





# THE WINNER'S CURSE

---

BEHAVIORAL ECONOMICS  
ANOMALIES, THEN AND NOW

---

**RICHARD H. THALER**  
**ALEX O. IMAS**

**SIMON & SCHUSTER PAPERBACKS**

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



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

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

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

Copyright © 1992 by Richard H. Thaler

Updated edition copyright © 2025 by Richard H. Thaler and Alex O. Imas  
Originally published in 1992 by Free Press, a division of Simon & Schuster, Inc.

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

This Simon & Schuster trade paperback edition July 2026

SIMON & SCHUSTER PAPERBACKS and colophon are registered trademarks of Simon & Schuster, LLC

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

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

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

*Interior design by Joy O'Meara*

Manufactured in the United States of America

1 3 5 7 9 10 8 6 4 2

Library of Congress Control Number: 2025944805

ISBN 978-1-9821-6511-6

ISBN 978-1-9821-6512-3 (pbk)

ISBN 978-1-9821-6513-0 (ebook)



Scan here to get book recommendations, exclusive offers, and more delivered to your inbox.

***From Richard:***

*To Greg (plus Olive, Julia, and Alexa),  
Maggie (plus Eric, Hallie, and Aidan),  
and Jessie (plus Evan, Athena, and Phoenix).*

***From Alex:***

*To my parents Dora and Shulim Imas,  
who have given me everything, and my sister Polina Imas.  
To my amazing wife Aislinn Bohren, who has been patient,  
kind, and always there with the best advice.  
To my wonderful boys Saul and Oran Imas.  
And to my advisor Uri Gneezy,  
and honorary advisor George Loewenstein,  
without whom this whole academia thing would have not been possible.*



# Contents

---

Preface ix

1. The Winner's Curse 1

2. Cooperation 26

*with Robyn M. Dawes*

3. The Ultimatum Game 46

*with Colin F. Camerer*

4. The Endowment Effect, Loss Aversion,  
and Status Quo Bias 65

*with Daniel Kahneman and Jack L. Knetsch*

5. A Primer on the Psychology of Risky  
Decision Making 88

6. Be Careful Before You Call Something Risk Aversion 103

*with Matthew Rabin*

7. Choosing Between Now and Later 130

*with George Loewenstein*

8. Savings, Fungibility, and Mental Accounts 156

9. Preference Reversals	184
<i>with Amos Tversky</i>	
10. Utility Maximization	199
<i>with Daniel Kahneman</i>	
11. A Brief Digression on the Efficient Market Hypothesis	221
12. The Law of One Price	227
<i>with Owen Lamont</i>	
Epilogue	247
Acknowledgments	265
Notes	267
Bibliography	285
Index	307

Names listed were co-authors of the original articles  
upon which these chapters are based.

# Preface

---

Notes from the Authors

## Part 1: By Richard Thaler

The origins of this book began in 1986. At home, the music in the air was “Brothers in Arms” by Dire Straits and “Born in the U.S.A.” by Bruce Springsteen. Ronald Reagan was president of the United States and pushed a major tax reform through Congress. Argentina won the World Cup aided by the “hand of God,” and the Chicago Bears won the Super Bowl. Gary Becker was the president of the American Economic Association.

Personally, I was about a decade into a high-risk professional adventure. I had become increasingly dissatisfied with the usual way economics was practiced, and I became convinced it was possible to enrich the field by adding insights from other social sciences, especially psychology. Having this conviction was one thing, figuring out how to convince my fellow economists was quite another. There was no well-established career path to follow. When I started, I was not sure I would be able to earn a living pursuing this idea.

But, to my surprise, my decision to do research in the field that eventually became known as “behavioral economics” had worked

out better than I could have ever expected. Not only did I manage to get an academic job at Cornell University's business school, but I had recently been granted tenure and had just returned from a productive year in Vancouver, British Columbia, visiting with my friend and mentor, Daniel Kahneman. One day during my visit, Danny (as he insisted everyone call him) had mused aloud that in academia there is an age at which someone can no longer be considered "promising." He thought it might be around age 40, not consciously aware that I was 39 at the time. Clearly, I needed to figure out what would constitute the "post-promising" stage of my career.

Sometime that year, I happened to have lunch with an economist friend, Hal Varian. Hal told me that he was joining the editorial board of a new academic journal being launched by the American Economic Association (AEA), called the *Journal of Economic Perspectives* (*JEP*). The profession was becoming increasingly specialized—so much so that it was hard for an economist from one field (say, labor economics) to understand academic papers in other fields (such as macroeconomics or finance). Articles in the *JEP* would be written in a way that any economist, or even undergraduate economics major, could understand them. Joseph Stiglitz was picked to be editor, with Carl Shapiro as co-editor (both stars in the profession), and in a stroke of genius, Timothy Taylor was picked to be the managing editor. It fell upon Tim to help economists translate their first drafts into readable prose. Luckily for the journal and the profession, Tim was very good at that task and has remained in the job to this day. In addition to Hal Varian, a leading theorist (who later became the chief economist at Google, a title he retained until recently), the editorial board included other luminaries such as Paul Krugman, Larry Summers, and Janet Yellen.

Hal told me that the *Journal* was considering including some recurring feature articles in every issue. The extraordinarily clever Yale economist Barry Nalebuff had agreed to write one on puzzles

and brain teasers—not normal material for an economics journal! At some point during our lunch, we hatched the idea of having a feature on economics anomalies, which would be defined as empirical observations that were inconsistent with standard economic theory. I told Hal that I was pretty sure I could come up with at least a dozen. He sold the idea to Joe and the rest of the editorial board, and I was off and running. But exactly to where I was running was decidedly unclear.

Despite the all-star editorial board, the future of the *Journal* was highly uncertain. Who would want to write articles for such a journal and how would they be received? Before my participation in this effort was widely known, I overheard a senior colleague at Cornell opine that articles in this journal certainly should not “count.” I was tempted to ask: Who was keeping score and what exactly were they counting? Writing articles that are easy to understand is not a skill that is well rewarded in academic economics, so this would be yet another risky career move. Still, tenure is supposed to give academics the freedom to take risks, and writing the columns would be fun—my principal criterion for deciding whether to undertake any task (including this book).

We will never know what tally my Cornell colleague was using when he said that articles in the *Journal of Economic Perspectives* would not “count,” but at least in hindsight, they did score highly on one important dimension: eyeballs reached. At one point early on, before anything had the possibility of going viral online, the AEA did a survey of members, and half of them (meaning thousands) reported reading my Anomalies columns “regularly,” an astonishing fact. Many academic articles are lucky to have readership reach three digits.

To define an anomaly, I offered this description in the first column: “Economics is distinguished from other social sciences by the belief that most (all?) behavior can be explained by assuming that agents have stable, well-defined preferences and make rational

choices consistent with those preferences in markets that (eventually) clear. An empirical result is anomalous if it is difficult to ‘rationalize,’ or if implausible assumptions are necessary to explain it within the paradigm.” Along with rationality, an additional component of economic theory as practiced was to assume that “agents” (the actual word economists use to refer to what everyone else calls “people”) are selfish, meaning they mostly care about payoffs (monetary or otherwise) to themselves, and possibly to close friends and (some) members of their family.

The journal came out four times a year, which meant I had to write a new column quarterly. That may not seem like a lot, but these were more like short academic articles than newspaper op-eds. They were typically about ten journal pages. Each column had to state the case for why the facts reported were a legitimate critical test of the standard theory, and therefore anomalous, and then provide some commentary on what it all meant. To qualify as a critical test, the result could not be (easily) explained without straining the usual assumptions of economic theory. It helped that I was able to recruit co-authors for many of the articles, including both Danny Kahneman and Amos Tversky (separately), as well as fellow economists who would make major contributions to the field we were trying to create: Colin Camerer, Owen Lamont, George Loewenstein, Andrei Shleifer, and more. The topics ranged from laboratory experiments about how people respond to risk or choices over time, to the workings of various markets. One thing that the columns had in common is that these anomalies were important for *economics*. The departures from the standard “rational” model were not marginal; in some cases, the core empirical predictions were violated (such as universal free riding—Chapter 2), and in others, the core assumptions of the models themselves were proven wrong—such as the case of Expected-Utility Theory (Chapter 5). So, while not everyone agreed with the columns, they did spur a heated discussion across the entire field.

Although writing the columns was indeed fun, after writing thirteen of them, I told the editors I wanted to get off the treadmill and submit new columns only when the spirit moved me. (Needless to say, this happened much less often than four times a year.) I took this break when I thought I might have enough material for a book. The Free Press agreed to publish the book and we decided to call it the *Winner's Curse*, which was the title of one of the columns.\* Soon after the book came out, my editor, Peter Dougherty of the Free Press, left to become the publisher at Princeton University Press, and he sold himself the rights to the paperback edition. He assured me this would be good for the book because it meant the book would never go out of print.

Imagine my surprise, then, when I was approached by Simon & Schuster (who had bought the Free Press) and was informed that *The Winner's Curse* was now going out of print. They asked whether I might be interested in publishing a new edition, maybe with a new preface? I told them I had written another half dozen columns since the book had been published, so maybe we could do something a bit more ambitious and add some of the new material? Of course, this is easier to say than do, but since the actual work would not begin for a year or so, I was not that worried (see Chapter 7 for why). I realized that it would help to add someone with a different perspective, and that someone is my co-author, Alex Imas.

To be clear about this, Alex was of no help whatsoever in crafting the original versions of these columns. He claims to have a legitimate excuse, as he was in kindergarten in Bender, Moldova, when I was getting started. If for some reason you are unfamiliar with Bender, it is about 600 kilometers due south of Chernobyl,

---

\* While visiting Kahneman in Vancouver, I had gotten to know his brilliant stepdaughter Deborah Treisman, who was then in high school. Deb went on to be an English major at UC Berkeley and was a recent graduate when the book was being finished. I hired her as a freelance copy editor. I am pretty sure this gig played a major role in her later becoming the fiction editor at *The New Yorker*, a job she has now held for decades.

which in April 1986 had the nuclear power plant disaster that is its unfortunate claim to fame. So, although Alex is a behavioral economics prodigy (despite, or perhaps because of, the radioactive fallout in the neighboring city), I guess he can be excused for not helping out with the original columns.

I met Alex when he was a graduate student at the University of California at San Diego in 2009. At that time, I used to spend my winters in that part of the world and had an office at the Rady School of Management. Alex had an office nearby, and we would frequently chat about research, including his thesis, which was on a topic near and dear to my heart: mental accounting (see Chapter 8). We found that we enjoyed talking to each other and kept in touch when he began his career at Carnegie-Mellon University, where he immediately started writing a slew of interesting papers on behavioral economics. In 2020, my home institution, the Booth School of Business, at the University of Chicago, managed to pry Alex away from CMU, and he and I were able to continue our conversations more often. I am happy to report that he was recently promoted to full professor.

## Part 2: By Alex Imas

While I had not contributed to the writing of the original Anomalies columns on account of being in grade school, Richard's writing had a huge impact on my life. In 2007, I was a newly minted college graduate seemingly on my way to medical school. As a first-generation immigrant, my parents were confident I was either going to be a doctor or—if I really screwed up—I'd do "something with computers." I discovered economics as a happy accident when an introductory class happened to fulfill a course requirement; I quickly fell in love with it. I also took quite a few psychology classes as part of my pre-med requirements, but I never saw the connection between the two.

Then, during a long drive helping my then-girlfriend move across the country from Chicago to LA, I happened to be listening to the radio. The host was interviewing an author of a new book that challenged the idea people were fully rational and highlighted the importance of bringing psychology to economics. The book was *Nudge* and the author was Richard Thaler. I bought the book, co-authored with Cass Sunstein, as soon as we arrived in Los Angeles. And I was hooked. Behavioral economics seemed to combine everything I loved about economics and psychology. It used rigorous formal models to understand complex behavior in markets, but put *humans*—yes, with all their follies and mistakes—front and center. I read as much as I could about it, and I decided to try to make a career studying this field. I thought the best way to go was to get a doctorate in economics, preferably at a place with at least a few faculty doing research in behavioral economics. Less than twelve months later, I was taking my first class as a graduate student at the University of California at San Diego.

I was lucky that UC San Diego had several prominent behavioral economists on the faculty, some of whom make it to the pages of this book. James Andreoni, an early pioneer of research on altruism and cooperation, was in the economics department, and Uri Gneezy—who ultimately became my advisor—was at the business school. Uri guided me through the first few years of the grueling classwork and exposed me to much of the modern work in behavioral economics in his graduate course. We review some of this modern work in the Updates to each chapter (discussed more later), as well as in the Epilogue of the book. Through him, I was also able to get an office at the brand-new business school—an office that happened to be adjacent to none other than Richard Thaler.

Although I was initially too intimidated to approach Richard, we ended up striking up a conversation, which eventually turned into a regular occurrence. I'd come up with an idea, he'd tell me

about the paper that already showed that in the 1980s, and I'd go back to the drawing board. And so it went for a couple of years. Eventually, I hit on an idea of using mental accounting to explain how people respond to prior gains and losses (more on this in Chapter 8). While previous work had studied this topic, the results were inconclusive. Sometimes studies found that people responded to losses by taking on more risk; other times they found the opposite: that people took on less. It was a mess. I got the idea that mental accounting (the way people encode economic outcomes) could explain this inconsistency by determining *how* the prior loss gets psychologically classified: whether people are anxious to recover from it (the account is still open) or whether they make peace with the loss and move on (the account is closed). Through several iterations (and endless conversations), this became my job-market paper, which landed me a position in the storied Social and Decision Sciences Department at Carnegie Mellon University, where George Loewenstein, Linda Babcock, and many of the other characters in this book wrote groundbreaking papers in behavioral economics.

Getting to Carnegie Mellon was one thing. But then I had to get ready for something that all academics have to do but that graduate school does little to prepare us for: I had to teach! Although undergraduates had their pick of courses in psychology, economics, and even decision science, one course was clearly missing from the curriculum: behavioral economics. Against my better judgment (and the advice of my advisors), I volunteered to design and teach this course in my first year as an assistant professor.

I quickly realized when planning the course that there was no commonly used textbook in behavioral economics; in fact, a textbook didn't seem to exist at all. This is where Richard Thaler's writing was a boon again. I organized a course around his original Anomalies columns. For the class on social preferences, the students read the columns on Cooperation and Ultimatum Games.

These columns quickly introduced them to the myriad studies showing that people are not in fact selfish money maximizers, as suggested in some other classes they had taken. For the class on risk preferences, the students read the columns on risk aversion and the equity premium puzzle. There, we learned that standard estimates of “risk aversion,” as normally defined in economic models, do a poor job of explaining actual behavior. For one thing, attitudes toward risk are highly context dependent: The same person may pay an excessive amount to buy insurance on a small purchase, like a toaster oven, and also buy lottery tickets. In many contexts, a more useful concept is “loss aversion”: Roughly speaking, losses hurt more than corresponding gains give pleasure. And so on. The columns were clear, easy to read, and communicated complicated concepts that may be difficult for a non-economist to extract from the arcane, byzantine writing style of the original scientific papers. Importantly, the students loved them. So when Richard approached me to help write the “now” part of the current book, I jumped at the opportunity.

As Richard often likes to say, the success of behavioral economics did not occur by changing the minds of old economists who were set in their ways but, rather, by “corrupting” young minds. This is the reason he, along with Colin Camerer and Danny Kahneman, started a two-week boot camp in 1994 for early-stage graduate students in economics and related disciplines. In recent years, the camp has been organized by David Laibson and Matthew Rabin, now professors at Harvard, who were attendees of that first camp. Many dozens of other alumni are now in leading universities around the world. I was part of a later cohort in the Laibson-Rabin era as well (although I didn’t need much indoctrination by that point; I had fully drunk the Kool-Aid). We hope that, with this book, we can continue the “corruption” of both young minds and those that are, uh, more mature, and that you enjoy reading it as much as we did writing it!

## How We Wrote This Book (from Both of Us)

There are many ways we could have gone about refreshing a book that was published thirty years ago. We had as raw material the thirteen chapters of the original book, plus the six new Anomalies columns that appeared in the years after the book came out. The first decision was what to include. Given that we were planning to add some entirely new material, and we did not want to publish an unwieldy doorstop, it was necessary to make some hard decisions about what topics to include. Our first criterion was to limit ourselves to topics that would be both accessible and inherently interesting to non-economists (costing us a chapter on foreign exchange markets and another on “inter-industry wage differentials”). The other factor we weighed was whether we felt the explanation for the anomaly was primarily “behavioral”: Did it illustrate something important about how people make decisions? This guideline cost us a chapter on some very odd seasonal patterns in the stock markets (markets go up on Fridays and down on Mondays, for example, and the turn of the year is also a good time for the stock market). For sure this was an anomaly, but no one had any explanation for it, behavioral or otherwise. The famous financial economist Richard Roll published an article on the so-called January effect that was titled “Vas Ist Das?” Similarly, some strange behavior of racetrack bettors and lottery ticket buyers did not make the cut. We can offer some free advice from that column: Don’t buy lottery tickets or bet on horse racing, but if you insist on doing so, avoid popular numbers for lotto tickets and don’t bet on long shots at the track. And not to worry, if you really want to read these or any of the other original columns, they are available for free at the *Journal of Economic Perspectives* website.

For the chapters we did decide to include, we did some combining, rearranging, and editing, but we chose to impose an unusual rule upon ourselves: For the primary body of the chapter,

we could rewrite and combine as much as we wanted, but *we could not introduce any new empirical facts or theoretical findings that did not exist at the time the original column was published*. One reason for imposing this highly unusual rule on ourselves was to maintain both the spirit and the timing of the original columns, which were written with the goal of capturing the attention of economists and getting them to consider some surprising evidence. In some places, you may feel we are arguing with straw men, but please believe us: There were very real economists who disagreed with the premise of each of these articles, many of whom wrote articles attempting to show why the alleged anomaly was an illusion or could easily be explained using standard economics concepts.

Another reason for this rule was to create a level playing field for the new material in this book, the writing of which Alex took the lead on, providing an Update to almost every chapter. The Updates answer these questions: Is this stuff still true? And if so, has it been shown to matter for the real world? It would not be fair to state we are testing the original claims about anomalies if we had the freedom to go back and change or remove passages that would later seem embarrassing. Again, feel free to check our work!

The addition of the Updates is the reason we felt this new version of the book was worth our time writing and your time reading. Attempting to synthesize thirty years of research on each of the topics took us much longer than we had initially planned (again, see Chapter 7). In the three years since we started working on the book, Alex has gone from “promising” to “can no longer be called promising.” But we believe that asking whether the original claims still hold up is an important question, for two reasons. First, the original columns were, after all, about *anomalies*. Anomalies are unexpected. Maybe some were chance observations, underpowered experiments, mismeasurements, or temporary phenomena that have since vanished. In other words, maybe behavioral economics is really much ado about nothing, a view then held by many econo-

mists.\* The Updates show this was not the case; the original lab experiments grew into a large body of work spanning different methods and subfields of economics. In fact, one way in which behavioral economics has evolved over time is that, after confirming the robustness of the laboratory findings, researchers went on to document that the same anomalies showed up in the field (often referred to as the “real world” by critics). It turns out that college students are not the only people whose behavior produces violations of the standard model; the behaviors of households, portfolio managers, business executives, and professional athletes are just as likely to produce anomalies.

Another reason for providing the Updates is that behavioral economics explicitly borrows some findings from psychology, and there are some subfields of psychology that have been experiencing what has been referred to as a “replication crisis.” For example, there are many reported results about a phenomenon called *priming*. The idea is that, if someone is exposed to some stimulus—say, a picture of a delicious-looking food item—that stimulus may influence some subsequent behavior (such as the choice of what to eat on some subsequent occasion) without the person ever being aware that their behavior has been altered. Undoubtedly, priming as a phenomenon certainly exists, but there were some articles that presented surprisingly impressive demonstrations of priming that other researchers were unable to reproduce. In light of these failures to replicate, Daniel Kahneman felt it necessary to disavow the attention he gave to these findings in his bestselling book, *Thinking, Fast and Slow*. Worse, recent investigations have even uncovered outright data fabrication in psychology papers on some topics that are at least adjacent to behavioral economics.

We are pleased to report there is no evidence of anything that

---

\* When Richard arrived at the University of Chicago in 1995, one of his new colleagues described the field as a “passing fad.”

could be called a replication crisis for the set of topics covered in the Anomalies columns upon which this book is based. To the contrary, many of the findings we discuss are so robust that they can be reliably reproduced as classroom demonstrations. Why are we so confident? We know that in part because we have done such demonstrations in our own classes for a decade (Alex) or many decades (Richard).

In Alex's case, he routinely replicated many of the experiments "live" in his classes. Here is how: For every topic, students would form groups and design replications of the main experiments we had been discussing. They would work together to incorporate the instructions into a survey platform like Qualtrics. The following class would begin with launching the experiments at the start of class, collecting data from participants via online platforms. Data would usually come in within the hour, and we would spend the rest of the class analyzing the results and seeing if the experiments indeed replicated; in nearly a decade of teaching the class, Alex has not seen a single "classic" study fail to replicate.

Of course, that does not mean every paper or classroom demonstration obtained exactly the same result, and it certainly does not mean that every researcher agreed on the appropriate interpretation of each finding. But we can say that in writing this book, we found that all the critical conclusions about anomalies have held up remarkably well, even after decades of subsequent research.

But don't take our word for it. Nearly every chapter in this book is built around evidence either from a "focal" laboratory experiment (e.g., Chapter 3) or from analysis of observational data (e.g., Chapter 12). The explosion of online crowdsourcing platforms such as Amazon's Mechanical Turk and Prolific has made it substantially easier to collect experimental data from human subjects. For the focal experiments, we have replicated the results on Prolific and include the relevant analyses in the Online Materials for the book, which are available at [www.TheWinnersCurse.org](http://www.TheWinnersCurse.org). What's

more, we also include the specific instructions so anyone can replicate the studies themselves! For the results using observational data, whenever possible, we have replicated the original analyses using just the time period since the original research was published, creating an out-of-sample test. These analyses are reported in the Updates, with further details included in the Online Materials as well. We hope you take the opportunity to do some behavioral economics yourself.

## How to Read This Book

Although we have tried to write this book in one coherent style, as you make your way through it you should be aware you are reading the thoughts and words of many economists and psychologists, and should remember that each chapter begins with a component initially written in the 1980s or 1990s, usually with a co-author. Although his memory is getting a bit hazy, Richard can pick out certain turns of phrase or conceptual ideas and know they were his co-author's thoughts. There is a fun riff on baseball umpires that Amos Tversky had proposed in spite of the fact that he was no fan of baseball. (He was a basketball nut.) In contrast, the Updates are clearly written in the twenty-first century. We have put these Updates in a different format just to nudge you to remember that this is the two of us talking now, as opposed to Richard and his friends pontificating three or four decades ago.

We should also make clear that, although Alex has used the original Anomalies columns in the class he teaches, this book is not intended to be a comprehensive treatise on behavioral economics. Far from it!\* For example, there is nothing in this book

---

\* The economist Sanjit Dhani (2020) has produced a thorough treatment that encompasses seven (thick) volumes.

about macroeconomics, because Richard could not come up with a sharp anomaly in that field. You might think that the Great Depression, in which economies around the world suffered huge stock market declines and enormous unemployment, might be a good candidate for an Anomalies column. After all, one can ask why wages did not fall enough to allow the market to clear at something like full employment (say, 5 percent rather than close to 25 percent). Rather than reducing wages enough to keep most workers employed, firms often chose to reduce the size of their workforce. (To be clear, some wages did fall, but not enough to prevent massive unemployment.) There was no column on this because these facts were well known to economists and there was no agreement about the cause.

We suggest that readers treat each of these chapters as a puzzle. A set of facts is described, and we (and the original authors) claim these facts are inconsistent with standard economic theory. Then, in the Update, we review the more recent evidence and highlight how the original anomalies—which were documented in the lab or in econometric analyses—are just as prevalent in current real-world settings. But one small spoiler alert: We do not think there has been a paradigm shift in the field of economics. We discuss in the Epilogue why this hasn't been the case. But as you are reading, we encourage you all, both economists and others, to try to think of new anomalies. It is fun!

One final note. The order of the chapters is somewhat arbitrary. If the topic of a chapter does not interest you, feel free to skip it and move on. But a hint: You don't want to miss the last chapter on the Law of One Price. It may be our favorite.



# THE WINNER'S CURSE



## CHAPTER 1

---

# The Winner's Curse<sup>\*</sup>

Next time you find yourself a little short of cash for a night on the town, try the following game in your neighborhood bar. First, take a jar and fill it with coins, noting the total value of the coins. Then, perhaps well into the evening, auction off the jar to the assembled tipplers, offering to pay the winning bidder in bills to control for coin aversion. (If the winner doubts your estimate, he can take possession of the jar and examine its contents.) Here is what is likely to happen:

1. The average bid will be significantly less than the value of the coins. (Bidders are risk averse.)
2. The winning bid will exceed the value of the jar.

In conducting this demonstration, you will have simultaneously obtained some funding for your evening's entertainment and enlightened the patrons of the bar about the perils of the "winner's curse." We advise buying the winner a drink with some of your profits.

---

<sup>\*</sup> This chapter is based on Thaler (1988).

The winner's curse is a concept that was first discussed in the scientific literature by engineers at the oil company Atlantic Richfield. Their company and many others had been participating in auctions run by the government, in which the winner obtained the right to drill in a specific geographical area. The engineers were puzzled about the fact that, although they always hired highly qualified experts to evaluate the attractiveness of each drilling site, when they won the rights to drill on a site, the average amount of oil at that site was typically less than the expert predicted.

Here was the firm's insight: Atlantic Richfield's bid only matters when they win the auction, *but winning is bad news about the amount of oil at the site*. When a firm wins the auction, it means that all the other firms probably had lower estimates on the amount of oil, so the winning firm's estimate is probably too high. A rational firm would account for this bad news at the time of bidding, shading their bid in anticipation. But if firms don't figure this out, there will be a winner's curse: Winners will, on average, find that there is less oil down there than they expected.

You might think that, given the amount of money at stake and the presence of highly qualified experts, over time the bidding firms would learn the optimal bidding strategy. However, figuring out the right thing to do is not easy. Because of the considerable difficulty of predicting the amount of oil that will be found in a given location, even the estimates of unbiased experts will vary substantially—some far too high and some too low. Companies may incorporate this uncertainty by bidding conservatively, well below the estimate their experts provided, but the firms whose experts provided high estimates will tend to bid more than the firms whose experts guessed lower. Indeed, it is likely that the firm that wins the auction will often be among those whose experts provided the highest estimates. If this happens, the winner of the auction is destined to be cursed, in that the value of the tract is less than the expert's estimate and the winning firm is (on average) disappointed.

It can be proven that a winner's curse will not occur if all the bidders are fully rational, so evidence of a winner's curse in market settings qualifies as an anomaly. But look at what's involved in acting "rationally" here. Rational bidding requires distinguishing between (1) the expected value of the object for sale, conditioned only on the prior information available, and (2) the expected value conditioned on winning the auction. Just estimating the expected value of the object is hard to do, but the real difficulty is trying to account for the behavior of the other bidders. Even if a bidder grasps this basic concept, the winner's curse can occur if winning bidders underestimate the magnitude of the adjustment necessary to compensate for the presence of other bidders.

One key ingredient for a smart bidder to consider is the number of other bidders. Here is a quick quiz you can try. Suppose you are in that bar thinking about bidding on the jar of coins, when a dozen folks in another part of the room hear about the fun and decide to bid as well. There is now more competition to win that auction. Should you change your bid? If so, do you want to bid more or less? Answering this question requires weighing two factors that work in opposite directions. An increase in the number of other bidders implies that to win the auction, you must bid more aggressively, but their greater presence also increases the chance that if you do win, you will have overestimated the value of the object for sale—suggesting that you should bid less aggressively. As the oil engineers so aptly put it: "If one wins a tract against two or three others, he may feel fine about his good fortune. But how should he feel if he won against 50 others? Ill."

As you can probably tell, solving for the optimal bid is not trivial. It is thus an empirical question whether bidders in various contexts get it right or are cursed. The evidence from both experimental and field studies suggests that the winner's curse appears to be a common and robust phenomenon.

## Experimental Evidence

The jar of coins example has in fact been carried out as an experiment by Max Bazerman and William Samuelson. Their subjects were MBA students taking microeconomics classes at Boston University—a population that should, at least in theory, be better at figuring this out than participants picked out from a random bar. The objects auctioned off were jars of coins or other objects such as paper clips valued at four cents each. Unknown to the subjects, each jar had a value of \$8.\* Subjects submitted sealed bids and were told the highest bidder would receive the value of the object, less his or her bid. A total of 48 auctions were conducted, 4 in each of 12 classes. No feedback was provided until the entire experiment was completed. Participants were also asked to estimate the value of each jar, and a prize was offered for the best guess in each class.

The estimates of the actual values turned out to be too low. The average estimate of the value of the jars was \$5.13—well below the true value of \$8. This makes sense: Any dislike of risk even at these low stakes would tend to work against observing a winner's curse. Nevertheless, the mean winning bid was \$10.01, producing an average loss to the winning bidder of \$2.01 and implying a profit to the auctioneer. Clearly these experiments do not require large experimental budgets. In auctions like this, it is better to be a seller than a bidder.

Samuelson and Bazerman have run another series of experiments about the winner's curse in a different context, using the following problem:

---

\* If you are worried that the amounts of money in the experiments we report seem too small to get excited about, note that we are not adjusting these for inflation (multiply by about 3 to do that). Also, the issue of stakes will be a recurring theme in this book, and the level of the stakes rarely plays a big role in understanding the phenomena. And in this case, remember that the phenomenon was discovered using evidence from the auction for oil leases. Big stakes!

In the following exercise, you will represent Company A (the acquirer) which is currently considering acquiring Company T (the target) by means of a tender offer. You plan to tender in cash for 100% of Company T's shares but are unsure how high a price to offer. The main complication is this: The value of the company depends directly on the outcome of a major oil exploration project it is currently undertaking.

The very viability of Company T depends on the exploration outcome. In the worst case (if the exploration fails completely), the company under current management will be worth nothing—\$0/share. In the best case (a complete success), the value under current management could be as high as \$100/share. Given the range of exploration outcomes, all share values between \$0 and \$100 per share are considered equally likely. By all estimates the company will be worth considerably more in the hands of Company A than under current management. In fact, whatever the value under current management, the company will be worth 50% more under the management of Company A than under Company T.

The board of directors of Company A has asked you to determine the price they should offer for Company T's shares. This offer must be made now, before the outcome of the drilling project is known.

Thus, you (Company A) will not know the results of the exploration project when submitting your offer, but Company T will know the results when deciding whether or not to accept your offer. In addition, Company T is expected to accept any offer by Company A that is greater than or equal to the (per share) value of the company under its own management.

As the representative of Company A, you are deliberating over price offers in the range \$0/share to \$150/share. What offer per share would you tender?

Think through this problem before continuing and decide what you would bid.

• • •

The typical subject considers this problem roughly as follows: “The firm has an expected value of \$50 to Company T, which makes it worth \$75 to my company. Therefore, if I suggest a bid somewhere between \$50 and \$75, we should make some money.” This analysis fails to take into consideration the asymmetric information built into the problem. The asymmetry is that the firm being sold (Company T) knows the true value, but the buying firm (Company A) does not. As is true of all the examples in this chapter, the correct analysis must calculate the expected value of winning the auction *conditional on the bid being accepted*. To see this, let’s work through an example.

Suppose you bid \$60. If the bid is rejected, meaning the firm was worth more than \$60 to the seller, you do not win or lose anything. But if the bid is accepted, the company must be worth no more than \$60 under the current management (or they would not have accepted your offer). Since all the values less than \$60 are equally likely (by assumption in this example), this implies that the firm, on average, will be worth \$30 to the current owners, or \$45 to you. That means if you bid \$60, you expect to lose \$15.

Make this calculation for a bunch of numbers. You’ll quickly realize a startling fact: Despite Company T being worth 50 percent more to Company A than its current owner, for any bid that is greater than zero, Company A will on average lose 25 percent of its bid. Thus, this problem produces an extreme form of the winner’s curse, in which any positive bid yields an expected loss to the bidder. It is a wicked environment!

This experiment was run in two conditions—one with monetary incentives and one without. The results are quite similar for

the two conditions, with the bids in the condition with monetary incentives somewhat lower. In both conditions, over 90 percent of the subjects make positive bids, and a majority are in the range between \$50 and \$75.

• • •

Economists often respond to examples like this by hypothesizing that although people can be fooled once or twice by such a problem, they will figure out the trap with experience. In this book, we call this the *confused-subjects hypothesis*. (It comes up again in the next two chapters that are based on experimental data.) One study investigated this hypothesis by giving the “buy-a-firm” problem to 69 Northwestern University MBA students. Each student repeated the experiment twenty times with financial incentives and feedback after each trial. The feedback included the “true” value of the company (which was publicly selected at random on each trial), whether their bid was accepted, and how much money they made or lost.

Everyone knows the expression: “Fool me once, shame on you. Fool me twice, shame on me.” We personally don’t know about any expression that covers twenty opportunities to repeat the same task in which any bid is a bad idea. Of the 69 subjects, only 5 managed to learn to bid \$1 or less by the end of the experiment. Even for these 5 subjects, the average trial in which they did so was Trial 8, so their learning was far from speedy. As for the rest of the participants, there was no sign of any learning whatsoever; in fact, the average bid drifted *up* over the last few trials. It may be possible to learn to avoid the winner’s curse in this problem, but the learning is neither easy nor fast.

Another series of experiments has been conducted by John Kagel and his colleagues to see whether people understand how to adjust their bid to the number of other participants. Many of

the experiments have the following structure: An object is to be sold using a sealed-bid auction. The value of the object varies from one trial to another, but is always between the announced low and high limits. Before the bidding, each bidder is given a private clue about the value of the object; the clue is picked at random from the possible values. You can think about this clue as corresponding to expert estimate in the oil bidding scenario.

An auction is then conducted, the bids are posted, and the winner has a profit or loss credited to his or her account. (Bidders were given some capital to start with, usually about \$10. Once their account reached zero, they were no longer permitted to bid.) The experiments varied on some technical details we will spare you, but the crucial one was the number of bidders. Typically, subjects first participated in small groups of three to five bidders, and then in “large” groups of six or seven. A nice feature of all the experiments is that for each trial, the authors compute the outcome that would be predicted by a bidding model in which everyone bids optimally.

The results varied with group size. In the small groups, bidders usually made money. But the profits were on average about two-thirds of what rational bidders would earn. However, in the larger groups, there were losses of \$0.88 per auction, in contrast to the \$4.68 profits predicted by the rational model. The winner’s curse emerged in the larger groups because the subjects bid more aggressively as the size of the group and the amount of perceived competition increased—just the opposite of what the rational strategy dictates.

This is the crucial feature of the winner’s curse, and it is counterintuitive. It is natural to think that the more bidders you are competing with, the higher you should bid. This would indeed be true if your goal was just to be the winner of the auction. However, if the goal is to only win auctions when the value of the object is more than the price you pay, you have to bid less as the number of other bidders goes up. Tricky!

Kagel's team also reports on a series of low-price auctions, meaning the winning bidder is the one who submits the lowest bid, as is common in auctions to find the low bidder for some construction project. (Think home improvement projects.) In these auctions, there were losses in both the small groups and the larger groups. However, the most interesting and innovative feature of the paper is that they included an experiment with a group of construction firm managers. (A common criticism of experimental economics, especially if the results of the experiment do not accord with the predictions of economic theory, is that the subjects are "only college students working on toy problems, and in the real world, experts wouldn't make these silly mistakes"—in other words, the *confused-subject hypothesis*.) So how did the construction managers do?

Although the experimenters worried that the experts might bust their research budget by bidding so cleverly, in fact the experts did no better or worse than the students. This is surprising, given that construction firms participate in low-bid auctions all the time and would soon go bankrupt if they fell prey to the winner's curse. The authors believe the result occurred because in their businesses, the managers have learned situation-specific rules of thumb rather than the relevant theory. For example, they might just include a fudge factor such as "construction markup" in all their bids that covers (some of) the unexpected costs that arise during the project. Firms that do this stay in business, but they have not internalized the deeper lesson, which is that the more bidders there are, the bigger the construction markup needs to be.

## Field Data

The laboratory evidence demonstrates that avoiding the winner's curse is not easy. Even experienced subjects who are given signifi-