


GAME THEORY

AN INTRODUCTION



STEVEN TADELIS

GAME THEORY

GAME THEORY

AN INTRODUCTION

Steven Tadelis

PRINCETON UNIVERSITY PRESS
Princeton and Oxford

Copyright © 2013 by Princeton University Press

Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540
In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock,
Oxfordshire OX20 1TW
press.princeton.edu

All Rights Reserved

Library of Congress Cataloging-in-Publication Data

Tadelis, Steve.

Game theory : an introduction / Steven Tadelis.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-691-12908-2 (hbk. : alk. paper)

1. Game theory. I. Title.

HB144.T33 2013

519.3—dc23

2012025166

British Library Cataloging-in-Publication Data is available

This book has been composed in Times Roman and Myriad using ZzT_EX by Princeton Editorial Associates Inc., Scottsdale, Arizona

Printed on acid-free paper. ∞

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

PART I Rational Decision Making

Chapter 1	The Single-Person Decision Problem	3
1.1	Actions, Outcomes, and Preferences	4
1.1.1	<i>Preference Relations</i>	5
1.1.2	<i>Payoff Functions</i>	7
1.2	The Rational Choice Paradigm	9
1.3	Summary	11
1.4	Exercises	11
Chapter 2	Introducing Uncertainty and Time	14
2.1	Risk, Nature, and Random Outcomes	14
2.1.1	<i>Finite Outcomes and Simple Lotteries</i>	15
2.1.2	<i>Simple versus Compound Lotteries</i>	16
2.1.3	<i>Lotteries over Continuous Outcomes</i>	17
2.2	Evaluating Random Outcomes	18
2.2.1	<i>Expected Payoff: The Finite Case</i>	19
2.2.2	<i>Expected Payoff: The Continuous Case</i>	20
2.2.3	<i>Caveat: It's Not Just the Order Anymore</i>	21
2.2.4	<i>Risk Attitudes</i>	22
2.2.5	<i>The St. Petersburg Paradox</i>	23
2.3	Rational Decision Making with Uncertainty	24
2.3.1	<i>Rationality Revisited</i>	24
2.3.2	<i>Maximizing Expected Payoffs</i>	24
2.4	Decisions over Time	26
2.4.1	<i>Backward Induction</i>	26
2.4.2	<i>Discounting Future Payoffs</i>	28
2.5	Applications	29
2.5.1	<i>The Value of Information</i>	29
2.5.2	<i>Discounted Future Consumption</i>	31
2.6	Theory versus Practice	32
2.7	Summary	33
2.8	Exercises	33

PART II Static Games of Complete Information

Chapter 3	Preliminaries	43
3.1	Normal-Form Games with Pure Strategies	46
3.1.1	<i>Example: The Prisoner's Dilemma</i>	48
3.1.2	<i>Example: Cournot Duopoly</i>	49
3.1.3	<i>Example: Voting on a New Agenda</i>	49
3.2	Matrix Representation: Two-Player Finite Game	50
3.2.1	<i>Example: The Prisoner's Dilemma</i>	51
3.2.2	<i>Example: Rock-Paper-Scissors</i>	52
3.3	Solution Concepts	52
3.3.1	<i>Assumptions and Setup</i>	54
3.3.2	<i>Evaluating Solution Concepts</i>	55
3.3.3	<i>Evaluating Outcomes</i>	56
3.4	Summary	57
3.5	Exercises	58
Chapter 4	Rationality and Common Knowledge	59
4.1	Dominance in Pure Strategies	59
4.1.1	<i>Dominated Strategies</i>	59
4.1.2	<i>Dominant Strategy Equilibrium</i>	61
4.1.3	<i>Evaluating Dominant Strategy Equilibrium</i>	62
4.2	Iterated Elimination of Strictly Dominated Pure Strategies	63
4.2.1	<i>Iterated Elimination and Common Knowledge of Rationality</i>	63
4.2.2	<i>Example: Cournot Duopoly</i>	65
4.2.3	<i>Evaluating IESDS</i>	67
4.3	Beliefs, Best Response, and Rationalizability	69
4.3.1	<i>The Best Response</i>	69
4.3.2	<i>Beliefs and Best-Response Correspondences</i>	71
4.3.3	<i>Rationalizability</i>	73
4.3.4	<i>The Cournot Duopoly Revisited</i>	73
4.3.5	<i>The "p-Beauty Contest"</i>	74
4.3.6	<i>Evaluating Rationalizability</i>	76
4.4	Summary	76
4.5	Exercises	76
Chapter 5	Pinning Down Beliefs: Nash Equilibrium	79
5.1	Nash Equilibrium in Pure Strategies	80
5.1.1	<i>Pure-Strategy Nash Equilibrium in a Matrix</i>	81
5.1.2	<i>Evaluating the Nash Equilibria Solution</i>	83
5.2	Nash Equilibrium: Some Classic Applications	83
5.2.1	<i>Two Kinds of Societies</i>	83
5.2.2	<i>The Tragedy of the Commons</i>	84
5.2.3	<i>Cournot Duopoly</i>	87
5.2.4	<i>Bertrand Duopoly</i>	88
5.2.5	<i>Political Ideology and Electoral Competition</i>	93
5.3	Summary	95
5.4	Exercises	95

Chapter 6	Mixed Strategies	101
6.1	Strategies, Beliefs, and Expected Payoffs	102
6.1.1	<i>Finite Strategy Sets</i>	102
6.1.2	<i>Continuous Strategy Sets</i>	104
6.1.3	<i>Beliefs and Mixed Strategies</i>	105
6.1.4	<i>Expected Payoffs</i>	105
6.2	Mixed-Strategy Nash Equilibrium	107
6.2.1	<i>Example: Matching Pennies</i>	108
6.2.2	<i>Example: Rock-Paper-Scissors</i>	111
6.2.3	<i>Multiple Equilibria: Pure and Mixed</i>	113
6.3	IESDS and Rationalizability Revisited	114
6.4	Nash's Existence Theorem	117
6.5	Summary	123
6.6	Exercises	123

PART III Dynamic Games of Complete Information

Chapter 7	Preliminaries	129
7.1	The Extensive-Form Game	130
7.1.1	<i>Game Trees</i>	132
7.1.2	<i>Imperfect versus Perfect Information</i>	136
7.2	Strategies and Nash Equilibrium	137
7.2.1	<i>Pure Strategies</i>	137
7.2.2	<i>Mixed versus Behavioral Strategies</i>	139
7.2.3	<i>Normal-Form Representation of Extensive-Form Games</i>	143
7.3	Nash Equilibrium and Paths of Play	145
7.4	Summary	147
7.5	Exercises	147
Chapter 8	Credibility and Sequential Rationality	151
8.1	Sequential Rationality and Backward Induction	152
8.2	Subgame-Perfect Nash Equilibrium: Concept	153
8.3	Subgame-Perfect Nash Equilibrium: Examples	159
8.3.1	<i>The Centipede Game</i>	159
8.3.2	<i>Stackelberg Competition</i>	160
8.3.3	<i>Mutually Assured Destruction</i>	163
8.3.4	<i>Time-Inconsistent Preferences</i>	166
8.4	Summary	169
8.5	Exercises	170
Chapter 9	Multistage Games	175
9.1	Preliminaries	176
9.2	Payoffs	177
9.3	Strategies and Conditional Play	178
9.4	Subgame-Perfect Equilibria	180
9.5	The One-Stage Deviation Principle	184
9.6	Summary	186
9.7	Exercises	186

Chapter 10	Repeated Games	190
10.1	Finitely Repeated Games	190
10.2	Infinitely Repeated Games	192
10.2.1	<i>Payoffs</i>	193
10.2.2	<i>Strategies</i>	195
10.3	Subgame-Perfect Equilibria	196
10.4	Application: Tacit Collusion	201
10.5	Sequential Interaction and Reputation	204
10.5.1	<i>Cooperation as Reputation</i>	204
10.5.2	<i>Third-Party Institutions as Reputation Mechanisms</i>	205
10.5.3	<i>Reputation Transfers without Third Parties</i>	207
10.6	The Folk Theorem: Almost Anything Goes	209
10.7	Summary	214
10.8	Exercises	215
Chapter 11	Strategic Bargaining	220
11.1	One Round of Bargaining: The Ultimatum Game	222
11.2	Finitely Many Rounds of Bargaining	224
11.3	The Infinite-Horizon Game	228
11.4	Application: Legislative Bargaining	229
11.4.1	<i>Closed-Rule Bargaining</i>	230
11.4.2	<i>Open-Rule Bargaining</i>	232
11.5	Summary	235
11.6	Exercises	236
PART IV	Static Games of Incomplete Information	
Chapter 12	Bayesian Games	241
12.1	Strategic Representation of Bayesian Games	246
12.1.1	<i>Players, Actions, Information, and Preferences</i>	246
12.1.2	<i>Deriving Posteriors from a Common Prior: A Player's Beliefs</i>	247
12.1.3	<i>Strategies and Bayesian Nash Equilibrium</i>	249
12.2	Examples	252
12.2.1	<i>Teenagers and the Game of Chicken</i>	252
12.2.2	<i>Study Groups</i>	255
12.3	Inefficient Trade and Adverse Selection	258
12.4	Committee Voting	261
12.5	Mixed Strategies Revisited: Harsanyi's Interpretation	264
12.6	Summary	266
12.7	Exercises	266
Chapter 13	Auctions and Competitive Bidding	270
13.1	Independent Private Values	272
13.1.1	<i>Second-Price Sealed-Bid Auctions</i>	272
13.1.2	<i>English Auctions</i>	275
13.1.3	<i>First-Price Sealed-Bid and Dutch Auctions</i>	276
13.1.4	<i>Revenue Equivalence</i>	279
13.2	Common Values and the Winner's Curse	282

13.3	Summary	285
13.4	Exercises	285

Chapter 14 Mechanism Design 288

14.1	Setup: Mechanisms as Bayesian Games	288
14.1.1	<i>The Players</i>	288
14.1.2	<i>The Mechanism Designer</i>	289
14.1.3	<i>The Mechanism Game</i>	290
14.2	The Revelation Principle	292
14.3	Dominant Strategies and Vickrey-Clarke-Groves Mechanisms	295
14.3.1	<i>Dominant Strategy Implementation</i>	295
14.3.2	<i>Vickrey-Clarke-Groves Mechanisms</i>	295
14.4	Summary	299
14.5	Exercises	299

PART V Dynamic Games of Incomplete Information

Chapter 15 Sequential Rationality with Incomplete Information 303

15.1	The Problem with Subgame Perfection	303
15.2	Perfect Bayesian Equilibrium	307
15.3	Sequential Equilibrium	312
15.4	Summary	314
15.5	Exercises	314

Chapter 16 Signaling Games 318

16.1	Education Signaling: The MBA Game	319
16.2	Limit Pricing and Entry Deterrence	323
16.2.1	<i>Separating Equilibria</i>	324
16.2.2	<i>Pooling Equilibria</i>	330
16.3	Refinements of Perfect Bayesian Equilibrium in Signaling Games	332
16.4	Summary	335
16.5	Exercises	335

Chapter 17 Building a Reputation 339

17.1	Cooperation in a Finitely Repeated Prisoner's Dilemma	339
17.2	Driving a Tough Bargain	342
17.3	A Reputation for Being "Nice"	349
17.4	Summary	354
17.5	Exercises	354

Chapter 18 Information Transmission and Cheap Talk 357

18.1	Information Transmission: A Finite Example	358
18.2	Information Transmission: The Continuous Case	361
18.3	Application: Information and Legislative Organization	365
18.4	Summary	367
18.5	Exercises	367

Chapter 19	Mathematical Appendix	369
19.1	Sets and Sequences	369
19.1.1	<i>Basic Definitions</i>	369
19.1.2	<i>Basic Set Operations</i>	370
19.2	Functions	371
19.2.1	<i>Basic Definitions</i>	371
19.2.2	<i>Continuity</i>	372
19.3	Calculus and Optimization	373
19.3.1	<i>Basic Definitions</i>	373
19.3.2	<i>Differentiation and Optimization</i>	374
19.3.3	<i>Integration</i>	377
19.4	Probability and Random Variables	378
19.4.1	<i>Basic Definitions</i>	378
19.4.2	<i>Cumulative Distribution and Density Functions</i>	379
19.4.3	<i>Independence, Conditional Probability, and Bayes' Rule</i>	380
19.4.4	<i>Expected Values</i>	382
	References	385
	Index	389

Preface

The study of economics, political science, and the social sciences more generally is an attempt to understand the ways in which people behave and make decisions, as individuals and in group settings. The goal is often to apply our understanding to the analysis of questions pertinent to the functioning of societies and their institutions, such as markets, governments, and legal institutions. Social scientists have developed frameworks and rigorous models that abstract from reality with the intent of focusing attention on the crux of the issues at hand, while ignoring details that seem less relevant and more peripheral. We use these models not only to shed light on what we observe but also to help us predict what we cannot yet see. One of the ultimate goals is to prescribe policy recommendations to the private and public sectors, based on the simplistic yet rigorous models that guide our analysis. In this process we must be mindful of the fact that the strength of our conclusions will depend on the validity of our assumptions, in particular those regarding human behavior and the environment in which people act.

Game theory provides a framework based on the construction of rigorous models that describe situations of conflict and cooperation between *rational* decision makers. Following the tradition of mainstream decision theory and economics, rational behavior is defined as choosing actions that maximize one's payoff (or some form of payoff) subject to the constraints that one faces. This is clearly a caricature of reality, but it is a useful benchmark that in many cases works surprisingly well. Game theory has been successfully applied to many relevant situations, such as business competition, the functioning of markets, political campaigning, jury voting, auctions and procurement contracts, and union negotiations, to name just a few. Game theory has also shed light on other disciplines, such as evolutionary biology and psychology.

This book provides an introduction to game theory. It covers the main ideas of the field and shows how they have been applied to many situations drawn mostly from economics and political science. Concepts are first introduced using simple motivating examples, then developed carefully and precisely to present their general formulation, and finally used in a variety of applications.

The Book's Origins and Intended Audience

As with many textbook authors, it was never my intention to write a textbook. This book grew out of lecture notes that I used for an advanced undergraduate game theory course that I taught at Stanford University from 1997 through 2004. Over the years I was convinced by some colleagues and a persistent executive editor to take those notes and expand them.

Given its origins, the book is aimed at more advanced undergraduates in economics. In writing the book I tried hard to be precise—as one ought to be with a more advanced textbook—while at the same time being reader friendly. Relative to the advanced course that I taught, I added many more examples, both easier and harder than the ones I had used in my class.

I was somewhat frustrated by the books that were available at the time I taught: some were too loose and others were too dense. As a consequence the ideas in this textbook are first presented in a way that most students with minimal mathematical training can follow; they are then further developed to meet the needs and address the curiosity of students who are more rigorously trained. Concepts are presented rigorously but illustrated using examples with varying degrees of complexity from easier to harder. I am therefore quite confident that first-year graduate students in economics and political science will find the book useful as well, especially as a backdrop to more demanding graduate-level textbooks.

This text is meant to be self-contained. Many examples should help the reader absorb the concepts of decision making and strategic analysis. Because precise logical reasoning is at the center of game theory, and of this textbook, a degree of mathematical maturity will be useful to comprehend the material fully. That said, it is not assumed that the reader has a strong mathematical background. A mathematical appendix covers most of what would be needed by someone with a relatively rigorous high school training. Calculus and knowledge of some basic probability theory are required to follow all of the material in this book, but even without this knowledge many of the basic examples and constructions can be appreciated.

The Book's Structure and Suggested Use

The book contains five parts:

- *Part I (Chapters 1–2)—Rational Decision Making:* This part presents the basic ideas of the *rational choice paradigm* used in economics and adopted by many other social science disciplines. Students who have had a basic microeconomics course can easily skip this part of the book.
- *Part II (Chapters 3–6)—Static Games of Complete Information:* The most fundamental aspects of game theory and “normal-form” games are developed in this part of the book. It starts with the notions of dominated and dominant strategies, moves on to consider the consequences of assuming rationality and common knowledge of rationality, and ends with the celebrated concept of Nash equilibrium. All the concepts are first introduced using “pure” (non-stochastic) strategies, and the more demanding concept of mixed strategies is introduced toward the end of this part.
- *Part III (Chapters 7–11)—Dynamic Games of Complete Information:* This part extends the static framework developed earlier to be able to deal with games that unfold over time. The “extensive-form” game is introduced, as well as the concepts of sequential rationality and subgame-perfect equilibrium. These concepts are then used to explore multistage and repeated games, as well as bargaining games. Applications include collusion between price-setting firms, the development of institutions that support anonymous trade, and legislative bargaining.

- *Part IV (Chapters 12–14)—Static Games of Incomplete Information:* This is where the analysis becomes more demanding. This part expands the concepts introduced in Part II to be able to tackle situations in which players are not exactly aware of the characteristics of the other players with whom they interact. The concepts of Bayesian games and Bayesian Nash equilibria are carefully developed and applied to such important contexts as adverse selection, jury voting, and auctions. Some of the more advanced treatments of auctions, and the chapter on mechanism design, are intended for graduate students and very rigorously trained undergraduates.
- *Part V (Chapters 15–18)—Dynamic Games of Incomplete Information:* The last part of the book extends Part IV to deal with games in which information unfolds over time. The idea of sequential rationality is extended to Bayesian games, and equilibrium concepts such as perfect Bayesian and sequential equilibrium are defined and illustrated. Applications include signaling games, the development of reputation, and information-transmission games.

As this outline illustrates, the book contains much more material than can be taught in a quarter or semester undergraduate course. Since it is not suited for an easy “math-free” course, I envision it being used in one or more of the following three courses:

- *Intermediate undergraduate game theory:* Such a course, most likely aimed at undergraduates who either have been exposed to intermediate microeconomics or are comfortable with logical analysis, would include Chapters 3–5 and 7–8. Depending on the instructor’s preferences, and the students’ level of preparation, some parts of Chapter 6, as well as Chapters 9–12 could be used, and possibly even parts of Chapters 15–16. Students who have not been exposed to rational choice theory will benefit from covering Chapters 1–2.
- *Advanced undergraduate game theory:* This is the course I taught at Stanford. It was aimed at undergraduates who had had intermediate microeconomics with calculus, and were familiar with probabilities and random variables. It included all of the material described for the intermediate course, as well as Chapter 6, most parts of Chapters 9–12, and parts of Chapters 13 and 15–16.
- *Graduate game theory:* This book would not be suited for an advanced graduate course on game theory, for which there are several excellent texts, such as Myerson (1991), Osborne and Rubinstein (1994), and most notably Fudenberg and Tirole (1991). It would be quite useful, however, for first-year Ph.D. students in economics and political science who are covering a first course on game theory. The topics would cover everything described in the advanced undergraduate course, as well as all of Chapters 12 and 15–16 and parts of Chapters 13–14 and 17–18.

Regardless of the way in which the book is used, all the topics are motivated with simple examples that illustrate the main concepts, many of which are used to slowly and carefully explain the main ideas. The easiest examples are tailored to appeal to students in an intermediate undergraduate class, while some examples will be suited only for the most advanced undergraduates as well as for graduate students. The book contains over 150 exercises, which should be more than enough to drive the various points home. About half the exercises have solutions that are freely accessible online, while solutions to the rest are available only to instructors. For more information, visit this book’s page on the Princeton University Press web site at <http://press.princeton.edu/titles/10001.html>.

The Book's Style

The book is casual yet precise in tone, and sometimes a bit demanding. Where mathematical concepts and notation are introduced, they are explained, and where applicable the reader is referred to the mathematical appendix. Given that I trust instructors to assign the book to its intended audience, I expect most students not to need the appendix. However, because my intention is for a curious reader to be able to use the book without an instructor, the mathematical appendix may come in handy.

Many authors struggle with the fact that the English language does not have a singular pronoun that is sex neutral. I have noticed that in recent years many authors have been careful to use a mix of male and female players, and some have decided to tip the balance toward the female pronoun. Alas, I don't trust myself to be careful enough to catch all the uses of "he" and "she" were I to attempt to use both, and I am sure that I would sometimes make an error. Hence I decided to be old fashioned and use the singular male pronoun. I hope that readers will not find this practice insensitive.

Acknowledgments

One person is responsible for me choosing an academic career: Yossi Greenberg. Yossi introduced me to game theory and planted the seed of academic curiosity deep in the soil of my mind. I continued learning game theory from Binyamin Shitovitz and Dov Monderer as an undergraduate, and in graduate school I had the pleasure and privilege of learning so much more game theory from Eric Maskin and Drew Fudenberg. I owe my passion for the subject to the wonderful professors I have had during the years of my training.

Teaching curious students at Stanford only fueled that passion further. Many wonderful students passed through the game theory course that I taught, some of whom have continued on to their own academic careers and from whom I am now learning new things. One student has had a particularly important impact on this book. In 2003 Wendy Sheu took her excellent written notes from my class, with all my examples, definitions, and tangent comments, and typed them up for me. These notes formed the skeleton of the lecture notes that I used over the next three years, and from those notes came this textbook—after many more hours than I am willing to admit to myself.

While I was developing the notes at Stanford, I had the good fortune of employing two excellent Ph.D. students as my teaching assistants, David Miller and Dan Quint, who have both offered many valuable suggestions that improved the notes. In addition, Victor Bennett, Peter Hammond, Igal Hendel, Matt Jackson, and Steve Matthews have kindly agreed to use parts of this book in its earlier stages and have provided valuable feedback. Several anonymous reviewers selected by Princeton University Press offered excellent direction and comments. The final manuscript of this textbook was read by Orie Shelef, whose careful reading is second to none and whose thoughtful comments proved to be invaluable.

Of course, the editors at Princeton University Press played an important role in making this book a reality. Tim Sullivan's relentless pursuits convinced me to embark on the journey of transforming class notes into a textbook. After Tim left the Press, Seth Ditchik provided the necessary encouragement to make me follow through on my commitment. There were many days, and even more nights, when I regretted that

decision, although I am now pleased with the outcome. I was once told that there are three things one should do in the course of one's life: have a child, plant a tree, and write a book.¹ I must thank Tim and Seth for helping me scratch the third and last item off that list! I am also thankful to Peter Strupp and his team at Princeton Editorial Associates for providing outstanding copyediting and production services that improved the book tremendously.

Last but not least, I am grateful to my wife, Irit, whose help and encouragement touch my work in so many ways. Her sharp mind, exceptional drive, and superb organizational skills are an inspiration. Without her constant support I surely would accomplish only a fraction of what I manage to do. And when I think of game theory, I can't deny that my two sons, Nadav and Noam, constantly teach me that when it comes to strategy, there is still so much more that I have to learn!

1. Some attribute this saying to the Cuban national hero, poet, and writer José Martí.

PART I

**RATIONAL
DECISION
MAKING**

The Single-Person Decision Problem

Imagine yourself in the morning, all dressed up and ready to have breakfast. You might be lucky enough to live in a nice undergraduate dormitory with access to an impressive cafeteria, in which case you have a large variety of foods from which to choose. Or you might be a less-fortunate graduate student, whose studio cupboard offers the dull options of two half-empty cereal boxes. Either way you face the same problem: what should you have for breakfast?

This trivial yet ubiquitous situation is an example of a *decision problem*. Decision problems confront us daily, as individuals and as groups (such as firms and other organizations). Examples include a division manager in a firm choosing whether or not to embark on a new research and development project; a congressional representative deciding whether or not to vote for a bill; an undergraduate student deciding on a major; a baseball pitcher contemplating what kind of pitch to deliver; or a lost group of hikers confused about which direction to take. The list is endless.

Some decision problems are trivial, such as choosing your breakfast. For example, if Apple Jacks and Bran Flakes are the only cereals in your cupboard, and if you hate Bran Flakes (they belong to your roommate), then your decision is obvious: eat the Apple Jacks. In contrast, a manager's choice of whether or not to embark on a risky research and development project or a lawmaker's decision on a bill are more complex decision problems.

This chapter develops a *language* that will be useful in laying out rigorous foundations to support many of the ideas underlying strategic interaction in games. The language will be formal, having the benefit of being able to represent a host of different problems and provide a set of tools that will lend structure to the way in which we think about decision problems. The formalities are a vehicle that will help make ideas precise and clear, yet in no way will they overwhelm our ability and intent to keep the more practical aspect of our problems at the forefront of the analysis.

In developing this formal language, we will be forced to specify a set of assumptions about the behavior of decision makers or players. These assumptions will, at times, seem both acceptable and innocuous. At other times, however, the assumptions will be almost offensive in that they will require a significant leap of faith. Still, as the analysis unfolds, we will see the conclusions that derive from the assumptions

that we make, and we will come to appreciate how sensitive the conclusions are to these assumptions.

As with any theoretical framework, the value of our conclusions will be only as good as the sensibility of our assumptions. There is a famous saying in computer science—“garbage in, garbage out”—meaning that if invalid data are entered into a system, the resulting output will also be invalid. Although originally applied to computer software, this statement holds true more generally, being applicable, for example, to decision-making theories like the one developed herein. Hence we will at times challenge our assumptions with facts and question the validity of our analysis. Nevertheless we will argue in favor of the framework developed here as a useful benchmark.

1.1 Actions, Outcomes, and Preferences

Consider the examples described earlier: choosing a breakfast, deciding about a research project, or voting on a bill. These problems all share a similar structure: an individual, or player, faces a situation in which he has to choose one of several alternatives. Each choice will result in some outcome, and the consequences of that outcome will be borne by the player himself (and sometimes other players too).

For the player to approach this problem in an intelligent way, he must be aware of three fundamental features of the problem: What are his possible choices? What is the result of each of those choices? How will each result affect his well-being? Understanding these three aspects of a problem will help the player choose his best action. This simple observation offers us a first working definition that will apply to *any decision problem*:

The Decision Problem A **decision problem** consists of three features:

1. **Actions** are all the alternatives from which the player can choose.
2. **Outcomes** are the possible consequences that can result from any of the actions.
3. **Preferences** describe how the player ranks the set of possible outcomes, from most desired to least desired. The **preference relation** $\succsim$ describes the player's preferences, and the notation $x \succsim y$ means “ x is at least as good as y .”

To make things simple, let's begin with our rather trivial decision problem of choosing between Apple Jacks and Bran Flakes. We can define the set of actions as $A = \{a, b\}$, where a denotes the choice of Apple Jacks and b denotes the choice of Bran Flakes.¹ In this simple example our actions are practically synonymous with the outcomes, yet to make the distinction clear we will denote the set of outcomes by $X = \{x, y\}$, where x denotes eating Apple Jacks (the consequence of *choosing* Apple Jacks) and y denotes eating Bran Flakes.

1. More on the concept of a set and the appropriate notation can be found in Section 19.1 of the mathematical appendix.

1.1.1 Preference Relations

Turning to the less familiar notion of a **preference relation**, imagine that you prefer eating Apple Jacks to Bran Flakes. Then we will write $x \succeq y$, which should be read as “ x is at least as good as y .” If instead you prefer Bran Flakes, then we will write $y \succeq x$, which should be read as “ y is at least as good as x .” Thus our preference relation is just a shorthand way to express the player’s ranking of the possible outcomes.

We follow the common tradition in economics and decision theory by expressing preferences as a “weak” ranking. That is, the statement “ x is at least as good as y ” is consistent with x being *better* than y or *equally as good as* y . To distinguish between these two scenarios we will use the **strict preference relation**, $x \succ y$, for “ x is strictly better than y ,” and the **indifference relation**, $x \sim y$, for “ x and y are equally good.”

It need not be the case that actions are synonymous with outcome, as in the case of choosing your breakfast cereal. For example, imagine that you are in a bar with a drunken friend. Your actions can be to let him drive home or to order him a cab. The outcome of letting him drive is a certain accident (he’s *really* drunk), and the outcome of ordering him a cab is arriving safely at home. Hence for this decision problem your actions are physically different from the outcomes.

In these examples the action set is *finite*, but in some cases one might have infinitely many actions from which to choose. Furthermore there may be infinitely many outcomes that can result from the actions chosen. A simple example can be illustrated by me offering you a two-gallon bottle of water to quench your thirst. You can choose how much to drink and return the remainder to me. In this case your action set can be described as the interval $A = [0, 2]$: you can choose any action a as long as it belongs to the interval $[0, 2]$, which we can write in two ways: $0 \leq a \leq 2$ or $a \in [0, 2]$.² If we equate outcomes with actions in this example then $X = [0, 2]$ as well. Finally it need not be the case that more is better. If you are thirsty then drinking a pint may be better than drinking nothing. However, drinking a gallon may cause you to have a stomachache, and you may therefore prefer a pint to a gallon.

Before proceeding with a useful way to represent a player’s preferences over various outcomes, it is important to stress that we will make two important assumptions about the player’s ability to think through the decision problem.³ First, we require the player to be able to rank *any two outcomes* from the set of outcomes. To put this more formally:

The Completeness Axiom The preference relation $\succeq$ is **complete**: any two outcomes $x, y \in X$ can be ranked by the preference relation, so that either $x \succeq y$ or $y \succeq x$.

At some level the completeness axiom is quite innocuous. If I show you two foods, you should be able to rank them according to how much you like them (including being indifferent if they are equally tasty and nutritious). If I offer you two cars, you should be able to rank them according to how much you enjoy driving them, their safety

2. The notation symbol $\in$ means “belongs to.” Hence “ $x, y \in X$ ” means “elements x and y belong to the set X .” If you are unfamiliar with sets and these kinds of descriptions please refer to Section 19.1 of the mathematical appendix.

3. These assumptions are referred to as “axioms,” following the language used in the seminal book by von Neumann and Morgenstern (1944) that laid many of the foundations for both decision theory and game theory.

specifications, and so forth. If I offer you two investment portfolios, you should be able to rank them according to the extent to which you are willing to balance risk and return. In other words, the completeness axiom *does not let you be indecisive between any two outcomes*.⁴

The second assumption we make guarantees that a player can rank *all* of the outcomes. To do this we introduce a rather mild consistency condition called *transitivity*:

The Transitivity Axiom The preference relation $\succsim$ is **transitive**: for any three outcomes $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.

Faced with several outcomes, completeness guarantees that any two can be ranked, and transitivity guarantees that there will be no contradictions in the ranking, which could create an indecisive cycle. To observe a violation of the transitivity axiom, consider a player who strictly prefers Apple Jacks to Bran Flakes, $a \succ b$, Bran Flakes to Cheerios, $b \succ c$, and Cheerios to Apple Jacks, $c \succ a$. When faced with any two boxes of cereal, say $A = \{a, b\}$, he has no problem choosing his preferred cereal a . What happens, however, when he is presented with all three alternatives, $A = \{a, b, c\}$? The poor guy will be unable to decide which of the three to choose, because for any given box of cereal, there is another box that he prefers. Therefore, by requiring that the player have complete and transitive preferences, we basically guarantee that among *any set of outcomes*, he will always have at least one *best outcome* that is as good as or better than any other outcome in that set.

To foreshadow what will be our premise for decision making, a preference relation that is complete and transitive is called a **rational preference relation**. We will be concerned only with players who have such rational preferences, for without such preferences we can offer neither predictive nor prescriptive insights.

Remark As noted by the Marquis de Condorcet in 1785, it is possible to have a group of rational individual players who, when put together to make decisions as a group, will become an “irrational” group. For example, imagine three roommates, called players 1, 2, and 3, who have to choose one box of cereal for their apartment kitchen. Player 1’s preferences are given by $a \succ_1 c \succ_1 b$, player 2’s are given by $c \succ_2 b \succ_2 a$, and player 3’s are given by $b \succ_3 a \succ_3 c$. Imagine that our three players make choices in a democratic way and use majority voting to reach a decision. What will be the resulting preferences of the group, $\succ_G$? When faced with the pair a and c , players 1 and 3 will vote for Apple Jacks, hence $a \succ_G c$. When faced with the pair c and b , players 1 and 2 will vote for Cheerios, hence $c \succ_G b$. When faced with the pair a and b , players 2 and 3 will vote for Bran Flakes, hence $b \succ_G a$. As a result, our three rational players will not be able to reach a conclusive decision using the group preferences that result from majority voting! This type of group indecisiveness resulting from majority voting is often referred to as the *Condorcet Paradox*. Because we will not be analyzing group decisions, it is not something we will confront, but it is useful to be mindful of such phenomena, in which imposing individual rationality does not imply “group rationality.”

4. In other words, this axiom prohibits the kind of problem referred to as “Buridan’s ass.” One version describes a situation in which an ass is placed between two identical stacks of hay, assuming that the ass will always go to whichever stack is closer. However, since the stacks are both the same distance from the ass, it will not be able to choose between them and will die of hunger.

1.1.2 Payoff Functions

When we restrict attention to players with rational preferences, not only do we get players who behave in a consistent and appealing way, but as an added bonus we can replace the preference relation with a much friendlier, and more operational, apparatus. Consider the following simple example. Imagine that you open a lemonade stand on your neighborhood corner. You have three possible actions: choose low-quality lemons (l), which imply a cost of \$10 and a revenue from sales of \$15; choose medium-quality lemons (m), which imply a cost of \$15 and a revenue from sales of \$25; or choose high-quality lemons (h), which imply a cost of \$28 and a revenue from sales of \$35. Thus the action set is $A = \{l, m, h\}$, and the outcome set is given by net profits and is $X = \{5, 10, 7\}$, where the action l yields a profit of \$5, the action m yields a profit of \$10, and the action h yields a profit of \$7. Assuming that obtaining higher profits is strictly better, we have $10 \succ 7 \succ 5$. Hence you should choose alternative m and make a profit of \$10.

Notice that we took a rather obvious profit-maximizing problem and fit it into our framework for a decision problem. We derived the preference relation that is consistent with maximizing profit, the objective of any for-profit business. Arguably it would be more natural and probably easier to comprehend the problem if we looked at the actions and their associated profits. In particular we can define the **profit function** in the obvious way: every action $a \in A$ yields a profit $\pi(a)$. Then, instead of considering a preference relation over profit outcomes, we can just look at the profit from each action directly and choose an action that maximizes profits. In other words, we can *use the profit function to evaluate actions and outcomes*.

As this simple example demonstrates, a profit function is a more direct way for a player to rank his actions. The question then is, can we find similar ways to approach decision problems that are not about profits? It turns out that we can do exactly that if we have players with rational preferences, and to do that we define a payoff function.⁵

Definition 1.1 A **payoff function** $u : X \rightarrow \mathbb{R}$ represents the preference relation $\succeq$ if for any pair $x, y \in X$, $u(x) \geq u(y)$ if and only if $x \succeq y$.

To put the definition into words, we say that the preference relation $\succeq$ is represented by the payoff function $u : X \rightarrow \mathbb{R}$ that assigns to each outcome in X a real number, if and only if the function assigns a higher value to higher-ranked outcomes.

It is important to notice that representing preferences with payoff functions is convenient, but that payoff values by themselves have no meaning whatsoever. Payoff is an *ordinal* construct: it is used to order the alternatives from most to least desirable. For example, if I like Apple Jacks more than Bran Flakes, then I can construct the payoff function $u(\cdot)$ so that $u(a) = 5$ and $u(b) = 3$. I can also use a different payoff function $\tilde{u}(\cdot)$ that represents the same preferences as follows: $\tilde{u}(a) = 100$ and $\tilde{u}(b) = -237$. Just as Fahrenheit and Celsius are two different ways to describe hotter and colder temperatures, there are many ways to represent preferences with payoff functions.

Using payoff functions instead of preferences will allow us to operationalize a theory of how decision makers with rational preferences ought to behave, and how they often will behave. They will choose actions that maximize a payoff function that

5. Recall that a function relates each of its inputs to exactly one output. For more on this see Section 19.2 of the mathematical appendix.

represents their preferences. One last question we need to ask is whether we know for sure that this method will work: is it true that players will surely have a payoff function representing their preferences? One case is easy and worth going through briefly. In what follows, we provide a formal proposition and a formal, yet fairly easy to follow, proof.

Proposition 1.1 *If the set of outcomes X is finite then any rational preference relation over X can be represented by a payoff function.*

Proof The proof is by construction. Because the preference relation is complete and transitive, we can find a least-preferred outcome $\underline{x} \in X$ such that all other outcomes $y \in X$ are at least as good as $\underline{x}$, that is, $y \succeq \underline{x}$ for all other $y \in X$. Now define the “worst outcome equivalence set,” denoted X_1 , to include $\underline{x}$ and any other outcome for which the player is indifferent between it and $\underline{x}$. Then, from the remaining elements of $X \setminus X_1$,⁶ define the “second worst outcome equivalence set,” X_2 , and continue in this fashion until the “best outcome equivalence set,” X_n , is created. Because X is finite and $\succeq$ is rational, such a finite collection of n equivalence sets exists. Now consider n arbitrary values $u_n > u_{n-1} > \dots > u_2 > u_1$, and assign payoffs according to the function defined by: for any $x \in X_k$, $u(x) = u_k$. This payoff function represents $\succeq$. Hence we have proved that such a function exists. ■

This proposition is useful: for many realistic situations, we can create payoff functions that work in a similar way as profit functions, giving the player a useful tool to see which actions are best and which ought to be avoided. We will not explore this issue further, but payoff representations exist in many other cases that include infinitely many outcomes. The treatment of such cases is beyond the scope of this textbook, but you are welcome to explore one of the many texts that offer a more complete treatment of the topic, which is referred to under the title “representation theorems.” (See, e.g., Kreps [1990a, pp. 18–37, and 1988] for an in-depth treatment of this topic.)

As we have seen so far, the formal structure of a decision problem offers a coherent framework for analysis. For decades, however, teachers, students, and practitioners have instead used the intuitive and graphically simple tool of *decision trees*.

Imagine that, in addition to Apple Jacks (a) and Bran Flakes (b), your breakfast options include a muffin (m) and a scone (s). Your preferences are given as $s \succ a \succ m \succ b$. (Recall that we now consider preferences over outcomes as directly over actions.) Consider the following payoff representation: $v(s) = 4$, $v(a) = 3$, $v(m) = 2$, and $v(b) = 1$. We can write down the corresponding decision tree, which is depicted in Figure 1.1.

To read this simple decision tree, notice that the player resides at the “root” of the tree on the left, and that the tree then branches off, each branch representing a possible action. In the example of choosing breakfast, each action results in a final payoff, and these payoffs are written to correspond to each of the action branches. Our rational decision maker will look down the tree, consider the payoff from each branch, and choose the branch with the highest payoff.

The node at which the player has to make a choice is called a **decision node**. The nodes at the end of the tree where payoffs are attached are called **terminal nodes**. As

6. The notation $A \setminus B$ means “the elements that are in A but are not in B ,” or sometimes “the set A less the set B .”




FIGURE 1.1 A simple breakfast decision tree.

the next chapter demonstrates, the structure of a decision tree will become slightly more involved and useful to capture more complex decision problems. We will return to similar trees in Chapter 7, where we consider the strategic interaction between many possible players, which is the main focus of this book.

1.2 The Rational Choice Paradigm

We now introduce *Homo economicus* or “economic man.” *Homo economicus* is “rational” in that he chooses actions that maximize his well-being as defined by his payoff function over the resulting outcomes.⁷ The assumption that the player is rational lies at the foundation of what is known as the **rational choice paradigm**. Rational choice theory asserts that when a decision maker is choosing between potential actions he will be guided by rationality to choose his best action. This can be assumed to be true for individual human behavior, as well as for the behavior of other entities, such as corporations, committees, or nation-states.

It is important to note, however, that by adopting the paradigm of rational choice theory we are imposing some implicit assumptions, which we now make explicit.

Rational Choice Assumptions The player fully understands the decision problem by knowing:

1. all possible actions, A ;
2. all possible outcomes, X ;
3. exactly how each action affects which outcome will materialize; and
4. his rational preferences (payoffs) over outcomes.

Perhaps at a first glance this set of assumptions may seem a bit demanding, and further contemplation may make you feel that it is impossible to satisfy for most decision problems. Still, it is a benchmark for a world in which decision problems are completely understood by the player, in which case he can approach the problems in a systematic and structured way. If we let go of any of these four knowledge

7. A naive application of the *Homo economicus* model assumes that our player knows what is best for his long-term well-being and can be relied upon to always make the right decision for himself. We take this naive approach throughout the book, though we will sometimes question how appropriate this approach is.

requirements then we cannot impose the notion of rational choice. If (1) is unknown then the player may be unaware of his best course of action. If (2) or (3) are unknown then he may not correctly foresee the actual consequences of his actions. Finally if (4) is unknown then he may incorrectly perceive the effect of his choice's consequence on his well-being.

To operationalize this paradigm of rationality we must choose among *actions*, yet we have defined preferences—and payoffs—over *outcomes* and not actions. It would be useful, therefore, if we could define preferences—and payoffs—over actions instead of outcomes. In the simple examples of choosing a cereal or how much water to drink, actions and outcomes were synonymous, yet this need not always be the case. Consider the situation of letting your friend drive drunk, in which the actions and outcomes are not the same. Still each action led to one and only one outcome: letting him drive leads to an accident, and getting him a cab leads to safe arrival. Hence, even though preferences and payoff were defined over outcomes, this *one-to-one correspondence*, or function, between actions and outcomes means that we can consider the preferences and payoffs to be over actions, and we can use this correspondence between actions and outcomes to define the payoff over actions as follows: if $x(a)$ is the outcome resulting from action a , then the payoff from action a is given by $v(a) = u(x(a))$, the payoff from $x(a)$. We will therefore use the notation $v(a)$ to represent the payoff from action a .⁸ Now we can precisely define a rational player as follows:

Definition 1.2 A player facing a decision problem with a payoff function $v(\cdot)$ over actions is rational if he chooses an action $a \in A$ that maximizes his payoff. That is, $a^* \in A$ is chosen if and only if $v(a^*) \geq v(a)$ for all $a \in A$.

We now have a formal definition of *Homo economicus*: a player who has rational preferences and is rational in that he understands all the aspects of his decision problem and always chooses an option that yields him the highest payoff from the set of possible actions.

So far we have seen some simple examples with finite action sets. Consider instead an example with a continuous action space, which requires some calculus. Imagine that you're at a party and are considering engaging in social drinking. Given your physique, you'd prefer some wine, both for taste and for the relaxed feeling it gives you, but too much will make you sick. There is a one-liter bottle of wine, so your action set is $A = [0, 1]$, where $a \in A$ is how much you choose to drink. Your preferences are represented by the following payoff function over actions: $v(a) = 2a - 4a^2$, which is depicted in Figure 1.2. As you can see, some wine is better than no wine (0.1 liter gives you some positive payoff, while drinking nothing gives you zero), but drinking a whole bottle will be worse than not drinking at all ($v(1) = -2$). How much should you drink? Your maximization problem is

$$\max_{a \in [0, 1]} 2a - 4a^2.$$

Taking the derivative of this function and equating it to zero to find the solution, we obtain that $2 - 8a = 0$, or $a = 0.25$, which is a bit more than two normal glasses of

8. To be precise, let $x : A \rightarrow X$ be the function that maps actions into outcomes, and let the payoff function over outcomes be $u : X \rightarrow \mathbb{R}$. Define the payoff over actions as the composite function $v = u \circ x : A \rightarrow \mathbb{R}$, where $v(a) = u(x(a))$.




FIGURE 1.2 The payoff from drinking wine.

wine.⁹ Thus, by considering how much wine to drink as a decision problem, you were able to find your optimal action.

1.3 Summary

- A simple decision problem has three components: actions, outcomes, and preferences over outcomes.
- A rational player has complete and transitive preferences over outcomes and hence can always identify a best alternative from among his possible actions. These preferences can be represented by a payoff (or profit) function over outcomes and the corresponding payoffs over actions.
- A rational player chooses the action that gives him the highest possible payoff from the possible set of actions at his disposal. Hence by maximizing his payoff function over his set of alternative actions, a rational player will choose his optimal decision.
- A decision tree is a simple graphic representation for decision problems.

1.4 Exercises

- 1.1 Your Decision:** Think of a simple decision you face regularly and formalize it as a decision problem, carefully listing the actions and outcomes without the preference relation. Then assign payoffs to the outcomes and draw the decision tree.
- 1.2 Going to the Movies:** There are two movie theaters in your neighborhood: Cineclass, which is located one mile from your home, and Cineblast, located three miles from your home. Each is showing three films. Cineclass is showing *Casablanca*, *Gone with the Wind*, and *Dr. Strangelove*, while Cineblast is showing *The Matrix*, *Blade Runner*, and *Aliens*. Your problem is to decide which movie to go to.

9. To be precise, we must also make sure that first, the second derivative is negative for the solution $a = 0.25$ to be a local maximum, and second, the value of $v(a)$ is not greater at the two boundaries $a = 0$ and $a = 1$. For more on maximizing the value of a function, see Section 19.3 of the mathematical appendix.

- a. Draw a decision tree that represents this problem without assigning payoff values.
 - b. Imagine that you don't care about distance and that your preferences for movies are alphabetic (i.e., you like *Aliens* the most and *The Matrix* the least). Using payoff values 1 through 6 complete the decision tree you drew in part (1). Which option would you choose?
 - c. Now imagine that your car is in the shop and that the cost of walking each mile is equal to one unit of payoff. Update the payoffs in the decision tree. Would your choice change?
- 1.3 **Fruit or Candy:** A banana costs \$0.50 and a piece of candy costs \$0.25 at the local cafeteria. You have \$1.25 in your pocket and you value money. The money-equivalent value (payoff) you get from eating your first banana is \$1.20, and that of each additional banana is half the previous one (the second banana gives you a value of \$0.60, the third \$0.30, and so on). Similarly the payoff you get from eating your first piece of candy is \$0.40, and that of each additional piece is half the previous one (\$0.20, \$0.10, and so on). Your value from eating bananas is not affected by how many pieces of candy you eat and vice versa.
 - a. What is the set of possible actions you can take given your budget of \$1.25?
 - b. Draw the decision tree that is associated with this decision problem.
 - c. Should you spend all your money at the cafeteria? Justify your answer with a rational choice argument.
 - d. Now imagine that the price of a piece of candy increases to \$0.30. How many possible actions do you have? Does your answer to (c) change?
- 1.4 **Alcohol Consumption:** Recall the example in which you needed to choose how much to drink. Imagine that your payoff function is given by $\theta a - 4a^2$, where θ is a parameter that depends on your physique. Every person may have a different value of θ , and it is known that in the population (1) the smallest θ is 0.2; (2) the largest θ is 6; and (3) larger people have higher θ s than smaller people.
 - a. Can you find an amount that no person should drink?
 - b. How much should you drink if your $\theta = 1$? If $\theta = 4$?
 - c. Show that in general smaller people should drink less than larger people.
 - d. Should any person drink more than one 1-liter bottle of wine?
- 1.5 **Buying a Car:** You plan on buying a used car. You have \$12,000, and you are not eligible for any loans. The prices of available cars on the lot are given as follows:

Make, model, and year	Price
Toyota Corolla 2002	\$9,350
Toyota Camry 2001	10,500
Buick LeSabre 2001	8,825
Honda Civic 2000	9,215
Subaru Impreza 2000	9,690

For *any given year*, you prefer a Camry to an Impreza, an Impreza to a Corolla, a Corolla to a Civic, and a Civic to a LeSabre. For *any given year*, you are willing to pay up to \$999 to move from any given car to the next preferred one. For example, if the price of a Corolla is \$ z , then you are willing to buy it rather than a Civic if the Civic costs more than $\$(z - 999)$, but you would prefer to buy the Civic if it costs less than this amount. Similarly you prefer the Civic at \$ z to a Corolla that costs more than $\$(z + 1000)$, but you prefer the Corolla if it costs less. For *any given car*, you are willing to move to a model a year older if it is cheaper by at least \$500. For example, if the price of a 2003 Civic is \$ z , then you are willing to buy it rather than a 2002 Civic if the 2002 Civic costs more than $\$(z - 500)$, but you would prefer to buy the 2002 Civic if it costs less than this amount.

- a. What is your set of possible alternatives?
- b. What is your preference relation between the alternatives in (a) above?
- c. Draw a decision tree and assign payoffs to the terminal nodes associated with the possible alternatives. What would you choose?
- d. Can you draw a decision tree with different payoffs that represents the same problem?

- 1.6 **Fruit Trees:** You have room for up to two fruit-bearing trees in your garden. The fruit trees that can grow in your garden are either apple, orange, or pear. The cost of maintenance is \$100 for an apple tree, \$70 for an orange tree, and \$120 for a pear tree. Your food bill will be reduced by \$130 for each apple tree you plant, by \$145 for each pear tree you plant, and by \$90 for each orange tree you plant. You care only about your total expenditure in making any planting decisions.

- a. What is the set of possible actions and related outcomes?
- b. What is the payoff of each action/outcome?
- c. Draw the associated decision tree. What will a rational player choose?
- d. Now imagine that the reduction in your food bill is half for the second tree of the same kind. (You like variety.) That is, the first apple tree still reduces your food bill by \$130, but if you plant two apple trees your food bill will be reduced by $\$130 + \$65 = \$195$, and similarly for pear and orange trees. What will a rational player choose now?

- 1.7 **City Parks:** A city's mayor has to decide how much money to spend on parks and recreation. City codes restrict this spending to no more than 5% of the budget, and the yearly budget of the city is \$20,000,000. The mayor wants to please his constituents, who have diminishing returns from parks. The money-equivalent benefit from spending \$ c on parks is $v(c) = \sqrt{400c} - \frac{1}{80}c$.

- a. What is the action set for the city's mayor?
- b. How much should the mayor spend?
- c. The movie *An Inconvenient Truth* has shifted public opinion, and now people are more willing to pay for parks. The new preferences of the people are given by $v(c) = \sqrt{1600c} - \frac{1}{80}c$. What now is the action set for the mayor, and how much spending should he choose to cater to his constituents?

Introducing Uncertainty and Time

Now that we have a coherent and precise language to describe decision problems, we move on to be more realistic about the complexity of many such problems. The cereal example was fine to illustrate a simple decision problem and to get used to our formal language, but it is certainly not very interesting.

Consider a division manager who has to decide on whether a research and development (R&D) project is worthwhile. What will happen if he does not go ahead with it? Maybe over time his main product will become obsolete and outdated, and the profitability of his division will no longer be sustainable. Then again, maybe profits will still continue to flow in. What happens if he does go ahead with the project? It may lead to vast improvements in the product line and offer the prospect of sustained growth. Or perhaps the research will fail and no new products will emerge, leaving behind only a hefty bill for the expensive R&D endeavor. In other words, both actions have uncertainty over what outcomes will materialize, implying that the choice of a best action is not as obvious as in the cereal example.

How should the player approach this more complex problem? As you can imagine, using language like “maybe this will happen, or maybe that will happen” is not very useful for a rational player who is trying to put some structure on his decision problem. We must introduce a method through which the player can compare uncertain consequences in a meaningful way. For this approach, we will use the concept of stochastic (random) outcomes and probabilities, and we will describe a framework within which payoffs are defined over random outcomes.

2.1 Risk, Nature, and Random Outcomes

Put yourself in the shoes of our division manager who is deciding whether or not to embark on the R&D project. Denote his actions as g for going ahead or s for keeping the status quo, so that $A = \{g, s\}$. To make the problem as simple as possible, imagine that there are only two final outcomes: his product line is successful, which is equivalent to a profit of 10 (choose your denomination), or his product line is obsolete, which is equivalent to a profit of 0, so that $X = \{0, 10\}$. However, as already explained, there is no one-to-one correspondence here between actions and outcomes. Instead

there is uncertainty about which outcome will prevail, and the uncertainty is tied to the choice made by the player, the division manager.

In order to capture this uncertainty in a precise way, we will use the well-understood notion of randomness, or risk, as described by a random variable. Use of random variables is the common way to precisely and consistently describe random prospects in mathematics and statistics. We will not use the most formal mathematical representation of a random variable but instead present it in its most useful depiction for the problems we will address. Section 19.4 of the mathematical appendix has a short introduction to random variables that you can refer to if this notion is completely new to you. Be sure to make yourself familiar with the concept: it will accompany us closely throughout this book.

2.1.1 Finite Outcomes and Simple Lotteries

Continuing with the R&D example, imagine that a successful product line is more likely to be created if the player chooses to go ahead with the R&D project, while it is less likely to be created if he does not. More precisely, the odds are 3 to 1 that success happens if g is chosen, while the odds are only 50-50 if s is chosen. Using the language of probabilities, we have the following description of outcomes following actions: If the player chooses g then the probability of a payoff of 10 is 0.75 and the probability of a payoff of 0 is 0.25. If, however, the player chooses s then the probability of a payoff of 10 is 0.5, as is the probability of a payoff of 0.

We can therefore think of the player as if he is choosing between two **lotteries**. A lottery is exactly described by a random payoff. For example, the state lottery offers each player either several million dollars or zero, and the likelihood of getting zero is extremely high. In our example, the choice of g is like choosing a lottery that pays zero with probability 0.25 and pays 10 with probability 0.75. The choice of s is like choosing a lottery that pays either zero or 10, each with an equal probability of 0.5.

It is useful to think of these lotteries as choices of another player that we will call “Nature.” The probabilities of outcomes that Nature chooses depend on the actions chosen by our decision-making player. In other words, Nature chooses a probability distribution over the outcomes, and the probability distribution is conditional on the action chosen by our decision-making player.

We can utilize a decision tree to describe the player’s decision problem that includes uncertainty. The R&D example is described in Figure 2.1. First the player takes an action, either g or s . Then, conditional on the action chosen by the player, Nature (denoted by N) will choose a probability distribution over the outcomes 10 and 0. The branches of the player are denoted by his actions, and the branches of Nature’s




FIGURE 2.1 The R&D decision problem.

choices are denoted by their corresponding probabilities, which are conditional on the choice made by the player.

We now introduce a definition that generalizes the kind of randomness that was demonstrated by the R&D example. Consider a decision problem with n possible outcomes, $X = \{x_1, x_2, \dots, x_n\}$.

Definition 2.1 A simple lottery over outcomes $X = \{x_1, x_2, \dots, x_n\}$ is defined as a probability distribution $p = (p(x_1), p(x_2), \dots, p(x_n))$, where $p(x_k) \geq 0$ is the probability that x_k occurs and $\sum_{k=1}^n p(x_k) = 1$.

By the definition of a probability distribution over elements in X , the probability of each outcome cannot be a negative number, and the sum of all probabilities over all outcomes must add up to 1. In our R&D example, following a choice of g , the lottery that Nature chooses is $p(10) = 0.75$ and $p(0) = 0.25$. Similarly, following a choice of s , the lottery that Nature chooses is $p(10) = p(0) = 0.5$.

Remark To be precise, the lottery that Nature chooses is conditional on the action taken by the player. Hence, given an action $a \in A$, the conditional probability that $x_k \in X$ occurs is given by $p(x_k|a)$, where $p(x_k|a) \geq 0$, and $\sum_{k=1}^n p(x_k|a) = 1$ for all $a \in A$.

Note that our trivial decision problem of choosing a cereal can be considered as a decision problem in which the probability over outcomes after any choice is equal to 1 for some outcome and 0 for all other outcomes. We call such a lottery a **degenerate lottery**. You can now see that decision problems with no randomness are just a very special case of those with randomness. Thus we have enriched our language to include more complex decision problems while encompassing everything we have developed earlier.

2.1.2 Simple versus Compound Lotteries

Arguably a player should care only about the probabilities of the various final outcomes that are a consequence of his actions. It seems that the exact way in which randomness unfolds over time should not be consequential to a player's well-being, but that only distributions over final outcomes should matter.

To understand this concept better, imagine that we make the R&D decision problem a bit more complicated. As before, if the player chooses not to embark on the R&D project (s) then the product line is successful with probability 0.5. If he chooses to go ahead with R&D (g) then two further stages will unfold. First, it will be determined whether the R&D effort was successful or not. Second, the outcome of the R&D phase will determine the likelihood of the product line's success. If the R&D effort is a failure then the success of the product is as likely as if no R&D had been performed; that is, the product line succeeds with probability 0.5. If the R&D effort is a success, however, then the probability of a successful product line jumps to 0.9. To complete the data for this example, we assume that R&D succeeds with probability 0.625 and fails with probability 0.375.

In this modified version of our R&D problem we have Nature moving once after the choice s and twice in a row after the choice g : once through the outcome of the R&D phase and then through the determination of the product line's success. This new decision problem is depicted in Figure 2.2.




FIGURE 2.2 The modified R&D decision problem.

It seems like the two decision problems in Figures 2.1 and 2.2 are of different natures (no pun intended). Then again, let's consider what a decision problem ought to be about: actions, distributions over outcomes, and preferences. It is apparent that the player's choice of s in both Figure 2.1 and Figure 2.2 leads to the *same distribution* over outcomes. What about the choice of g ? In Figure 2.2 this is followed by two random stages. However, the outcomes are still either 10 or 0. What are the probabilities of each outcome?

There are two ways that 10 can be obtained after the choice of g : First, with probability 0.625 the R&D project succeeds, and then with probability 0.9 the payoff 10 will be obtained. Hence the probability of "R&D success followed by 10" is equal to $0.625 \times 0.9 = 0.5625$. Second, with probability 0.375 the R&D project fails, and then with probability 0.5 the payoff 10 will be obtained. Hence the probability of "R&D failure followed by 10" is equal to $0.375 \times 0.5 = 0.1875$. Thus if the player chooses g then the probability of obtaining 10 is just the sum of the probabilities of these two *exclusive events*, which equals $0.5625 + 0.1875 = 0.75$. It follows that if the player chooses g then the probability of obtaining a payoff of 0 is 0.25, the complement of the probability of obtaining 10 (you should check this).

What then is the difference between the two decision problems? The first, simpler, R&D problem has a simple lottery following the choice of g . The second, more complex, problem has a *simple lottery over simple lotteries* following the choice of g . We call such lotteries over lotteries **compound lotteries**. Despite this difference, we impose on the player a rather natural sense of rationality. In his eyes the two decision problems are the same: he has the same set of actions, each one resulting in the same probability distributions over final outcomes. This innocuous assumption will make it easier for the player to evaluate and compare the benefits from different lotteries over outcomes.

2.1.3 Lotteries over Continuous Outcomes

Before moving on to describe how the player will evaluate lotteries over outcomes, we will go a step further to describe random variables, or lotteries, over continuous-outcome sets. To start, consider the following example. You are growing 10 tomato vines in your backyard, and your crop, measured in pounds, will depend on two inputs. The first is how much you water your garden per day and the second is the weather. Your action set can be any amount of water up to 50 gallons (50 gallons will completely

flood your backyard), so that $A \in [0, 50]$, and your outcome set can be any amount of crop that 10 vines can yield, which is surely no more than 100 pounds, hence $X = [0, 100]$. Temperatures vary daily, and they vary continuously. This implies that your final yield, given any amount of water, will also vary continuously.

In this case we will describe the uncertainty not with a discrete probability, as we did for the R&D case, but instead with a **cumulative distribution function** (CDF) defined as follows:¹

Definition 2.2 A simple lottery over an interval $X = [\underline{x}, \bar{x}]$ is given by a cumulative distribution function $F : X \rightarrow [0, 1]$, where $F(\hat{x}) = \Pr\{x \leq \hat{x}\}$ is the probability that the outcome is less than or equal to $\hat{x}$.

For those of you who have seen continuous random variables, this is not new. If you have not, Section 19.4 of the mathematical appendix may fill in some of the gaps.² The basic idea is simple. Because we have infinitely many possible outcomes, it is somewhat meaningless to talk about the probability of growing a certain exact weight of tomatoes. In fact it is correct to say that the probability of producing any particular predefined weight is zero. However, it is meaningful to talk about the probability of being below a certain weight x , which is given by the CDF $F(x)$, or similarly the probability of being above a certain weight x , which is given by the complement $1 - F(x)$.

Remark Just as in the case of finite outcomes, we wish to consider the case in which the distribution over outcomes is conditional on the action taken. Hence, to be precise, we need to use the notation $F(x|a)$.

Now that we have concluded with a description of what randomness is, we can move along to see how our decision-making player evaluates random outcomes.

2.2 Evaluating Random Outcomes

From now on we will consider the choice of an action $a \in A$ as the *choice of a lottery* over the outcomes in X . If the decision problem does not involve any randomness, then these lotteries are degenerate. This implies that we can stick to our notation of defining a decision problem by the three components of actions, outcomes, and preferences. The novelty is that each action is a lottery over outcomes.

The next natural question is: how will a player faced with the R&D problem in Figure 2.1 choose between his options of going forward or staying the course? Upon reflection, you may have already reached a conclusion. Despite the fact that his different choices lead to different lotteries, it seems that the two lotteries that follow g and s are easy to compare. Both have the same set of outcomes, a profit of 10 or a profit of 0. The choice g has a higher chance at getting the profit of 10, and hence we would expect anyone in their right mind to choose g . This implicitly assumes, however, that there are no costs to launching the R&D project.

1. The definition considers the outcome set to be a finite interval $X = [\underline{x}, \bar{x}]$. We can use the same definition for any subset of the real numbers, including the real line $(-\infty, \infty)$. An example of a lottery over the real line is the normal “bell-shape” distribution.

2. You are encouraged to learn this material since it will be useful, but one can continue through most of Parts I–III of this book without this knowledge.




FIGURE 2.3 The R&D problem with costs.

Let's consider a less obvious revision of the R&D problem, and imagine that there is a real cost of pursuing the R&D project equivalent to 1. Hence the outcome of success yields a profit of 9 instead of 10, and the outcome of failure yields a profit of -1 instead of 0. This new problem is depicted in Figure 2.3. Now the comparison is not as obvious: is it better to have a coin toss between 10 and 0, or to have a good shot at 9, with some risk of losing 1?

2.2.1 Expected Payoff: The Finite Case

To our advantage, there is a well-developed methodology for evaluating how much a lottery is worth for a player, how different lotteries compare to each other, and how lotteries compare to “sure” payoffs (degenerate lotteries). This methodology, called “expected utility theory,” was first developed by John von Neumann and Oskar Morgenstern (1944), two of the founding fathers of game theory, and explored further by Leonard Savage (1951). It turns out that there are some important assumptions that make this method of evaluation valid. (The foundations that validate expected payoff theory are beyond the scope of this text, and are rather technical in nature.)³

The intuitive idea is about averages. It is common for us to think of our actions as sometimes putting us ahead and sometimes dealing us a blow. But if *on average* things turn out on the positive side, then we view our actions as pretty good because the gains will more than make up for the losses. We want to take this idea, with its intuitive appeal, and use it in a precise way to tackle a single decision problem. To do this we introduce the following definition:

Definition 2.3 Let $u(x)$ be the player's payoff function over outcomes in $X = \{x_1, x_2, \dots, x_n\}$, and let $p = (p_1, p_2, \dots, p_n)$ be a lottery over X such that $p_k = \Pr\{x = x_k\}$. Then we define the player's **expected payoff from the lottery** p as

$$E[u(x)|p] = \sum_{k=1}^n p_k u(x_k) = p_1 u(x_1) + p_2 u(x_2) + \dots + p_n u(x_n).$$

The idea of an expected payoff is naturally related to the intuitive idea of averages: if we interpret a lottery as a list of “weights” on payoff values, so that numbers that appear with higher probability have more weight, then the expected payoff of a lottery is nothing other than the weighted average of payoffs for each realization of the lottery.

3. The key idea was introduced by von Neumann and Morgenstern (1944) and is based on the “Independence Axiom.” A nice treatment of the subject appears in Kreps (1990a, Chapter 3).

That is, payoffs that are more likely to occur receive higher weight while payoffs that are less likely to occur receive lower weight.

Using the definition of expected payoff we can revisit the R&D problem in Figure 2.3. First assume that the payoff to the player is equal to his profit, so that $u(x) = x$. By choosing g , the expected payoff to the player is

$$v(g) = E[u(x)|g] = 0.75 \times 9 + 0.25 \times (-1) = 6.5.$$

In contrast, by choosing s his expected payoff is

$$v(s) = E[u(x)|s] = 0.5 \times 10 + 0.5 \times 0 = 5.$$

Hence his best action using expected profits as a measure of preferences over actions is to choose g . You should be able to see easily that in the original R&D game in Figure 2.1 the expected payoff from s is still 5, while the expected payoff from g is 7.5, so that g was also his best choice, as we intuitively argued earlier.

Notice that we continue to use our notation $v(a)$ to define the expected payoff of an action given the distribution over outcomes that the action causes. This is a convention that we will use throughout this book, because the object of our analysis is what a player should do, and this notation implies that his ranking should be over his actions.

2.2.2 Expected Payoff: The Continuous Case

Consider the case in which the outcomes can be any one of a continuum of values distributed on some interval X . The definition of expected utility will be analogous, as follows:

Definition 2.4 Let $u(x)$ be the player's payoff function over outcomes in the interval $X = [\underline{x}, \bar{x}]$ with a lottery given by the cumulative distribution $F(x)$, with density $f(x)$. Then we define the player's expected payoff as⁴

$$E[u(x)] = \int_{\underline{x}}^{\bar{x}} u(x) f(x) dx.$$

To see an example with continuous actions and outcomes, recall the tomato growing problem in Section 2.1.3, in which your choice is how much water to use in the set $A = [0, 50]$ and the outcome is the weight of your crop that will result in the set $X = [0, 100]$. Imagine that given a choice of water $a \in A$, the distribution over outcomes is uniform over the quantity support $[0, 2a]$. (Alternatively the distribution of x conditional on a is given by $x|a \sim U[0, 2a]$.) For example, if you use 10 gallons of water, the output will be uniformly distributed over the weight interval $[0, 20]$, with the cumulative distribution function given by $F(x|a = 10) = \frac{x}{20}$ for $0 \leq x \leq 20$,

4. More generally, if there are continuous distributions that do not have a density because $F(\cdot)$ is not differentiable, then the expected utility is given by

$$E[u(x)] = \int_{x \in X} u(x) dF(x).$$

This topic is covered further in Section 19.4 of the mathematical appendix.

and $F(x|a = 10) = 1$ for all $x > 20$. More generally the cumulative distribution function is given by $F(x|a) = \frac{x}{2a}$ for $0 \leq x \leq 2a$, and $F(x|a) = 1$ for all $x > 2a$. The density is given by $f(x|a) = \frac{1}{2a}$ for $0 \leq x \leq 2a$, and $f(x|a) = 0$ for all $x > 2a$. Thus if your payoff from quantity x is given by $u(x)$ then your expected payoff from any choice $a \in A$ is given by

$$v(a) = E[u(x)|a] = \int_0^{2a} u(x) f(x|a) dx = \frac{1}{2a} \int_0^{2a} u(x) dx.$$

Given a specific function to replace $u(\cdot)$ we can compute $v(a)$ for any $a \in [0, 50]$. As a concrete example, let $u(x) = 18\sqrt{x}$. Then we have

$$v(a) = \frac{1}{2a} \int_0^{2a} 18x^{\frac{1}{2}} dx = \frac{9}{a} \left[\frac{2}{3} x^{\frac{3}{2}} \right]_0^{2a} = \frac{6}{a} (2a)^{\frac{3}{2}} = 12\sqrt{2a}.$$

2.2.3 Caveat: It's Not Just the Order Anymore

Recall that when we introduced the idea of payoffs in Section 1.1.2, we argued that any payoff function that preserves the order of outcomes as ranked by the preference relation $\succsim$ will be a valid representation for the preference relation $\succsim$. It turns out that this statement is no longer true when we step into the realm of expected payoff theory as a paradigm for evaluating random outcomes.

Looking back at the R&D problem in Figure 2.3, we took a leap when we equated the player's payoff with profit. This step may seem innocuous: it is pretty reasonable to assume that, other things being equal, a rational player will prefer more money to less. Hence for the player in the R&D problem we have $10 \succ 9 \succ 0 \succ -1$, a preference relation that is indeed captured by our imposed payoff function where $u(x) = x$.

What would happen if payoffs were not equated with profits? Consider a different payoff function to represent these preferences. In fact, consider only a slight modification as follows: $u(10) = 10$, $u(9) = 9$, $u(0) = 0$, and $u(-1) = -8$. The order of outcomes is unchanged, but what happens to the expected payoffs? $E[u(s)] = 5$ is unchanged, but now

$$v(g) = E[u(x)|g] = 0.75 \times 9 + 0.25 \times (-8) = 4.75.$$

Thus even though the order of preferences has not changed, the player would now prefer to choose s instead of g , just because of the different payoff number we assigned to the profit outcome of -1 .

The reason behind this reversal of choice has important consequences. When we choose to use expected payoff then the *intensity of preferences* matters—something that is beyond the notion of simple order. We can see this from our intuitive description of expected payoff. Recall that we used the intuitive notion of “weights”: payoffs that appear with higher probability have more weight in the expected payoff function. But then, if we change the number value of the payoff of some outcome without changing its order in the payoff representation, we are effectively changing its weight in the expected payoff representation.

This argument shows that, unlike payoff over *certain* outcomes, which is meant to represent *ordinal preferences* $\succsim$, the expected payoff representation involves a *cardinal ranking*, in which values matter just as much as order. At some level this implies that we are making assumptions that are not as innocuous about decision making

when we extend our rational choice model to include preferences over lotteries and choices among lotteries. Nevertheless we will follow this prescription as a benchmark for putting structure on decision problems with uncertainty. We now briefly explore some implications of the intensity of preferences in evaluating random outcomes.

2.2.4 Risk Attitudes

Any discussion of the evaluation of uncertain outcomes would be incomplete without addressing a player's attitudes toward risk. By treating the value of outcomes as "payoffs" and by invoking the expected payoff criterion to evaluate lotteries, we have effectively circumvented the need to discuss risk, because by assumption all that people care about is their expected payoff.

To illustrate the role of risk attitudes, it will be useful to distinguish between monetary rewards and their associated payoff values. Imagine that a player faces a lottery with three monetary outcomes: $x_1 = \$4$, $x_2 = \$9$, and $x_3 = \$16$ with the associated probabilities p_1 , p_2 , and p_3 . If the player's payoff function over money x is given by some function $u(x)$ then his expected payoff is

$$E[u(x)|p] = \sum_{k=1}^3 p_k u(x_k) = p_1 u(x_1) + p_2 u(x_2) + p_3 u(x_3).$$

Now consider two different lotteries: $p' = (p'_1, p'_2, p'_3) = \left(\frac{7}{12}, 0, \frac{5}{12}\right)$ and $p'' = (p''_1, p''_2, p''_3) = (0, 1, 0)$. That is, the lottery p' randomizes between \$4 and \$16 with probabilities $\frac{7}{12}$ and $\frac{5}{12}$, respectively, while the lottery p'' picks \$9 for sure. Which lottery should the player prefer? The obvious answer will depend on the expected payoff of each lottery. If $\frac{7}{12}u(4) + \frac{5}{12}u(16) > u(9)$, then p' will be preferred to p'' , and vice versa. This answer, by itself, tells us nothing about risk, but taken together with the special way in which p' and p'' relate to each other, it tells us a lot about the player's risk attitudes.

The lotteries p' and p'' were purposely constructed so that the average payoff of p' is equal to the sure payoff from p'' : $\frac{7}{12} \times 4 + \frac{5}{12} \times 16 = 9$. Hence, *on average*, both lotteries offer the player the same amount of money, but one is a sure thing while the other is uncertain. If the player chooses p' instead of p'' , he faces the risk of getting \$5 less, but he also has the chance of getting \$7 more. How then do his choices imply something about his attitude toward risk?

Imagine that the player is indifferent between the two lotteries, implying that $\frac{7}{12}u(4) + \frac{5}{12}u(16) = u(9)$. In this case we say that the player is **risk neutral**, because replacing a sure thing with an uncertain lottery that has the same expected monetary payout has no effect on his well-being. More precisely we say that a player is risk neutral if he is willing to exchange any sure payout with any lottery that promises *the same expected* monetary payout.

Alternatively the player may prefer not to be exposed to risk for the same expected payout, so that $\frac{7}{12}u(4) + \frac{5}{12}u(16) < u(9)$. In this case we say that the player is **risk averse**. More precisely a player is risk averse if he is *not* willing to exchange a sure payout with any (nondegenerate) lottery that promises *the same expected* monetary payout. Finally a player is **risk loving** if the opposite is true: he *strictly prefers* any lottery that promises the same expected monetary payout.

Remark Interestingly risk attitudes are related to the fact that the payoff representation of preferences matters above and beyond the rank order of outcomes, as

discussed in Section 2.2.3. To see this imagine that $u(x) = x$. This immediately implies that the player is risk neutral: $\frac{7}{12}u(4) + \frac{5}{12}u(16) = 9 = u(9)$. In addition it is obvious from $u(\cdot)$ that the preference ranking is $\$16 \succ \$9 \succ \$4$. Now imagine that we use a different payoff representation for the same preference ranking: $u(x) = \sqrt{x}$. Despite the fact that the ordinal ranking is preserved, we now have $\frac{7}{12}u(4) + \frac{5}{12}u(16) = \frac{17}{6} < 3 = u(9)$. Hence a player with this modified payoff function, which *preserves the ranking* among the outcomes, will exhibit different risk attitudes.

2.2.5 The St. Petersburg Paradox

Some trace the first discussion of risk aversion to the St. Petersburg Paradox, so named in Daniel Bernoulli's original presentation of the problem and his solution, published in 1738 in the *Commentaries of the Imperial Academy of Science of Saint Petersburg*. The decision problem goes as follows.

You pay a fixed fee to participate in a game of chance. A “fair” coin (each side has an equal chance of landing up) will be tossed repeatedly until a “tails” first appears, ending the game. The “pot” starts at \$1 and is doubled every time a “head” appears. You win whatever is in the pot after the game ends. Thus you win \$1 if a tail appears on the first toss, \$2 if it appears on the second, \$4 if it appears on the third, and so on. In short, you win 2^{k-1} dollars if the coin is tossed k times until the first tail appears. (In the original introduction, this game was set in a hypothetical casino in St. Petersburg, hence the name of the paradox.)

The probability that the first “tail” occurs on the k th toss is equal to the probability of the “head” appearing $k - 1$ times in a row and the “tail” appearing once. The probability of this event is $(\frac{1}{2})^k$, because at any given toss the probability of any side coming up is $\frac{1}{2}$. We now calculate the expected *monetary value* of this lottery, which takes expectations over the possible events as follows: You win \$1 with probability $\frac{1}{2}$; \$2 with probability $\frac{1}{4}$; \$4 with probability $\frac{1}{8}$, and so on. The expected value of this lottery is

$$\sum_{k=1}^{\infty} \frac{1}{2^k} \times 2^{k-1} = \sum_{k=1}^{\infty} \frac{1}{2} = \infty.$$

Thus the expected monetary value of this lottery is infinity! The reason is that even though large sums are very unlikely, when these events happen they are huge. For example, the probability that you will win more than \$1 million is less than one in 500,000!

When Bernoulli presented this example, it was very clear that no reasonable person would pay more than a few dollars to play this lottery. So the question is: where is the paradox? Bernoulli suggested a few answers, one being that of decreasing marginal payoff for money, or a concave payoff function over money, which is basically risk aversion. He correctly anticipated that the value of this lottery should not be measured in its expected monetary value, but instead in the monetary value of its *expected payoff*.

Throughout the rest of this book we will make no more references to risk preferences but instead assume that every player's preferences can be represented using expected payoffs. For a more in-depth exposition of attitudes toward risk, see Chapter 3 in Kreps (1990a) and Chapter 6 in Mas-Colell et al. (1995).