This chapter consists of two parts. In the first, we describe the three types of models. In the second, we cover the uses of models: to reason, explain, design, communicate, act, predict, and explore. These form the acronym REDCAPE, a notso-subtle reminder that many-model thinking endows us with superpowers.<sup>1</sup>

## Types of Models

When constructing a model, we take one of three approaches. We can aim for realism and follow an *embodiment approach*. Such models include the important parts and either strip away unnecessary dimensions and attributes or lump them together. Models of ecological glades, legislatures, and traffic systems take this approach, as do climate models and models of the brain. Or we can take an *analogy approach* and abstract from reality. We can model crime spreading like a disease and the taking of political positions as choices on a left-right continuum. The spherical cow is a favorite classroom example of the analogy approach: to make an estimate of the amount of leather in a cowhide, we assume a spherical cow. We do so because the integral tables in the back of calculus textbooks include  $\tan(x)$  and  $\cos(x)$  but not  $\text{cow}(x)$ .<sup>2</sup>

While the embodiment approach stresses realism, the analogy approach tries to capture the essence of a process, system, or phenomenon. When a physicist assumes away friction but otherwise makes realistic assumptions, she takes the embodiment approach. When an economist represents competing firms as different species and defines product niches, she makes an analogy. She does so using a model developed to embody a different system. No bright line differentiates the embodiment approach from the analogy approach. Psychological models of learning that assign weights to alternatives lump together dopamine responses and other factors; they also invoke the analogy of a scale on which we balance alternatives.

A third approach, the *alternative reality approach*, purposely does not represent or capture reality. These models function as analytic and computational playgrounds in which we can explore possibilities. This approach allows us to discover general insights that apply outside our physical and social world. They help us to understand the implications of real-world constraints: What if energy could be sent safely and efficiently through the air? And they allow us to run

impossible experiments: What if we tried to evolve a brain? This book contains a few such models, notably the Game of Life, which consists of a checkerboard whose squares are classified as either alive (black) or dead (white) that switch between alive and dead according to fixed rules. Though unrealistic, the model produces insights into self-organization, complexity, and, some argue, even life itself.

Whether embodying a more complex reality, creating an analogy, or building a made-up world for exploring ideas, a model must be *communicable* and *tractable*. We should be able to write the model in a formal language such as mathematics or computer code. When describing a model, we cannot toss out terms like *beliefs* or *preferences* without providing a formal description. Beliefs can be represented as a probability distribution over a set of events or priors. Preferences can be represented in several ways such as a ranking over a set of alternatives or as a mathematical function.

How tractable something is means how amenable it is to analysis. In the past, analysis relied on mathematical or logical reasoning. A modeler had to be able to prove each step in an argument. This constraint led to an aesthetic that valued stark models. English friar and theologian William of Ockham (1287–1347) wrote, “Plurality must never be posited without necessity.” Einstein summed up this principle, known as *Ockham’s Razor*, as follows: *everything should be made as simple as possible, but not simpler*. Today, when we run up against the constraint of analytic tractability, we can turn to computation. We can build elaborate models with many moving parts without concern for analytic tractability. Scientists take this approach when constructing models of the global climate, the brain, forest fires, and traffic. They still pay heed to Ockham’s advice, but recognize that “as simple as possible” might require a lot of moving parts.

## The Seven Uses of Models

The academic literature describes dozens of uses of models. Here, we focus on seven categories of uses: to *reason*, *explain*, *design*, *communicate*, *act*, *predict*, and *explore*.

## The Uses of Models (REDCAPE)

**Reason:** To identify conditions and deduce logical implications.

**Explain:** To provide (testable) explanations for empirical phenomena.

**Design:** To choose features of institutions, policies, and rules.

**Communicate:** To relate knowledge and understandings.

**Act:** To guide policy choices and strategic actions.

**Predict:** To make numerical and categorical predictions of future and unknown phenomena.

**Explore:** To investigate possibilities and hypotheticals.

### REDCAPE: Reason

When constructing a model, we identify the most important actors and entities along with relevant characteristics. We then describe how those parts interact and aggregate, enabling us to derive what follows from what, and why. In doing so, we improve our reasoning. While what we can derive depends upon what we assume, we uncover more than tautologies. Rarely can we infer the full range of implications of our assumptions from inspection alone. We need formal logic. Logic also reveals impossibilities and possibilities. With

it, we can derive precise and sometimes unexpected relationships. We can discover the conditionality of our intuitions.

*Arrow's theorem* provides an example of how logic reveals impossibilities. The model addresses the question of whether individual preferences aggregate to form a collective preference. This model represents preferences as ordinal rankings over alternatives. If applied to five Italian restaurants, denoted by the letters A through E, the model allows any of the 120 orderings. Arrow required that the collective ordering be *monotonic* (if everyone ranks A above B, then so does the collective), *independent of irrelevant alternatives* (if no person's relative rankings of A and B are unchanged but rankings of other alternatives change, then the order of A and B in the collective ranking does not change), and *nondictatorial* (no single person should decide the collective ordering). Arrow then proved that if any preferences are allowed, then no collective ordering necessarily exists.<sup>3</sup>

Logic can also reveal paradoxes. Using models we can show the possibility of each subpopulation containing a larger percentage of women than men but the total population containing a larger percentage of men, a phenomenon (*Simpson's paradox*). This actually happened: 1973, the University of California, Berkeley, accepted a larger percentage of women in most departments. Overall, it accepted men at a higher rate. Models also show that it is possible for two losing bets, when played alternately, to produce a positive expected return (*Parrondo's paradox*). With models, we can show that it is possible to add a node to a network and reduce the total length of the edges needed to connect all the nodes.<sup>4</sup>

We should not dismiss these examples as mathematical novelties. Each has practical applications: efforts to increase the population of women could backfire, combinations of losing investments could win, and the total length of a network of electric lines, pipelines, ethernet lines, or roads



could be reduced by adding more nodes.

Logic also uncovers mathematical relationships. Given Euclid's axioms, a triangle can be uniquely determined by any two angles and a side, or by any two sides and an angle. With standard assumptions about consumer and firm behaviors, in markets with a large number of competing firms, price equals marginal cost. Some results are unexpected: among them the *friendship paradox*, which states that in any friendship network, on average, people's friends have more friends than they do.

The paradox arises because highly popular people have more friends. Figure 2.1 shows Zachary's Karate Network. The person represented by the dark circle has six friends, denoted by gray circles. His friends have nine friends on average. These people are represented by white circles. Over the entire network, twenty-nine of the thirty-four people have friends who are more popular than they are.<sup>5</sup> Later we show that if we make a few more assumptions, most people's friends will also be, on average, better-looking, kinder, richer, and smarter than they are.




Figure 2.1: The Friendship Paradox: A Person's Friends Have More Friends

Last, and most important of all, logic reveals the conditionality of truths. A politician may claim that lowering income taxes increases government revenue by spurring economic growth. A rudimentary model in which revenue equals the tax rate times the income level proves that revenue increases only if the percentage growth in income exceeds the percentage cut in taxes.<sup>6</sup> Thus, a 10% cut in income taxes increases revenue only if it causes income to grow by more than 10%. The politician's logic only holds

given certain conditions. Models identify those conditions.

The power of conditionality becomes evident when we contrast claims derived from models with narrative claims, even when the latter have empirical support. Consider the management proverb *first things first*: the idea that when facing multiple tasks, you should do the most important task first. This rule is also known as *big rocks first*, because when filling a bucket with rocks of various sizes, you should put the big rocks in first—if you put the little rocks in first, the big rocks will not fit.

The rule *big rocks first*, inferred from expert observation, may be a good rule most of the time, but it is unconditional. A model-based approach would make specific assumptions about the task and then derive an optimal rule. In the *bin packing problem*, a set of objects of various sizes (or weights) must be allocated into bins of finite capacity. The objective is to use as few bins as possible. Imagine, for example, you are packing up your apartment and putting everything into 2-foot-by-2-foot boxes. Ordering your possessions by size and putting each object in the first box with sufficient space (known as the *first fit algorithm*) turns out to be quite effective. Big rocks first works well. However, suppose that we consider a more complex task: allocating space on the International Space Station for research projects. Each project has a payload weight, a size, and power requirements along with demands on the astronauts' time and cognitive abilities. Each also makes a potential scientific contribution. Even if we came up with some measure of bigness as a weighted average of these attributes, big rocks first would prove a poor rule given the dimensionality of interdependencies. More sophisticated algorithms and possibly market mechanisms would perform much better.<sup>7</sup> Thus, under some conditions, big rocks first is a good rule. Under other conditions, it is not. With models, we can trace the boundaries of when we should place the big rocks first and when we should not.

Critics of formalism claim that models repackage what we already know, that they pour old wine into shiny

mathematical bottles, that we do not need a model to know that two heads are better than one or that he who hesitates is lost. We can learn the value of commitment from reading of Odysseus tying himself to the mast. That criticism fails to recognize that inferences drawn from models take conditional forms: if condition *A* holds, then result *B* follows (e.g., if you are packing bins and size is the only constraint, pack the biggest objects first). Lessons drawn from literature or proverbial advice from great thinkers often provide no conditions. If we try to lead our lives or manage others by unconditional rules, we find ourselves lost in a sea of *opposite proverbs*. Are two heads better than one? Or, do too many cooks spoil the broth?

**Proverb:** Two heads are better than one

**Opposite:** Too many cooks spoil the broth

**Proverb:** He who hesitates is lost

**Opposite:** A stitch in time saves nine

**Proverb:** Tie yourself to the mast

**Opposite:** Keep your options open

**Proverb:** The perfect is the enemy of the good

**Opposite:** Do it well or not at all

**Proverb:** Actions speak louder than words

**Opposite:** The pen is mightier than the sword

While opposite proverbs abound, opposite theorems cannot. Within models, we make assumptions and prove theorems. Two theorems that disagree on the optimal action, make different predictions, or offer distinct explanations must make different assumptions.

**REDCAPE: Explain**

Models provide clear logical explanations for empirical phenomena. Economic models explain price movements and market shares. Physics models explain the rate of falling objects and the shape of trajectories. Biological models explain the distributions of species. Epidemiological models explain the speed and patterns of disease spread. Geophysical models explain the size distribution of earthquakes.

Models can explain point values and changes in their values. A model can explain the current price of pork belly futures and why prices rose over the past six months. A model can explain why a president appoints a moderate Supreme Court justice and why a candidate moves to the left or right. Models also explain shape: models of the diffusion of ideas, technologies, and diseases produce an S-shaped curve of adoption (or contagion).

The models we learn in physics, such as Boyle's Law (a model stating that the pressure of oxygen times the volume equals a constant ( $PV = k$ )), explain phenomena unreasonably well.<sup>8</sup> If we know the volume, we can estimate the constant  $k$ , and then explain or predict pressure  $P$  as a function of  $V$  and  $k$ . The model owes its accuracy to the fact that gases consist of simple parts that exist in large numbers and follow fixed rules: any two oxygen molecules placed in the identical situation follow the same physical laws. They exist in such large numbers that statistical averaging cancels out any randomness. Most social phenomena share none of these three attributes: social actors are heterogeneous, interact in small groups, and do not follow fixed rules. People also think. Even more problematic, people respond to social influences, meaning that behavioral variations may not cancel out. As a result, social phenomena are much less predictable than physical phenomena.<sup>9</sup>

The most effective models explain both straightforward outcomes and puzzling ones. Textbook models of markets can explain why an unanticipated increase in the demand for a normal good like shoes or potato chips increases the price in the short run, an intuitive result. These same models explain

why in the long run, demand increases have less of an effect on price than the marginal cost of producing the good. Increases in demand can even produce reductions in price that result from increased returns to scale in production, a more surprising result. The same models can explain paradoxes such as why diamonds, which have little practical value, have high prices, but water, a necessity for survival, costs little.

As for the claim that models can explain anything: it is true, they can. However, a model-based explanation includes formal assumptions and explicit causal chains. Those assumptions and causal chains can be taken to data. A model that claims that high levels of criminal behavior can be explained by low probabilities of being caught can be tested.

## **REDCAPE: Design**

Models aid in design by providing frameworks within which we can contemplate the implications of choices. Engineers use models to design supply chains. Computer scientists use models to design web protocols. Social scientists used models to design institutions.

In July 1993, a group of economists met at Caltech in Pasadena, California, to design an auction to allocate the electronic spectrum for cellular phones. In the past, the government had allocated spectrum rights to large companies for modest fees. A provision within the Omnibus Budget Reconciliation Act of 1993 allowed for auctioning the spectrum to raise money.

The radio signal from a tower covers a geographic range. Therefore, the government sought to sell licenses for specific regions: Western Oklahoma, Northern California, Massachusetts, Eastern Texas, and so on. This created a design challenge. The value of any given license for a company depended on the other licenses that company won. The license for Southern California would be worth more to a company that also owned the license for Northern California,

for example. Economists refer to these interdependent valuations as *externalities*. The externalities had two main sources: construction and advertising. Holding neighboring licenses meant lower construction costs and the potential to exploit overlapping media markets.

The externalities created a problem with holding simultaneous auctions. A company trying to win a bundle of licenses might lose one license to another bidder and therefore lose the externalities. That company might then want to back out of its bids on other licenses. Sequential auctions had a different shortcoming. Bidders would underbid in early auctions to hedge against losing subsequent licenses.

A successful auction design had to be immune to strategic manipulation, generate efficient outcomes, and be comprehensible to participants. The economists used game theory models to analyze whether features could be exploited by strategic bidders, computer simulation models to compare the efficiency of various designs, and statistical models to choose parameters for experiments with real people. The final design, a multiple-round auction that allowed participants to back out of bids and prohibited sitting out early periods to mask intentions, proved successful. Over the past thirty years, the FCC has raised nearly \$60 billion using this type of auction.<sup>10</sup>

## **REDCAPE: Communicate**

By creating a common representation, models improve communication. Models require formal definitions of the relevant features and their relationships that we can then communicate with precision. The model  $F = MA$  relates three measurable quantities, force, mass, and acceleration, and does so in equation form. Each term is expressed in measurable units that can be communicated without fear of mis-interpretation. By comparison, the claim that “bigger, faster things generate more power” offers far less precision.

Much can get lost in translation. Does bigger mean weight or size? Does faster mean velocity or acceleration? Does power mean energy or force? And how do bigger and faster combine to produce power? Attempts to formalize the claim could result in any of several forms: power could be written incorrectly as weight plus velocity ( $P = W + V$ ), weight times velocity ( $P = WV$ ), or weight plus acceleration ( $P = W + A$ ).

When we formally define an abstract concept like political ideology using a reproducible methodology, those concepts take on some of the same features as physical qualities such as mass and acceleration. We can use a model to say that one politician is more liberal than another based on their voting records. We can then communicate that claim with precision. Liberalness is well defined and measurable. Someone can use the same method to compare other politicians. Of course, voting records may not be the only measure of liberalness. We might construct a second model that assigns ideologies based on textual analysis of speeches. With that model as well, we can communicate with clarity what we mean by more liberal.

Many underappreciate the impact of communication on progress. An idea that cannot be communicated is like a tree falling in a forest with no one around to notice it. The remarkable economic growth in the Age of Enlightenment was due in no small part to the transferability of knowledge, often in model form. In fact, the evidence suggests that the transferability of ideas may have contributed more to economic growth during that time than did levels of education: city-level growth in eighteenth-century France correlates more strongly with the number of subscriptions to Diderot's *Encyclopédie* than with literacy rates.<sup>11</sup>

## **REDCAPE: Act**

Francis Bacon wrote, “The great end of life is not knowledge but action.” Good actions require good models. Governments, corporations, and nonprofits all use models to guide actions.

Whether it be raising or lowering prices, opening a new location, acquiring a company, offering universal health care, or funding an after-school program, decision-makers rely on models. On the most important actions, decision-makers use sophisticated models. Models are linked to data.

In 2008, as part of the Troubled Asset Relief Program (TARP), the Federal Reserve gave \$182 billion in financial assistance to bail out the multinational insurance company American International Group (AIG). According to the US Department of the Treasury, the government chose to stabilize AIG “because its failure during the financial crisis would have had a devastating impact on our financial system and the economy.”<sup>12</sup> The purpose of the bailout was not to save AIG but to prop up the entire financial system. Businesses fail every day, and the government does not intervene.<sup>13</sup>

The particular choices made within TARP were based on models. Figure 2.2 shows a version of a network model produced by the International Monetary Fund. The nodes (circles) represent financial institutions. The edges (the lines between the circles) represent correlations between the values of the holdings of those institutions. The color and width of an edge corresponds to the strength of the correlation between the institutions, with darker and thicker lines implying greater correlation.<sup>14</sup>

AIG occupies a central position in the network because it sold insurance to other firms. AIG held promises to pay other firms if those firms’ assets lost value. If prices fell, then AIG owed those firms money. By implication, if AIG failed, so too would the firms connected to AIG. A cascade of failures might ensue. By stabilizing AIG’s position, the government could prop up the market values of other firms in the network.<sup>15</sup>



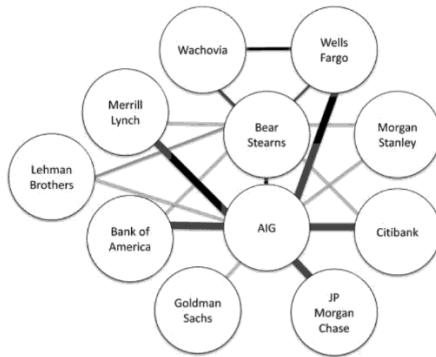


Figure 2.2: Correlation Graph Between Financial Institutions

Figure 2.2 also helps to explain why the government let Lehman Brothers fail. Lehman did not occupy a central position in the network. We cannot rerun history, so we cannot know if the Federal Reserve took the correct action. We do know that the financial industry did not collapse as a result of Lehman’s failure. We also know that the government earned a \$23 billion profit on its loan to AIG. So, we can infer that the policy choices—based on many-model thinking—were not a failure.

Models that guide action, such as policy models, often rely on data, but not all do. Most policy models also use mathematics, though that was not always true. In the past, policymakers built physical models as well. Phillips’s hydraulic model of the British economy was used to think through policy choices in the mid-twentieth century, and a physical model of San Francisco Bay was instrumental in the decision not to dam the bay for fresh water.<sup>16</sup> The Mississippi River Basin Model Waterways Experiment Station, which covers nearly 200 acres near Clinton, Mississippi, is a miniature replica of the river’s basin built on a horizontal scale of 1:100. The model can test the upstream and downstream effects of building new dams and reservoirs. The released water follows the laws of physics within the physical structure. In these physical models, the entities themselves are analogs of the real world. The models are logical because

Each of these hypotheticals can be explored with models. We can also use models to explore unrealistic environments. What if Lamarck had been correct and acquired traits could be passed on to our offspring, so the children of parents with orthodontically corrected teeth would not need braces? What happens in such a world? Asking that question and exploring its implications can help to reveal the limits of evolutionary processes. Abandoning the constraints of reality can spur creativity. For this reason, advocates of the critical design movement engage in speculative fictions to generate new ideas.<sup>19</sup>

Exploration sometimes consists of comparing common assumptions across domains. To understand network effects, a modeler might begin a collection of stylized network structures and then ask whether and how network structure affects cooperation, disease spread, or social uprisings. Or a modeler might apply a collection of learning models to decisions, two-person games, and multiperson games. The purpose of these exercises is not to explain, predict, act, or design. It is to explore and learn.

When we apply a model in practice, we may use it in any of several ways. The same model may explain, predict, and guide action. As an example, on August 14, 2003, sagging trees leaning on power lines near Toledo, Ohio, created a localized power outage that spread when a software failure prevented an alarm from alerting technicians to redistribute power. Within a day, more than 50 million people in the northeastern United States and Canada had lost power. That same year, a storm knocked out a power line between Italy and Switzerland, leaving 60 million Europeans without power. Engineers and scientists turned to models that represent the power grid as a network. The models helped to explain how the failures occurred, offered predictions of regions where future failures might be likely, and also guided actions by identifying locations where new lines, transformers, and power supplies would enhance the robustness of the network. Putting one model to many uses

will be a recurrent theme in this book. As we see next, one-to-many is a necessary complement to our central theme of applying many models to make sense of complex phenomena.

### 3. *The Science of Many Models*

*Nothing is less real than realism. Details are confusing. It is only by selection, by elimination, by emphasis that we get to the real meaning of things.*

—Georgia O’Keeffe

In this chapter, we take a scientific approach to motivate the many-model approach. We begin with the Condorcet jury theorem and the diversity prediction theorem, which make quantifiable cases for the value of many models in helping us act, predict, and explain. These theorems may overstate the case for many models. To show why, we introduce categorization models, which partition the world into boxes. Using categorization models shows us that constructing many models may be harder than we expect. We then apply this same class of model to discuss model granularity—how specific our models should be—and help us decide whether to use one big model or many small models. The choice will depend on the use. When predicting, we often want to go big. When explaining, smaller is better.

The conclusion addresses a lingering concern. Many-model thinking might seem to require learning a lot of models. While we must learn some models, we need not learn as many as you might think. We do not need to master a hundred models, or even fifty, because models possess a one-to-many property. We can apply any one model to many cases by reassigning names and identifiers and modifying assumptions. This property of models offers a counterpoise to the demands of many-model thinking. Applying a model in a

new domain requires creativity, an openness of mind, and skepticism. We must recognize that not every model will be appropriate to every task. If a model cannot explain, predict, or help us reason, we must set it aside.

The skills required to excel at one-to-many differ from the mathematical and analytic talents many people think of as necessary for being a good modeler. The process of one-to-many involves creativity. It is to ask: *How many uses can I think of for a random walk?* To provide a hint of the forms that creativity takes, at the end of the chapter we apply the geometric formula for area and volume as a model and use it to explain the size of supertankers, to criticize the body mass index, to predict the scaling of metabolisms, and to explain why we see so few women CEOs.

## **Many Models as Independent Lies**

We now turn to formal models that help reveal the benefits of many-model thinking. Within those models, we describe two theorems: the Condorcet jury theorem and the diversity prediction theorem. The *Condorcet jury theorem* is derived from a model constructed to explain the advantages of majority rule. In the model, jurors make binary decisions of guilt or innocence. Each juror is correct more often than not. In order to apply the theorem to collections of models instead of jurors, we interpret each juror's decision as a classification by a model. These classifications could be actions (buy or sell) or predictions (Democratic or Republican winner). The theorem then tells us that by constructing multiple models and using majority rule we will be more accurate than if we used one of the constituent models. The model relies on the concept of a *state of the world*, a full description of all relevant information. For a jury, the state of the world consists of the evidence presented at trial. For models that measure the social contribution of a charitable project, the state of the world might correspond to the project's team, the

organizational structure, the operational plan, and the characteristics of the problem or situation the project would address.

## Condorcet Jury Theorem

Each of an odd number of people (models) classifies an unknown state of the world as either true or false. Each person (model) classifies correctly with a probability  $p > \frac{1}{2}$ , and the probability that any person (model) classifies correctly is statistically independent of the correctness of any other person (model).

**Condorcet jury theorem:** A majority vote classifies correctly with higher probability than any person (model), and as the number of people (models) becomes large, the accuracy of the majority vote approaches 100%.

Ecologist Richard Levins elaborates on how the logic of the theorem applies to the many-model approach: “Therefore, we attempt to treat the same problem with several alternative models each with different simplifications but with a common biological assumption. Then, if these models, despite their different assumptions, lead to similar results, we have what we can call a robust theorem, which is relatively free of the details of the model. Hence our truth is the intersection of independent lies.”<sup>1</sup> Note that here he aspires to a unanimity of classification. When many models make a common classification, our confidence should soar.

Our next theorem, the *diversity prediction theorem*, applies to models that make numerical predictions or valuations. It quantifies the contributions of model accuracy and model diversity to the accuracy of the average of those models.<sup>2</sup>

and “summer vacations.” We use categories to guide actions. We categorize restaurants by ethnicity—Italian, French, Turkish, or Korean—to decide where to have lunch. We categorize stocks by their price-to-earnings ratios and sell stocks with low price-to-earnings ratios. We use categories to explain, as when we claim that Arizona’s population has grown because the state has good weather. We also use categories to predict: we might forecast that a candidate for political office with military experience has an increased chance of winning.

We can interpret the contributions of categorization models within the wisdom hierarchy. The objects constitute the data. Binning the objects into categories creates information. The assigning of valuations to categories requires knowledge. To critique the Condorcet jury theorem, we rely on a *binary categorization model* that partitions the objects or states into two categories, one labeled “guilty” and one “innocent.” The key insight will be that the number of relevant attributes constrains the number of distinct categorizations, and therefore the number of useful models.

## Categorization Models

There exists a set of objects or states of the world, each defined by a set of attributes and each with a value. A **categorization model**,  $M$ , partitions these objects or states into a finite set of categories  $\{S_1, S_2, \dots, S_n\}$  based on the object’s attributes and assigns **valuations**  $\{M_1, M_2, \dots, M_n\}$  for each category.

Imagine we have one hundred student loan applications, half of which were paid back and half of which were defaulted. We know two pieces of information for each loan: whether the loan amount exceeded \$50,000, and whether the

recipient majored in engineering or the liberal arts. These are the two attributes. With two attributes we can distinguish between four types of loans: large loans to engineers, small loans to engineers, large loans to liberal arts majors, and small loans to liberal arts majors.

A binary categorization model classifies each of these four types as either repaid or defaulted. One model might classify small loans as repaid and large loans as defaulted. Another model might classify loans to engineers as repaid and loans to liberal arts majors as defaulted. It seems plausible that each of these models could be correct more than half the time, and that the two models might be approximately independent of each other. A problem arises when we try to construct more models. There exist only sixteen unique models that map four categories into two outcomes. Two of those models classify all loans as repaid or defaulted. Each of the remaining fourteen has an exact opposite. Whenever the model classifies correctly, its opposite model classifies incorrectly. Thus, of the fourteen possible models, at most seven can be correct more than half the time. And if any model happens to be correct exactly half of the time, then so must its opposite.

The dimensionality of our data limits the number of models we can produce. At most we can have seven models. We cannot construct eleven independent models, much less seventy-seven. Even if we had higher-dimensional data—say, if we knew the recipient's age, grade point average, income, marital status, and address—the categorizations that relied on those attributes must yield accurate predictions. Each subset of attributes would have to be relevant to whether the loan was repaid and be uncorrelated with the other attributes. Both are strong assumptions. For example, if address, marital status, and income are correlated, then models that swap those attributes will be correlated as well.<sup>4</sup> In the stark probabilistic model, independence seemed reasonable: different models make independent mistakes. When we unpack that logic with categorization models, we see the difficulty of constructing multiple independent



models.

Attempts to construct a collection of diverse, accurate models encounter a similar problem. Suppose that we want to build an ensemble of categorization models that predict unemployment rates across five hundred mid-size cities. An accurate model must partition cities into categories such that within a category the cities have similar unemployment rates. The model must also predict unemployment accurately for each category. For two models to make diverse predictions, they must categorize cities differently, predict differently, or do both. Those two criteria, though not in contradiction, can be difficult to satisfy. If one categorization relies on average education level and a second relies on average income, they may categorize similarly. If so, the two models will be accurate but not diverse. Creating twenty-six categories using the first letter of each city's name will create a diverse categorization but probably not an accurate model. Here as well, the takeaway is that in practice "many" may be closer to five than fifty.

Empirical studies of prediction align with that inference. While adding models improves accuracy (they have to, given the theorems), the marginal contribution of each model falls off after a handful of models. Google found that using one interviewer to evaluate job candidates (instead of picking at random) increases the probability of an above-average hire from 50% to 74%, adding a second interviewer increases the probability to 81%, adding a third raises it to 84%, and using a fourth lifts it to 86%. Using twenty interviewers only increases the probability to a little over 90%. That evidence suggests a limit to the number of relevant ways of looking at a potential hire.

A similar finding holds for an evaluation of tens of thousands of forecasts by economists regarding unemployment, growth, and inflation. In this case, we should think of the economists as models. Adding a second economist improves the accuracy of the prediction by about 8%, two more increase it by 12%, and three more by 15%. Ten

economists improve the accuracy by about 19%. Incidentally, the best economist is only about 9% better than average—assuming you knew which economist was best. So three random economists perform better than the best one.<sup>5</sup> Another reason for averaging many and not relying on the economist who has been best historically is that the world changes. The economist who performs at the top today may be middling tomorrow. That same logic explains why the US Federal Reserve relies on an ensemble of economic models rather than just one: the average of many models will typically be better than the best model.

The lesson should be clear: if we can construct multiple diverse, accurate models, then we can make very accurate predictions and valuations and choose good actions. The theorems validate the logic of many-model thinking. What the theorems do not do, and cannot do, is construct the many models that meet their assumptions. In practice, we may find that we can construct three or maybe five good models. If so, that would be great. We need only read back one paragraph: adding a second model yields an 8% improvement, while adding a third gets us to 15%. Keep in mind, these second and third models need not be better than the first model. They could be worse. If they are a little less accurate, but categorically (in the literal sense) different, they should be added to the mix.

## **One Big Model and the Granularity Question**

Many models work in theory and in practice. That does not mean that they are always the correct approach. Sometimes we are better off constructing a single large model. In this section, we put some thought into when we should use each approach and along the way take up the *granularity question* of how finely we should partition our data.

To take on the first question, of whether to use one big model or many small ones, recall the uses of models: to

*reason, explain, design, communicate, act, predict, and explore.* Four of these uses—to reason, explain, communicate, and explore—require simplification. By simplifying, we can apply logic allowing us to explain phenomena, communicate our ideas, and explore possibilities.

Think back to the Condorcet jury theorem. Within it, we could unpack logic, explain why an approach that uses many models was more likely to produce a correct result, and communicate our findings. Had we constructed a model of jurors with personality types and described the evidence as vectors of words, we would have been lost in a mangle of detail. Borges elaborates on this point in an essay on science. He describes mapmakers who make ever more elaborate maps: “The Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it. The following Generations, who were not so fond of the Study of Cartography as their Forebears had been, saw that this vast Map was useless.”

The three other uses of models—to *predict, design, and act*—can benefit from high-fidelity models. If we have BIG data, we should use it. As a rule of thumb, the more data we have, the more granular we should make our model. This can be shown by using categorization models to structure our thinking. Suppose first that we want to construct a model to explain variation in a data set. To provide context, suppose that we have an enormous data set from a chain of grocery stores detailing monthly spending on food for several million households. These households differ in the amount they spend, which we measure as variation: the sum of the squared differences between what each family spends and average spending across all households. If average spending is \$500 a month and a given family spends \$520, that family contributes 400, or 20 squared, to the total variation. Statisticians call the proportion of the variation that a model explains the model’s  $R^2$ .

If the data had a total variation of 1 billion and a model explains 800 million of that variation, then the model has an

# Model Error Decomposition Theorem

The Bias-Variance Trade-off

Model Error = Categorization Error + Valuation Error

$$\sum_{x \in X} (M(x) - V(x))^2 = \sum_{i=1}^n \sum_{x \in S_i} (V(x) - V_i)^2 + \sum_{i=1}^n (M_i - V_i)^2$$

where  $M(x)$  and  $M_i$  denote the model's values for data point  $x$  and category  $S_i$  and  $V(x)$  and  $V_i$  denote their true values.<sup>6</sup>

## One-to-Many

Learning models takes time, effort, and breadth. To reduce those demands, we take a *one-to-many* approach. We advocate mastering a modest number of flexible models and applying them creatively. We use a model from epidemiology to understand the diffusion of seed corn, Facebook, crime, and pop stars. We apply a model of signaling to advertising, marriage, peacock feathers, and insurance premiums. And we apply a rugged-landscape model of evolutionary adaption to explain why humans lack blowholes. Of course, we cannot take any model and apply it to any context, but most models are flexible. We gain even when we fail because attempts at creative uses of models reveal their limits. And it is fun.

The one-to-many approach is relatively new. In the past, models belonged to specific disciplines. Economists had models of supply and demand, monopolistic competition, and economic growth; political scientists had models of electoral competition; ecologists had models of speciation and replication; and physicists had models describing laws of motion. All of these models were developed with specific purposes in mind. One would not apply a model from physics to the economy or a model from economics to the brain any more than one would use a sewing machine to repair a leaky

pipe.

Taking models out of their disciplinary silos and practicing one-to-many has produced notable successes. Paul Samuelson reinterpreted models from physics to explain how markets attain equilibria. Anthony Downs applied a model of ice cream vendors competing on a beach to explain the positioning of political candidates competing in ideological space. Social scientists have applied models of interacting particles to explain poverty traps, variation in crime rates, and even economic growth across countries. And economists have taken models of self-control based on economic principles to understand the functioning of the brain.<sup>7</sup>

### One-to-Many: Higher Powers ( $X^N$ )

Creatively applying models requires practice. To provide a preview of the potential of the *many-to-one* principle, we take the familiar formula of a variable raised to a power,  $X^N$ , and apply it as a model. When the power equals 2, the formula gives the area of a square, when the power equals 3, it gives the volume of a cube. When raised to higher powers, it captures geometric expansion or decay.

**Supertankers:** Our first application considers a cubic supertanker whose length is eight times its depth and width, which we denote by  $S$ . As shown in figure 3.1, the supertanker has a surface area of  $34S^2$  and a volume of  $8S^3$ . The cost of building a supertanker depends primarily on its surface area, which determines the amount of steel used. The amount of revenue a supertanker generates depends on its volume. Computing the ratio of volume to surface area,  $\frac{8S^3}{34S^2} \approx \frac{S}{4}$ , reveals a linear gain in profitability from increasing size.




Figure 3.1: A Cubic Supertanker: Surface Area =  $34S^2$ , Volume =  $8S^3$

Shipping magnate Stavros Niarchos, who knew this ratio, built the first modern supertankers and made billions during the period of rebuilding that followed World War II. To give some sense of scale: the T2 oil tanker used during World War II measured 500 feet long, 25 feet deep, and 50 feet wide. Modern supertankers such as the *Knock Nevis* measure 1,500 feet long, 80 feet deep, and 180 feet wide. Imagine tipping the Willis (Sears) Tower in Chicago on its side and floating it in Lake Michigan. The *Knock Nevis* resembles a T2 oil tanker scaled up by a factor of a little over three. The *Knock Nevis* has about ten times the surface area as a T2 oil tanker and over thirty times the volume. A question arises as to why supertankers are not even larger. The short answer is that tankers must pass through the Suez Canal; the *Knock Nevis* squeezes through with a gap of a few feet on each side.<sup>8</sup>

**Body mass index:** Body mass index (BMI) is used by the medical profession to define weight categories. Developed in England, BMI equals the ratio of a person's weight (in kilograms) to her height in meters squared.<sup>9</sup> Holding height constant, BMI increases linearly with weight. If one person weighs 20% more than another person of the same height, the first person's BMI will be 20% higher.

We first apply our model to approximate a person as a perfect cube made up of some mixture of fat, muscle, and bone. Let  $M$  denote the weight of one cubic meter of our cubic person. The human cube's weight equals its volume times the weight per cubic meter, or  $H^3 \cdot M$ . Our cube's BMI equals  $H \cdot M$ . Our model reveals two flaws: BMI increases linearly with height, and given that muscle weighs more than fat, fit people have higher  $M$  and therefore higher BMIs. Height should be unrelated to obesity, and muscularity is the *opposite* of fatness. These flaws remain if we make the model more realistic. If we make a person's depth (thickness front to back) and width proportional to height using parameters  $d$

and  $w$ , then BMI can be written as follows:  $BMI = \frac{H \cdot (dH) \cdot (wH) \cdot M}{H^2} = dwHM$ . The BMIs of many NBA stars and other athletes place them in the overweight category (BMI > 25), along with many of the world's top male decathletes.<sup>10</sup> Given that even moderately tall, physically fit people will likely have high BMIs, we should not be surprised that a meta-analysis of nearly a hundred studies with a combined sample size in the millions found that slightly overweight people live longest.<sup>11</sup>

**Metabolic rates:** We now apply our model to predict an inverse relationship between an animal's size and its metabolic rate. Every living entity has a metabolism, a repeated sequence of chemical reactions that breaks down organic matter and transforms it into energy. An organism's metabolic rate, measured in calories, equals the amount of energy needed to remain alive. If we construct cubic models of a mouse and an elephant, figure 3.2 shows that the smaller cube has a much larger ratio of surface area to volume.




Figure 3.2: The Exploding Elephant

We can model the mouse and the elephant as composed of cells 1 cubic inch in volume, each with a metabolism. Those metabolic reactions produce heat that must dissipate through the surface of the animal. Our mouse has a surface area of 14 square inches and a volume of 3 cubic inches, a surface-to-volume ratio of roughly 5:1.<sup>12</sup> For each cubic-inch cell in its volume, the mouse has five square inches of surface area through which it can dissipate heat. Each heat-producing cell

in the elephant has only one-fifteenth of a square inch of surface area. The mouse can dissipate heat at seventy-five times the rate of the elephant.

For both animals to maintain the same internal temperature, the elephant must have a slower metabolism. It does. An elephant with a mouse's metabolism would require 15,000 pounds of food per day. The elephant's cells would also produce too much heat to be dissipated through its skin. As a result, elephants would smolder and then explode. The reason elephants do not blow up is that they have a metabolism roughly twenty times lower than that of mice. The model does not predict the rate at which metabolism scales with size, only the direction. More elaborate models can explain the scaling laws.<sup>13</sup>

**Women CEOs:** For our last application, we increase the exponent in the formula and use the model to explain why so few women become CEOs. In 2016, fewer than 5% of Fortune 500 companies had women CEOs. To become a CEO a person must receive multiple promotions. We can model those promotion opportunities as probabilistic events: a person has some probability of receiving a promotion. We further assume that to become CEO, a person must be promoted at each opportunity.


We assume fifteen promotion opportunities as a benchmark, as that corresponds to a promotion every two years on a thirty-year path to CEO. The weight of evidence reveals modest biases in favor of men, which we can model as men having a higher probability of being promoted.<sup>14</sup> We model this as a man's probability of promotion,  $P_M$ , being slightly larger than a woman's,  $P_W$ . If we benchmark these probabilities at 50% and 40%, respectively, then a man is nearly thirty times more likely than a woman to become CEO.<sup>15</sup> The model reveals how modest biases accumulate. A 10% difference in promotion rates becomes a 30-fold bias at the top. This same model provides a novel explanation for why a much larger percentage (about 25%) of college and



## Bagging and Many Models

Often we fit a model to a sample from an existing data set and then test that same model against the remainder of the data. Other times we fit a model to existing data and use that model to predict future data. This type of modeling creates a tension: the more parameters we include in our model, the better we can fit data and the more we risk overfitting. Good fit does not imply a good model. Physicist Freeman Dyson tells of Enrico Fermi's reaction to a piece of Dyson's research that had exceptional model fit. "In desperation I asked Fermi whether he was not impressed by the agreement between our calculated numbers and his measured numbers. He replied, 'How many arbitrary parameters did you use for your calculations?' I thought for a moment about our cut-off procedures and said, 'Four.' He said, 'I remember my friend Johnny von Neumann used to say, with four parameters I can fit an elephant, and with five I can make him wiggle his trunk.' With that, the conversation was over."<sup>16</sup>

The estimates used to "wiggle the trunk" often include higher-order terms: squares, cubes, and fourth powers. This introduces a risk of large errors, because higher-order terms amplify. While 10 is twice as large as 5,  $10^4$  is 16 times as large as  $5^4$ . The figure below shows an example of overfitting.




Overfitting and Out-of-Sample Error

The graph on the left shows (hypothetical) sales data from a company that manufactures industrial 3-D printers as a function of the number of site visits made (on average) per month by their sales team. The graph on the left shows a nonlinear best fit that includes nonlinear terms up to the fifth power. The graph on the right shows that the model predicts sales of 100 printers if sales visits reach 30. That cannot be correct if customers buy at most one 3-D printer. By overfitting, the model makes a huge error out of the sample.

To prevent overfitting, we could avoid higher-order terms. A more sophisticated solution known as *bootstrap aggregation* or *bagging* constructs many models. To bootstrap a data set, we create multiple data sets of equal size by randomly drawing data points from the original data. The points are drawn with replacement—after we draw a data point, we put it back in the “bag” so that we might draw it again. This technique produces a collection of data sets of equal size, each of which contains multiple copies of some data points and no copies of others.

We then fit (nonlinear) models to each data set, resulting in multiple models.<sup>17</sup> We can then plot all the models on the same set of axes, creating a *spaghetti graph* (see below). The dark line shows the average of the different models.



Bootstrapping and a Spaghetti Graph

Bagging will capture robust nonlinear effects, as they will be evident in multiple random samples of the data, while avoiding fitting idiosyncratic patterns in any

single data set. By building diversity through random samples and then averaging the many models, bagging applies the logic that underpins the diversity prediction theorem. It creates diverse models, and as we know, the average of those models will be more accurate than the models themselves.

## 4. Modeling Human Actors

*It is not possible yet to point to a single theory of human behavior that has been successfully formulated and tested in a variety of settings.*

—Elinor Ostrom

In this chapter, we address a question that lies at the core of this book: How do we model people? In many of the models that follow, people will be the fundamental unit of analysis. We will construct models of people who vote, cooperate, start uprisings, participate in fads, invest in retirement accounts, and become addicted to drugs. Within each model, we will have to make assumptions about people. What are their objectives? Do they only care about themselves or are they altruistic? What are their potential actions? How do they choose what action to take, or do they not even have a choice?

We could make ad hoc assumptions for each model. But to do so would invite confusion and miss an opportunity. We would be left with an idiosyncratic set of constructions. Each new model would require new thinking about how people act. The resulting heterogeneity would limit our ability to think across and combine models. We could not be effective many-model thinkers.

The approach we follow stresses coherence along with variety. We will model people as either *rule-based actors* or *rational actors*. Within the set of rule-based actors, we consider those who act based on *simple fixed rules* and those who act based on *adaptive rules*. Someone acting based on an adaptive

rule can change her behavior based upon information or past success or because she watches others. As we shall discuss, no bright lines distinguish these cases from one another; an adaptive rule can sometimes be interpreted as a fixed rule, and sometimes rational actions will take the form of simple rules.

How we choose to model people will depend on the context and on our goals. Are we predicting or explaining? Are we evaluating policy actions? Are we trying to design an institution? Or are we exploring? In low-stakes environments, such as modeling what color coat people buy or whether they stand for an ovation after a show, we will most often assume that people apply fixed rules. When people decide whether or not to cooperate in a venture or to trust another person, we will assume that people learn and adapt. Finally, in high-stakes environments, we will assume that informed, sophisticated people make optimal choices.

Before describing our approach in more detail, we first address some common misconceptions. Many people first encounter formal models of social phenomena in introductory economics courses. Those models often rely on a rudimentary rational-actor model in which everyone is self-interested and capable of optimizing. The model may also assume everyone has the same preferences and income level. Economists then solve for equilibria within these models, enabling them to evaluate the effects of shocks to a market or policy changes. These models, though based on incorrect assumptions, are useful. They help economists to communicate and students to understand.

Based on this experience, many people infer that formal modeling requires a narrow, unrealistic view of human nature, in which people are self-interested and never make mistakes. That is not the case. In fact, not even economists think that to be true. The frontiers of economics consist of models with imperfectly informed, heterogeneous actors who adapt in response to what they learn, and who sometimes, though not always, care about the payoffs to others. The

## The Challenge of Modeling People

Modeling people presents a challenge because while models require low-dimensional representations, people defy simple characterizations. People are *diverse*, we are *socially influenced*, we are *error-prone*, we are *purposive*, and we *learn*. In addition, people *possess agency*—we have the capacity to act.

By way of contrast, physical objects, such as carbon atoms and billiard balls, exhibit none of these six properties. Carbon atoms lack diversity (though they can occupy heterogeneous positions within compounds, such as in propane). Carbon atoms never violate the laws of physics nor do they lead purposive lives. They do not change their behaviors based on past experiences. They lack agency; they do not decide to lead uprisings or switch careers. Hence the oft-repeated quip by social scientists: how difficult physics would be if electrons could think. Physics would be even harder if electrons could write models.

We can start with the problems created by diversity. People differ in our preferences, in our capacity to act, in the social networks we form, in our levels of altruism, and in the level of cognitive attention we allocate to actions. Modeling would be easier if everyone were the same. Sometimes we rely on statistical logic and assume that the behavioral diversity cancels. For example, we might construct a model that predicts charitable donations as a function of income. For a given income level and tax rate, some people may be more altruistic than we assume and others may be less. If the deviations from the model average out (and in Chapter 5 we cover models of distributions that explain why they might), then our model may be accurate. This canceling out of diversity will not occur unless actions are independent. When behavior is socially influenced, extreme actions can create spillovers. This occurs when political activists energize voters. We will encounter this effect of diversity when we model riots.

Whether or not mistakes cancel in the aggregate depends

on the context. Errors that result from a lack of cognitive attachment may be random and independent. Errors that arise from cognitive biases may be systematic and correlated. For example, people may overweight recent events and recall narratives better than statistics. A shared bias like this will not cancel out.

The next challenge relates to what people desire. A central challenge in writing models of people will be making an accurate assessment of their goals and objectives. Some people desire wealth and fame. Others want to contribute to the betterment of their communities and the world. In the rational-actor model, we represent a person's payoff directly in the form of a function. In rule-based models, purposes are more implicit. A behavioral rule in which people seek to live in an integrated neighborhood but move out of a neighborhood if the percentage of people who share their racial identity falls below 10% embeds certain beliefs about what people desire.

The final challenge to modeling people results from the fact that people have agency: the ability to take action, to change what we do, and to learn. That said, in some contexts, people may be better characterized as creatures of habit. Actions may be outside our control. Few people choose to be addicted to opioids or to be poor. Yet people take actions that produce those outcomes.

Often, when people take actions that produce bad outcomes, they adapt their behavior. We can capture this by including learning in our models. How people learn varies by context. When learning how many hours they need to study for an exam in order to get a good grade, or how many times a week they need to exercise, people may learn based on individual experiences and introspection. When learning what grocery store to visit or whether to contribute to a charity, people may learn through observing others. In Chapter 26, we show how in non-strategic contexts, learning generally works. People learn the best action. We also show that in strategic contexts, which we model as games, all bets

are off. Neither individual nor social learning necessarily produces good outcomes.

Each of these six characteristics are potential model features. If we include a feature, we must decide how much of it to include. How diverse do we make our actors? How much social influence do we include? Do people learn from others? How do we define objectives? How much agency do people possess? We may possess less agency than we believe. Jonathan Haidt describes our lack of agency with his metaphor of the rider and the elephant. “The image I came up with for myself, as I marveled at my weakness, was that I was a rider on the back of an elephant. I’m holding the reins in my hands, and by pulling one way or the other I can tell the elephant to turn, to stop, or to go. I can direct things, but only when the elephant doesn’t have desires of his own. When the elephant really wants to do something, I’m no match for him.”<sup>1</sup> Sometimes we do ride the elephant. Sometimes we do not. No single approach to modeling humans will be appropriate in all settings, so we model humans in a variety of ways.

## The Rational-Actor Model

The *rational-actor model* assumes that people make optimal choices given a payoff or utility function. These actions can be *decisions*, where the payoff depends only on the individual’s own action, or they can take place within a *game*, where payoffs depend on what others do. In a game with simultaneous choices or with incomplete information, the rational-actor model also specifies *beliefs* about what the other actors will do.

### Rational-Actor Model

An individual’s preferences are represented by a



mathematical utility or payoff function defined over a set of possible actions. The individual chooses the action that maximizes the function's value. In a game, that choice may require beliefs about the actions of other players.

As an example, we construct a primitive rational-actor model of an individual's decision for how much income to allocate to housing. The model characterizes her utility as a function of housing and all other consumption, with the latter including food, clothing, and entertainment (see box). The model assumes a price for housing and a price for all other goods. The model is far from realistic. It treats all housing the same. And, it lumps all other goods into one category called consumption and prices them all the same. We can set those inaccuracies aside for the moment, as the purpose of the model is to explain the proportion of income spent on housing.

## A Rational-Actor Model of Consumption

**Assumption:** An individual's utility from general consumption,  $C$ , and housing,  $H$ , can be written as follows:

$$U(C, H) = C^{\frac{2}{3}} H^{\frac{1}{3}}$$

**Result:** A utility-maximizing individual (a rational actor) spends exactly one-third of her income on housing.<sup>2</sup>

In the model, the proportion of income a person spends on housing does not depend on the price of housing or on