

Markets, Games, and Strategic Behavior: Recipes for Interactive Learning

Charles A. Holt
University of Virginia
Comments welcome: holt@virginia.edu

August 2004

This draft includes 9 new chapters on Extensive-Form Games, Monopoly/Cournot, Market Institutions and Power, Collusion, Lemons Markets, Asset Markets and Bubbles, Multi-Unit Auctions, Trust and Reciprocity, Voting (unfinished), and Common Pool Resources. A discussion of the Guessing Game has been added to Chapter 3. The more advanced chapters on Bayes' Rule and Information Cascades have been moved to the final part on Information. The appendices contain many new sets of instructions for non-computerized classroom experiments.

As before, the book is organized so that one can cover the first 5 chapters and then skip around to particular topics of special interest. Each chapter is designed to be a reading assignment for a single class, and each chapter is matched with one or more suggested experiments, either run on the Veconlab website or by hand, with instructions provided in the appendices.

Markets, Games, and Strategic Behavior – Charles A. Holt

Charles A. Holt
University of Virginia

© Copyright, All Rights Reserved
All rights reserved.

Markets, Games and Strategic Behavior

Charles A. Holt

© Copyright, All Rights Reserved

Table of Contents

<i>Preface</i>	7
Part I. Basic Concepts: Decisions, Game Theory, and Market Equilibrium	11
<i>Chapter 1. Introduction</i>	13
<i>Chapter 2. A Pit Market</i>	25
<i>Chapter 3. Some Simple Games: Competition, Coordination, and Guessing</i> ..	39
<i>Chapter 4. Risk and Decision Making</i>	51
<i>Chapter 5. Randomized Strategies</i>	63
Part II. Individual Decision Experiments	75
<i>Chapter 6. Probability Matching</i>	77
<i>Chapter 7. Lottery Choice Anomalies</i>	87
<i>Chapter 8. ISO (In Search of ...)</i>	97
Part III. Game Theory Experiments: Treasures and Intuitive Contradictions	107
<i>Chapter 9. Multi-Stage Games in Extensive Form</i>	109
<i>Chapter 10. Generalized Matching Pennies</i>	119
<i>Chapter 11. The Traveler's Dilemma</i>	131
<i>Chapter 12. Coordination Games</i>	145
Part IV. Market Experiments	157
<i>Chapter 13. Monopoly and Cournot Markets</i>	159
<i>Chapter 14. Vertical Market Relationships</i>	171
<i>Chapter 16. Collusion and Price Competition</i>	193
<i>Chapter 17. Product Quality, Asymmetric Information, and Market Failure</i>	201

Markets, Games, and Strategic Behavior – Charles A. Holt

<i>Chapter 18. A Limit Order Asset Market</i>	207
Part V. Auctions	221
<i>Chapter 19. Private Value Auctions</i>	223
<i>Chapter 20. The Takeover Game</i>	239
<i>Chapter 21. The Winner's Curse</i>	245
<i>Chapter 22. Multi-Unit Auctions: Irrigation Reduction and FCC Bandwidth</i> <i>(in progress)</i>	255
Part VI. Bargaining and Fairness.....	265
<i>Chapter 23. Ultimatum Bargaining</i>	267
<i>Chapter 24. Trust, Reciprocity, and Principal-Agent Games</i>	279
Part VII. Public Choice	285
<i>Chapter 25. Voting (in progress)</i>	287
<i>Chapter 26. Voluntary Contributions</i>	289
<i>Chapter 27. The Volunteer's Dilemma</i>	303
<i>Chapter 28. Externalities, Congestion, and Common Pool Resources</i>	313
<i>Chapter 29. Rent Seeking</i>	325
Part VIII. Information and Learning	333
<i>Chapter 30. Bayes' Rule</i>	335
<i>Chapter 31. Information Cascades</i>	351
<i>Chapter 32. Statistical Discrimination</i>	358
<i>Chapter 33. Signaling (unfinished)</i>	369
Appendices: Instructions for Class Experiments	371
<i>Pit Market Instructions (Chapter 2)</i>	371
<i>Game Instructions (Chapter 3)</i>	374
<i>Lottery Choice Instructions (Chapter 4)</i>	376
<i>BS Game Instructions (Chapter 5)</i>	378
<i>Binary Prediction Game (Chapter 6)</i>	380

<i>Lottery Choice Instructions (Chapter 7)</i>	381
<i>Search Instructions (Chapter 8)</i>	383
<i>Instructions: Traveler's Dilemma (Chapter 11)</i>	386
<i>Instructions: Coordination Game (Chapter 12)</i>	388
<i>Instructions: Market Game (Chapter 13)</i>	390
<i>Instructions: Price/Quality Market (Chapter 17)</i>	392
<i>Private Value Auction (Chapter 19)</i>	395
<i>Instructions for Takeover Game (Chapter 20)</i>	397
<i>Common Value Auction (Chapter 21)</i>	399
<i>Multi-Unit Auction (Chapter 22)</i>	401
<i>Ultimatum Bargaining (Chapter 23)</i>	404
<i>Play-or-Keep Game (Chapter 26)</i>	405
<i>Volunteer's Dilemma Game (Chapter 27)</i>	407
<i>Lobbying Game Instructions (Chapter 29)</i>	409
<i>Bayes' Rule Instructions (Chapter 30)</i>	410
References	415

Preface

Economics is enjoying a resurgence of interest in behavioral considerations, i.e. in the study of how people actually make decisions. Laboratory experiments are increasingly used to study behavior in markets, games, and other strategic situations. The rising interest in experiments is reflected in the 2002 Nobel Prize, which went to an experimental economist and an experimental psychologist. Whole new sub-disciplines are arising in the literature (e.g., behavioral game theory, behavioral law and economics, behavioral finance), and experiments provide a discipline guides these developments.

This book is designed to combine a behavioral approach with active classroom learning exercises. This book is a collection of chapters, each containing a single game and the associated analysis of the results. The games set up simple economic situations, e.g. a market or auction, which highlight several related economic ideas. Each chapter provides the reading for a particular class, in a one-a-day approach. The reading can serve either as a supplement to other material, or (preferably) it can be assigned in conjunction with an in-class “experiment” in which students play the game with each other. Doing the experiments *before* the assigned reading enhances the teaching value of these experiments. Many of the games can be run in class “by hand,” with dice or playing cards. The appendix contains instructions for hand-run games in many cases.

For those with student computer access, all games are available online at the Veconlab site:

<http://veconlab.econ.virginia.edu/admin.htm> (for instructor setup),
<http://veconlab.econ.virginia.edu/login.htm> (for participant login).

These web-based games can be set up and run from any standard browser (Internet Explorer or Netscape) that is connected to the internet, without loading any software. For instructions, see the link on the “admin” page above. The web-based programs have fully integrated instructions that automatically conform to the features selected by the instructor in the setup. Web-based games are quicker to administer, and the instructor data displays can provide records of decisions, earnings, round-by-round data averages, and in some cases, theoretical calculations. These displays can be printed or projected for post-experiment discussions. There is an extensive menu of setup options for each game that lets the instructor select parameters, e.g. the numbers of buyers, sellers, decision rounds, fixed payments, payoffs, etc. For example, the private-value auction setup menu allows one to choose the range of randomly determined private

values, the number of bidders, the number of rounds, the pricing rule (first-price or second-price, and “winner-pays” or “all-pay”). I have even taught classes where students design their own experiments and run them on the others in the class, with a formal (Power-Point) presentation of results in the following class.

The first several chapters contain an example of an individual decision, a simple two-person game, and a market with buyers and sellers. These initial chapters raise a few methodological issues, such as whether and when financial incentives are needed for research and/or teaching experiments. In addition, some basic notions of decision making and equilibrium are introduced. The focus in these chapters is on the basics; discussion of anomalies and alternative theories is deferred until later. After these chapters have been covered, there is a lot of flexibility in terms of the order of coverage of the remaining chapters, which are divided into groups: individual decisions, games, auctions, markets, bargaining, public choice, and asymmetric information. It is possible to pick and choose, based on the level and subject matter of the course.

This book provides an organizing device for courses in game theory, topics in microeconomics, and introductory experimental economics. For an upper-level or graduate course in experimental economics, Davis and Holt's (1993) *Experimental Economics* has the advantage of being organized around the main classes of experiments, with presentations of the associated theory and methodological concepts. This book, in contrast, presents a sequence of particular games, one per chapter, with a carefully measured amount of theory and methodology that is mainly presented in the context of specific examples. These chapters are relatively self-contained, which makes it possible to choose a selection tailored for a particular course, e.g. public economics. The book could also be integrated into courses in microeconomics, managerial economics, or strategy at the M.B.A. level. Many of the experimental designs may be of interest to non-economists, e.g. students of political science, anthropology, and psychology, as well as anyone interested in behavioral finance or behavioral law and economics.

Mathematical arguments are simple, since experiments are typically based on parametric cases that distinguish alternative theories. The mathematical calculations are sometimes illustrated with spreadsheet programs that are constructed in a step-by-step process. Then a process of copying blocks of cells results in iterative calculations that converge to equilibrium solutions. Calculus is avoided, except in a couple of places, e.g., in the chapters on auctions. These optional sections are preceded by discrete examples that provide the intuition behind the more general results.

I have tried to keep the text simple. There are no footnotes. References to other papers are very limited, and most are confined to an “extensions and further reading” section at the end of each chapter. The book contains no extensive

surveys of related literature on research experiments. Such surveys can be found in Davis and Holt (1993), Kagel and Roth's (1995) *Handbook of Experimental Economics*, and Hey's (1994) *Experimental Economics*, which are pitched at a level appropriate for advanced undergraduates, graduate students, and researchers in the field. For more discussion of methodology, see Friedman and Sunder's (1994) *Experimental Methods*. The references in all of these books are somewhat out of date, since the number of published economics papers using laboratory methods has almost doubled since 1995, when the last of these books was published. This is an increase of about 1300 publications, and there are many more unpublished working papers (see Figure 1.1 in Chapter 1). These new publications are listed and categorized by keyword on the bibliography of experimental economics and social science that is available on line at <http://www.people.virginia.edu/~cah2k/y2k.htm>

A friend once asked me who my intellectual hero was, and without hesitation I mentioned Vernon Smith (who was then at Arizona and is currently at George Mason). The work that he has done with his colleagues and coauthors has always served as an inspiration for me. His 2002 Nobel Prize in Economics (together with Danny Kahneman) is richly deserved, and the effects of his work pervade many parts of this book.

I would especially like to thank my coauthor Lisa Anderson and her William and Mary students for many helpful suggestions on this book and on the software that it utilizes. Much of what I know about these topics is due to joint research projects with Jacob Goeree (UVA - University van Amsterdam), and with others: Lisa Anderson (William and Mary), Simon Anderson (Virginia), Sheryl Ball (Virginia Tech), Jordi Brandts (Barcelona), Monica Capra (Washington and Lee), Catherine Eckel (Virginia Tech), Jean Ensminger (Caltech), Roland Fryer (Chicago/Harvard), Rosario Gomez (Malaga), Susan Laury (Georgia State), David Reiley (Arizona), Tom Palfrey (Cal Tech), Al Roth (Harvard), Roger Sherman (Houston), and Rick Wilson (Rice). I also received useful suggestions from James Cox (Arizona). Annie Talman suggested the relevance of the film *Wall Street* to the material in Chapter 20. (She is an actress who played the role of Michael Douglas's secretary in the film.) In addition, I was fortunate to have an unusually talented and enthusiastic group of Virginia students who read parts of the manuscript: A.J. Bostian, Jeanna Composti, Kari Ellasson, Erin Golub, Katie Johnson, Shelley Johnson, Kurt Mitman, Angela Moore, and Loren Pitt. Another student, Joe Monaco, set up the Linux server and procedures for running the programs on a network of hand-held, wireless "pocket" PCs with color, touch-sensitive screens. Imagine a game theory class, outside on the University of Virginia "lawn," with students competing in an auction via wireless pocket PC's!

Part I. Basic Concepts: Decisions, Game Theory, and Market Equilibrium

The next several chapters provide an introduction to the three main types of experiments: individual decisions, games, and markets. Each chapter introduces some key concepts, such as expected value, risk aversion, Nash equilibrium, and market efficiency. Knowledge of these concepts makes it possible to pick and choose among the remaining chapters based on the goals and level of the course. All chapters in this section should be completed before proceeding to the sections that follow.

The first part of Chapter 1 provides an introduction to the use of experimental games for teaching and research, which is followed by an overview of the topics covered in the book. The final two sections of this chapter contain a very brief discussion of methodology and a history of the development of experimental economics. These final two sections are optional for students in most courses (other than experimental economics).

Chapter 2 introduces students to a “pit market” that corresponds to the trading of futures contracts in a trading pit, with a free intermingling of prospective buyers and sellers. The chapter describes a design that illustrates the role of prices in achieving an efficient set of trades. The best procedure is to run a class pit market experiment *before* discussing the chapter material. Instructions for this are provided in the appendix, and the instructor is free to use playing cards to design a setup that will complement or supplement the lessons that can be derived from the market design presented in the chapter. The hand-run experiment is a recommended way to stimulate a discussion of how markets work. However, some with computer access may prefer to set up the experiment using the Double Auction program, which is listed as DA on the Veconlab instructor “admin” menu under Markets.

Chapter 3 presents some simple games of competition, coordination and guessing that serve to introduce the notion of a Nash equilibrium. The Prisoner’s Dilemma and Coordination games can be done by giving students playing cards and using the instructions provided in the appendix, or these games can be run from the Matrix Game program (MG) on the Veconlab menu. The Guessing Game, which is easily run by hand or with the program GG, can be used to focus a discussion on how behavior may converge to a Nash equilibrium.

The fourth chapter shifts attention to individual decisions in situations with random elements that affect money payoffs. This permits a discussion of expected money value for someone who is “neutral” towards risk. Non-neutral attitudes (risk aversion or risk seeking) may also arise in some situations.

The notions of decision making in the presence of random elements are used in Chapter 5 to analyze simple matrix games in which the equilibrium may involve randomization. We consider games of chicken, battle-of-sexes, and matching pennies. The techniques used to find equilibria in these simple games will be applied in more complex situations encountered later, such as the choice of price when there is danger of being undercut slightly by competitors.

Chapter 1. Introduction

I. Origins

Like other scientists, economists observe naturally occurring data patterns and then try to construct explanations. Then the resulting theories are evaluated in terms of factors like plausibility, generality, and predictive success. As with other sciences, it is often difficult to sort out cause and effect when many factors are changing at the same time. Thus, there may be several reasonable theories that are roughly consistent with the same observations. As Keynes noted, without a laboratory to control for extraneous factors, economists often “test” their theories by gauging reactions of colleagues. In such an environment, theories may gain support on the basis of mathematical elegance, persuasion, and focal events in economic history like the Great Depression. Theories may fall from fashion, but the absence of sharp empirical tests leaves an unsettling clutter of plausible alternatives. For example, economists are fond of using the word equilibrium preceded by a juicy adjective (e.g. proper, perfect, divine, or universally divine). This clutter is often not apparent in refined textbook presentations.

The development of sophisticated econometric methods has added an important discipline to the process of devising and evaluating theoretical models. Nevertheless, any statistical analysis of naturally occurring economic data is typically based on a host of auxiliary assumptions. Economics has only recently moved in the direction of becoming a laboratory science in the sense that key theories and policy recommendations are suspect if they cannot provide intended results in controlled experiments. This book provides an introduction to the study of economic behavior, organized around games that can be played in class.

The first classroom market games were conducted in a Harvard class by Edward Chamberlin (1948). He had proposed a new theory of “monopolistic” competition, and he used experiments to highlight failures of the standard model of perfect competition. Students were given buyer and seller roles and instructions about how trades could be arranged. For example, a seller would be given a card with a “cost” in dollars. If the seller were to find a buyer who agreed to pay a price above this cost, the seller would earn the difference. Similarly, a buyer would be given a card with a “resale value,” and the buyer could earn the difference if a purchase could be arranged at a price below this resale value. Different sellers may be given cards with different cost numbers, and likewise, buyers may receive different values. These values and costs are the key elements that any theory of market price determination would use to derive predictions, as explained in Chapter 2. Without going into detail, it should be clear that it is possible to set up a laboratory market, and to provide strong financial incentives

by using “large” value and cost numbers and by paying subjects their earnings in cash.

Chamberlin’s classroom markets produced some inefficiencies, which he attributed to the tendency for buyers and sellers in a market to break off and negotiate in small groups. Vernon Smith, who attended Chamberlin’s class, later began running classroom markets with an enforced central clearinghouse for all offers to buy and sell. This trading institution is called a “double auction” since sellers’ ask prices tend to decline at the same time as buyers’ bid prices rise. A trade occurs when the bid-ask spread closes and someone accepts another’s offer to buy or sell. Smith observed efficient competitive outcomes, even with as few as 6-10 traders. This result was significant, since the classical “large numbers” assumptions were not realistic approximations for most market settings. Smith’s early work on the double auction market figured prominently in his 2002 Nobel Prize in Economics.

These classroom markets can be quite useful, because students learn what an economic situation is like “from the inside” before standard presentations of the relevant economic theory. A classroom game (followed by structured question-and-answer discussion) can let students discover the relevant economic principle for themselves, which enhances the credibility of seemingly abstract economic models. Each of the chapters that follow will be built around a single game that can be run in class, using a web-based software suite and/or simple props like dice and playing cards.

II. Overview

Individual Decisions

A market typically involves a relatively complex set of interactions between multiple buyers and sellers over time. Sometimes it is useful to study key aspects of behavior of individuals in isolation. For example, stock markets involve major risks of gains and losses, but such risks depend on anticipated behavior of others in the market. It is straightforward to set up a simple decision experiment by giving a person a choice between gambles or “lotteries,” e.g. between a sure \$10 and a coin flip that yields \$25 in the event of heads. Many of the earlier chapters pertain to such individual decisions, where payoffs may depend on random events but do not depend on the decisions made by others. Some of the topics in these early chapters, like expected values and risk aversion, are useful in the analysis of interactive situations in subsequent chapters. In addition, the chapters in part II of the book are focused on a series of specific decision making situations, e.g. prediction, search, and information processing.

Game Theory

The simplest strategic interactions among economic agents have been modeled as games. In a “matching-pennies” game, for example, each player chooses heads or tails with the prior knowledge that one will win a sum of money when the coins match, and the other will win when the coins do not match. Each person’s optimal decision in such a situation depends on what the other player is expected to do. The systematic study of such situations began with John von Neumann and Oscar Morgenstern’s (1944) *Theory of Games and Economic Behavior*. They asserted that standard economic theory of competitive markets did not apply to the bilateral and small-group interactions that make up a significant part of economic activity. Their “solution” was incomplete, except for the case of “zero-sum” games in which one person’s loss is another’s gain. While the zero-sum assumption may apply to some extremely competitive situations, like sports contests or matching pennies games, it does not apply to situations where players might all prefer some outcomes to others. In particular, economists and mathematicians at the RAND Corporation in Santa Monica, California were trying to apply game-theoretic reasoning to military tactics at the dawn of the Cold War. In many nuclear scenarios it is easy to imagine that the “winner” may be much worse off than would be the case in the absence of war, which results in a non-zero-sum game. At about this time, a young graduate student at Princeton entered von Neumann’s office with a notion of equilibrium that applies to a wide class of games, including the special case of those that satisfy the zero-sum property. John Nash’s notion of equilibrium and the half-page proof that it generally exists were recognized by the Nobel Prize committee about 50 years later. With the Nash equilibrium as its keystone, game theory has recently achieved the central role that von Neumann and Morgenstern envisioned. Indeed, with the exception of supply and demand, the “Nash equilibrium” is probably used as often today as any other construct in economics.

Intuitively speaking, a Nash equilibrium is a set of strategies, one for each player, with the property that nobody would wish to deviate from their planned action given the strategies being used by the other players. Consider a “coordination” game in which two sellers must decide independently whether to sell in a black market, or in an alternative (“red”) market. There is not enough room for two entrants in either market; they both earn zero for any red/red or black/black outcome. With a red/black outcome, the entrant in the black market earns 10 while the entrant in the red market only earns 5. In this game, the congested outcomes (red/red or black/black) that yield zero payoffs are not Nash equilibria, since each person would wish to deviate from their decision given the other’s decision to enter the same market. By this time you should have realized that the Nash equilibrium is a red/black outcome, since each person earns a

positive amount (5 or 10), so neither would want to switch to the other's market and end up earning 0.

This analysis does not settle the issue of which person gets to enter the more profitable black market, and many developments in game theory are focused on the issue of how to predict which of several Nash equilibria will have more predictive power in a particular game. Such games are easy to set up as experiments by letting each person play a red card (Hearts or Diamonds) or a black card (Clubs or Spades). When existing theory makes no prediction, observed behavioral tendencies in such experiments can provide important guideposts for the development of new theories. Indeed, the mathematicians and economists at the RAND Corporation began running game experiments at about the same time as Chamberlin's classroom market experiments. The game theory chapters in part III pertain to simple games, with a careful considerations of conditions under which behavior does or does not conform to equilibrium predictions.

Markets

The chapters in part IV cover several ways that economists have modeled market interactions, with firms choosing prices, production quantities, or entry decisions. Some markets have distinct groups of buyers and sellers, and others more closely resemble stock markets in which purchase and resale is common. Market experiments can be used to assess the antitrust implications of mergers, contracts, and other market conditions. One goal of such experiments is to assess factors that increase the extent to which markets achieve all possible gains from trade, i.e. the *efficiency* of the market. Markets with enough price flexibility and good information about going prices tend to be highly efficient.

Auctions

The advent of web-based communications has greatly expanded the possibilities for setting up auctions that connect large numbers of geographically dispersed buyers and sellers. For example, several economists were recently involved in designing and running an auction whereby farmers made offers to suspend irrigation on designated tracts of land in southwest Georgia during the 2001 growing season. Bids were collected from 8 sites and were ranked and processed by a web-based program. The interim results were viewed online hundreds of miles away by Georgia Environmental Protection Division officials who had to decide when to terminate bidding (Cummings, Holt, and Laury, forthcoming 2004). Web-based auctions have also been used with great success in the sale of communications bandwidth in the U.S. and in Europe.

There are numerous decisions that must be made in setting up an auction, for example, whether to allow bid revisions during the auction. Experiments can

be used to “testbed” alternative sets of rules. In the Georgia Irrigation Reduction Auction, for example, economists at Georgia State began running experiments as soon as the law passed. State officials observed some of these experiments before drafting the actual auction procedures. A large-scale dry run with over a hundred bidders at 5 southwest Georgia locations preceded the actual auction, which involved about 200 farmers and \$5 million dollars in irrigation reduction payments. The auction had much of the look and feel of an experiment, with the reading of instructions, a round-by-round collection of bids, and web-based bid collection and processing. The chapters in part V pertain to auctions, including a web-based irrigation reduction auction.

Bargaining

Economic decisions in small-group settings often raise issues of fairness since earnings may be inequitable. In a simple “ultimatum bargaining game,” one person proposes a way to divide a fixed amount of money, say \$100, and the other may either agree to the division, which is implemented, or may reject the division, which results in zero earnings for both. Attitudes about inequity are more difficult to model than the simpler selfish money-seeking motives that may dominate in impersonal market situations. In this case, research experiments can provide insights in areas where theory is silent or less developed. Many topics in Law and Economics involve bargaining (e.g. bankruptcy, pre-trial settlements), and experiments can provide useful insights. Experiments such as these also have been used by anthropologists (Ensminger, 2001) to study attitudes about fairness in primitive societies in Africa and South America. The chapters in part VI pertain to bargaining and other classes of games where fairness, equity, and other interpersonal factors seem to matter.

Public Choice

Inefficiencies can occur when some costs and values are not reflected in prices. For example, it may be difficult to set up a market that allows public goods like national defense to be provided by decentralized contributions. Another source of potential inefficiency occurs when one person’s activity has a negative impact on others’ well being, as is the case with pollution or the over-use of a freely available, shared resource. Non-price allocation mechanisms often involve the dedication of real resources to lobbying. No individual contestant would want the value of their effort to exceed the value of the object being sought, but the aggregate lobbying costs for a number of individuals may be large relative to the value of the prize. The public choice chapters in part VII pertain to such situations; topics include voluntary contributions and the use of a “common-pool” resource.

Information

Markets may also fail to generate efficient results when prices do not convey private information. With limited information, individuals may rely on “signals” like educational credentials or ethnic background. Informational asymmetries can produce interesting patterns of conformity that may have large effects on stock prices, hiring patterns, etc. Experiments are particularly useful in these cases, since the effect of informational disparities is to produce many Nash equilibria. The information chapters in part VIII include experiments on signaling and discrimination.

III. A Brief Comment on Methodology

This section and the next are optional for students who are primarily interested in microeconomics and game theory, as opposed to research in experimental economics.

Treatments

In an experiment, a treatment is a completely specified set of procedures, which includes instructions, incentives, rules of play, etc. Just as scientific instruments need to be calibrated, it is useful to calibrate economics experiments, which typically involves establishing a “baseline treatment” for comparisons. For example, suppose that individuals are given sums of money that can either be invested in an “individual” or a “group” account, where investments in the group account have a lower return to the individual, but a higher return to all group members. If the typical pattern of behavior is to invest half in each account, then this might be attributed either to the particular investment return functions used in the experiment or to “going fifty-fifty” in an unfamiliar situation. In this case, a pair of treatments with differing returns to the individual account may be used. Suppose that the investment rate for the individual account is fifty percent in one treatment, which could be attributed to confusion or uncertainty, as indicated above. The importance of the relevant economic incentives could be established if this investment rate falls, say to ten percent, when the return for investing in the individual account is reduced.

Next, consider a market example. High prices may be attributed to small numbers of sellers or to the way in which buyers are constrained from requesting private discounts. These issues could be investigated by changing the number of sellers, holding buyers’ shopping rules constant, or by changing buyers’ shopping opportunities by holding the number of sellers constant. Many economics experiments involve a “2x2” design with treatments in each of the four cells: low numbers with buyer shopping; low numbers with no shopping; high numbers with no shopping; high numbers with shopping.

Treatment Structure: “Between-Subjects” Versus “Within-Subjects” Designs

Each of the following chapters is based on a single experiment. In order to preserve time for discussion, the experiments typically involve a pair of treatments. One issue is whether half of the people are assigned to each treatment, which is called a “between-subjects” design. The alternative is to let each person make decisions in both treatments, which results in a larger sample size but less time to complete multiple rounds of decision making in each treatment. This is called a “within-subjects” design, since each person’s behavior in one treatment is compared with the same person’s behavior in the other treatment. Each method has its advantages. If behavior is slow to converge or if many observations are required to measure what is being investigated, then the between-subjects design may be preferred since it will generate the most observations in the limited time available. Moreover, subjects are only exposed to a single treatment, which avoids sequence effects. For example, a market experiment that lets sellers discuss prices may result in some successful collusive arrangements, with high prices that may carry over even if communication is not allowed in a second treatment. These sequence effects may be avoided with a between-subjects design. The advantage of a within-subjects design is that individual differences are controlled for by letting each person serve as their own control. For example, suppose that you have a group of eight adults, and you want to determine whether they run as fast wearing blue jeans as with running shorts. In any group of eight adults, running speeds may vary by factors of 2 or 3, depending on weight, age, health, etc. In this case it would be desirable to time running speeds for each person under both conditions unless the distance were so great that fatigue would cause major sequence effects. In general, a within-subjects design is more appealing if there is a high behavioral variability across individuals.

Incentives

Economics experiments typically involve monetary decisions like prices, costly efforts, etc. Most economists are suspicious of results of experiments done with hypothetical incentives, and therefore real cash payments are generally used in research experiments. As we shall see, sometimes incentives matter a lot and sometimes not much at all. For example, people have been shown to be more generous in offers to others when such offers are hypothetical than when generosity has a real cost. For purposes of teaching, it is typically not possible or even desirable to use monetary payments. Nevertheless, the results of class experiments can provide a useful learning experience as long as the effects of payments in research experiments are provided when important incentive effects have been documented. Therefore, the presentations in the chapters that follow will be based on a mixture of class and research experiments. When the term “experiment” or “research experiment” is used, this will mean that all earnings

were at reasonable levels and were paid in cash. The term “classroom experiment” indicates that payoffs were basically hypothetical. However, for non-market classroom experiments discussed in this book, I would pick one person at random *ex post* and pay a small fraction of earnings, usually several dollars. This procedure is not generally necessary, but it was followed here to reduce unexpected differences between classroom and research data.

Replication

One of the main advantages of experimental analysis is the ability to repeat the same setup numerous times in order to determine average tendencies that are relatively insensitive to individual or group effects. Replication requires that instructions and procedures be carefully documented. For example, it is essential that instructions to subjects be written as a script that is followed in exactly the same manner with each cohort that is brought to the laboratory. Having a set of written instructions helps ensure that “biased” terminology is avoided, and it permits other researchers to replicate the reported results. The general rule is that enough detail should be reported so that someone else could redo the experiment in a manner that the original author(s) would accept as being valid, even if the results turned out to be different. For example, if the experimenters provide a number of examples of how prices determine payoffs in a market experiment, and if these examples are not contained in the written instructions, the different results in a replication may be due to differences in the way the problem is presented to the subjects.

Control

A second main advantage of experimentation is the ability to control the factors that may be affecting observed behavior, so that extraneous factors are held constant (controlled) as the treatment variable changes. The most common cause of loss of control is changing more than one factor at the same time, so that it is difficult to determine the cause of any observed change. For example, a well-known study of lottery choice compared the tendency of subjects to choose a sure amount of money with a lottery that may yield higher or lower payoffs. In one treatment, the payoffs were in the zero to \$10 range, and in another treatment the payoffs were in the thousands of dollars. Then the author concluded that there were no incentive effects, since there was no observed difference in the tendency to choose the safe payoff. The trouble with this conclusion is that the high-payoff treatment was conducted under hypothetical conditions, so two factors were being changed: the scale of the payoffs, and whether or not they were real or hypothetical. This conclusion turns out to be questionable, given results of experiments that change one factor at a time. For example, Holt and Laury (2001) report that scaling up hypothetical choices has little effect on the tendency to

select the safer option in their experiments. Thus, choice patterns with low real payoffs do look like choice patterns with both low and high hypothetical payoffs. However, scaling up the payoffs in the real (non-hypothetical) treatments to hundreds of dollars causes a sharp increase in the tendency to choose the safer option.

Control is also lost when procedures make it difficult to determine the incentives that participants actually faced in an experiment. If people are trading physical objects like university sweatshirts, differences in individual valuations make it hard to reconstruct the nature of demand in a market experiment. There are, of course, situations where non-monetary rewards are desirable, such as experiments designed to test whether ownership of a physical object makes it more desired (“the endowment effect”). Thus control should always be judged in the context of the purpose of the experiment.

Fatal Errors

Professional economists often look to experimental papers for data patterns that support existing theories or that suggest desirable properties of new theories. Therefore, the researcher needs to be able to distinguish between results that will replicate from those that are artifacts of improper procedures. Even students in experimental sciences should be sensitive to procedural matters so that they can evaluate others’ results critically. Experiments also provide a rich source of topics for term papers, senior dissertations, etc. Students and others who are new to experimental methods in economics should be warned that there are some fatal errors that can render useless the results of economics experiments. As the above discussion indicates, these include: inappropriate incentives, non-standardized instructions and procedures, loss of control due to the failure to provide a calibrated baseline treatment or the change in more than one design factor at the same time.

A more detailed treatment of methodology can be found in Chapters 1 and 9 of Davis and Holt (1993) *Experimental Economics* and Friedman and Sunder (1994) *Experimental Methods*.

IV. A Brief History of Experimental Economics

Figure 1.1 below shows the trends in published papers in experimental economics. The first papers by Chamberlin and some of the game theorists at RAND were written in the late 1940’s and early 1950’s. In addition, some of the early interest in experimental methods was generated by the work of Fouraker and Siegel (1962, 1963). Siegel was a psychologist with high methodological standards; some of his work on “probability matching” will be discussed in a later chapter. In the late 1950’s, some business school faculty at places like Carnegie-Mellon became interested in business games, both for teaching and research. And

Vernon Smith's early market experiments were published in 1962. Even so, there were less than 10 publications per year before 1965, and less than 30 per year before 1975. Some of the interesting work during this period was being done by Reinhard Selten and other Germans, and there was an international conference on experimental economics held in Germany in 1973. During the late 1970's, Vernon Smith was a visitor at Caltech, where he began working with Charles Plott, who had studied at Virginia under James Buchanan and was interested in public choice issues. Plott's (1979) first voting experiments stimulated work on voting and agendas by political scientists in the early 1980's. Other interesting work included the Battalio, Green, and Kagel (1981) experiments with rats and pigeons, and Al Roth's early bargaining experiments, e.g. Roth and Malouf (1979). There were still fewer than 50 publications per year in the area before 1985. At this time, my former thesis advisor and colleague, Ed Prescott, told me that "Experimental economics was dead end in the 1960's and it will be dead end in the 1980's."

In the 1980's, Vernon Smith and his colleagues and students at Arizona established the first large laboratory and began the process of developing computerized interfaces for experiments. In particular, Arlington Williams wrote the first posted-offer program. After a series of conferences in Tucson, the Economic Science Association was founded in 1986, and the subsequent presidents constitute a partial list of key contributors (Smith, Plott, Battalio, Hoffman, Holt, Forsythe, Palfrey, Cox, Schotter, Camerer, and Fehr).




Figure 1.1. Numbers of Published Papers in Experimental Economics

Vernon Smith and Mark Isaac begin editing *Research In Experimental Economics*, a series of collected papers appearing about every other year since 1979. The first comprehensive books in this area were published in the 1990's, e.g. Davis and Holt's (1993) and Hey (1994). The 1995 *Handbook of Experimental Economics*, edited by Kagel and Roth, contains survey papers on key topics like auctions, bargaining, public goods, etc, and these are still good sources for reference. The first specialty journal, *Experimental Economics*, was started in 1998. The strong interest among Europeans is indicated by the fact that one of the founding coeditors is Arthur Schram from Amsterdam. The Economic Science Association now has an annual international meeting and additional regional meetings in the U.S. and Europe. These developments have resulted in over a hundred publications in this area every year since 1990, with highs of over 200 per year since 1999.

These numbers indicate the growing acceptance of experimental methods in economics. Since the total number of publications in all areas of economics has increased, it is natural to consider papers published in the top journals, where the total article count is roughly constant. In the *American Economic Review*, for example, there was one experimental economics article in the 1950's, 3 articles in the 1960's, 11 in the 1970's, 40 in the 1980's, and 86 articles in the 1990's. Similarly, the number of publications in *Econometrica* went from 2 in the 1970's to 12 in the 1980's to 27 in the 1990's. A searchable bibliography of over 4,000 papers in experimental economics and related social sciences can be found in the author's *Y2K Bibliography of Experimental Economics*: <http://www.people.virginia.edu/~cah2k/y2k.htm> (which also contains HTML reading lists for specific topics).

There are lots of exciting developments in this field in recent years. Economics experiments are being integrated into introductory courses and the workbooks of some major texts. Theorists are looking at laboratory results for applications and tests of their ideas, and policy makers are increasingly willing to look at how proposed mechanisms perform in controlled tests. Experimental methods have been used to design large auctions (e.g. the FCC auction for bandwidth and the Georgia Irrigation Auction) and systems for matching people with jobs (e.g. medical residents and hospitals). Two of the recipients of the 1994 Nobel Prize in Economics, Nash and Selten, were game theorists who had run their own experiments. Most significantly, the 2002 Nobel Prize in Economics was awarded to an experimental economist (Smith) and to an experimental psychologist (Kahneman) who is widely cited in the economics literature. Economics is well on its way to becoming an experimental science.

Chapter 2. A Pit Market

When buyers and sellers can communicate freely and openly as in a trading “pit,” the price and quantity outcomes can be predictable and efficient. Deviations tend to be relatively small and may be due to informational imperfections. The discussion should be preceded by a class experiment, either using playing cards and the instructions provided in the Appendix, or using the Vecon web-based double auction program (“DA”).

I. A Simple Example

Chamberlin (1948) set up the first market experiment by letting students with buyer or seller roles negotiate trading prices. The purpose was to illustrate systematic deviations from the standard theory of perfect competition. Ironically, this experiment is most useful today in terms of what factors it suggests are needed to promote efficient, competitive market outcomes.

Each seller was given a card with a dollar amount or “cost” written on it. For example, one seller may have a cost of \$2, and another may have a cost of \$8. The seller earns the difference between the sale price and the cost, so the low-cost seller would be searching for a price above \$2 and the high-cost seller would be searching for a price above \$8. The cost is not incurred unless a sale is made, i.e. the product is “made to order.” The seller would not want to sell below cost, since the resulting loss is worse than the zero earnings from no sale.

Similarly, each buyer was given a card with a dollar amount or “value” written on it. A buyer with a value of \$10, for example, would earn the difference if a price below this amount could be negotiated. A buyer with a lower value, say \$4, would refuse prices above that level, since a purchase above value would result in negative earnings.

The market is composed of groups of buyers and sellers who can negotiate trades with each other, either bilaterally or in larger groups. For example, suppose that the market structure is shown in Table 2.1. There are four buyers with values of \$10, and four buyers with values of \$4. Similarly, there are four sellers with costs of \$2, and four sellers with costs of \$8.

In addition to the “structural” elements of the market (numbers of buyers and sellers, and their values and costs), we must consider the nature of the market price negotiations. A *market institution* is a full specification of the rules of trade. For example, one might let sellers “post” catalogue prices and then let buyers contact sellers if they wish to purchase at a posted price, with discounts not permitted. This institution, known as a “posted-offer auction,” is sometimes used in laboratory studies of retail markets. The asymmetry, with one side posting and the other responding, is common when there are many people on the

Table 2.1. A Market Example

	Values		Costs
Buyer 1	\$10	Seller 9	\$2
Buyer 2	\$10	Seller 10	\$2
Buyer 3	\$10	Seller 11	\$2
Buyer 4	\$10	Seller 12	\$2
Buyer 5	\$4	Seller 13	\$8
Buyer 6	\$4	Seller 14	\$8
Buyer 7	\$4	Seller 15	\$8
Buyer 8	\$4	Seller 16	\$8

responding side and few on the posting side. Posting on the thin side may conserve on information costs, and agents on the thin side may have the “power” to impose prices on a “take-it-or-leave-it” basis. In contrast, Chamberlin used an institution that was symmetric and less structured; he let buyers and sellers mix together and negotiate bilaterally or in small groups, much as traders of futures contracts interact in a trading “pit.” Sometimes he announced transactions prices as they occurred, much as the market officials watching from a “pulpit” over the trading pit would key in contract prices that are posted electronically and flashed to other markets around the world. At other times, Chamberlin did not announce prices as they occurred, which may have resulted in more decentralized trading negotiations.

II. A Classroom Experiment

Figure 2.1 shows the results of a classroom pit market experiment using the setup in Table 2.1. Participants were University of Virginia Education School students who were interested in new approaches to teaching economics at the secondary school level. When a buyer and a seller agreed on a price, they came together to the recording desk, where the price was checked to ensure that it was no lower than the seller’s cost and no higher than the buyer’s value, as required by the instructions given in the appendix of this chapter. (The prohibition of trading at a loss was used since the earnings in this classroom experiment were only hypothetical.) Each dot in the figure corresponds to a trade. Notice that the quantity of trades begins at 5, goes up to 6, and then declines to four. The prices are variable in the first two periods and then stabilize in the \$5-\$7 range.

Consider the question of why the prices converged to the \$5-\$7 range, with a transactions quantity of four units per period. Notice that the quantity could have been as high as eight. For example, suppose that the four high-value (\$10) buyers negotiated with high-cost (\$8) sellers and agreed on prices of \$9. Similarly, suppose that the negotiated prices were \$3 for sales from the low-cost (\$2) sellers to the low-value (\$4) buyers. In this scenario, all of the sellers’ units

sell, the quantity is eight, and each person earns \$1 for the period. Prices, however, are quite variable. These patterns are not observed in the price sequence of trades shown in Figure 2.1. There is high variability initially, as some of the high-cost units are sold at high prices, and some of the low-value units are purchased at low prices. By the third period, prices have converged to the \$5-\$7 range that prevents high-cost sellers from selling at a profit, and that prevents low-value buyers from buying at a profit. The result is that only the four low-cost units are sold to the four high-value buyers.




Figure 2.1. Contract Price Sequence

The intuitive reason for the low price dispersion in later periods is fairly obvious. At a price of \$9, all eight sellers would be willing (and perhaps eager) to sell, but only the four high-value buyers would be willing to buy, and perhaps not so eager given the low buyer earnings at that high price. This creates a competitive situation in which sellers may try to lower prices to get a sale, causing a price decline. Conversely, suppose that prices began in the \$3 range. At these low prices, all eight buyers would be willing to buy, but only the four low-cost sellers would be willing to sell. This gives sellers the power to raise price without losing sales.

Finally, consider what happens with prices in the intermediate range, from \$4 to \$8. At a price of \$6, for example, each low cost seller earns \$4 ($= \$6 - \2), and each high-value buyer earns \$4 ($= \$10 - \6). Together, the eight people who make trades earn a total of \$32. These earnings are much higher than \$1 per person that was earned with price dispersion, for a total of \$1 times 16 people =

\$16. In this example, the effect of reduced price dispersion is to reduce quantity by half and to double total earnings. The reduced price dispersion benefits buyers and sellers as a group, since total earnings go up, but the excluded high-cost sellers and low-value buyers are worse off. Nevertheless, economic efficiency has increased by these exclusions. The sellers have high costs because the opportunity costs of the resources they use are high, i.e. the value of the resources they would employ (perhaps inefficiently) is higher in alternative uses. Moreover, the buyers with low values are not willing to pay an amount that covers the opportunity cost of the extra production needed to serve them.

One way to measure the efficiency of a market is to compare the actual earnings of all participants with the maximum possible earnings. It is a simple calculation to verify that \$32 is the highest total earnings level that can be achieved by any combination of trades in this market. The efficiency was 100% for the final two periods shown in Figure 2.1, since the four units that traded in each of those periods involved low costs and high values. If a fifth unit had traded, as happened in the first period, the aggregate earnings must go down since the high cost (\$8) for the fifth unit exceed the low value (\$4) for this unit. Thus the earnings total goes down by the difference \$4, reducing earnings from the maximum of \$32 to \$28. In this case, the outcome with price dispersion and five units traded only has an efficiency level of $28/32 = 87.5$ percent. The trading of a sixth unit in the second period reduced efficiency even more, to 75 percent.

The operation of this market can be illustrated with the standard supply and demand graph. First consider a seller with a cost of \$2. This seller would be unwilling to supply any units at prices below this cost, and would offer the entire capacity (1 unit) at higher prices. Thus the seller's individual supply function has a "step" at \$2. The total quantity supplied by all sellers in the market is zero for prices below \$2, but market supply jumps to four units at slightly higher prices as the four low-cost sellers offer their units. The supply function has another step at \$8 when the four high-cost sellers offer their units at prices slightly above this high-cost step. This resulting market supply function, with steps at each of the two cost levels, is shown by the solid line in Figure 2.2. The market demand function is constructed analogously. At prices above \$10, no buyer is willing to purchase, but the quantity demanded jumps to four units at prices slightly below this value. The demand function has another step at \$4, as shown by the dashed line in Figure 2.2. Demand vertical at prices below \$4, since all buyers will wish to purchase at any lower price. The supply and demand functions overlap at a quantity of 4 in the range of prices from \$4 to \$8. At any price in this region, the quantity supplied equals the quantity demanded. At lower prices, there is excess demand, which would tend to drive prices up. At prices above the region of overlap, there is excess supply that would tend to drive prices back down.

The fact that the maximum aggregate earnings are \$32 for this design can be seen directly from Figure 2.2. Suppose that the price is \$6 for all trades. The value of the first unit on the left is \$10, and since the buyer pays \$6, the surplus value is the difference, or \$4. The surplus on the second, third, and fourth units is also \$4, so the “consumers’ surplus” is the sum of the surpluses on individual consumers’ units, or \$16. Notice that consumer’s surplus is the area under the




Figure 2.2. A Simple Market Design

demand curve and above the price paid. Since sellers earn the difference between price and cost, the “producers’ surplus” is the area (also \$16) above the supply curve and below the price. Thus the “total surplus” is the sum of these areas, which equals the area under demand and above supply (to the left of the intersection). This total surplus area is four times the vertical distance (\$10 - \$2), or \$32. The total surplus is actually independent of the particular price at which the units trade. For example, a higher price would reduce consumers’ surplus and increase producers’ surplus, but the total area would remain fixed at \$32. Adding a fifth unit would reduce surplus since the cost (\$8) is greater than the value (\$4).

These types of surplus calculations do not depend on the particular form of the demand and supply functions in Figure 2.2. Individual surplus amounts on each unit are the difference between the value of the unit (the height of demand) and the cost of the unit (the height of supply). Thus the area between demand and supply to the left of the intersection represents the maximum possible total surplus, even if demand and supply have more steps than the example in Figure 2.2.

Figure 2.3 shows the results of a classroom experiment with 9 buyers and 9 sellers, using a design with more steps and an asymmetric structure. Notice that demand is relatively “flat” on the left side, and therefore that the competitive price range (from \$7 to \$8) is relatively high. For prices in the competitive range, the consumers’ surplus will be much smaller than the producers’ surplus. The right side shows the results of two periods of pit market trading, with the prices plotted in the order of trade. Prices start at about \$5, in the middle of the range between the lowest cost and the highest value. The prices seem to be converging to the competitive range from below, which could be due to buyer resistance to increasingly unequal earnings. In both periods, all seven of the higher value (\$10 and \$9) units were purchased, and all seven of the lower cost (\$2 and \$7) units were sold. As before, prices stayed in a range needed to exclude the high-cost and low-value units, and efficiency was 100% in both periods.

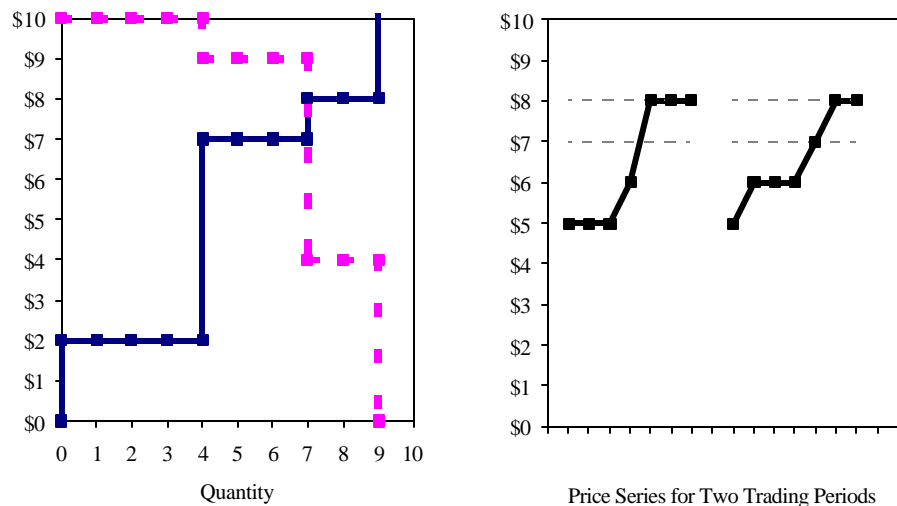


Figure 2.3. Demand, Supply, and Transactions Prices for a Pit Market

The asymmetric structure in Figure 2.3 was used to ensure that prices did not start in the equilibrium range, in order to illustrate some typical features of price adjustments. The first units that traded in each period were the ones with \$10 values and \$2 costs on the left side of the supply and demand figure. After these initial transactions in the \$5-\$7 range, the remaining traders were closer to the competitive “margin.” These marginal buyers were those with values of \$9 and \$4. The marginal sellers were those with costs of \$7 and \$8. Clearly these units will have to sell for prices above \$7, which forces the prices closer to the competitive prediction at the end of the period. When sellers who sold early in

the period see these higher prices, they may hold out for higher prices in the next period. Similarly, buyers will come to expect prices to rise later in the period, so they will scramble to buy early, which will tend to drive prices up earlier in each subsequent period. To summarize, the convergence process is influenced by the tendency for the highly profitable units (on the left side of demand and supply) to trade early, leaving traders “at the margin” where price negotiations tend to be near competitive levels. Price dispersion is narrowed in subsequent periods as traders come to expect the higher prices at the end of the period.

III. Chamberlin’s Results and Vernon Smith’s Reaction

The negotiations in the class experiments discussed above took place in a central area that served as the trading pit, although some people tended to break off in pairs to finalize deals. Even so, the participants could hear the public offers being made by others, a process which should reduce price dispersion. A high-value buyer, who is better off paying up to \$10 instead of not trading, may not be willing to pay such high prices when some sellers are making lower offers. Similarly, a low-cost seller will be less willing to accept a low price when other buyers are observed to pay more. In this manner, good market information about the going prices will tend to reduce price dispersion. A high dispersion is needed for high-cost sellers and low-value buyers to be able to find trading partners, so less dispersion will tend to exclude these “extra-marginal” traders.

Chamberlin did report some tendency for the markets to yield “too many” trades relative to competitive predictions, which he attributed to the dispersion that can result from small group negotiations. In order to evaluate this conjecture, he took the value and cost cards from his experiment and used them to *simulate* a *decentralized* trading process. These simulations were not laboratory experiments with student traders; they were mechanical, the way one would do computer simulations today.

For groups of size two, Chamberlin would shuffle the cost cards and the value cards, and then he would match one cost with one value and “make” a trade at an intermediate price if the value exceeded the cost. Cards for trades that were not made were returned to the deck to be reshuffled and re-matched. It is useful to see how this random matching would work in a simple example with only four buyers and four sellers, as shown in Table 2.2. The random pairing process would result in only two trades if the two low-cost units were matched with the two high-value units. It would result in four trades when both low-cost units were matched with the low-value units, and high-cost units are matched with high-value units. The intermediate cases would result in three trades, for example when the value/cost combinations are: **\$10/\$2**, **\$10/\$8**, **\$4/\$2**, and **\$4/\$8**. The three value/cost pairs that result in a trade are shown in bold. To summarize, random matches in this example produce a quantity of trades of either two, three, or four,

and some simulations should convince you that on average there will be three units traded.

Table 2.2. An Eight-Trader Example

	Values		Costs
Buyer 1	\$10	Seller 5	\$2
Buyer 2	\$10	Seller 6	\$2
Buyer 3	\$4	Seller 7	\$8
Buyer 4	\$4	Seller 8	\$8

For groups larger than two traders, Chamberlin would shuffle and deal the cards into groups and would calculate the competitive equilibrium quantity for each group, as determined by the intersection of supply and demand for that group alone. In the four-buyer/four-seller example in Table 2.2, there is only one group of size 4, i.e. all four traders, and the equilibrium for all four is the competitive quantity of 2 units. This illustrates Chamberlin's general finding that quantity tended to decrease with larger group size in his simulations.

Notice the relationship between the use of simulations and laboratory experiments with human participants. The experiment provided the empirical regularity (excess quantity) that motivated a theoretical model (competitive equilibrium for subgroups), and the simulation confirmed that the same regularity would be produced by this model. In general, computer simulations can be used to derive properties of models that are too complex to solve analytically, which is often the case for models of out-of-equilibrium behavior and dynamic adjustment. The methodological order can, of course, be reversed, with computer simulations being used to derive predictions that can be tested with laboratory experiments using human participants.

The simulation analysis given above suggests that market efficiency will be higher when price information is centralized, so that all traders know the "going" levels of bid, ask, and transactions prices. Vernon Smith, who attended some of Chamberlin's experiments when he was a student at Harvard, used this intuition to design a trading institution that promoted efficiency. After some classroom experiments of his own, he began using a "double auction" in which buyers made bids, sellers made offers (or "asks"), and all could see the highest outstanding bid and the lowest outstanding ask. Buyers could raise the current best bid at any time, and sellers could undercut the current best ask at any time. In this manner, the bid/ask spread would typically diminish until someone closed a contract by accepting the terms from the other side of the market, i.e. until a seller accepted a buyer's bid or a buyer accepted a seller's ask. This is called a *double auction*, since it involves both buyers (bidding up, as in an auction for

antiques) and sellers (bidding down, as would occur if contractors undercut each others' prices). A trade occurs when these processes meet, i.e. when a buyer accepts a seller's ask or when a seller accepts a buyer's bid.

Table 2.3 shows a typical sequence of bids and asks in a double auction. Buyer 2 bids \$3.00, and seller 5 offers to sell for \$8.00. Seller 6 comes in with an ask price of \$7.50, and buyer 1 bids \$4.00 and then \$4.50. At this point seller 7 offers to sell at \$6.00, which buyer 2 accepts.

Table 2.3. A Price Negotiation Sequence

	<i>Bid</i>	<i>Ask</i>	
Buyer 2	\$3.00		
		\$8.00	Seller 5
		\$7.50	Seller 6
Buyer 1	\$4.00		
Buyer 1	\$4.50		
		\$6.00	Seller 7
Buyer 2	Accepts \$6.00		

In a double auction, there is always good public information about the bid/ask spread and about past contract prices. This information creates a large-group setting that tends to diminish price variability and increase efficiency. Smith (1962, 1964) reported that this double auction resulted in efficiency measures of over ninety percent, even with relatively small numbers of traders (4-6) and with no information about others' values and costs. The double auction tends to be somewhat more centralized than a pit market, which does not close off the possibility of bilateral negotiations that are not observed by others. Double auction trading more closely resembles trading for securities on the New York Stock Exchange, where the specialist collects bids and asks, and all trades come across the "ticker tape."

Besides introducing a centralized posting of all bids, asks, and trading prices, Smith introduced a second key feature into his early market experiments: the repetition of trading in successive market "periods" or "trading days." The effects of repetition are apparent in Figure 2.4, which was done with the Veconlab web-based interface in a classroom setting. The supply and demand curves for the first two periods are shown in blue on the left side; the values and costs that determine the locations of the steps were randomly generated. There were four trading periods which are delineated by the vertical lines that separate the dots that show the transactions price sequence. The prices converge to near competitive levels in the second period. The demand and supply shift in period 3

(higher red lines on the left) causes a quick rise in prices, and convergence is observed by the end of period 4.

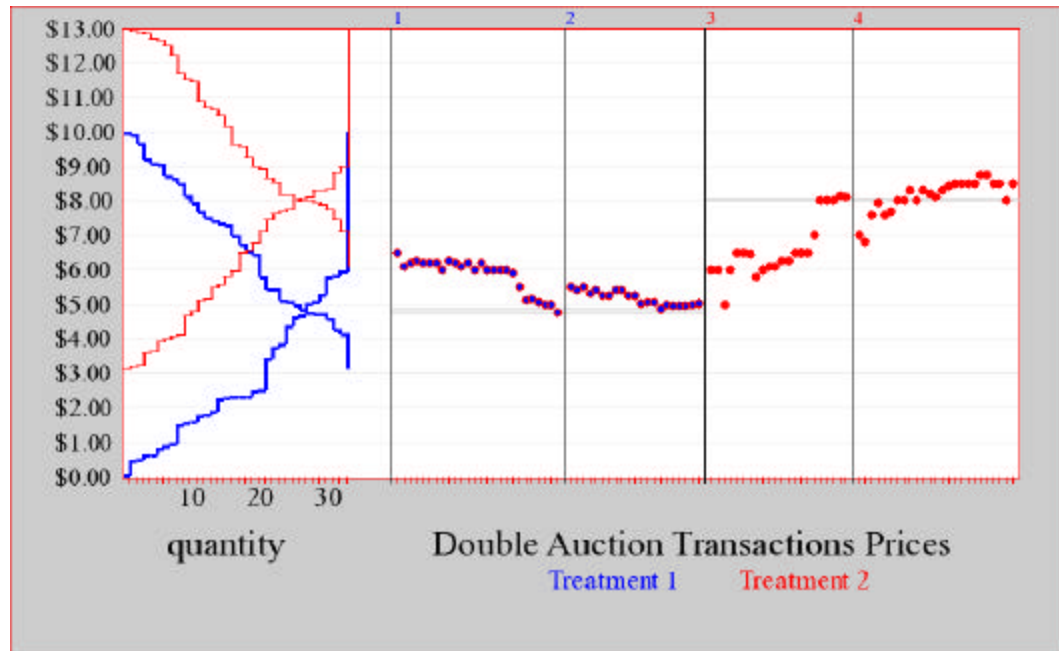


Figure 2.4. Demand, Supply, and Transactions Prices for a Double Auction with a Structural Change after the Second Period

Figure 2.5 shows the results of a double auction market conducted under research conditions with a much more extreme configuration of values and costs that was used to provide a stress test of the convergence properties of a double auction. The four buyers begin with four units each, for a total of 16 units, all with values of \$6.80. The four sellers have a total of 11 units, all with costs of \$5.70. In addition, each person received a “commission” of 5 cents for each unit that they bought or sold. The left side of the figure shows the resulting supply function, which is vertical at a quantity of 11 where it crosses demand. The intersection is at a price of \$6.80, and at lower prices there is excess demand. The period-by-period price sequences are shown on the right side of the figure. Notice that first-period prices start in the middle of the range between values and costs, and then rise late in the period. Prices in period 2 start higher and rise again. This upward trend continues until prices reach the competitive level of about \$6.80 in period 4. At this point, buyers are only earning pennies (e.g. the commission) on each transaction. A second treatment, not shown, began in period 6 with the total

number of buyer units reduced to 11 and the number of seller units increased to 16. Buyers, who were already earning very little on each unit, were disappointed at having fewer units. But prices started to fall immediately; the very first trading price in period 6 was only \$6.65. Prices declined steadily, reaching the new competitive price of \$5.70 by period 10. This session illustrates the strong convergence properties of the double auction, where excess demand or supply pressures can push prices to create severe earnings inequities that are unlikely to arise in bilateral bargaining situations.




Figure 2.5. Demand, Supply, and Transactions Prices for a Double Auction

IV. Extensions

The simple experiment presented in this chapter can be varied in numerous useful directions. For example, an upward shift in demand, accomplished by increasing buyers' values, should raise prices. An increase in the number of sellers would tend to shift supply outward, which has the effect of lowering prices. The imposition of a \$1 per-unit tax on buyers would (in theory) shift demand down by \$1. To see this, note that if one's value is \$10, but a tax of \$1 must be paid upon purchase, then the net value is only \$9. All demand steps would shift down by \$1 in this manner. Alternatively, a \$1 tax per unit imposed on sellers would shift supply up by \$1, since the tax is analogous to a \$1 increase in cost. Some extensions will be considered in the chapter on double auctions that is in the market experiments section of this book.

We have discussed several alternative market institutions here. Chamberlin's *pit market* corresponds most closely to trading of futures contracts in a trading pit, whereas Smith's *double auction* is more like the trading of securities on the New York Stock Exchange. A *posted-offer auction* (where sellers set prices on a take-it-or-leave-it basis) is more like a retail market with many buyers and few sellers. These different trading institutions have different properties and applications, and outcomes need not match competitive predictions. Nevertheless, it is useful to consider outcomes in terms of market efficiency, measured as the percentage of maximum earnings achieved by the trades that are made. In considering alternative designs for ways to auction off some licenses, for example, one might want to consider the efficiency of the allocations along with the amounts of revenue generated for the seller(s).

The theoretical predictions discussed in this chapter are derived from an analysis of supply and demand. In the competitive model, all buyers and sellers take price as given. This model may not provide good predictions when some traders perceive themselves as being price makers, with "power" to push prices in their favor. A single-seller monopolist may be able to raise prices above competitive levels, and a small group of sellers with enough of a concentration of market capacity may be able to raise prices as well. The exercise of this type of market power is discussed in the chapter on posted-offer auctions, in which sellers post prices on a take-it-or-leave-it basis. The exercise of market power by price-setting sellers is a strategic decision, and the relevant theoretical models are those of *game theory*, which is the analysis of interrelated strategic decisions. The word "interrelated" is critical here, since the amount that one firm may wish to raise price depends on how much others are expected to raise prices. We turn to a simple introduction of game theory in the next chapter. The discussion will be in terms of simple games with two players and two decisions, but the same principles will be applied later in the book to the analysis of more complex games with many sellers and many possible price levels.

Questions:

1. Suppose that all of the numbered Diamonds and Spades from a deck of cards (excluding Ace, King, Queen, and Jack) are used to set up a market. The Diamonds determine Demand, e.g. a 10 represents a buyer with a redemption value of \$10. Similarly, the Spades determine Supply. There are 9 buyers and 9 sellers, each with a single card.
 - a) Graph supply and demand and derive the competitive price and quantity predictions.
 - b) What is the predicted effect of removing the 3, 4, and 5 of Spades?
2. When people are asked to come up with alternative theories about why prices in a pit market stabilize at a particular level observed in class, they

sometimes suggest taking the average of all buyers' values and sellers' costs.

- a) Devise and graph two experimental designs (list all buyers' values and all sellers' costs) that have the same competitive price prediction but different averages of values and costs.
 - b) Devise and graph two experimental designs that have the same average of all values and costs, but which have different competitive price predictions.
3. Next consider the following price-determination theory, which was suggested in a recent experimental economics class: "Rank all buyer values from high to low and find the median (middle) value. Then rank all seller costs from low to high and find the median cost. The price should be the average of the median value and the median cost." By giving buyers a lot more units than sellers, it is possible to create a design where the median cost and the median value are both equal to the lowest of all values and costs. Use this idea to set up a supply and demand array where the competitive price prediction is considerably higher than the prediction based on medians. (If your design has a range of possible competitive prices, you may assume for purposes of discussion that the average price prediction is at the mid-point of this range.)
4. Large earnings asymmetries between buyers and sellers may slow convergence to theoretical predictions. This tendency will be especially pronounced if the traders on one side of the market are predicted to have *zero* earnings. If either of the experiment designs that you suggested in your answer to question 3 involve zero earnings for traders on one side of the market, how could you alter the design to ensure that all trades made by those on the "disadvantaged" side of the market in each treatment generate earnings of about 50 cents per unit?

Chapter 3. Some Simple Games: Competition, Coordination, and Guessing

A game with two players and two decisions can be represented by a 2x2 payoff table or “matrix.” Such games often highlight the conflict between incentives to compete or cooperate. This chapter introduces two such games, the prisoner’s dilemma and the coordination game. Discussion can be preceded with an experiment, either using playing cards (with instructions provided in the Appendix) or using the Veconlab web-based version (game MG). The main purpose of this chapter is to develop the idea of a Nash equilibrium, which is a somewhat abstract and mathematical notion. A simple Guessing Game, discussed last, serves to highlight the extent to which ordinary people might come to such an equilibrium.

I. Game Theory and the Prisoner’s Dilemma

The Great Depression, which was the defining economic event of the 20th century, caused a major rethinking of existing economic theories that represented the economy as a system of self-correcting markets needing little in the way of active economic policy interventions. On the macroeconomic side, John Maynard Keynes’ *The General Theory of Employment, Interest, and Money* focused on psychological elements (“animal spirits”) that could cause a whole economy to become mired in a low-employment equilibrium, with no tendency for self-correction. An equilibrium is a state where there are no net forces causing further change, and Keynes’ message implied that such a state may not necessarily be good. On the microeconomics side, Edward Chamberlin argued that markets may not yield efficient, competitive outcomes, as noted in the previous chapter. At about the same time, von Neumann and Morgenstern (1944) published *Theory of Games and Economic Behavior*, which was motivated by the observation that a major part of economic activity involves bilateral and small-group interactions, where the classical assumption of non-strategic, price-taking behavior (used in the previous chapter) is not realistic. In light of the protracted Depression, these new theories met with a quick acceptance. Chamberlin’s models of “monopolistic competition” were quickly incorporated into textbooks as alternative models, and von Neumann and Morgenstern’s book on game theory received front-page coverage in the *New York Times*.

A game is a mathematical model of a strategic situation in which players’ payoffs depend on their own and others’ decisions. A *game* is characterized by the players, their sets of feasible decisions, the information available at each

decision point, and the payoffs (as functions of all decisions and random events). A key notion is that of a *strategy*, which is essentially a complete plan of action that covers all contingencies. For example, a strategy in an auction could be an amount to bid for each possible estimate of the value of the prize. Since a strategy covers all contingencies, even those that are unlikely to be faced, it could be given to a hired employee to be played out on behalf of the player in the game. An equilibrium is a set of strategies that is stable in some sense, i.e. with no inherent tendency for change. Economists are interested in notions of equilibrium that will provide good predictions after behavior has had a chance to “settle down.”

As noted in Chapter 1, a Princeton graduate student, John Nash, corrected a major shortcoming in the von Neumann and Morgenstern analysis by developing a notion of equilibrium for non-zero-sum games. Nash showed that this equilibrium existed under general conditions, and this proof caught the attention of researchers at the RAND Corporation headquarters in Santa Monica. In fact, two RAND mathematicians immediately conducted a laboratory experiment designed to test Nash’s new theory. Nash’s thesis advisor was in the same building when he noticed the payoffs for the experiment written on a blackboard. He found the game interesting and made up a story of two prisoners facing a dilemma of whether or not to make a confession. This story was used in a presentation to the psychology department at Stanford, and the “prisoner’s dilemma” became the most commonly discussed paradigm in the new field of game theory.

In a prisoner’s dilemma, the two suspects are separated and offered a set of threats and rewards that make it best for each to confess and essentially “rat” on the other person, whether or not the other person confesses. For example, the prosecutor may say: “if you do not confess and the other person rats on you, then I can get a conviction and I will throw the book at you, so you are better off ratting on them.” When the prisoner asks what happens if the other does not confess, the prosecutor replies: “Even without a confession, I can frame you both on a lesser charge, which I will do if nobody confesses. If you do confess and the other person does not, then I will reward you with immunity and book the holdout on the lesser charge.” Thus each prisoner is better off implicating the other person even if the other remains silent. Both prisoners, aware of these incentives, decide to “rat” on the other, even though they would both be better off if they could somehow form a binding code of silence. This analysis suggests that two prisoners might be bullied into confessing to a crime that they did not commit, which is a scenario from *Murder at the Margin*, written by two mysterious economists under the pseudonym Marshall Jevons (1977).

A game with a prisoner’s dilemma structure is shown in the table below. The row player’s Bottom decision corresponds to confession, as does the column

player's Right decision. The Bottom/Right outcome, which yields payoffs of 3 for each, is worse than the Top/Left outcome, where both receive payoffs of 8. (A high payoff here corresponds to a light penalty.) The dilemma is that the "bad" confession outcome is an equilibrium in the following sense: if either person expects the other to talk, then their own best response to this belief is to confess as well. For example, consider the Row player, whose payoffs are listed on the left side of each outcome cell in the table. If Column is going to choose Right, then Row either gets 0 for playing Top or 3 for playing Bottom, so Bottom is the best response to Right. Similarly, column's Right decision is the best response to a belief that Row will play Bottom. There is no other cell in the table with this stability property. For example, if Row thinks column will play Left, then Row would want to play Bottom, so Top/Left is not stable. It is straightforward to show that the diagonal elements, Bottom/Left and Top/Right are also unstable.

Table 3.1 A Prisoner's Dilemma (Row's payoff, Column's payoff)

	Column Player:	
	Left	Right
Row Player:		
Top	8, 8	0, 10
Bottom	10, 0	3, 3

II. A Prisoner's Dilemma Experiment

The payoff numbers in the previous prisoner's dilemma experiment can be derived in the context of a simple example in which both players must choose between a low effort (0) and a higher effort (1). The cost of exerting the higher effort is \$10, and the benefits to each person depend on the total effort for the two individuals combined. Since each person can choose an effort of 0 or 1, the total effort must be 0, 1, or 2, and the benefit per person is given:

Total effort	0	1	2
Benefit per person	\$3	\$10	\$18

To see the connection between this "production" function and the prisoner's dilemma payoff matrix, let Bottom and Right correspond to zero effort. Thus the Bottom/Right outcome results in 0 total effort and payoffs of 3 for each person, as shown in the payoff matrix in Table 3.1. The Top/Left cell in the matrix is relevant when both choose efforts of 1 (at a cost of 8 each). The total effort is 2, so each earns $18 - 10 = 8$, as shown in the Top/Left cell of the matrix. The Top/Right and Bottom/Left parts of the payoff table pertain to the case where one person receives the benefit of 10 at no cost, the other receives the benefit of 10 at

a cost of 10, for a payoff of 0. Notice that each person has an incentive to “free ride” on the other’s effort, since one’s own effort is more costly (\$10) than the marginal benefit of effort, which is 7 ($= 10 - 3$) for the first unit of effort and 8 ($= 18 - 10$) for the second unit.

Cooper, DeJong, Forsythe, and Ross (1996) conducted an experiment using the prisoner’s dilemma payoffs in the above matrix, with the only change being that the payoffs were (3.5, 3.5) at the Nash equilibrium. Their matching protocol prevented individuals from being matched with the same person twice, or from being matched with anyone who had been matched with them or one of their prior partners. This “no-contageon” protocol is analogous to going down a receiving line at a wedding and telling everyone the same bit of gossip about the bride and groom. Even if this story is so curious that everyone you meet in line repeats it to everyone they meet, you will never encounter anyone who has heard the story, since all of the people who have heard the story will be behind you in the line. In an experiment, the elimination of any type of repeated matching, either direct or indirect, means that nobody is able to send a “message” or to punish or reward others for cooperating. Even so, cooperative decisions (Top or Left) were fairly common in early rounds (43%), and the incidence of cooperation declined to about 20% in rounds 15-20. A recent classroom experiment using the Veconlab program with the Table 3.1 payoffs and random matching produced cooperation rates of about 33 %, with no downward trend. A second experiment with a different class produced cooperation rates starting at about 33% and declining to zero by the fourth period. This latter group behaved quite differently when they were matched with the same partner for 5 periods; cooperation rates stayed level in the 33-50% range until the final period. The “end-game effect” is not surprising since cooperation in earlier periods may consist of an effort to signal good intentions and stimulate reciprocity. These forward-looking strategies are not available in the final round.

The results of these prisoner’s dilemma experiments are representative. There is typically a mixture of cooperation and defection, with the mix being somewhat sensitive to payoffs and procedural factors, with some variation across groups. In general, cooperation rates are higher when individuals are matched with the same person in a series of repeated rounds. In fact, the first prisoner’s dilemma experiment run at the RAND Corporation over fifty years ago lasted for 100 periods, and cooperative phases were interpreted as evidence against the Nash equilibrium.

A strict game-theoretic analysis would have people realizing that there is no reason to cooperate in the final round, and knowing this, nobody would try to cooperate in the next-to-last round with the hope of stimulating final-round cooperation. Thus, there should be no cooperation in the next-to-last round, and hence there is no reason to try to stimulate such cooperation. Reasoning

“backwards” from the end in this manner, one might expect no cooperation even in the very first round, at least when the total number of rounds is finite and known. Nash responded to the RAND mathematicians with a letter maintaining that it is unreasonable to expect people to engage in this many levels of iterative reasoning in an experiment with many rounds. There is an extensive literature on related topics, e.g., the effects of punishments, rewards, adaptive behavior, and various “tit-for-tat” strategies in repeated prisoner’s dilemma games.

Finally, it should be noted that there are many ways to present the payoffs for an experiment, even one as simple as the prisoner’s dilemma. Both the Cooper et al. and the Veconlab experiment used a payoff matrix presentation. Alternatively, the instructions could be presented in terms of the cost of effort and the table showing the benefit per person for each possible level of total effort. This presentation is perhaps less neutral, but the economic context makes it less abstract and artificial than a matrix presentation. The Appendix for this chapter takes a different approach by setting up the prisoner’s dilemma game as one where each person chooses which of two cards to play (Capra and Holt, 2000). For example, suppose that each person has playing cards numbered 8 and 6. Playing the 6 “pulls” \$6 (from the experimenter’s reserve) into one’s own earnings, and playing the 8 “pushes” \$8 (from the experimenter’s reserve) to the other person’s earnings. If they both pull the \$6, earnings are \$6 each. Both would be better off if they played 8, yielding \$8 for each. Pulling \$6, however, is better from a selfish perspective, regardless of what the other person does. A consideration of the resulting payoff matrix (Question 2) indicates that this is clearly a prisoner’s dilemma. The card presentation is not neutral enough for a research experiment, but it is quick and easy to implement in class, where students hold the cards played against their chests, and the instructor picks people in pairs to reveal their cards. A Nash equilibrium survives an “announcement test” in that neither would wish they had played a different card given the card played by the other. In this case, the Nash equilibrium is for each to play 6.

III. A Coordination Game

Many production processes have the property that one person’s effort increases the productivity of another’s effort. For example, a mail order company must take orders by phone, produce the goods, and ship them. Each sale requires all three services, so if one process is slow, the effort of those in other activities is to some extent wasted. In terms of our two-person production function, recall that 1 unit of effort produced a benefit of \$10 per person and 2 units produced a benefit of \$18. Suppose that the second unit of effort makes the first one more productive in the sense that the benefit is more than doubled when a second unit of effort is added. In the table that follows, per-person benefit is \$30 for 2 units of effort:

Total effort	0	1	2
Benefit per person	\$3	\$10	\$30

If the cost of effort remains at \$10 per unit, then each person earns $\$30 - \$10 = \$20$ in the case where both supply a unit of effort, as shown in the revised payoff matrix below:

Table 3.2 A Coordination Game
(Row's payoff, Column's payoff)

		Column Player:	
		Left	Right
Row Player:	Top	↑ 20, 20 ←	0, 10
	Bottom	10, 0	↓ 3, ⇒ 3

As before, the zero effort outcome (Bottom/Right) is a Nash equilibrium, since neither person would want to change from their decision unilaterally. For example, if Row knows that Column will choose Right, then Row can get \$0 from Top and \$3 from Bottom, so Row would want to choose Bottom, as indicated by the downward pointing arrow in the lower-right box. Similarly, Right is a best response to Row's decision to use Bottom, as indicated by the right arrow in the lower-right box.

There is a second Nash equilibrium in the Top/Left box of this modified game. To verify this, think of a situation in which we “start” in that box, and consider whether either person would want to deviate, so there are two things to check. Step 1: if Row is thought to choose Top, then Column would want to choose Left, as indicated by the left arrow in the Top/Left box. Step 2: if Row expects Column to choose Left, then Row would want to choose Top, as indicated by the up arrow. Thus Top/Left survives an “announcement test” and would be stable. This second equilibrium yields payoffs of 20 for each, far better than the payoffs of 3 each in the Bottom/Right equilibrium outcome. This is called a “coordination game,” since the presence of multiple equilibria raises the issue of how players will guess which one will be relevant. (In fact, there is a third Nash equilibrium in which each player chooses randomly, but we will not discuss such strategies until Chapter 5.)

It used to be common for economists to *assume* that rational players would somehow coordinate on the best equilibrium, if all could agree which is the best equilibrium. While it is apparent that the Top/Left outcome is likely to be the most frequent outcome, one example does not justify a general assumption that players will always coordinate on an equilibrium that is better for all. This

assumption can be tested with laboratory experiments, and it has been shown to be false. In fact, players sometimes get driven to an equilibrium that is *worst* for all concerned (Van Huyck, Battalio, and Beil; 1990). Examples of data from such coordination games will be provided in a subsequent chapter on coordination problems. For now, consider the intuitive effect of increasing the cost of effort from \$10 to \$19. With this increase in the cost of effort, the payoffs in the high-effort Top/Left outcome are reduced to $\$30 - \$19 = \$11$. Moreover, the person who exerts effort alone only receives a \$10 benefit and incurs a \$19 cost, so the payoffs for this person are $-\$9$, as shown:

Table 3.3 A Coordination Game with High Effort Cost
(Row's payoff, Column's payoff)

		Column Player:	
		Left	Right
Row Player:	Top	↑ 11, 11 ←	-9, 10
	Bottom	10, -9	↓ 3, ⇒ 3

Notice that Top/Left is still a Nash equilibrium: if Column is expected to play Left, then Row's best response is Top and vice versa, as indicated by the directional arrows in the upper-left corner. But Top and Left are very risky decisions, with possible payoffs of \$11 and $-\$9$, as compared with the \$10 and \$3 payoffs associated with the Bottom and Right strategies that yield the other Nash equilibrium. While the absolute payoff is higher for each person in the Top/Left outcome, each person has to be very sure that the other one will not deviate. In fact, this is a game where behavior is likely to converge to the Nash equilibrium that is worst for all. These payoffs were used in a Veconlab experiment that lasted 5 periods with random matching of 12 participants. Participants were individuals or pairs of students at the same computer. By the fifth period, only a quarter of the decisions were Top or Left. The same people then played the less risky coordination game shown in Table 3.2, again with random matching. The results were quite different; all but one of the pairs ended up in the good (Top, Left) equilibrium outcome in all periods.

To summarize, it is not appropriate to assume that behavior will somehow converge to the best outcome for all, even when this is a Nash equilibrium as in the coordination games considered here. With multiple equilibria, the one with the most "attraction" may depend on payoff features (e.g., the effort cost) that determine which equilibrium is riskier. These issues will be revisited in the chapter on coordination.

IV. A Guessing Game

A key aspect of most games is the need for players to guess what others will do in order to determine their own best decision. This aspect is somewhat obscured in a 2x2 game since a wide range of beliefs may lead to the same decision, and therefore, it is not possible to say much about the beliefs that stand behind any particular observed decision. In this section, we turn to a game with a continuum of decisions, from 0 to 100, and any number of players, N . Each player selects a number in this interval, with the advance knowledge that the person whose decision is closest to $1/2$ of the average of all N decisions will win a money prize. (The prize is divided equally in the event of a tie.) Thus each person's task is to make a guess about the average, and then submit a decision that is half of that amount.

If the game is only played once, then there is no way to learn from past decisions, and each person must “learn” by thinking *introspectively* about what others might do. For example, someone might reason that since the decisions must be between 0 and 100, without further knowledge it might be ok to guess that the average would be at the midpoint, 50. This person would submit a decision of 25 in order to be at half of the anticipated average. But then the person might realize that others could be thinking in the same way, and also submitting decisions of 25, so that a decision of 12.5 would be closest to the target level of $1/2$ of the average. A third level of iteration along these lines suggests a decision of 6.25, and even higher levels of iterated thinking will lead the person to choose a decision that is closer and closer to 0. So the best decision to make depends on one's beliefs about the extent to which others are thinking iteratively in this manner, i.e. on beliefs about the others' rationality. Since the best response to any common decision is to go down to one half of that level, it follows that no common decision can be a Nash equilibrium. Clearly a common decision of 0 is a Nash equilibrium, since each person gets a positive share ($1/N$) of the prize, and to deviate and choose a higher decision would result in a zero payoff. It turns out that this is the only Nash equilibrium for this game (problem 5).

Figure 3.1 shows data for a classroom guessing game run with the Veconlab software. There were 5 participants, who made first round decisions of 55, 29, 22, 21, and 6.25, which average to about 27. Notice that the lowest person is one who engaged in three levels of iterated reasoning about what the others might do, and this person was closest to the target level of about $17/2 = 13.5$. This person submitted the same decision in round 2 and won again; the other decisions were 20, 15, 11.5, and 10. The other people do not seem to be engaging in any precise iterated reasoning, but it is clear the person with the lowest decision is being rewarded, which will tend to pull decisions down in subsequent rounds.

These decisions converge to near-Nash levels by the fifth round, although nobody selected 0 exactly.




Figure 3.1 Data from a Classroom Guessing Game

The second treatment changes the calculation of the target from $1/2$ of the average to $20 + 1/2$ of the average. The instructor (a.k.a. author) who ran this in class had selected the parameters quickly, with the conjecture that adding 20 to the target would raise the Nash equilibrium from 0 to 20. When the average came in at 28, the instructor thought that it would then decline toward 20, just as it had declined toward 0 in the first treatment. But the average rose to 33 in the next round, and continued to rise to a level that is near 40. Notice that if all choose 40, then the average is 40, half of the average is 20, and $20 + \text{half of the average}$ is $20 + 20 = 40$, which is the Nash equilibrium for this treatment (question 6).

This game illustrates a couple of things. Some people think iteratively about what others will do, but this is not uniform, and the effect of such introspection is not enough to move decisions to near-Nash levels in a single round of play. When people are able to learn from experience, behavior does converge to the Nash equilibrium in this game, even in some cases when the equilibrium is not apparent to the person who designed the experiment.

V. Extensions

The prisoner's dilemma game discussed in this chapter can be given an alternative interpretation, i.e. that of a "public good." This is because each person's benefit depends on the total effort, including that of the other person. In effect, neither person can appropriate more than half of the benefit of their own effort, and in this sense the benefit is publicly available to both, just as national defense or police protection are freely available to all. We will consider the public goods provision problem in more detail in a later chapter in the Public Choice section of this book. The setup allows a large number of effort levels, and the effects of changes in costs and other incentives will be considered.

The prisoner's dilemma has a Nash equilibrium that is worse than the outcome that results when individuals cooperate and ignore their private incentives to defect. The dilemma is that the equilibrium is the bad outcome, and the good outcome is not an equilibrium. In contrast, the coordination game considered above has a good outcome that is also an equilibrium. Even so, there is no guarantee that the better equilibrium will be realized. A more formal analysis of the "risk" associated with alternative equilibria is presented in a subsequent chapter on coordination. The main point is that there may be multiple Nash equilibria, and which equilibrium is observed may depend on intuitive factors like the cost of effort. In particular, a high effort cost may cause players to "get stuck" in a Nash equilibrium that is worse for all concerned, as compared with an alternative Nash equilibrium. The coordination game provides a paradigm in which an equilibrium may not be desirable. Such a possibility has concerned macroeconomists like Keynes, who worried that a market economy may become mired in a low-employment, low-effort state.

The guessing game is examined in Nagel (1995, 1997), who shows that there is a widespread failure of behavior to converge to the Nash equilibrium in a one-round experiment, regardless of subject pool. Nagel discusses the degree to which some of the people go through iterations of thinking that may lead them to choose lower and lower decisions (e.g., 50, 25, 12.5, etc.). The game was originally motivated by one of Keynes' remarks, that investors are typically in a position of trying to guess what stock others will find attractive (first iteration), or what stock that others think other investors will find attractive (second iteration), etc.

In all of the equilibria considered thusfar, each person chooses a decision without any randomness. It is easy to think of games where it is not good to be perfectly predictable. For example, a soccer player in a penalty kick situation should sometimes kick to the left and sometimes to the right, and the goalie should also avoid a statistical preference for diving to one side or the other. Such behavior in a game is called a "randomized strategy." Before considering randomized strategies (in Chapter 5), it is useful to introduce the notions of

probability, expected value, and other aspects of decision making in uncertain situations. This is the topic of the next chapter, where we consider the choice between simple lotteries over cash prizes.

For a brief, non-technical discussion of John Nash's original 1948 *Proceedings of the National Academy of Sciences* paper, its content and subsequent impact and policy applications, see Holt and Roth (2004).

Questions:

1. Suppose that the cost of a unit of effort is raised from \$10 to \$25 for the example based on the effort-value table shown below. Is the resulting game a coordination game or a prisoner's dilemma? Explain, and find the Nash equilibrium (or equilibria) for the new game.

Total effort	0	1	2
Benefit per person	\$3	\$10	\$30

2. What is the payoff table for the prisoner's dilemma game described in section II of the text based on "pulling" \$6 or "pushing" \$8?
3. Each player is given a 6 of Hearts and an 8 of Spades. The two players select one of their cards to play. If the suit matches, then they each are paid \$6 for the case of matching Hearts and \$8 for the case of matching Spades. Earnings are zero in the event of a mismatch. Is this a coordination game or a prisoner's dilemma? Show the payoff table and find all Nash equilibria (that do not involve random play).
4. Suppose that the payoffs for the original prisoner's dilemma are altered as follows. The cost of effort remains at \$10, but the per-person benefits are given in the table below. Recalculate the payoff matrix, find all Nash equilibria (that do not involve random play), and explain whether the game is a prisoner's dilemma or a coordination game.

Total effort	0	1	2
Benefit per person	3	5	18

5. Suppose that 2 people are playing a guessing game with a prize going to the person closest to $1/2$ of the average. Guesses are required to be between 0 and 100. Show that none of the following are Nash equilibria:
 - a) Both choose 1. (Hint, consider a deviation to 0 by one person, so that the average is $1/2$, and half of the average is $1/4$.)
 - b) One person chooses 1 and the other chooses 0.
 - c) One chooses x and the other chooses y ,

where $x > y > 0$. (Hint: If they choose these decisions, what is the average, what is the target, and is the target less than the midpoint of the range of guesses between x and y ? Then use these calculations to figure out which person would win, and whether the other person would have an incentive to deviate.)

6. Consider a guessing game with N people, and the person closest to 20 plus $1/2$ of the average is awarded the prize, which is split in the event of a tie. The range of guesses is from 0 to 100. Show that 40 is a Nash equilibrium. (Hint: Suppose that $N-1$ people choose 40 and one person deviates to $x < 40$. Calculate the target as a function of the deviant's decision, x . The deviant will lose if the target is greater than the midpoint between x and 40, i.e. if the target is greater than $(40 + x)/2$. Check to see if this is the case.)
7. Consider a guessing game with N people, and the person closest to 10 plus $2/3$ of the average is awarded the prize. Find the Nash equilibrium?
8. A more general formulation of the guessing game might be for the target level to be: $A + B \cdot (\text{average decision})$, where $A > 0$ and $0 < B < 1$, and decisions must be between 0 and an upper limit $L > 0$. Find the Nash equilibrium. When is it equal to L ?

Chapter 4. Risk and Decision Making

In this chapter, we consider decisions in risky situations. Each decision has a set of money consequences or “prizes” and the associated probabilities. In such cases, it is straightforward to calculate the expected money value of each decision. A person who is neutral to risk will select the decision with the highest expected payoff. A risk-averse person is willing to accept a lower expected payoff in order to reduce risk. A simple lottery choice experiment is used to illustrate the concepts of expected value maximization and risk aversion. Variations on the experiment can be conducted prior to class discussions, either with the instructions included in the appendix or with the Veconlab software (select experiment LC).

I. *Who Wants to Be a Millionaire?*

Many decision situations involve consequences that cannot be predicted perfectly in advance. For example, suppose that you are a contestant on the television game show *Who Wants to Be a Millionaire?* You are at the \$500,000 point and the question is on a topic that you know nothing about. Fortunately, you have saved the “fifty-fifty” option that rules out two answers, leaving two that turn out to be unfamiliar. At this point, you figure you only have a one-half chance of guessing correctly, which would make you into a millionaire. If you guess incorrectly, you receive the safety level of \$32,000. Or you can fold and take the sure \$500,000. In thinking about whether to take the \$500,000 and fold, you decide to calculate the expected money value of the guess option. With probability 0.5 you earn \$32,000, and with probability 0.5 you earn \$1,000,000, so the expected value of a guess is:

$$0.5(\$32,000) + 0.5(\$1,000,000) = \$16,000 + \$500,000 = \$516,000.$$

This expected payoff is greater than the sure \$500,000 one gets from stopping, but the trouble with guessing is that you either get 32K or one million dollars, not the average. So the issue is whether the \$16,000 increase in the average payoff is worth the risk, which is considerable if you guess. If you love risk, then there is no problem; take the guess. If you are neutral to risk, the extra 16K in the average payoff should cause you to take the risk. Alternatively, you might reason that the 32K would be gone in 6 months, and only a prize of 500K or greater would be large enough to bring about a change in your lifestyle (new SUV, tropical vacation, etc.), which may lead you to fold. To an economist, the person who folds would be classified as being *risk averse* in this case, because the risk for this

person is sufficiently bad that it is not worth the extra 16K in average payoff associated with the guess.

Now consider an even more extreme case. You have the chance to secure a sure 1 million dollars. The alternative is to take a coin flip, which provides a prize of 3 million dollars for Heads, and nothing for Tails. If you believe that the coin is fair, then you are choosing between two lotteries:

<i>Safe Lottery:</i>	1 million for sure;
<i>Risky Lottery:</i>	3 million with probability 0.5, 0 with probability 0.5.

As before, the expected money value for the risky lottery can be calculated by multiplying the probabilities with the associated payoffs, yielding an expected payoff of 1.5 million dollars. When asked about this (hypothetical) choice, most students will select the safe million. Notice that for these people, the extra \$500,000 in expected payoff is not worth the risk of ending up with nothing. An economist would call this risk aversion, but to a layman the reason is intuitive: the first million provides a major change in lifestyle. A second million is an equally large money amount, but it is hard for most of us to imagine how much *additional* benefit this would provide. Roughly speaking, the additional (marginal) utility of the second million dollars is much lower than the utility of the first million, and the marginal utility of the third million is likely to be even lower. A utility function with a diminishing marginal utility is one with a curved “uphill” shape, as shown in Figure 4.1.

This utility function is the familiar square root function, where utility (in millions) is the square root of the payoff (in millions). Thus the utility of 0 is 0, the utility of 1 million dollars is 1, the utility of 4 million dollars is 2, the utility of 16 million is 4, etc. In fact, each time you multiply the payoff by 4, the utility only increases by a factor of 2, which is indicative of the diminishing marginal utility. This feature is also apparent from the fact that the slope of the utility function near the origin is high, and it diminishes as we move to the right. The more curved the utility function, the faster the utility of an additional million diminishes.




Figure 4.1. A "Square Root" Utility Function

The diminishing-marginal-utility hypothesis was first suggested by Daniel Bernoulli (1738). There are many functions with this property. One is the class of power functions: $U(x) = x^{1-r}$, where x is money income and r is a measure of risk aversion that is less than 1. Notice that when $r = 0$, we have the linear function: $U(x) = x^1 = x$, which has no curvature, and hence no diminishing marginal utility for income. For a person whose choices are represented by this function, the second million is just as good as the first, etc., so the person is neutral to risk. Alternatively, if the risk aversion measure r is increased from 0 to 0.5, we have $U(x) = x^{0.5}$, which is the square root function in the figure. Further increases in r result in more curvature, and in this sense r is a measure of the extent to which marginal utility of additional money income decreases. This measure is often called the coefficient of *relative risk aversion*.

In all of the examples considered thusfar in this chapter, you have been asked to think about what you would do if you had to choose between gambles involving millions of dollars. The payoffs were (unfortunately) hypothetical. This raises the issue of whether what you say is what you would actually do if you faced the real situation, e.g. if you were really in the final stage of the series of questions on *Who Wants to be a Millionaire*? It would be fortunate if we did not really have to pay large sums of money to find out how people would behave in high-payoff situations. The possibility of a *hypothetical bias*, i.e. the proposition that behavior might be dramatically different when high hypothetical stakes become real, is echoed in the film *Indecent Proposal*:

John (a.k.a. Robert): Suppose I were to offer you one million dollars for one night with your wife.

David: I'd assume you were kidding.

John: Let's Pretend I'm not. What would you say?

Diana (a.k.a. Demi): He'd tell you to go to hell.

John: I didn't hear him.

David: I'd tell you to go to hell.

John: That's just a reflex answer because you view it as hypothetical. But let's say there were real money behind it. I'm not kidding. A million dollars. Now, the night would come and go, but the money could last a lifetime. Think of it – a million dollars. A lifetime of security for one night. And don't answer right away. But consider it – seriously.

In the film, John's proposal was ultimately accepted, which is the Hollywood answer to the incentives question. On a more scientific note, incentive effects are an issue that can be investigated with experimental techniques. Before returning to this issue, it is useful to consider the results of a lottery choice experiment designed to evaluate risk attitudes, which is the topic of the next section.

II. A Simple Lottery-Choice Experiment

In all of the cases discussed above, the safe lottery is a sure amount of money. In this section we consider a case where both lotteries have random outcomes, but one is riskier than the other. In particular, suppose that Option A pays either \$40 or \$32, each with probability one half, and that Option B pays either \$77 or \$2, each with probability one half. First, we calculate the expected values:

$$\text{Option A: } 0.5(\$40) + 0.5(\$32) = \$20 + \$16 = \$36.00$$

$$\text{Option B: } 0.5(\$77) + 0.5(\$2) = \$38.50 + \$1 = \$39.50.$$

In this case, option B has a higher expected value, by \$3.50, but a lot more risk since the payoff spread from \$77 to \$2 is almost ten times as large as the spread from \$32 to \$40.

This choice was on a menu of choices used in an experiment (Holt and Laury, 2001). There were about 200 participants from several universities, including undergraduates, about 80 MBA students and about 20 business school faculty. Even though Option B had a higher expected payoff when each prize is equally likely, eighty-four percent of the participants selected the safe option, which indicates some risk aversion.

Table 4.1. A Menu of Lottery Choices Used to Evaluate Risk Preferences

	Option A	Option B	Your Choice A or B
Decision 1	\$40.00 if throw of die is 1 \$32.00 if throw of die is 2-10	\$77.00 if throw of die is 1 \$2.00 if throw of die is 2-10	_____
....			
Decision 4	\$40.00 if throw of die is 1-4 \$32.00 if throw of die is 5-10	\$77.00 if throw of die is 1-4 \$2.00 if throw of die is 5-10	_____
Decision 5	\$40.00 if throw of die is 1-5 \$32.00 if throw of die is 6-10	\$77.00 if throw of die is 1-5 \$2.00 if throw of die is 6-10	_____
Decision 6	\$40.00 if throw of die is 1-6 \$32.00 if throw of die is 7-10	\$77.00 if throw of die is 1-6 \$2.00 if throw of die is 7-10	_____
....			
Decision 10	\$40.00 if throw of die is 1-10	\$77.00 if throw of die is 1-10	_____

The payoff probabilities in the Laury and Holt experiment were determined by the throw of a ten-sided die. This allowed the researchers to alter the probability of the high payoff in one-tenth increments. A part of the menu of choices is shown in Table 4.1, where the probability of the high payoff (\$40 or \$77) is one tenth in Decision 1, four tenths in Decision 4, etc. Notice that Decision 10 is a kind of rationality check, where the probability of the high payoff is 1, so it is a choice between \$40 for sure and \$77 for sure. The subjects indicated a preference for all ten decisions, and then one decision was selected at random, *ex post*, to determine earnings. In particular, after all decisions were made, we threw a 10-sided die to determine the relevant decision, and then we threw the ten-sided die again to determine the subject's earnings for the selected decision. This procedure has the advantage of providing data on all ten decisions without any "wealth effects." Such wealth effects could come into play, for example, if a person wins \$77 on one decision and this makes them more willing to take a risk on the subsequent decision. The disadvantage is that incentives are diluted, which was compensated for by raising payoffs.

The expected payoffs associated with each possible choice are shown under "Risk Neutrality" in the second and third columns of the Table 4.2. First, look in the fifth row where the probabilities are 0.5. This is the choice previously discussed in this chapter, with expected values of \$36 and \$39.50 as calculated above. The other expected values are calculated in the same manner, by multiplying probabilities by the associated payoffs, and adding up these products.

Table 4.2. Optimal Decisions for Risk Neutrality and Risk Aversion

Probability of High Payoff	Risk Neutrality		Risk Aversion ($r = 0.5$)	
	Expected Payoffs for $U(x) = x$		Expected Utilities for $U(x) = x^{1/2}$	
	Safe \$40 or \$32	Risky \$77 or \$2	Safe \$40 or \$32	Risky \$77 or \$2
0.1	\$32.80	\$9.50	5.72	2.15
0.2	\$33.60	\$17.00	5.79	2.89
0.3	\$34.40	\$24.50	5.86	3.62
0.4	\$35.20	\$32.00	5.92	4.36
0.5	\$36.00	\$39.50	5.99	5.09
0.6	\$36.80	\$47.00	6.06	5.83
0.7	\$37.60	\$54.50	6.12	6.57
0.8	\$38.40	\$62.00	6.19	7.30
0.9	\$39.20	\$69.50	6.26	8.04
1.0	\$40.00	\$77.00	6.32	8.77

The best decision when the probability of the high payoff is only 0.1 (in the top row of Table 4.2) is obvious, since the safe decision also has a higher expected value, or \$32.80 (versus \$9.50 for the risky lottery). In fact, 98 percent of the subjects selected the safe lottery (Option A) in this choice. Similarly, the last choice is between sure amounts of money, and all subjects chose Option B in this case. The expected payoffs are higher for the top four choices, as shown by the bolded numbers. Thus a risk-neutral person, who by definition only cares about expected values regardless of risk, would make 4 safe choices in this menu. In fact, the average number of safe choices was 6, not 4, which indicates some risk aversion. As can be seen from the 0.6 row of the table, the typical safe choice for this option involves giving up about \$10 in expected value in order to reduce the risk. About two-thirds of the people made the safe choice for this decision, and 40 percent chose safe in Decision 7.

The percentages of safe choices are shown by the thick solid line in Figure 4.2, where the decision number is listed on the horizontal axis. The behavior predicted for a risk-neutral person is represented by the dashed line, which stays at the top (100 percent safe choices) for the first four decisions, and then shifts to the bottom (no safe choices) for the last six decisions. The thick line representing actual behavior is generally above this dashed line, which indicated the tendency to make more safe choices. There is some randomness in actual choices, however, so that choice percentages do not quite get to 100 percent on the left side of the figure.




Figure 4.2. Percentages of Safe Choices with Real Incentives and Hypothetical Incentives
(Source: Holt and Laury (2001)).

This real-choice experiment was preceded by a *hypothetical* choice task, in which subjects made the same ten choices with the understanding that they would not be paid their earnings for that part. The data averages for the hypothetical choices are plotted as points on the thin line in the figure. Several differences between the real and hypothetical choice data are clear. The thin line lies below the thick line, indicating less risk aversion when choices have no real impact. The average number of safe choices was about 6 with real payments, about 5 with hypothetical payments. Even without payments, subjects were a little risk averse as compared with the risk-neutral prediction of 4 safe choices, but they could not imagine how they really would behave when they had to face real consequences. Second, with hypothetical incentives, there may be a tendency for people to think less carefully, which may produce “noise” in the data. In particular, for decision 10, the thin line is a little higher, corresponding to the fact that 2 percent of the people chose the sure \$40 over the sure (but hypothetical) \$77.

The issue of whether or not to pay subjects is one of the key issues that divides research in experimental economics from some work on similar issues by psychologists (see Hertwig and Ortmann, 2001, for a provocative survey on practices in psychology, with about 30 comments and an authors’ reply). One justification for using high hypothetical payoffs is realism. Two prominent psychologists, Kahneman and Tversky, justify the use of hypothetical incentives:

Experimental studies typically involve contrived gambles for small stakes, and a large number of repetitions of very similar problems. These features of laboratory gambling complicate the interpretation of the results and restrict their generality. By default, the method of hypothetical choices emerges as the simplest procedure by which a large number of theoretical questions can be investigated. The use of the method relies on the assumption that people often know how they would behave in actual situations of choice, and on the further assumption that the subjects have no special reason to disguise their true preferences. (Kahneman and Tversky, 1979, p. 265)

Of course, there are many documented cases where hypothetical and real-incentive choices coincide quite closely (one such example will be presented in the next section). But in the absence of a widely accepted theory about when they do and do not coincide, it is dangerous to assume that real incentives are not needed.

III. Risk Aversion: Incentive Effects, Demographics, and Order Effects

The intuitive effect of risk aversion is to diminish the utility associated with higher earnings levels, as can be seen from the curvature for the “square root” utility function in Figure 4.1. With nonlinear utility, the calculation of a person’s expected utility is analogous to the calculation of expected money value. For example, the safe option A for Decision 5 is a 1/2 of \$40 and a 1/2 chance of \$32. The expected payoff is found by adding up the products of money prize amounts and the associated probabilities:

$$\text{Expected payoff (safe option)} = 0.5(\$40) + 0.5(\$32) = \$20 + \$16 = 36.$$

The expected *utility* of this option is obtained by replacing the money amounts, \$40 and \$32, with the utilities of these amounts, which we will denote by $U(40)$ and $U(32)$. If the utility function is the square-root function, then $U(40) = (40)^{1/2} = 6.32$ and $U(32) = (32)^{1/2} = 5.66$. Since each prize is equally likely, we take the average of the two utilities:

$$\begin{aligned}\text{Expected utility (safe option)} &= 0.5 U(40) + 0.5 U(32) \\ &= 0.5(6.32) + 0.5(5.66) = 5.99.\end{aligned}$$

This is the expected utility for the safe option when the probabilities are 0.5, as shown in the 0.5 row of Table 4.2 in the column under “Risk Aversion” for the Safe Option. Similarly, it can be shown that the expected utility for the risky option, with payoffs of \$77 and \$2, is 5.09. The other expected utilities for all ten

decisions are shown in the two right-hand columns of the Table 4.2. The safe option has the higher expected utility for the top six decisions, so the theoretical prediction for someone with this utility function is that they choose six safe options before crossing over to the risky option. Recall that the analogous prediction for risk neutrality (see the left side of the table) is 4 safe choices, and that the data for this treatment exhibit 6 safe choices on average. In this sense, the square-root utility function provides a better fit to the data than the linear utility function that corresponds to risk neutrality.

One interesting question is whether the order in which a choice is made matters. In particular, the Holt and Laury (2002) experiment involved subjects who first went through a lottery choice trainer involving a choice between a safe \$3 and a menu of gambles with payoffs of \$1 or \$6, to familiarize them with the dice-throwing procedure for picking one decision at random and then for determining the payoff for that decision. Then all participants made 10 decisions for a low real-payoff choice menu, where all money amounts were 1/20 of the level shown in Table 4.1. Thus the payoffs for the risky option are \$3.85 and \$0.10, and the possible payoffs for the safer option are \$2.00 and \$1.60. These low payoffs will be referred to as the “1x” treatment, and the payoffs in Table 4.1 will be called the 20x payoffs. Other treatments are designated similarly as multiples of the low payoff level. The initial low-payoff choice was followed by a choice menu with high *hypothetical* payoffs (20x, 40x, or 90x), followed by the same menu with high *real* payoffs (20x, 40x, or 90x), and ending with a second 1x real choice. The average numbers of safe choices, shown in the top two rows of Table 4.3, indicate that the number of safe choices increases steadily as real payoffs are scaled up, but this incentive effect is not observed as hypothetical payoffs are scaled up.


Table 4.3. Average Numbers of Safe Choices: Order and Incentive Effects
Key: Superscripts indicate Order (^a = 1st, ^b = 2nd, ^c = 3rd, ^d = 4th)

Experiment	Incentives	1x	10x	20x	50x	90x
Holt and Laury (2002) U.C.F., Ga. St., U. Miami 208 subjects	Real	5.2 ^a		6.0 ^c	6.8 ^c	7.2 ^c
	Hypothetical	5.3 ^d		4.9 ^b	5.1 ^b	5.3 ^b
Harrison et al. (2003) South Carolina 178 subjects	Real	5.3 ^a	6.0 ^a 6.4 ^b			
	Real	5.7 ^a		6.7 ^a		
Holt and Laury (2004) U. of Virginia 168 subjects	Hypothetical	5.6 ^a		5.7 ^a		

Harrison et al. (2003) point out this design mingles order and payoff scale effects, producing a possible fatal flaw (loss of control) as discussed in Chapter 1. The letter superscripts in the table indicate the order in which the decision was made (*a* first, *b* second, etc.), so a comparison of the 1x decisions with those of higher scales is called into question by the change in order, as are the comparisons of high hypothetical and high real payoffs (done in orders 2 and 3 respectively). The presence of order effects is supported by evidence summarized in the third row of the table, where risk aversion seems to be higher when the 10x scale treatment follows the 1x treatment than when the 10x treatment is done first. The order effect produces a difference of 0.4 safe choices. Such order effects call into question the real versus hypothetical comparisons, although it is still possible to make inferences about the scale effects from 20x to 40x to 90x in the Holt and Laury experiment, since these were made in the same order.

In response to these issues of interpretation, Holt and Laury (2004) ran a followup experiment with a 2x2 design: real or hypothetical payoffs and 1x or 20x payoffs. Each participant began with a lottery choice trainer (again with choices between \$3 and gambles involving \$1 or \$6). Then each person made a single menu of choices from one of the 4 treatment cells, so all decisions are made in the *same* order. Again, the scaling up of real payoffs causes a sharp increase in the average number of safe choices (from 5.7 to 6.7), whereas this incentive effect is not observed with a scaling up of hypothetical choices. Finally, the data from the lottery-choice trainers can be used to check whether observed differences in the 4 treatments are somehow due to unanticipated demographic differences or other unobserved factors that may cause people in one cell to be intrinsically more risk averse than those in other cells. A check of data from the trainers shows virtually no differences across treatment cells, and if anything, the trainer data indicate slightly lower risk aversion for the group of people who were subsequently given the high-real payoff menu.

The main conclusions from the second study are that participants are risk averse, the risk aversion increases with higher real payoffs, and that looking at high hypothetical payoffs may be very misleading. These conclusions are apparent from the graph of the distributions of the number of safe choices for each of the 10 decisions, as shown in Figure 4.3. In this context, it does not seem to matter much whether or not one used real money when the scale of payoffs is low, but observing no difference and then inferring that money payoffs do not matter in other contexts would be incorrect. Finally, note that payoff-scale effects may not be apparent in classroom experiments where payoffs are typically small or hypothetical.



**Figure 4.3. Results for New Experiment with Single Decision with No Order Effects:
Percentages of Safe Choices for High (20x) and Low (1x) Payoffs,
Under Real and Hypothetical Incentives (36 Subjects per Treatment)
Source: Holt and Laury (2004)**

Another interesting question is whether there are systematic demographic effects. The subjects in the original Holt and Laury (2002) experiment included about 60 MBA students, about thirty business school faculty (and a Dean), and over a hundred undergraduates from three universities (University of Miami, University of Central Florida, and Georgia State University). There is a slight tendency for people with higher incomes to be less risk averse. The women were more risk averse than men in the low-payoff (1x) condition, as observed in some previous studies, but this effect disappeared for higher payoff scales. All of that bravado went away with high stakes (e.g. 20x scale). There was no white/non-white difference, but Hispanics in the sample were a little less risk averse. It is not appropriate to make broad inferences about demographic effects from a single study. For example, the Hispanic effect could be an artifact of the fact that most Hispanic subjects were M.B.A. students in Miami, many of whom are from families that emigrated from Cuba.

IV. Extensions

There are, of course, other possible values for the risk aversion coefficient of the power function, which correspond to more or less curvature. And there are functions with the curvature of Figure 4.1 that are not power functions, e.g. the natural logarithm. Finding the function that provides the best fit to the data involves a statistical analysis based on specifying an “error” term that

incorporates un-modeled random variations due to emotions, perceptions, interpersonal differences, etc. (These differences might be controlled by including demographic variables.) Various functional forms for the utility function will be discussed in later chapters.

Another issue that is being deferred is whether utility should be a function of final wealth or of gains from the current wealth position. Here we treat utility as a function of gains only. There is some experimental and theoretical evidence for this perspective that will also be discussed in a subsequent chapter.

Questions

1. For the square root utility function, find the expected utility of the risky lottery for Decision 6, and check your answer with the appropriate entry in Table 4.2.
2. Consider the quadratic utility function $U(x) = x^2$. Sketch the shape of this function in a figure analogous to Figure 4.1. Does this function exhibit increasing or decreasing marginal utility? Is this shape indicative of risk aversion or risk seeking behavior? Calculate the expected utilities for the safe and risky options in decision 4 for Table 4.1, i.e. when the probabilities are equal to 0.4 and 0.6.
3. Suppose that a deck with all face cards (Ace, King, Queen, and Jack) removed is used to determine a money payoff, e.g. a draw of a 2 of Clubs would pay \$2, a draw of a 10 of Diamonds would pay \$10, etc. Write down the nine possible money payoffs and the probability associated with each. What is the expected value of a single draw from the deck, assuming that it has been well shuffled? How much money would you pay to play this game?

Chapter 5. Randomized Strategies

In many situations there is a strategic advantage associated with being unpredictable, much as a tennis player does not always lob in response to an opponent's charge to the net. This chapter discusses randomized strategies in the context of simple matrix games (Matching Pennies and Battle of Sexes). The associated class experiments can be run using the instructions in the Appendix or with the Veconlab software (MG from the Games menu).

I. Symmetric Matching Pennies Games

In a matching pennies game, each person uncovers a penny showing either Heads or Tails. By prior agreement, one person can take both coins if the pennies match (two Heads or two Tails), and the other can take the coins if the pennies do not match. In this case, a person cannot afford to have a reputation of always choosing Heads, or of always choosing Tails, because being predictable will result in a loss every time. Intuition suggests that each person play Heads half of the time.

A similar situation may arise in a soccer penalty kick, where the goalie has to dive to one side or the other, and the kicker has to kick to one side or the other. The goalie wants a "match" and the kicker wants a "mismatch." Again, any tendency to go in one direction more often than another can be exploited by the other player. If there are no asymmetries in kicking and diving ability for one side versus the other, then each direction should be selected about half of the time.

It is easy to imagine economic situations where people would not like to be predictable. For example, a lazy manager only wants to prepare for an audit if such an audit is likely. On the other side, the auditor is rewarded for discovering problems, so the auditor would only want to audit if it is likely that the manager is unprepared. The qualitative structure of the payoffs for this situation may be represented by Table 5.1.

Table 5.1. A Matching Pennies Game
(Row's payoff, Column's payoff)

Row Player (manager):	Column Player (auditor):	
	Left (audit)	Right (not audit)
Top (prepare)	1, -1 \Rightarrow	\Downarrow -1, 1
Bottom (not prepare)	-1, 1 \Uparrow	\Leftarrow 1, -1

First suppose that the manager expects an audit, so the payoffs on the left side of the table are more likely. Then the manager prefers to prepare for the anticipated

audit to obtain the good outcome with a payoff of 1, which is greater than the payoff of -1 for getting caught unprepared. This preferred deviation is represented by the “up” arrow in the lower-left box. Alternatively, if the manager expects no audit (the right side of the table), then the manager prefers not to prepare, which again yields a payoff of 1 (notice the down arrow in the top-right box). Conversely, the auditor prefers to audit when the manager is not prepared, and to skip the audit otherwise.


Given the intuition about being unpredictable, it is not surprising that we typically observe a near-equal split on aggregate decisions for the symmetric matching pennies game in Table 5.1, although there can be considerable variation in the choice proportions from round to round. (In laboratory experiments, the payoffs are scaled up, and a fixed payment is added to eliminate the possibility of losses, but the essential structure of the game is unchanged, as will be seen in Chapter 11.) Despite the intuitive nature of “fifty-fifty” splits for the game in Table 5.1, it is useful to see what behavior is stable in the sense of being a Nash equilibrium. Up to now, we have only considered strategies without random elements, but such strategies will not constitute an equilibrium in this game. First consider the Top/Left box in the table, which corresponds to audit/prepare. This would be a Nash equilibrium if neither player has an incentive to change unilaterally, which is not the case, since the auditor would not want to audit if the manager is going to prepare. In all cells, the player with the lower payoff would prefer to switch unilaterally.

As mentioned in Chapter 1, Nash (1948) proved that an equilibrium always exists (for games in which each person has a finite number of strategies). Since there is no equilibrium in non-random strategies in the matching pennies game, there must be an equilibrium that involves random play. The earlier discussion of matching pennies already indicated that this equilibrium involves using each strategy half of the time. Think about the “announcement test.” If one person is playing Heads half of the time, then playing Tails will win half of the time; playing Heads will win half of the time, and playing a “fifty-fifty” mix of Heads and Tails will win half of the time. In other words, when one player is playing randomly with probabilities of one half, the other person cannot do any better than using the same probabilities. If each person were to announce that they would use a coin flip to decide which side to play, the other could not do any better than using the coin flip. This is the Nash equilibrium for this game. It is called a “mixed equilibrium” since players use a probabilistic mix of each of their decisions. In contrast, an equilibrium in which no strategies are random is called a “pure strategy” equilibrium, since none of the strategies are probability mixes.

Another perspective on the mixed equilibrium is based on the observation that a person is only willing to choose randomly if no decision is any better than another. So the row player’s choice probabilities must keep the column player

indifferent, and vice versa. The only way one person will be indifferent is if the other is using equal probabilities of Heads and Tails, which is the equilibrium outcome.

Even though the answer is obvious, it is useful to introduce a graphical device that will help clarify matters in more complicated situations. This graph will show each person's best response to any given beliefs about the other's decisions. In Figure 5.1, the **thick solid line** shows the best response for the row player (manager). The horizontal axis represents what Row expects Column to do. These beliefs can be thought of as a probability of Right, going from 0 on the left to 1 on the right. If Column is expected to choose Left, then Row's best response is to choose Top, so the best response line starts at the top-left part of the figure, as shown by the thick line. If Column is expected to choose Right, then Row should play bottom, so the best response line ends up on the bottom-right side of the graph. The crossover point is where the Column's probability is exactly 0.5, since Row does better by playing Top whenever Column is more likely to choose Left.



**Figure 5.1. Row's Best Response to Beliefs about Column's Decision
Column's Best Response to Beliefs about Row's Decision**

A mathematical derivation of the crossover point (where Row is indifferent and willing to cross over) requires that we find the probability of Right for which Row's expected payoff is exactly equal for each decision. Let p denote

Row's beliefs about the probability of Right, so $1-p$ is the probability of Left. Recall from the previous chapter that expected payoffs are found by adding up the products of payoffs and associated probabilities. From the top row of Table 5.1, we see that if Row chooses Top, then Row earns 1 with probability $1-p$ and Row earns -1 with probability p , so the expected payoff is:

$$\text{Row's Expected Payoff for Top} = 1(1-p) - 1(p) = 1 - 2p.$$

Similarly, by playing Bottom, Row earns -1 with probability $1-p$ and 1 with probability p , so the expected payoff is:

$$\text{Row's Expected Payoff for Bottom} = -(1-p) + p = -1 + 2p.$$

These expected payoffs are equal when:

$$1 - 2p = -1 + 2p.$$

Solving, we see that $p = 2/4 = 0.5$, which confirms the earlier conclusion that Row is indifferent when Column is choosing each decision with equal probability.

A similar analysis shows that Column is indifferent when Row is using probabilities of one half. Therefore the best response line will “cross over” when Row's probability of Top is 0.5. To see this graphically, change the interpretation of the axes in Figure 5.1 to let the vertical axis represent Column's *beliefs* about what Row will do. And instead of interpreting the horizontal axis as a probability representing Row's beliefs about Column's action, now interpret it in terms of Column's actual best response. With this change, the dashed line that crosses at a height of one half is Column's best response to beliefs on the vertical axis. If Column thinks Row will play Bottom, then Column wants to play Left, so this line starts in the bottom/left part of the figure. This is because the high payoff of 1 for Column is in the bottom/left part of the payoff table in Table 5.1. There is another payoff of 1 for Column in the top/right part of the table, i.e. when Column thinks Row will play Top, then Right is the best response. For this reason, the dashed best response line in Figure 5.1 ends up in the top/right corner.

In a Nash equilibrium, neither player can do better by deviating, so a Nash equilibrium must be on the best response lines for both players. The only intersection of the solid and dashed lines in Figure 5.1 is at probabilities of 0.5 for each player. If both players think the other's move is like a coin flip, they would be indifferent themselves, and hence willing to decide by a flip of a coin. This equilibrium prediction works well in the symmetric matching pennies game, but we will consider qualifications in a later chapter. The focus here is on calculating

and interpreting the notion of an equilibrium in randomized strategies, not on summarizing all behavioral tendencies in these games.

II. Battle of the Sexes

Thus far, we have considered two types of games with unique Nash equilibria, the prisoner's dilemma (Chapter 3) and the matching pennies game. In contrast, the coordination game discussed in Chapter 3 had *two* equilibria in non-random strategies, one of which was preferred by both players. Next we consider another game with two equilibria in non-random strategies, the payoffs for which are shown in Table 5.2. Think of this as a game where two friends who live on opposite sides of Central Park wish to meet at one of the entrances, i.e. on the East side or on the West side. It is obvious from the payoffs that Column wishes to meet on the East side, and Row prefers the West side. But notice the zero payoffs in the mis-matched (West and East) outcomes, i.e. each person would rather be with the other than to be on the preferred side of the park alone. Games with the structure shown in Table 5.2 are generally known as "battle-of-the-sexes" games.

Table 5.2. A Battle-of-Sexes Game
(Row's payoff, Column's payoff)

Row:	Column:	
	West	East
West	1, 4	0, 0
East	0, 0	4, 1

Think of what you would do in a repeated situation. Clearly most people would take turns. This is exactly what tends to happen when two people play the game repeatedly in controlled experiments with the same partner. In fact, coordinated switching often arises even when explicit communication is not permitted. Table 5.3 shows a decision sequence for a pair who were matched with each other for 6 periods, using payoffs from Table 5.2. There is a "match" on East in the first period, and the Column player then switches to West in period 2, which results in a mismatch and earnings of \$0 for both. Row switches to West in period 3, and they alternate in a coordinated manner in each subsequent period, which maximizes their joint earnings. Four of the six pairs alternated in this manner, and the other two pairs settled into a pattern where one earned \$4 in all rounds.

Table 5.3. Alternating Choices for a Pair of Subjects
in a Battle-of-Sexes Game with Fixed Partners

	Round 1	Round 2	Round 3	Round 4	Round 5	Round 6
Row Player	East (\$4)	East (\$0)	West (\$1)	East (\$4)	West (\$1)	East (\$4)
Column Player	East (\$1)	West (\$0)	West (\$4)	East (\$1)	West (\$4)	East (\$1)

The problem is harder if the battle-of-sexes game is played only once, without communication, or if there is repetition with random matchings. Consider the battle-of-sexes game shown in Table 5.4.

Table 5.4. A Battle-of-Sexes Game
(Row's payoff, Column's payoff)

Row:	Column:	
	Left	Right
Top	2, 1	0, 0
Bottom	0, 0	1, 2

The payoffs in this table were used a recent classroom experiment conducted at the College of William and Mary. There were 30 students located in several different computer labs, with random matchings between those designated as row players and those designated as column players. Table 5.5 shows the percentage of times that players chose the preferred location (Top for Row and Bottom for Column). Intuitively, one would expect that the percentage of preferred-location choices to be above one half, and in fact this percentage converges to 67%. This mix of choices does not correspond to either equilibrium in pure strategies.

The remarkable feature of Table 5.4 is that both types seem to be choosing their preferred location about two-thirds of the time. It is not surprising that this fraction is above one-half, but why two thirds? If you look at the payoff table, you will notice that Row gets either 2 or 0 for playing Top, and either 0 or 1 for playing Bottom, so in a loose sense Top is more attractive unless Column is expected to play Right with high probability. Similarly, from Column's point of view, Right (with payoffs of 0 or 2) is more attractive than Left (with payoffs of 1 and 0), unless Row is expected to play Top with high probability. This intuition only provides a qualitative prediction, that each person will choose their preferred

decision more often than not. The remarkable convergence of the frequency of preferred decisions to $2/3$ cries out for some mathematical explanation, especially considering that each person's 2 and 1 payoffs add up to 3, and two thirds of sum can only be obtained with the preferred decision.

Table 5.5. Percentage of Preferred-Location Decisions
for the Battle-of-Sexes Game in Table 5.4

Round	Row Players	Column Players	All Players
1	80	87	83.5
2	87	93	90
3	87	60	73.5
4	67	67	67
5	67	67	67
Nash	67	67	67

Instead of looking for mathematical coincidences, let us calculate some expected payoff expressions as was done in the previous section. First consider Row's perspective when Column is expected to play Right with probability p . Consider the top row of the payoff matrix in Table 5.4, where Row thinks the right column is relevant with probability p . If Row chooses Top, Row gets 2 when Column plays Left (expected with probability $1-p$) and Row gets 0 when Column plays Right (expected with probability p). When Row chooses Bottom, these payoffs are replaced by 0 and 1. Thus Row's expected payoffs are:

$$(5.1) \quad \begin{aligned} \text{Row's expected payoff for Top} &= 2(1-p) + 0(p) = 2 - 2p \\ \text{Row's expected payoff for Bottom} &= 0(1-p) + 1(p) = 0 + p. \end{aligned}$$

The expected payoff for Top is higher when $2 - 2p > p$, or equivalently, when $p < 2/3$. Obviously, the expected payoffs are equal when $p = 2/3$, and Top provides a lower expected payoff when $p > 1/3$.

Since Row's expected payoffs are equal when column's probability of choosing Right is $2/3$, Row would not have a preference between the two decisions and would be willing to choose randomly. At this time, you might just guess that since the game looks symmetric, the equilibrium involves each player choosing their preferred decision with probability $2/3$. This guess would be correct, as we shall show with Figure 5.2. As before, the solid line represents Row's best response that we analyzed above. If the horizontal axis represents Row's beliefs about how likely it is that Column will play Right, then Row would want to "go to the top of the graph" (play Top) as long as Column's probability p

is less than $2/3$. Row is indifferent when $p = 2/3$, and row would want to “go to the bottom” when $p > 2/3$.

So far, we have been looking at things from Row’s point of view, but an equilibrium involves both players, so let’s think about column’s decision where the vertical axis now represents Column’s belief about what Row will do. At the top, Row is expected to play Top, and Column’s best response is to play Left in order to be in the same location as Row. Thus Column’s dashed best response line starts in the upper-left part of the figure. This line also ends up in the lower-right part, since when Row is expected to play Bottom (coming down the vertical axis), Column would want to switch to the Right location that is preferred. It can be verified by simple algebra that the switchover point is at $2/3$, as shown by the horizontal segment of the dashed line.




Figure 5.2. Best Responses for the Game in Table 5.4

Since a Nash equilibrium is a pair of strategies such that each player cannot do better by deviating, each player has to be making a best response to the other’s strategy. In the figure, a Nash equilibrium will be on both Row’s (solid) best-response line and on Column’s (dashed) best-response line. Thus the final

step is to look for equilibrium points at the intersections of the best response lines. There are three intersections. There is one in the upper-left corner of the figure where both coordinate on Row's preferred outcome (Row earns 2, Column earns 1). Similarly, the lower-right intersection is an equilibrium at Column's preferred outcome (Row earns 1, Column earns 2). We already found these equilibria by looking at the payoff matrix directly, but the third intersection point in the interior of the figure is new. At this point, players choose their preferred decisions (Top for Row and Bottom for Column) with probability $2/3$, which is what we see in the data.

The graph shows the equilibria clearly, but it is useful to see how the random strategy equilibrium would be found with only simple algebra, since the graphical approach will not be possible with more players or more decisions.

Step 1. First we need to summarize notation:

p = probability that Column chooses Right

q = probability that Row chooses Top

Step 2. Calculate expected payoffs for each decision.

Row's Expected Payoff for Top = $2(1-p) + 0(p) = 2 - 2p$

Row's Expected Payoff for Bottom = $0(1-p) + 1(p) = 0 + p$

Column's Expected Payoff for Right = $0(q) + 2(1-q) = 2 - 2q$

Column's Expected Payoff for Left = $1(q) + 0(1-q) = q + 0$

Step 3. Calculate the equilibrium probabilities.

Equate Row's expected payoffs to determine p .

Equate Column's expected payoffs to determine q .

We already showed that the first part yields $p = 2/3$, and it is easy to verify that $q = 2/3$.

The idea behind these calculations is that, in order to randomize willingly, a person must be indifferent between the decisions, and indifference is found by equating expected payoffs. The tricky part is that equating *Row's* expected payoffs pins down *Column's* probability, and vice versa.

Extensions

The data in Table 5.5 are atypical in the sense that such sharp convergence to an equilibrium in randomized strategies is not always observed. Often there is a little more bouncing around the predictions, due to noise factors (see Question 7). Remember that each person is seeing a series of other people, so people have different experiences, and hence different beliefs. This raises the issue of how

people learn after observing others' decisions, which will be discussed in a later chapter on Bayesian learning. Second, the games discussed in this chapter are symmetric in some sense; payoff asymmetries may cause biases, as will be discussed in a later chapter. Finally, the battle-of-sexes game discussed here was conducted under very low payoff conditions, with only one person of 30 being selected *ex post* to be paid their earnings in cash. High payoffs might cause other factors like risk aversion to become important, especially when there is a lot more variability in the payoffs associated with one decision than with another. If risk aversion is a factor, then the expected payoffs would have to be replaced with expected utility calculations.

Questions

1. Use the expected payoff calculations in Step 2 above to solve for an equilibrium level of q for the battle-of-sexes game.
2. The appendix to this chapter contains instructions for a game with playing cards. a) Write out the payoff matrix for round 1 of this game. b) Find all equilibria in non-random strategies and say what kind of game this is. c) Find expressions for Row's expected payoffs, one for each of Row's decisions. d) Find expressions for Column's expected payoffs. e) Find the equilibrium with randomized strategies, using algebra. f) Illustrate your answer with a graph.
3. Suppose that the payoffs in Table 5.4 are changed by raising the 2 payoff to a 3, for both players. Answer all parts c), d) and e) of question 2 for this game.
4. Find the Nash equilibrium in mixed strategies for the coordination game in Table 3.2, and illustrate your answer with a graph.
5. Graph the best-response lines for the prisoner's dilemma game in Table 3.1, and indicate why there is a unique Nash equilibrium.
6. The data in the table below were for the battle-of-sexes game shown in Table 5.2. The 12 players were randomly matched, and the final 5 periods out of a 10-period sequence are shown. The game was run with an experimental economics class at the University of Virginia. Each "player" consisted of one or two people at the same PC, and one player was selected at random *ex post* to be paid a third of earnings. Calculate the percentage of times that Row players chose East in all 5 periods, the percentage of times that Column players chose West, and the average of these two numbers. Then calculate the mixed-strategy Nash equilibrium.

Round	6	7	8	9	10
Row	6 East	3 East, 3 West	5 East, 1 West	5 East, 1 West	6 East
Column	2 East, 4 West	2 East, 4 West	1 East, 5 West	3 East, 3 West	2 East, 4 West

Part II. Individual Decision Experiments

The next part of the book contains a series of chapters on games involving decision making when the payoff outcomes cannot be known for sure in advance. Uncertainty about outcomes is represented by probabilities, which can be used to calculate the expected value of some payoff or utility function. As indicated in Chapter 4, a risk-neutral person will choose the decision with the highest expected money value. Non-neutral attitudes towards risk can be represented as the maximization of a nonlinear utility function.

The notion of expected utility became widely accepted in economics after von Neumann and Morgenstern (1944) provided a set of plausible assumptions or “axioms” which ensure that decisions will correspond to the maximization of the expected value of some utility function. This is, of course, an “as if” claim; you may sometimes see people calculate expected values, but it is rare to see a person actually multiplying out utilities and the associated probabilities. Today, expected utility is widely used in economics (and related areas like finance), despite some well-known situations where behavior is sensitive to biases and “behavioral” factors.

Chapter 6 pertains to the simplest two-way prediction task, e.g. rain or shine, when the underlying probabilities are fixed but unknown in repeated rounds. Each new observation provides more information about the relative likelihood of the two events, and the issue of “probability matching” concerns the relationship between the predictions and the underlying frequency of each event, and this matching represents a deviation from optimal behavior. This topic is also used to introduce some simple learning rules, e.g. “reinforcement learning.”

Anomalies in lottery-choice situations are discussed in Chapter 7, e.g. the “Allais paradox,” which is pattern of choices that cannot be explained by standard expected utility theory.

All decision problems considered up this point will have been “static,” i.e. without any time-dependent elements. The eighth chapter is organized around a dynamic situation involving costly search (e.g., for a high wage or a low price). Despite the complex nature of sequential search problems, observed behavior is often surprisingly consistent with search rules derived from an analysis of optimal decision-making.

Chapter 6. Probability Matching

Perhaps the simplest prediction problem involves guessing which of two random events will occur. The probabilities of the two events are fixed but not known, so a person can learn about these probabilities by observing the relative frequencies of the two events. One of the earliest biases recorded in the psychology literature was the tendency for individuals to predict each event with a frequency that approximately matches the fraction of times that the event is observed. This bias, known as “probability matching,” has been widely accepted as being evidence of irrationality despite Siegel’s experiments in the 1960’s. These experiments provide an important methodological lesson for how experiments should be conducted. The results are also used to begin a discussion of learning models that may explain paths of adjustment to some “steady state” where systematic changes in behavior have ceased. Binary prediction tasks can be run by hand (with instructions in the Appendix) or using the Veconlab program (PM listed under Decisions).

I. Being Treated Like a Rat

Before the days of computers, the procedures for binary prediction tasks in psychology experiments sometimes seemed like a setup for rats or pigeons that had been scaled up for human subjects. I will describe the setup used by Siegel and Goldstein (1959). The subject was seated at a desk with a screen that separated the working area from the experimenter on the other side. The screen contained two light bulbs, one on the left and one on the right, and a third, smaller light in the center that was used to signal that the next decision must be made. When the signal light went on, the subject recorded a prediction by pressing one of two levers (left and right). Then one of the lights was illuminated, and the subject would receive reinforcement (if any) based on whether or not the prediction was correct. When the next trial was ready, the signal light would come on, and the process would be repeated, perhaps hundreds of times. No information was provided about the relative likelihood of the two events (Left and Right), although sometimes people were told how the events were generated, e.g. by using a printed list. In fact, one of the events would be set to occur more often, e.g. 75 percent of the time. The process generating these events was not always random, e.g. sometimes the events were rigged so that in each block of 20 trials, exactly 15 would result in the more likely event.

By the time of the Siegel and Goldstein experiments in the late 1950’s, psychologists had already been studying probability matching behavior for over twenty years. The results seemed to indicate a curious pattern: the proportion of times that subjects predicted each event roughly matched the frequency with

which the events occurred. For example, if Left occurred three-fourths of the time, then subjects would come to learn this by experience and then would tend to predict Left three-fourths of the time.

II. Are Hungry Rats Really More Rational than Humans ?

The astute reader may have already figured out what is the best thing to do in such a situation, but a formal analysis will help ensure that such a conclusion does not rely on hidden assumptions like risk neutrality. In order to evaluate the rationality of this matching behavior, let us assume that the events really were independent random realizations, and that the more likely event occurred with probability $p > 1/2$. Let U_C denote the utility of the reward for a correct prediction, and let U_I denote the utility of the reward for an incorrect prediction. The rewards could be “external” (money payments, food), “internal” (psychological self-reinforcement), or some combination. The only assumption is that there is some preference for making a correct prediction: $U_C > U_I$. These utilities may even be changing over time, depending on the rewards received thus far; the only assumption is that an additional correct prediction is preferred.

Once a number of trials have passed, the person will have figured out which event is more likely, so let p denote the subjective probability that represents those beliefs, with $p > 1/2$. There are two decisions: predict the more likely event and predict the less likely event. Each decision yields a lottery:

Predict more likely event: U_C with probability p
 U_I with probability $1-p$

Predict less likely event: U_I with probability p
 U_C with probability $1-p$

Thus the expected utility for predicting the more likely event is higher if

$$pU_C + (1-p)U_I > pU_I + (1-p)U_C,$$

or equivalently,

$$(2p-1)U_C > (2p-1)U_I,$$

which is always the case since $p > 1/2$ and $U_C > U_I$. Although animals may become satiated with food pellets and other physical rewards, economists have no trouble with a non-satiation assumption for money rewards. Note that this argument does not depend on any assumption about risk attitudes. The

implication is that the more likely event should always be predicted, i.e. with probability 1. In this sense, probability matching is irrational as long as there is no satiation, so $U_C > U_I$.

Despite the clear prediction that one should predict the more likely event 100 per cent of the time after it becomes clear which event that is, this behavior is often not observed in laboratory experiments. For example, in a recent summary of the probability matching literature, the psychologist Fantino (1998, pp. 360-361) concludes: "human subjects do not behave optimally. Instead they match the proportion of their choices to the probability of reinforcement.... This behavior is perplexing given that non-humans are quite adept at optimal behavior in this situation." As evidence for the higher degree of rational behavior in animals, he cites a 1996 study that reported choice frequencies for the more likely event to be well above the probability matching predictions in most treatments conducted with chicks and rats. Before concluding that the animal subjects are more rational than humans, it will be instructive to review a particularly well-designed probability matching experiment.

III. Siegel and Goldstein's Experiments

Sidney Siegel is perhaps the psychologist who has had the largest direct and indirect impact on experiments in economics. His early work provides a high standard of careful reporting and procedures, appropriate statistical techniques, and the use of financial incentives where appropriate. His experiments on probability matching are a good example of this work. In one experiment, 36 male Penn State students were allowed to make predictions for 100 trials, and then 12 of these were brought back on a later day to make predictions in 200 more trials. The proportions of predictions for the more likely event are graphed in Figure 6.1, with each point being the average over 20 trials.




Figure 6.1. Prediction Proportions for the Event with Frequency 0.75
Source: Siegel, Siegel, and Andrews (1964) for Dark Lines and Holt (1992) for Light Lines

The 12 subjects in the “no-pay” treatment were simply told to “do your best” to predict which light bulb would be illuminated. These averages are plotted as the heavy dashed line, which begins at about 0.5 as would be expected in early trials with no information about which event is more likely. Notice that the proportion of predictions for the more likely event converges to the level of 0.75 (shown by a horizontal line on the right) predicted by probability matching, with a leveling off at about trial 100.

In the “pay-loss” treatment, 12 participants received 5 cents for each correct prediction, and they lost 5 cents for each incorrect decision. The 20-trial averages are plotted as the dark solid line in the figure. Notice that the line converges to a level of about 0.9, as shown by the upper horizontal line on the right. A third “pay” treatment offered a 5-cent reward but no loss for an incorrect prediction, and the results (not shown) are in between the other two treatments, and clearly above 0.75. Clearly, incentives matter, and probability matching is not observed with incentives in this context.

It would be misleading to conclude that incentives always matter, or that probability matching will be observed in no-pay treatments using different procedures. The two thin lines in Figure 6.1 show 20-trial averages for an experiment (Holt, 1992) in which 6 University of Virginia students in no-pay and pay treatments made decisions using computers, where the events were determined by a random number generator. The instructions matched those reported in Siegel, Siegel, and Andrews (1964), except that colored boxes on the screen were used instead of light bulbs. Notice that probability matching is not observed in either the pay/loss treatment (with a reward of 10 cents and a penalty of 10 cents, shown by the thin solid line) or the no-pay treatment (thin dashed line). The results for a third treatment with a 20-cent reward and no penalty (not shown) were similar. The reason that matching was not observed in the Holt no-pay treatment is unclear. One conjecture is that the computer interface makes the situation more anonymous and less like a matching pennies game. In the Siegel setup with the experimenter on one side of the screen, the subject might incorrectly perceive the situation as having some aspects of a game against the experimenter. Recall that the equilibrium in a matching pennies game involves equal probabilities for each decision. This might explain the lower choice percentages for the more likely event reported in the non-computerized setup, but this is only a guess.

Finally, note that Siegel's findings suggest a resolution to the paradoxical finding that rats are smarter than humans in binary prediction tasks. Since you cannot just tell a rat to "do your best," animal experiments are always run with food or drink incentives. As a result, the observed choice proportions are closer to those of financially motivated human subjects. In a recent survey of over fifty years of probability matching experiments, Vulkan (1998) concluded that probability matching is generally not observed with real payoffs, although humans can be surprisingly slow learners in this simple setting.

IV. A Simple Model of Belief Learning

Although the probability matching bias is not taken seriously these days, at least by economists, the experiments provide a useful data set for the study of learning behavior. Given the symmetry of the problem, a person's initial beliefs ought to be that each event is equally likely, but the first observation should raise the probability associated with the event that was just observed. One way to model this learning process is to let initial beliefs for the probability of events L and R be calculated as:

$$(6.1) \quad \Pr(L) = \frac{a}{a+a} \quad \text{and} \quad \Pr(R) = \frac{a}{a+a} \quad (\text{priors}),$$

where α is a positive parameter to be explained below. Of course, α has no role yet, since both of the above probabilities are equal to $1/2$.

If event L is observed, then $\Pr(L)$ should increase, so let us add 1 to the numerator for $\Pr(L)$. To make the two probabilities sum to 1, we must add 1 to the denominators for each probability expression:

$$(6.2) \quad \Pr(L) = \frac{\mathbf{a}+1}{\mathbf{a}+1+\mathbf{a}} \quad \text{and} \quad \Pr(R) = \frac{\mathbf{a}}{\mathbf{a}+1+\mathbf{a}} \quad (\text{after observing L}).$$

Note that α determines how quickly the probabilities respond to the new information; a large value of α will keep these probabilities close to $1/2$. Continuing to add 1 to the numerator of the probability for the event just observed, and to add 1 to the denominators, we have a formula for the probabilities after N_L observations of event L and N_R observations of event R. Let N be the total number of observations to date. Then the resulting probabilities are:

$$(6.3) \quad \Pr(L) = \frac{\mathbf{a}+N_L}{2\mathbf{a}+N} \quad \text{and} \quad \Pr(R) = \frac{\mathbf{a}+N_R}{2\mathbf{a}+N} \quad (\text{after } N \text{ observations}).$$

where $N = N_L + N_R$.

In the early periods, the totals, N_L and N_R , might switch in terms of which one is higher, but the more likely event will soon dominate, and therefore $\Pr(L)$ will be greater than $1/2$. The prediction of this learning model (and the earlier analysis of perfect rationality) is that all people will eventually start to predict the more likely event every time. Any “unexpected” prediction switches would have to be explained by adding some randomness to the decision making (Goeree and Holt, 2003), or by adding “recency effects” which make probability assessments more sensitive to the most recently observed outcomes. These kinds of modeling changes will be deferred until a later chapter. The point here is that a simple and intuitive learning model can be constructed, and that this model can explain a high proportion of predictions of the more likely event.

V. Reinforcement Learning

In a psychology experiment, the rewards and punishments are referred to as “reinforcements.” One prominent theory of learning associates changes in behavior to the reinforcements actually received. For example, suppose that the person earns a reinforcement of x for each correct prediction, nothing otherwise. If one predicts event L and is correct, then the probability of choosing L should

increase, and the extent of the behavioral change may depend on the size of the reinforcement. One way to model this is to let the choice probability be:

$$(6.4) \quad \Pr(\text{choose L}) = \frac{\mathbf{a} + x}{\mathbf{a} + x + \mathbf{a}} \quad \text{and} \quad \Pr(\text{choose R}) = \frac{\mathbf{a}}{\mathbf{a} + x + \mathbf{a}} .$$

Despite the similarity with equation (6.2), there are two important differences. The left side of (6.4) is a choice probability, not a probability that represents beliefs. With reinforcement learning, beliefs are not explicitly modeled, as is the case for the “belief learning” models of the type discussed in the previous section. The second difference is that the x in (6.4) represents a reinforcement, not an integer count as in (6.2).

Of course, reinforcement is a broad term, which can include both physical things like food pellets given to rats, as well as psychological feelings associated with success or failure. One way that this model is implemented for experiments with money payments is to make the simplifying assumption that reinforcement is measured by earnings. Suppose that event L has been predicted N_L times and that the predictions have sometimes been correct and sometimes not. Then the total earnings for predicting L, denoted e_L , would be less than xN_L . Similarly, let e_R be the total earnings from the correct R predictions. The choice probabilities would then be:

$$(6.5) \quad \Pr(\text{choose L}) = \frac{\mathbf{a} + e_L}{2\mathbf{a} + e_L + e_R} \quad \text{and} \quad \Pr(\text{choose R}) = \frac{\mathbf{a} + e_R}{2\mathbf{a} + e_L + e_R} .$$

Notice that the α parameters again have the role of determining how quickly learning responds to the stimulus, which is the money reinforcement in this case.

This kind of model might also explain some aspects of behavior in probability matching experiments with financial incentives. The choice probabilities would be equal initially, but a prediction of the more likely event will be correct 75% of the time, and the resulting asymmetries in reinforcement would tend to raise prediction probabilities for that event, and the total earnings for this event would tend to be larger than for the other event. Then e_L would be growing faster, so that e_R/e_L would tend to get smaller as e_L gets larger. Thus the probability of choosing L in (6.5) would tend to converge to 1. This convergence may be quite slow, as evidenced by some computer simulations reported by Goeree and Holt (2003). (See question 6 for some details on how such simulations might be structured.) The simulations normalized the payoff x to be 1, and used a value of 5 for α . Figure 6.2 shows the 20-period averages for simulations run for 1000 people, each making a series of 300 predictions. The

simulated averages start somewhat above the observed averages for Siegel's subjects (more like the Holt data in Figure 6.1), but similar qualitative patterns are observed.




Figure 6.2. Simulated and Observed Prediction Proportions for the More Likely Event
Source: Goeree and Holt (2003) for the simulated data.

VI. Extensions

Both of the learning models discussed here are somewhat simple, which is part of their appeal. The reinforcement model incorporates some randomness in behavior and has the appealing feature that incentives matter. But it has less of a cognitive element; there is no reinforcement for decisions not made. For example, suppose that a person chooses L three times in a row (by chance) and is wrong each time. Since no reinforcement is received, the choice probabilities stay at 0.5, which seems like an unreasonable prediction. Obviously, people learn something in the absence of previously received reinforcement, since they realize that making a good decision may result in higher earnings in the next round. Camerer and Ho (1999) have developed a generalization of reinforcement

learning that contains some elements of belief learning. Roughly speaking, observed outcomes receive partial reinforcement even if nothing is earned.

The belief-learning model in section IV can be derived from statistical principles (Bayes's rule to be discussed in the next chapter). Then beliefs determine the expected payoffs (or utilities) for each decision, which in turn determine the decisions made. In theory, the decision with the highest expected payoff is selected with certainty. In an experiment, however, some randomness in decision making might be expected if the expected payoffs for the two decisions are not too different. This randomness may be due to random emotions, calculation errors, selective forgetting of past experience, etc. Building some randomness into models that use belief learning (as in Capra, et al., 1999) is a useful way to go about explaining data from laboratory experiments, as we shall see in later chapters.

These learning models can be enriched in other ways to obtain better predictions of behavior. For example, the sums of event observations in the belief-learning model weigh each observation equally. It may be reasonable to allow for "forgetting" in some contexts, so that the observation of an event like L in the most recent trial may carry more weight than something observed a long time ago. This is done by replacing sums with weighted sums. For example, if event L were observed three times, N_L in (6.2) would be 3, which can be thought of as $1+1+1$. If the most recent observation (listed on the right in this sum) is twice as prominent as the one before it, then the prior event would get a weight of one half, and the one before that would get a weight of one-fourth, etc. These and other enrichments will be discussed in later chapters.

Questions

1. The initial beliefs implied by (6.1) are that each event is equally likely. How might this equation be altered in a situation where a person has some reason to believe that one event is more likely than another, even before any draws are observed?
2. A recent class experiment used the probability matching (PM) *Veconlab* software with payoffs of \$0.20 for each correct prediction, and in-class earnings averaged several dollars. There were 6 teams of 1-2 students, who made predictions for 20 trials only. (Participants were University of Virginia undergraduates in an experimental economics class who had not seen a draft of this chapter, but who had read drafts of the earlier chapters.) The more likely event was programmed to occur with probability 0.75. Calculate the expected earnings per trial for a team that follows perfect probability matching. How much more would a team earn per trial by being perfectly rational after learning which event is more likely?

3. Answer the two parts of question 2 for the case where a correct answer results in a gain of \$0.10 and an incorrect answer results in a loss of \$0.10.
4. For the gain treatment described in question 2, the more likely event actually occurred with probability 0.77 in the first 20 trials, averaged over all 6 teams. This event was predicted with a frequency of 0.88. Where would a data point representing this average be plotted in Figure 6.1?
5. For the gain/loss treatment described in question 3, the more likely event was observed with a frequency of about 0.75 in the first ten trials and 0.78 in the second 10 trials. Predictions were made by 6 teams of 1-2 students, and their earnings averaged about 25 cents per team. (To cover losses, each team began with a cash balance of \$1, as did the 6 other teams in the parallel gains treatment.) In the gain/loss treatment, the more likely event was predicted with a frequency of 0.58 in the first 10 trials and 0.70 in trials 11-20. To what extent do these results provide evidence in support of probability matching?
6. (A 10-sided die is required.) The discussion of long-run tendencies for the learning models was a little loose, since there are random elements in these models. One way to proceed is to simulate the learning processes implied by these models. Consider the reinforcement learning model, with initial choice probabilities of one half each. You can simulate the initial choice by throwing the 10-sided die twice, where the first throw determines the “tens” digit and the second determines the “ones” digit. For example, throws of a 6 and a 2 would determine a 62. If the die is marked with numbers 0, 1, ...9, then any integer from 0 to 99 is equally likely. The simulation could proceed by letting the person predict L if the throw is less than 50 (i.e., one of the 50 outcomes: 1, 2, 3, ... 49), which would occur with probability 0.5. The determination of the random event, L or R, could be done similarly, with the outcome being L if the next two throws determine a number that is less than 75. Given the prediction, the event, and the reward, say 10 cents, you can use the formula in (6.4) to determine the choice probabilities for the next round. This whole process can be repeated for a number of rounds, and then one could start over with a new simulation, which can be thought of as a simulation of the decision pattern of a second person. Simulations of this type are easily done with computers, or even with the random number features of a spreadsheet program. However you choose to generate the random numbers, your task is to simulate the process for four rounds, showing the choice probabilities, the actual choice, the event observed, and the reinforcement for each round.

Chapter 7. Lottery Choice Anomalies

Choices between lotteries with money payoffs may produce anxiety and other emotional reactions, especially if these choices may result in significant monetary gains and losses. Expected utility theory implies that such choices, even the difficult and stressful ones, can be modeled as the maximization of a mathematical function (a sum of products of probabilities and utilities). As would be expected, actual decisions sometimes deviate from these mathematical predictions, and this chapter begins with one of the more common anomalies: the Allais paradox. Other biases, such as the misperception of large and small probabilities are discussed. The web program PLC can be used to evaluate many of these anomalies; the default setup provides paired lottery choices that are based on Allais paradox and reflection effects.

I. Introduction

The predominant approach to the study of individual decision making in risky situations is the expected utility model introduced in Chapter 4. Expected utility calculations are sums of probabilities times (possibly nonlinear) utility functions of the prize amounts. In contrast, the probabilities enter linearly, which precludes over-weighting of low probabilities, for example. The nonlinearities in utility permit an explanation of risk aversion. This model, which dates to Bernoulli (1738), received a formal foundation in von Neumann and Morgenstern's (1944) book on game theory. This book specified a set of assumptions ("axioms") that imply behavior consistent with the maximization of the expected value of a utility function.

Almost from the very beginning, economists were concerned that behavior in some situations seemed to contradict the predictions of this model. The most famous contradiction, the Allais paradox, is presented in the next section. Such anomalies have stimulated a lot of work, theoretical and experimental, on developing alternative models of choice under risk. The most commonly mentioned alternative, "Prospect Theory," is presented and discussed in the sections that follow.

A reader looking for a resolution of the key modeling issues will be disappointed; some progress has been made, but much of the evidence is mixed. The purpose of this chapter is to introduce the issues, and to help the reader interpret seemingly contradictory results obtained with different procedures. For example, it is not uncommon for the estimates of the environmental benefits of some policy to differ by a factor of 2, which may be attributed to the way the questions were asked and to a "willingness-to-pay/willingness-to-accept bias." A familiarity with this and other biases is crucial for anyone interested in

interpreting the results of experimental and survey studies of situations where the outcomes are unknown in advance.

II. The Allais Paradox

Consider a choice between a sure 3,000 and a 0.8 chance of winning 4,000. This choice can be thought of as a choice between two “lotteries” that yield random earnings:

Table 7.1. A Lottery Choice Experiment (Kahneman and Tversky, 1979)

Lottery S (selected by 80 %)	Lottery R (selected by 20%)
3,000 with probability 1.0	4,000 with probability 0.8 0 with probability 0.2

A lottery is an economic item that can be owned, given, bought, or sold. People may prefer some lotteries to others, and economists assume that these preferences can be represented by a utility function, i.e. that there is some mathematical function with an expected value that is higher for the lottery selected than for the lottery not selected. This is not an assumption that people actually think about utility or do such calculations, but rather, that choices can be represented by (and are consistent with) rankings provided by the utility function.

The simplest utility function is the expected money value, which would represent the preferences of someone who is risk neutral. An expected payoff comparison would favor Lottery R, since $0.8(4,000) = 3,200$, which is higher than the 3,000 for the safe option. In this situation, Kahneman and Tversky reported that 80% of the subjects chose the safe option, which indicates some risk aversion. (Payoffs, in Israeli pounds, were hypothetical.) A person who is not neutral to risk would have preferences represented by a utility function with some curvature. The decision of an expected utility maximizer who prefers the safe option could be represented:

$$(7.1) \quad U(3,000) > 0.8U(4,000) + 0.2U(0),$$

where the utility function represents preferences over money income.

Next consider some simple mathematical operations that can be used to obtain a prediction for how such a person would choose in a different situation. For example, suppose that there is a three-fourths chance that all gains from either lottery will be confiscated, i.e. that is a 0.75 chance of earning \$0. To analyze this possibility, multiply both sides of (7.1) by 0.25, to obtain:

$$(7.2) \quad 0.25U(3,000) > 0.2U(4,000) + 0.05U(0).$$

In order to make the probabilities on each side sum to 1, we need to add the $0.75U(0)$ that corresponds to confiscation. This amount is added to both sides of (7.2) to obtain:

$$(7.3) \quad 0.25U(3,000) + 0.75U(0) > 0.2U(4,000) + 0.8U(0).$$

The inequality implies that the same person (who initially preferred Lottery S to Lottery R in Table 7.1) would prefer a one-fourth chance of 3,000 to a one-fifth chance of 4,000. Any reversal of this preference pattern, e.g. preferring the sure 3,000 in the first choice and the 0.2 chance of 4,000 in the second choice would violate expected utility theory. A risk-neutral person, for example, would prefer the lottery with the possibility of winning 4,000 in both cases (we already checked the first case, and the second is assigned in question 1).

The intuition underlying these predictions can be seen by reexamining equation (7.3). The left side is the expected utility of a one-fourth chance of 3,000 and a three-fourths chance of 0. Equivalently, we can think of the left side as a one-fourth chance of Lottery S (which gives 3,000) and a three-fourths chance of 0. Although it is not so transparent, the right side of (7.3) can be expressed analogously as a one fourth chance of Lottery R and a three-fourths chance of 0. To see this, note that Lottery R provides 4,000 with probability 0.8, and one fourth of 0.8 is 0.2, as indicated on the right side of (7.3). Thus the implication of the inequality in (7.3) is that a one-fourth chance of Lottery S is preferred to a one-fourth chance of Lottery R. What the mathematics of expected utility implies is that if you prefer Lottery S to Lottery R as in (7.1), then you prefer a one-fourth chance of Lottery S to a one-fourth chance of Lottery R as in (7.3). The intuition for this prediction is that the “extra” three-fourths chance of 0 that was added to both sides of the equation in going from (7.2) to (7.3) is the same for both sides. This extra probability of 0 dilutes the chances of winning in both Lottery S and Lottery R, but this added probability of 0 is a common, and hence “irrelevant” addition. Indeed, one of the basic axioms used to motivate the use of expected utility is the assumption of “independence with respect to irrelevant alternatives.”

As intuitive as the argument in the previous paragraph may sound, a significant fraction of the Kahneman and Tversky subjects violated this prediction. Eighty percent chose Lottery S over Lottery R, but sixty-five percent chose the diluted version of Lottery R over the diluted version of Lottery S. This behavior is inconsistent with expected utility theory, and is known as an Allais Paradox, named after the French economist, Allais (1953), who first proposed these types of paired lottery choice situations. In particular, this is the “common-ratio” version of the Allais paradox, since probabilities of positive payoffs for

both lotteries are diluted by a common ratio. Anomalous behavior in Allais paradox situations has also been reported for experiments in which the money prizes were paid in cash (e.g., Starmer and Sugden, 1989, 1991). Battalio, Kagel, and Macdonald (1985) even observed similar choice patterns with rats that could choose between levers that provided food pellets on a random basis.

III. Prospect Theory: Probability Misperception

The Allais paradox results stimulated a large number of studies that were intended to develop alternatives to expected utility theory. The alternatives are typically more complicated to use, and none of them show a clear advantage over expected utility in terms of predictive ability. The possible exception to this conclusion is “prospect theory” proposed by Kahneman and Tversky (1979). Prospect theory is really a collection of elements that specify how a person evaluates a risky prospect in relation to some *status quo* position for the individual. Roughly speaking, there is a reference point, e.g. current wealth, from which gains and losses are evaluated, and gains are treated differently from losses. People are assumed to be averse to losses, but when making choices where payoffs are all losses, they are thought to be risk seeking. On the other hand, when payoffs are all gains, people are assumed to be risk averse in general. A final element is a notion that probabilities are not always correctly perceived, i.e. that low probabilities are over-weighted and high probabilities are under-weighted. These elements can be unbundled and evaluated one at a time, and modifications of expected utility theory that include one or more of these elements can be considered, which is the plan of the remainder of this chapter.

The possibility of probability misperception is suggested by Figure 7.1 in the previous chapter. Recall that this was a graph with a 45-degree line representing the Bayesian prediction from which deviations could be observed. There is a “reverse S-shaped” pattern to the deviations (over-weighting of low probabilities and under-weighting of high probabilities), although the biases there are somewhat small. Prospect theory predicts a probability weighting function which has the general shape observed in Figure 7.1, starting at the origin, rising above the 45-degree line for low probabilities, falling below for high probabilities, and ending up on the 45-degree line in the upper-right corner. The notion that the probability weighting function should cross the 45-degree line at the upper-right corner is based on the intuition that it is difficult to misperceive a probability of one.

To see how prospect theory may explain the Allais paradox, note that Lottery S on the left side of Table 7.1 is a sure thing, so no misperception of probability is possible. The 0.8 chance of a 4,000 payoff on the right, however, may be affected. The largest deviation between Bayesian predictions and average reported probabilities is at a posterior of 0.8 on the horizontal axis in

Figure 7.1. Here, the average report was about seven-tenths. Now reconsider Lottery R on the right side of Table 7.1, when the 0.8 probability of 4,000 is treated as if it were 0.7. This enhances the attractiveness of Lottery S, so that even a risk neutral person would prefer S if the probability was misperceived in this manner (exercise 2). Next consider the choice that results when both lotteries are diluted by a three-fourths chance of a 0 payoff. It is apparent from (7.3) that the 0.25 probability of the 3,000 gain on the left is about the same as the 0.2 probability of the 4,000 on the right. In other words, 0.25 and 0.2 are located close to each other, so that a smooth probability weighting function will overweight them more or less by the same amount. Here a probability weighting function will have little effect, and a risk-neutral person will prefer the diluted version of Lottery R, even though the same person might prefer the undiluted version of Lottery S when the high probability of the 4,000 payoff for lottery R is under-weighted. Put differently, the certain payoff in the original Lottery S is not misperceived, but the 0.8 probability of 4,000 in Lottery R is under-weighted. So probability weighting tends to bias choices towards S. But the diluted S and R lotteries have probabilities in the same range (0.25 and 0.2) so no such bias would be introduced by a “smooth” probability weighting function.

This explanation of the Allais paradox is plausible, but it is not universally accepted. First, the evidence on probability weighting functions is mixed. Recall that the “reverse-S” pattern of deviations in Figure 7.1 could also be due to the tendency for random errors in the elicitation process to spread out on the side where there is “more room” for error. A number of other recent studies have failed to find the “reverse-S” pattern of probability misperceptions (Goeree, Holt, and Plafrey, 2000, 2002). Harbaugh, Krause, and Vesterlund (2002) found this “reverse S” pattern when probabilities were elicited by asking for subjects to assign prices to the lotteries, but the opposite pattern (S-shaped) was observed when the probabilities were elicited by giving people choices between lotteries. As intuitive as the probability weighting explanation for the Allais paradox seems, one would have to say that the evidence is inconclusive at this time.

IV. Prospect Theory: Gains, Losses, and “Reflection Effects”

A second part of prospect theory pertains to the notion of a reference point from which gains and losses are evaluated. In experiments with money payments, the most obvious candidate for the reference point is the current level of wealth, which includes earnings up to the present. The reference point is the basis for the notion of “loss aversion,” which implies that losses are given more weight in choices with outcomes that involve both gains and losses. A second property of the reference point is known as the *reflection effect*, which pertains to cases where positive payoffs are multiplied by minus one in a manner that “reflects” them around 0. The reflection effect postulates that the risk aversion exhibited by

choices when all outcomes are gains will be transformed into a preference for risk when all outcomes are losses.

A comparison of behavior in the gain and loss domains is difficult in a laboratory experiment for a number of reasons. First, the notion of a reference point is not precisely defined. For example, would “paper” earnings recorded up to the present point (but not actually paid) in cash be factored into the current wealth position? A second problem is that human subjects committees do not allow researchers to collect losses from subjects in an experiment, who cannot walk out with less money that they started with. One solution is to give people an initial stake of cash before they face losses, but it is not clear that this process will really change a person’s reference point. The notion that an initial stake is treated differently than hard-earned cash is called the “house-money effect.” There is some evidence that gifts (e.g. candy) tend to make people more willing to take risks in some contexts and less willing in others (Arkes et al. 1988, 1994). It is at least possible that the warm glow of a house-money effect may cause people to appear risk seeking for losses when this may not ordinarily be the case with earned cash. One solution is to give people identical stakes before both gain and loss treatments, which holds the house-money effect constant. And making people earn the initial stake through a series of experimental tasks is probably more likely to change the reference point.

Kahneman and Tversky (1979) presented strong experimental evidence for a reflection effect. The design involved taking all gains in a choice pair like those in Table 7.1 and reflecting them around zero to get losses, as in Table 7.2. Now the “safe” lottery, S*, involves a sure loss, whereas the risky lottery, R*, may yield a worse loss or no loss at all. The choice pattern in Table 7.1, with 80% safe choices, is reversed in Table 7.2, with only 8% safe choices.

Table 7.2. A Reflection Effect Experiment (Kahneman and Tversky, 1979)

Lottery S* (selected by 8%)	Lottery R* (selected by 92%)
<i>minus 3,000 with probability 1.0</i>	<i>minus 4,000 with probability 0.8</i>
	<i>0 with probability 0.2</i>

The Kahneman and Tversky experiments used hypothetical payoffs, which raises the issue of whether this reflection effect will persist with economic incentives. (Recall from Chapter 4 that risk aversion was strongly affected by the use of high economic incentives, as compared with hypothetical payoffs).

Holt and Laury (2002) evaluate the extent of the hypothetical bias in a reflection experiment. They took a menu of paired lottery choices similar to that in Table 4.1 and reflected all payoffs around 0. Recall that this menu has safe lotteries on one side and risky lotteries on the other, and that the probability of the

higher payoff number increases as one moves down the menu. Risk aversion is inferred by looking at the number of safe choices relative to the number of safe choices that would be made by a risk-neutral person (in this case 5). All participants made decisions in both the gains menu and in the losses menu, with the order of menu presentation alternated in half of the sessions. In their hypothetical payoff treatment, subjects were paid a fixed amount \$45 in exchange for participating in a series of tasks (search, public goods) in a different experiment, and afterwards they were asked to indicate their decisions for the lottery choice menus with the understanding that all gains and losses would be hypothetical. When all payoffs were hypothetical gains, about half of the subjects were risk averse, and slightly more than 50% of those who showed risk aversion for gains were risk seeking for losses. The modal pattern in this treatment was reflection, although other patterns (e.g. risk aversion for gains and losses) were also observed with some frequency. The choice frequencies for the hypothetical choices are shown in the left panel of Figure 7.1. The modal pattern of reflection is represented by the tall spike in the back-right corner of the right panel.

The real-incentive treatments for gains and losses were run in a parallel manner with the same choice menus. Participants were allowed to build up earnings of about \$45 in a different experiment using the same tasks used under the hypothetical treatment. In contrast with earlier results, the modal pattern of behavior with real incentives did not involve reflection. The most common pattern was for people to exhibit risk aversion for both gains and losses. The real-payoff choice frequencies are shown in the right panel of Figure 7.1. The modal pattern of risk aversion in both cases is represented by the spike in the back-left side of this panel. There is a little more risk aversion with real payoffs than with hypothetical payoffs; sixty percent of the subjects exhibit risk aversion in the gain condition, and of these only about a fifth are risk seeking for losses. The rate of reflection with real payoffs is less than half of the reflection rate observed with hypothetical payoffs.

Despite the absence of a clear reflection effect, there is some evidence that gains and losses are treated differently. On average, people tended to be essentially risk neutral in the loss domain, but they were generally risk averse in the gain domain. This result provides some support for the notion of a reference point, around which gains and losses are evaluated, which suggests that laboratory data should be analyzed using utility as a function of earnings (gains and losses), not final wealth. In other words, if the net worth of a person's assets is w , and if a decision may produce earnings or losses of x , then the analysis of expected utility (with or without the probability weighting of prospect theory) should be expressed in terms of $U(x)$, not $U(w+x)$.




Figure 7.1. Inferred Risk Aversion for Hypothetical Payoffs (Left) and Real Payoffs (Right)
Source: Holt and Laury (2002)

The absence of a clear reflection effect in the Holt and Laury data is a little surprising given the results of several other studies that found reflection with real money incentives (Camerer, 1989; Battalio et al. 1990). One difference is that instead of holding initial wealth constant in both treatments (at a level high enough to cover losses), these studies provided a high initial stake in the loss treatment, so the final wealth position is constant across treatments. For example, a lottery over gains of 4,000 and 0 could be replaced with an initial payoff of 4,000 and a choice involving *losses* of 4,000 and 0. Each presentation or “frame” provides the same possible final wealth positions (0 or 4,000), but the framing is in terms of gains in one treatment and in terms of losses in the other. A setup like this is exactly what is needed to document a “framing effect.” Such an effect is present since both studies report a tendency for subjects to be risk averse in the gain frame and risk seeking in the loss frame. Whether these results indicate a reflection effect is less clear, since the higher stake provided in the loss treatment may itself have induced more risk seeking behavior, just as gifts of candy and money tend to increase risk seeking in experiments reported by psychologists.

V. Extensions and Further Reading

The dominant method of modeling choice under risk in economics and finance involves expected utility, either applied to gains and losses or to final wealth. The final wealth approach involves a stronger type of rationality in the sense that people can see past gains and losses and focus on the variable that determines consumption opportunities (final wealth). The Camerer (1989) and Battalio et al. (1989) experiments provide strong evidence that decisions are

framed in terms of gains and losses, and that people do not “integrate” gains and losses into a final asset position. Indeed, there is little if any experimental evidence for such “asset integration.” Rabin (2000) and Thaler and Rabin (2001) also provide a theoretical argument against the use of expected utility as a function of final wealth; the argument being that the risk aversion needed to explain choices involving small amounts of money implies absurd levels of risk aversion for choices involving large amounts of money. Most analyses of risk aversion in laboratory experiments are, in fact, already done in terms of gains and losses (Binswanger, 1980; Kachelmeier and Shehata, 1992; Goeree, Holt, and Pfaffrey, 2000, 2002).

Even if expected utility is modeled in terms of gains and losses, there is the issue of whether to incorporate other elements like nonlinear probability weighting that represents systematic misperceptions of probabilities. As noted above, the evidence on this method is mixed, as is the evidence for the reflection effect. Some, like Camerer (1995), have urged economists to give up on expected utility theory in favor of prospect theory and other alternatives. More recently, Rabin and Thaler (2001) have expressed the hope that they have written the final paper that discusses the expected utility hypothesis, referring to it as the “ex-hypothesis” with the same tone that is sometimes used in talking about an ex-spouse. Other economists like Hey (1995) maintain that the expected-utility model outperforms the alternatives, especially when decision errors are explicitly modeled in the process of estimation. In spite of this controversy, expected utility continues to be widely used, either implicitly by assuming risk neutrality or explicitly by modeling risk aversion in terms of either gains and losses or in terms of final wealth.

Some may find the mixed evidence on some of these issues to be worrisome, but to an experimentalist it provides an exciting area for further research. For example, there remains a lot of work to be done in terms of finding out how people behave in high-stakes situations. One way to run such experiments is to go to countries where using high incentives would not be so expensive. For example, Binswanger (1980) studied the choices of farmers in Bangladesh when the prize amounts sometimes involved more than a month’s salary. (He observed considerable risk aversion, which was more pronounced with very high stakes.) Similarly, Kachelmeier and Shehata (1992) performed lottery choice experiments in rural China. They found that the method of asking the question has a large impact on the way people value lotteries. For example, if you ask for a selling price (the least amount of money you would accept to sell the lottery), people tend to give a high answer, which would seem to indicate a high value for the risky lottery, and hence a preference for risk. But if you ask them for the most they would be willing to pay for a risky lottery, they tend to give a much lower number, which would seem to indicate risk aversion. The incentive

structure was such that the optimal decision was to provide a “true” money value in both treatments (much as the instructions for the previous chapter provided an incentive for people to tell the truth about their probabilities in a Bayes’ rule experiment). It seems that people go into a bargaining mode when presented with a pricing task, demanding high selling prices and offering low buying prices. The nature of this “willingness-to-pay/willingness-to-accept bias” is not well known, at least beyond the simple bargaining mode intuition provided here. Nevertheless, it is important for policy makers to be aware of the WTP/WTB bias, since studies of non-market goods (like air and water quality) may have estimates of environmental benefits that vary by 100% depending on how the question is asked. Given the strong nature of this WTP/WTB bias, it is usually advisable to avoid using pricing tasks to elicit valuations. For more discussion of this topic, see Shogren, et al. (1994).

A number of additional biases have been documented in the psychology literature on judgment and decision making. For example, there may be a tendency for people to be overconfident about their judgments in some contexts. Some of the systematic types of judgmental errors will be discussed at length in later chapters, such as the “winner’s curse” in auctions for prizes of unknown value. For further discussion of these and other anomalies, see Camerer (1995).

Questions

1. Show that a risk-neutral person would prefer a 0.8 chance of winning 4,000 to a sure payment of 3,000, and that the same person would prefer a 0.2 chance of winning 4,000 to a 0.25 chance of winning 3,000.
2. If the probability of the 4,000 payoff for lottery R in Table 7.1 is replaced by 0.7, show that a risk neutral person would prefer Lottery S.

Chapter 8. ISO (In Search of ...)

When you go out to make a purchase, you probably do not check prices at all possible locations. In fact, you might stop searching as soon as you find an acceptable offer, even if you know that a better offer would probably turn up eventually. Such behavior may be optimal if the search process is costly, which makes it worthwhile to compare the costs and benefits of additional searching. The game discussed in this chapter is a search problem, where each additional offer costs a fixed amount of money. Observed behavior in such situations is often surprisingly rational, and the classroom game can be used to introduce a discussion of optimal search. As usual, the game can be done using ten-sided dice to generate the random offers, following the instructions in the Appendix, or it can be run on the Veconlab software (select SR).

I. Introduction

Many economists believed that the rise of e-commerce would lead to dramatically lower search costs and a consequent reduction in price levels and dispersion. Evidence to date, however, suggests that there is still a fair amount of price variability on the internet, and most of us can attest to the time costs of searching for low prices. Of course, a reduction in the cost of obtaining a price quote would cause one to get more quotes before deciding on a major purchase, but the total time spent searching may go either way, since this total is the product of the cost per search and the number of searches.

Issues of search and price dispersion are central to the study of how markets promote efficiency by connecting buyers and sellers. Search is also important on a macroeconomic level, since much of what is called “frictional” unemployment is due to workers searching for an acceptable wage offer. This chapter presents a particularly simple search problem, in which a person pays a constant cost for obtaining each new observation (e.g. a wage offer or price quote). Observations are independent draws from a probability distribution that is known. From a methodological point of view, the search problem is interesting because it is dynamic, i.e. decisions are made in sequence.

II. Search from a Uniform Distribution

The setup used here is based on a particular probability framework, the uniform distribution. Many situations of interest involve equal probabilities for the relevant events. For example, consider an airport with continuously circulating buses numbered from 1 to 20, which all pass by a certain central pickup point. If there is no particular pattern, the next bus to pass might be modeled as a uniform random variable. The uniform distribution is easy to

explain since all probabilities are equal, and the theoretical properties of models with uniform distributions are often quite simple. Therefore, the uniform distribution will be used repeatedly in other parts of this book, for example, in the chapters on auctions.

A uniform distribution is easily generated with a random device, like drawing ping pong balls with different numbers written on them, so that each number is equally likely to be drawn. Even if a computer random number generator is used, it is important to provide a physical example to explain what the distribution implies. One way to explain a uniform distribution on some interval, say from 0 to 99, is to imagine a roulette wheel with 100 stops, labeled 0, 1, 2, ... 99, so that a hard spin is just as likely to stop at one point as at any other. A convenient and inexpensive way to generate the realization of a uniform distribution is with throws of a 10-sided die, which is generally labeled 0, 1, 2, ... 9. The first throw can determine the “tens” digit, and the second throw can determine the “ones” digit. In this manner, each of the 100 integers on the interval from 0 to 99 is equally likely.

Besides being easy to explain and implement, the uniform distribution has the useful property that the expected value is the midpoint of the interval. This midpoint property is due to the absence of any asymmetry around the midpoint that would pull expected value in one direction or the other. To see this, consider the simplest case, i.e. with two equally likely outcomes, 0 and 1. Since each occurs with probability $1/2$, the expected value is $(1/2)(0) + (1/2)(1) = 1/2$, which is the midpoint of the interval. For a distribution from 0 to 99 determined by the die throws, each integer in the interval has the same probability, 0.01, and the expected value is $0.01(0) + 0.01(1) + 0.01(2) \dots + 0.01(99)$, which is the midpoint, 49.5.

The search instructions for the experiment discussed in this chapter present subjects with an opportunity to purchase draws from a distribution that is uniform on the interval from 0 to 90 (pennies). This distribution is generated by a computer random number generator for the web version (SR), and it can be determined by the throws of dice by ignoring all outcomes above 90. Each draw costs 5 cents, and there is no limit on the number of draws. The subject may decline to search, thereby earning zero. If the first draw is D_1 pennies, then stopping at that point would result in earnings of $D_1 - 5$ pennies. Suppose that the second draw is D_2 . Then the options are: 1) pay another nickel and search again, 2) stop and earn $D_2 - 10$, or if going back is allowed, 3) accept the first draw and earn $D_1 - 10$. We will say that there is “recall” if going back to take any previously rejected draw is permitted; otherwise we say that there is “no recall.” Recall may not be possible in some situations, e.g. when going through the classified (or personal ISO) ads.

Some typical search sequences for one person are shown in Table 8.1. The round number is in the left column, and the draws are listed in sequence in the middle column, with the accepted draw shown in boldface type. The search cost was 5 cents, and the resulting earnings are shown in the right column. This person, who had no prior practice, seems to have stopped as soon as a draw of about 45 cents or above was obtained.

Table 8.1. Individual Search with a 5 Cent Cost

Round	Draw Sequence (cents)	Earnings
1	70	65
2	31, 43, 63	48
3	21, 8, 43, 43, 51	26
4	53	48
5	87	82
6	69	64
7	3, 8, 35, 43	23
8	85	80
9	46	41
10	22, 12, 65	50

The analysis of the optimal way to search is simplified if we assume that the person is risk neutral, which lets us determine a benchmark from which the effects of risk aversion can be evaluated. In addition, assume that the subject faces no cash constraint (wouldn't it be nice) on the number of nickels (well that's not so unreasonable) available to pay the search costs. In this case, the future always looks the same because the horizon is infinite and the payoff parameters (search cost, prize distribution) are fixed over time. Since the future opportunities are the same regardless of how many draws have been rejected thus far, any draw rejected at one point in time should never be taken at a later point. Thus each possible draw in the interval from 0 to 90 should either be always accepted or always rejected, and the boundary that separates the acceptable and unacceptable draws is a goal or *reservation value*. The person who made the choices shown in Table 8.1 seems to have a reservation value level of about 45 cents. It is useful to see a broader sample of how people behave in this situation before further analysis of the optimal manner of search.

III. Experimental Data

Figure 8.1 shows the results of a classroom experiment with two different search cost treatments. Each "search sequence" was a series of draws made until the subject stopped and accepted a draw, and there were ten such search sequences in each treatment. (The data are from the last 5 periods for the second

and third sequences, so some learning has already occurred. Using the final 5 periods takes out the inertia effects caused by slow convergence that may show up when the search cost is changed between treatments. As a result, the data in the figure are a little “cleaner” than is normally the case.)




Figure 8.1. A Classroom Sequential Search Experiment

Consider the 5-cent search cost treatment, shown on the left side of the figure, where the draw number is indicated on the horizontal axis. The marks directly over the “1st” label on the axis represent the first-round draws, one per search sequence, which are fairly uniformly distributed from 0 to 90 as would be expected. Those accepted in the first round are labeled as “+” signs, and the rejected draws are labeled as dots. Each of the dots in the “1st” column becomes a mark in the “2nd” column, either a “+” if the second draw is accepted or a dot if not. The dots in the second column become marks (plus or dot) in the third column. The majority of search sequences ended by the fourth draw, so all data that ended in the fourth period and afterwards are shown in the “4-12th” column


(nobody made more than 12 draws). As before, this last column contains plus marks for accepted draws and dots for rejected draws. Hence a draw that is rejected in period 4 and accepted in period 5 will show up as two marks in the 4-12th column. In addition, the 4-12th column contains several asterisk marks, which represent accepted draws that were initially rejected. Since the draw was both rejected and accepted, think of the asterisk as a combination of a plus sign and a minus sign. Such recalls are occasionally observed, even though they are irrational under the assumptions made above.

The divide between accepted and rejected draws for the 5-cent cost treatment is in the range from 55 to 60 cents for most people in this group. In contrast, the dramatic increase in search cost to 20 cents results a much larger acceptance region, as shown on the right side of the figure. With high search costs, there are fewer searches, and the divide is somewhere between 25 and 30 cents. The infrequency of recalls (asterisks) and the relatively clean divide between acceptance and rejection regions is roughly consistent with the notion of a reservation value or goal level. The next step is to consider whether the observed divides can be explained with simple expected value calculations of benefits and costs.

IV. Optimal Search

First, consider the low-search-cost treatment. Obviously any amount above 85 cents should be accepted, since the potential gain from searching again is at most 5 cents (if a very lucky draw of 90 is obtained) and the cost of the new draw is 5 cents. Thus there is only a very low chance ($1/91$) that the new draw will cover the cost. At the opposite extreme, a low draw, say 0, would be rejected since the cost of another draw is only 5 cents and the expected value of the next draw is 45 cents (the midpoint of the distribution from 0 to 90). Notice the use of an expected value here, which is justified by the assumption of risk neutrality. To summarize, the benefits of further search exceed the cost when the best current draw is very low, and the cost exceeds the expected benefit when the best current draw is near the top of the range. The optimal reservation prize level is found by locating the point at which the expected benefits of another search are equal to the search cost.

Suppose that the best current draw is 60. There is essentially a $2/3$ chance that the next draw is at 60 or below, in which case the net gain is 0. This situation is represented in Figure 8.2, where the flat horizontal line represents the fact that each of the draws will be observed with the same probability. The area below this line represents probability. For example, two-thirds of the area below the dashed line is to the left of 60, and one third is to the right of 60. In other words, starting from 60, there is a $1/3$ chance of obtaining better draw.



**Figure 8.2. A Uniform Distribution on the Interval [0, 90]
(With One Third of the Probability to the Right of 60)**

Next consider the expected gains from search when the best current draw is 60 (with recall and a 5-cent search cost). A draw below 60 will not produce any gain, so the expected value of the gain will have a term that is the product: $(2/3)0 = 0$. The region of gain is to the right of the vertical line at 60 in Figure 8.2, and the area below the dashed line is one third of the total area. Thus, there is essentially a $1/3$ chance of an improvement, which on average will be half of the distance from 60 to 90, i.e. half of 30. The expected value of an improvement, therefore, is 15 cents. An improvement occurs with a probability about $1/3$, so the expected improvement is approximately $(1/3)(15) = 5$ cents. Any lower current draw will produce an expected improvement from further search that is above 5 cents, and any higher current draw will produce a lower expected improvement. It follows that 60 is the reservation draw level for which the cost of another draw equals the expected improvement. The same reasoning applied to the 20-cent search cost treatment implies that the optimal reservation draw is 30 cents (see question 2).

The predicted reservation values of 60 and 30 are graphed as horizontal lines in Figure 8.1. The power of these predictions is quite clear. There is a break between the plus signs indicating acceptance and the dots indicating rejection, with little overlap. This break is slightly below the 30 and 60 levels predicted for a risk-neutral person. Actually, adding risk aversion lowers the predicted cutoff. The intuition behind risk aversion effects is clear. While a risk-neutral person is approximately indifferent between searching again at a cost of 5 cents and stopping with a draw of 60, a risk-averse person would prefer to stop with 60, since it is a sure thing. In contrast, drawing again entails the significant risk (two thirds) of having to pay the cost without getting any improvement.

One implication of optimal behavior is that the reservation value is essentially the value of having the opportunity to go through this search process. This is because one rejects all draws below the reservation value, thereby indicating that the value of playing the search game is greater than those draws.

Conversely, draws above the reservation draw level are accepted, indicating a preference for those sure amounts of money over a continuation of the search process.

To summarize, the value of continuing to search is greater than any number below the reservation value, and it is less than any number above the reservation value, so the value of continuing to search must be equal to the reservation value. This implication was tested by Schotter and Braunstein (1981), who elicited a price at which individuals would be willing to sell the option to search. In one of their treatments, the draws were from a uniform distribution on the interval from 0 to 200 with a search cost of 5 cents. The optimal reservation value for a risk neutral person is 155 (see question 3). This should be the value of being able to play the search game, assuming that there is no enjoyment derived from doing an additional search sequence after several have already been completed. The mean reported selling price was 157, and the average accepted draw was 170, which is about halfway between the theoretical reservation value of 155 and the upper bound of 200. These and other results lead to the conclusion that the observed behavior is roughly consistent with the predictions of optimal search theory in this experiment.

Further Reading and Extensions

The search game discussed here is simple, but it illustrates the main intuition behind the determination of the reservation draw level in a sequential search problem. There are many interesting and realistic variations of this problem. The planning horizon may not be infinite, e.g. when you have a deadline for finding a new apartment before you have to move out of the current one. In this case, the reservation value will tend to fall as the deadline nears and desperation takes over. In the Cox and Oaxaca (1989) experiments with a finite horizon, subjects stopped at the predicted point about three-fourths of the time, and the deviations were in the direction of stopping too soon, which would be consistent with risk aversion.

In all situations considered thusfar, the person searching has been assumed to know the probability distribution from which draws were being made. This is a strong assumption, and therefore, it is interesting to consider cases in which people learn about the distribution of draws as they search. Suppose that you think the draws will be in the range from 0 to 100, and you see a draw of 180. This would have been well above your reservation value if the distribution from which draws were made had an upper limit of 100, but now you realize that you do not know what the upper limit really is. In this case, a high draw may be rejected in order to find out whether even higher draws are possible (Cox and Oaxaca, 2000).

Although the aggregate data discussed here are roughly consistent with theoretical predictions, this leaves open the issue of how people learn to behave in this manner. Certainly they do not generally do any mathematical analysis. Hey (1981) has analyzed individual behavior in search of heuristics and adaptive patterns that may explain behavior at the individual level. Also, see Hey (1987, when he is “still searching”). He specified a number of rules of thumb, such as:

One-Bounce Rule:

Buy at least two draws, and stop if a draw received is less than the previous draw.

Modified One-Bounce Rule:

Buy at least two draws, and stop if a draw received is less than the sum of the previous draw and the search cost.

With a search cost of 5 and initial draws of 80 and 50, for example, the modified one-bounce rule would predict that the person would stop if the third draw turned out to be less than 55. The experiment involved treatments with and without recall, and with and without information about the distribution (a truncated normal distribution) from which draws were drawn. Behavior of some subjects in some rounds corresponded to one or more of these rules, but by far the most common pattern of behavior was to use a reservation draw level. In the fifth and final round, about three-fourths of the participants exhibited behavior that conformed to the optimal reservation-value rule alone.

Questions

1. The computerized experiments discussed in this chapter were set up so that draws were equally likely to be any amount on the interval from 0 to 90. When throwing 10-sided dice with numbers from 0 to 9, the throws will be 0, 1, 2, ... 99. In order to truncate this distribution at 90, is it is convenient to ignore the first throw if it is a 9, so that all of the 90 integers from 0 through 89 are equally likely, i.e. each has a probability of $1/90$. (This is the approach taken for the hand-run experiment instructions in the appendix to this chapter.) If the current draw is 60, then the chances that the next draw will be as good or better are: a $1/90$ chance of a draw of 60, for a gain of 0, a $1/90$ chance of a draw of 61 for a gain of 1 cent; a $1/90$ chance of a draw of 62 for a gain of 2; etc. Use a spreadsheet to find the expected value of the *gain* over 60 (not the expected value of the draw), and compare this with the search cost of 5 cents. Would a draw of 60 be rejected? Would a draw of 59 be rejected? Hint: make a column of integers from 60 up to 89, and then make a second column of gains over

the current draw (of either 59 or 60). Finally, weight each gain by the probability ($1/90$) of that draw, and add up all of the weighted gains to get the expected gain, which can then be compared with the search cost of 5 cents.

2. With a search cost of 20 cents per draw from a distribution that is uniform from 0 to 90, show that the person is approximately indifferent between searching again and stopping when the current best draw is 30.
3. With a search cost of 5 cents for draws from a uniform distribution on the interval from 0 to 200, show that the reservation value is about 155.
4. Evaluate the data in Table 8.1 in terms of the “One-Bounce Rule” and the “Modified One-Bounce Rule.” Do these rules work well, and if not, what seems to be going wrong?
5. (Open-ended) Rules of thumb have more appeal in limited-information situations where rational behavior is less likely to be observed. Can you think of other rules of thumb that might be good when the distribution of draws is not known and the person searching has not had time to learn much about it?

Part III. Game Theory Experiments: Treasures and Intuitive Contradictions

As one might expect, human behavior will not always conform tightly to simple mathematical models, especially in interactive situations. These models, however, can be extremely useful if deviations from the Nash equilibrium are systematic and predictable. This part presents several games in which behavior is influenced by intuitive economic forces in ways that are not captured by basic game theory. Many of the applications are taken from Goeree and Holt (2001), “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions.” The “treasures” are treatments where data conform to theory, and the contradictions are produced by payoff changes with strong behavioral effects, even though these changes do not affect the Nash predictions. Since game theory is so widely used in the study of strategic situations like auctions, mergers, and legal disputes, it is fortunate that there is progress in understanding these anomalies. We will consider a modification of the Nash equilibrium that injects some randomness or “noise” into players’ best response functions. You might think of the randomness as being due to un-modeled effects of emotions, attention lapses, partial calculations, and other factors that may vary from person to person and from time to time for the same person.

Many games involve decisions by different players made in sequence. In these situations, players who make initial decisions must predict what those with subsequent decisions will do. These sequential games, which are represented by an extensive form or “decision tree” structure, will be considered in Chapter 9.

The generalized matching pennies game in Chapter 10 shows that behavioral deviations can be strongly influenced by payoff asymmetries. Existing data indicate that the Nash equilibrium only provides good, unbiased predictions by coincidence, i.e. when the payoffs for each decision exhibit a kind of balance or symmetry. This effect is even more dramatic in the Traveler’s Dilemma game discussed in Chapter 11, where the data patterns and the Nash equilibrium may be on *opposite* ends of the range of possible decisions. Finally, the Coordination Game in chapter 12 has a whole series of Nash equilibria, and the issue is which one will have more drawing power. This game is important in developing an understanding of how people may become “stuck” in an equilibrium that is bad for all concerned.

There has been a considerable amount of research devoted to explaining why behavior in experiments might be responsive to payoff conditions that do not affect the Nash equilibrium (or equilibria). The resulting models often relax the strong game-theoretic assumptions of perfect rationality and perfect predictions of

others' decisions, assumptions which imply that one always chooses the decision with the highest expected payoff, even if payoff differences are small. Such rationality leaves no room for computation errors, emotional or impulsive actions, and other random factors that are not embodied in the formal model of the game. In a simple decision, such noise elements tend to spread decisions around the optimal choice, with the most likely choice being near the optimal level, but with some probability of error in either direction. This would create a kind of "bell curve" pattern. With strategic interaction, however, even small amounts of noise in one person's decisions may cause large biases in others' decisions, especially if there is a lot of risk. Imagine a person trying to find the highest point on a hill, where the peak is on the edge of a steep cliff. The effect of this asymmetry in risk is that any slight chance of wind gusts may cause the person to move off of the peak. In a game, the payoff peaks depend on others' decisions, i.e. the peaks *move*. In such cases, the effects of noise or small un-modeled factors may have a "snowball" effect, moving all decisions well away from the Nash predictions. One subplot in the following chapters is an analysis of the effects of introducing some randomness or noise into the analysis; this approach is presented in a series of optional sections at the end of each chapter.

Chapter 9. Multi-Stage Games in Extensive Form

The games considered in Chapters 3 and 5 each involved simultaneous decisions. Many interesting games, however, are sequential in nature, so that the first mover must try to anticipate how the subsequent decision maker(s) will react. For example, the first person may make a take-it-or-leave-it proposal for how to split a sum of money, and the second person must either accept the proposed split or reject. This is an example of an ultimatum bargaining game, which will be considered in Chapter 23. Another example is a labor-market interaction, where the employer first chooses a wage or other set of contract terms, and the worker, seeing the wage, selects an effort level. These two-stage principal-agent games are considered in Chapter 24 below. There are several common principles that are used in the analysis of such games, and these will be covered in the current chapter. Somewhat paradoxically, the first person's decision is sometimes more difficult, since the optimal decision may depend on a forecast of the second person's response, whereas the person who makes the final decision does not need to forecast the other's action if it has already been observed. In this case, we begin by considering the final decision-maker's choice, and then we typically work backward to consider the first person's decision, in a process that is known as "backward induction." Of course, the extent to which this method of backward induction yields good predictions is a behavioral question, which will be evaluated in the context of laboratory experiments.

Several of the two-stage games discussed below are the default settings for the Veconlab two-stage game program (TS). There is a separate program (CPG) for the Centipede game program, with setup options that allow any number of stages (even 100). A hand-run version of the Centipede game can be run with a roll of quarters and a tray (question 6 at the end of the chapter).

I. Extensive Forms and Strategies

This section will present a simple two-stage bargaining game, in which the first move is a proposal of how to divide \$4, which must either be \$3 for the proposer and \$1 for the other, or \$2 for each. The other player (responder) sees this proposal and either accepts, which implements the split, or rejects, which results in \$0 for each. As defined in Chapter 3, a *strategy* for such a game is a complete plan of action that covers each possible contingency. In other words, a strategy for a player tells the player what decision to make, and these instructions are so detailed that a strategy could be carried out by an employee or agent, who would never need ask questions about what to do. In this particular game, a

strategy for the proposer is which proposal to offer, i.e. (\$3, \$1) or (\$2, \$2), where the proposer's payoff is listed on the left. A responder's strategy must specify a reaction to each of these proposals. Thus the responder had a number of options: a) accept both, b) reject both, c) accept (\$3, \$1) and reject (\$2, \$2), or d) reject (\$3, \$1) and accept (\$2, \$2). These strategies will be referred to as: AA, RR, AR, and RA respectively, where the first letter indicates the response, Accept or Reject, to the unequal proposal. Similarly, the proposer's strategies will be referred to as Equal (equal split) or Unequal (unequal, favoring the proposer). A Nash equilibrium is a pair of strategies, one for each player, with the property that neither can increase their own payoff by deviating to a different strategy under the assumption that the other player stays with the equilibrium strategy.

One Nash equilibrium for this game is (Equal, RA), where the proposal is for an equal split, and the responder rejects the unequal proposal and accepts the equal proposal. Deviations are not profitable, since the responder is already receiving the proposal that offers the higher of the two possible payoffs, and the proposer would not want to switch to the unequal proposal given that the responder's strategy requires that it be rejected. This Nash equilibrium, however, involves a threat by the responder to reject an unequal proposal (\$3 for the proposer, \$1 for the responder), thereby causing both to earn \$0. Such a rejection would reduce the responder's payoff from \$1 to \$0, so it violates a notion of sequential rationality that requires play to be rational in all parts of the game ("subgames"). The reader might wonder how this kind of irrationality can be part of a Nash equilibrium, and the answer is that the simple notion of a Nash equilibrium only requires rationality in terms of considering unilateral deviations from the equilibrium. In the equilibrium being considered, (Equal, RA), the unequal proposal is not made, so the responder never has to carry out the threat to reject this proposal. Selten (1965) proposed that the rationality requirement be expanded to cover all parts of the game, which he defined as "subgames." He called the resulting equilibrium a *subgame perfect* Nash equilibrium. Obviously, not all Nash equilibria are subgame perfect; as the equilibrium (Equal, RA) demonstrates.

Next consider the idea of starting in the final stage, where the responder is considering a specific proposal. Each of the possible proposals involve strictly positive amounts of money for the responder, so a rational responder who only cares about own payoff would accept either proposal. This analysis requires that the responder accept either proposal, so the only responder strategy that satisfies rationality in all subgames is AA. Next consider the first stage. If the proposer predicts that the responder is rational and will accept either offer, then the proposer will demand the larger share, so the equilibrium will be (Unequal, AA). This is a Nash equilibrium, and from the way it was constructed, we know that it satisfies sequential rationality, i.e. is subgame perfect.

The reasoning process that was used in the previous paragraph is an example of backward induction, which simply means starting at the end and working backwards toward the beginning when analyzing a problem. The concepts, backward induction, sequential rationality, and subgame perfection, can all be defined more precisely, but the goal here is for the reader to obtain an intuitive understanding of the principles involved, based on specific examples. This understanding is useful in constructing predictions for outcomes of simple multi-stage games, which can serve as benchmarks for evaluating observed behavior in experiments. As we shall see, these benchmarks do not always yield very good predictions, for a variety of reasons to be discussed below.

Before proceeding with examples, it is instructive to show how the two-stage bargaining game would be represented as a decision tree, which is known as the “extensive form” of a game. In Figure 9.1, the game begins at the top, at the node that is labeled “Proposer.” This player either chooses an Equal proposal (\$2,\$2), or an Unequal proposal (\$3,\$1). The Responder has a decision node following each of these proposals. On the left, the Responder may either choose A (accept) or R (reject) in response to the Equal proposal. The same options are also available on the right. Notice that each decision node is labeled with the name of the player who makes a decision at that node, and each branch emanating from a node is labeled with the name of one of the feasible decisions. The branches end at the bottom (terminal nodes), which show the payoffs, with the payoff of the Proposer listed on the left.




Figure 9.1. An Extensive Form Representation for the Bargaining Game

II. Two-Stage Trust Games

The first two experiments to be considered are based on the two games shown in Figure 9.2. Consider the top part, where the first player must begin by making a safe choice (S) or a risky choice (R). Decision S is safe in the sense that the payoffs are deterministic: 80 for the first player and 50 for the second (the first

player's payoff will be listed on the left in each case). The payoffs that may result from choosing R depend on the second player's response, P or N. For the game shown in the top part of Figure 9.2, this is an easy choice, since the second player would earn 10 from choosing P and 70 from choosing N. The first player should make the risky choice R as long as the first player trusts the second player to be rational. This game was played by once (no repetition) by pairs of subjects, where the payoffs were in pennies (see Goeree and Holt, 2001, for details). As indicated in the top part of the figure, 84 percent of the first movers were confident enough to choose R, and all of these received the anticipated N response. Note that the second movers would have reduced their earnings by 60 cents if any of them had selected P in response to R.

The game shown in the bottom part of Figure 9.2 is almost identical, except that the second mover only loses 2 cents by making a P response, which generates a payoff of 68 for the second mover instead of 70. In this case, more than half of the first movers were sufficiently concerned about the possibility of a P response that they chose decision the safe decision, S. This fear was well founded, since a fourth of the first movers who chose R encountered the P response. In fact, the first players who chose S in the bottom game in Figure 9.2 earned more on average than those who chose R (question 1).




Figure 9.2. A Two-Stage Game Where Mistakes Matter

Source: Goeree and Holt (2001)

The standard game-theoretic analysis of the two games in Figure 9.2 would yield the same prediction, since this analysis is based on the assumption that both players are rational and are concerned with maximizing own payoff. This assumption implies that the second player will choose N, regardless of whether it increases earnings by 60 cents, as in the top game, or by only 2 cents, as in the bottom game. In each of these games, we begin by analyzing the second player's best choice (choose N), and knowing this, we can calculate the best choice for the first player (choose R). The pair of choices, N and R, constitutes a Nash equilibrium, since neither player can increase earnings by deviating.

The Nash outcome (R and N) that was identified for the games in Figure 9.2 involves little tension in the sense that it maximizes the payoffs for each player. There is another Nash outcome, (S and P), which can be verified by showing that neither person has an incentive to deviate unilaterally (question 2). In particular, the second player earns 50 in this equilibrium, and a unilateral deviation from P to N would not change the outcome, since the outcome is fully determined by the first player's S decision. This second equilibrium is typically "ruled out" because it implies a type of behavior for the second player that is not rational in a sequential sense. Despite these arguments, the S outcome predicted by second equilibrium is the most commonly observed outcome for the bottom game in Figure 9.2.

There is more tension for the games shown in Figure 9.3. As before, there are two Nash equilibria: (S and P) and (R and N). First consider the (R and N) outcome, which yields 90 for the first player and 50 for the second. If the second player were to deviate to P, then the outcome would be (R and P) with lower payoffs for the second player, since we are considering a unilateral deviation, i.e. a deviation by one player given that the other continues to use the equilibrium decision. Similarly, given that the second player is using N, the first player cannot deviate and increase the payoff above 90, which is the maximum for the first player in any case. This equilibrium, R and N, is the one preferred by the first player. There is another equilibrium, S and P (question 3). Here, you can think of P as indicating "punishment," since a P decision in the second stage reduces the first player's payoff from 90 to 60. The reason that such a punishment might be enacted is that the second player actually prefers the outcome on the left side of the figure for each of the games in Figure 9.3. The difference between the two games in the figure is that the cost of punishment is 40 cents for the top game, and only 2 cents for the bottom game.

The standard game-theoretic analysis of the games in Figure 9.3 would be the same, i.e. that the (S and P) equilibrium is implausible since it implies that the second mover be ready to use a punishment that is costly, and hence not credible.

So the predicted outcome would be the equilibrium (R and N) in each case. This prediction is accurate for the game shown in the top part of the figure, where 88 percent of the outcomes are as predicted. The outcomes deviate sharply from this prediction for the game shown in the bottom part of the figure, since only about a third of the observed outcomes are at the (R and N) node.




Figure 9.3. A Two-Stage Game With a Threat That Is “Not Credible”
Source: Goeree and Holt (2001)

To summarize, the notion of sequential rationality rules out the possibility that the second mover will make a mistake in either of the games in Figure 9.2, or that the second mover will carry out a costly (and hence non-credible) threat for either of the games in Figure 9.3. Thus the predicted outcome would be the same (R and N) in all four games. This prediction works well for one of the games in each figure, but not for the other one. The problem with the notion of sequential rationality being used is that it does not allow room for mistakes, which are more likely if the cost of the mistake is small (2 cents for each of the bottom games in the figures). A more systematic analysis of the costs of mistakes and the effects on behavior will be undertaken in the three chapters that follow.

III. The Centipede Game

In the two-stage games considered thusfar, the first mover typically has to consider how the second mover will react to the initial decision. This thought process may involve the first mover thinking introspectively: “What would I do on the second stage, having just seen this particular initial decision?” This process of seeing a future situation through the eyes of someone else may not be easy or natural, and the difficulty will increase greatly if there are more than two stages, so an initial decision maker might have to think about reactions several stages later. In such cases, it is still often straightforward, but tedious, to begin with the last decision-maker, consider what is rational for that person, and work backwards to the second-to-last decision maker, and then to the third-to-last decision maker, etc. in a process of backward induction. The “noise” observed in some the two-stage games already considered in this chapter makes it plausible that players in multi-stage games may not be very predictable, at least in terms of their ability to reason via backward induction. Rosenthal (1982) proposed a sequential game with 100 stages, which serves as an extreme “stress test” of the backward induction reasoning process. Each of the 100 stages has a terminal node in the extensive form that hangs down, so the extensive form looks like an insect with 100 legs, which is why this game is known as the “centipede game.”

Figure 9.4 shows a truncated version of this game, with only 4 legs. At each node, the player (Red or Blue) can either move down, which stops the game, or continue to the Right, which passes the move to the other player, except in the final stage. The play begins on the left side, where the Red player must decide whether to stop the game (the downward arrow) or continue (the right arrow). A stop/down move in this first stage results in payoffs of 40 cents for Red and 10 cents for the other player, Blue. Note that the payoffs for this first terminal node are listed as (40, 10), with payoffs for the Red player on the left. If Red decides to continue to the right, then the next move is made by Blue, who can either continue or stop, yielding payoffs of (20, 80), where the left payoff is for Red, as before. Blue might reason that stopping gives 80 for sure, whereas a decision to continue will either yield payoffs of 40 (if Red stops in the next stage), 320 (if Blue stops by moving down in the final stage), or 160 (if Blue moves right in the final stage). Obviously, it is better for Blue to move down in the final stage, getting 320 instead of 160. Having figured out the best decision in the final stage, we can begin the process of backward induction. If Blue is expected to go down in the final stage, then Red will expect a payoff of 80 instead of 640 if play reaches the final stage. Thus it would be better for Red to take the 160 and stop the game in the third stage, rather than pass to Blue and end up with only 80 in the final stage.

To summarize thusfar, we have concluded that Blue will stop in the final stage, and that Red, anticipating this, will stop in the third stage. Similar reasoning can be used to show that Blue will want to stop in the second stage, if play reaches this point, and hence that Red will want to stop in the first stage, yielding payoffs of 40 for Red and only 10 for Blue. Thus the process of backward induction yields a very specific prediction, i.e. that the game will stop in the initial stage, with relatively low payoffs for both.




Figure 9.4. A Centipede Game
Source: McKelvey and Palfrey (1992)

The experiment, conducted by McKelvey and Palfrey (1992), involved groups of 20 participants, with 10 people in each role (Red or Blue). Each participant was matched with all 10 people in the other role, in a series of 10 games, each game being the centipede game shown in Figure 9.4. The aggregate percentages for each outcome are shown in the figure, below the payoffs at each node. Only 8% of the Red players stopped in the first stage, so the experiment provides a sharp rejection of predictions based on backward induction. Rates of continuation were somewhat lower toward the end of the session, but the incidence of stopping in the initial stage remained relatively low throughout. Most of the games ended in the second and third stages, but 12% of the games went on until the final stage.

The failure of backward induction predictions in this game is probably due to the fact that even a small amount of unpredictability of others' decisions in later stages may make it optimal to continue in early stages, since the high payoffs are concentrated at the terminal nodes on the right side of the figure. McKelvey suggested that a small proportion of people were somewhat concerned with others' payoffs ("altruists"), and that even if these types are not predominant, it would be optimal for selfish individuals to continue in early stages. In fact, any type of noise or unpredictability, regardless of its source, will have similar effects. Notice that about one in six Blue players continue in the final stage, which cuts their payoff in half but raises the Red player's payoff by a factor of 8, from 80 to 640. Whether this is due to altruism, to randomness from another cause, or to

miscalculation is an open question. For an analysis of the effects of randomness in centipede games, see Zauner (1999).

IV. Extensions

One possible reaction to the centipede game experiment is that the failure of backward induction may be due to low incentives, and that behavior would be more “rational” in high stakes experiments. Recent experiments with potential stakes of thousands of dollars indicate that this is not the case. Other recent work has considered how subjects in experiments learn and adjust in centipede games, see Nagel and Tang (1998) and Ponti (2000). Bornstein, Kugler, and Ziegelmeyer (2002) compare behavior of individuals and groups in centipede games. Fey, McKelvey and Palfrey (1996) report results of a variation of the centipede game, where the payoffs sum to a constant.

Questions

1. Show that the first players who chose R earned more on average than those who chose S in the game shown in the bottom part of Figure 9.2. (Hint: you have to use the response percentages shown below each outcome.)
2. Show that (S for the first player, P for the second player) is a Nash equilibrium for the games in Figure 9.2. To do this, you have to check to see whether a unilateral deviation by either player will not increase that players' payoff, under the assumption that the other player stays with the equilibrium decision.
3. Show that (S for the first player, P for the second player) is a Nash equilibrium for each of the games in Figure 9.3; i.e. check the profitability of unilateral deviations for each player.
4. Consider a game in which there is \$4 to be divided, and the first mover is only permitted to make one of three proposals: a) \$3 for the first mover and \$1 for the second mover, b) \$2 for each, or c) \$1 for the first mover and \$3 for the second mover. The second mover is shown the proposal and can either accept, in which case it is implemented, or reject and cause each to earn \$0. Show this game in extensive form. (Hint, the decision node for the first mover has three arrows, one for each decision, and each of these arrows leads to a node with two arrows. Be sure and show the payoffs for each person, with the first mover listed on the left, for each of the 6 terminal nodes.)
5. Show that proposal (c) in question 4 is a part of a Nash equilibrium for the game in question 4. To finish specifying the decisions for this equilibrium, you have to say what the second mover's response is to each

of the three proposals. Does behavior in this equilibrium satisfy sequential rationality?

6. Consider a game in which the instructor divides the class into two groups and takes a \$10 roll of quarters. The instructor then puts a quarter into a collection tray and allows those on one side of the room to take the quarter (and somehow divide or allocate it among themselves) or to pass. A pass results in doubling the money (to 2 quarters) and giving those on the other side of the room the option to take or pass. This process of doubling the number of quarters in the tray continues until one side takes the quarters, or until one side's pass forces the instructor to put all remaining coins into the plate, which are then given to the side of the room whose decision would be the next one. Draw a figure that shows the extensive form for this game. Label the players A and B, and show payoffs for each terminal node as an ordered pair: (\$ for A, \$ for B). What is the predicted outcome of the game, on the basis of backward induction, perfect rationality, and selfish behavior?

Chapter 10. Generalized Matching Pennies

There are many situations in which a person does not want to be predictable. The equilibrium in such cases involves randomization, but to be willing to randomize, each player must be indifferent between the decisions over which they are randomizing. In particular, each player's decision probabilities have to keep the *other* player indifferent. For this reason, changes in a player's own payoffs are not predicted to affect the player's own decisions. This counter-intuitive feature is contradicted by data from matching games with payoff imbalances, where "own-payoff effects" are systematic. These games can be run with minor modifications to the instructions used for Chapter 5, or with the program MG.

I. The Case of Balanced Payoffs

In a classic game of matching pennies, each person places a coin on a table, covering it so that the other person cannot see which side is up. By prior agreement, one person takes both pennies if the pennies match, and the other takes the pennies if they do not match. This is analogous to a soccer penalty kick, where the goalie must dive to one side or another before it is clear which way the kick will go, and the kicker cannot see which way the goalie will dive at the time of the kick. In this case, the goalie wants a match and the kicker wants a mismatch. Table 10.1 shows a game where the Row player prefers to match:


Table 10.1. A Modified Matching Pennies Game
(row's payoff, column's payoff)

Row Player:	Column Player:	
	Left (heads)	Right (tails)
Top (heads)	72, 36 \Rightarrow	\Downarrow 36, 72
Bottom (tails)	36, 72 \Uparrow	\Leftarrow 72, 36

(It can be shown that this game is really equivalent to a matching pennies game in which there is always one player with a penny gain and another with a penny loss; see question 1). In each of the cells of the payoff table, there is one person who would gain by altering the placement of their penny unilaterally, as indicated by the arrows. For example, if they were both going to choose Heads, then the column player would prefer to switch to Tails, as indicated by the arrow pointing to the right in the upper/left box of the payoff table. The arrows in each box indicate the direction of a unilateral payoff-increasing move, so there is no equilibrium in non-random strategies. (Non-random strategies are commonly

called “pure strategies” because they are not probability-weighted mixtures of other strategies.)

The game in Table 10.1 is *balanced* in the sense that each possible decision for each player has the same set of possible payoffs, i.e. 36 and 72. In games with this type of balance, the typical result is for subjects to play each decision with an approximately equal probability (Ochs, 1994; Goeree and Holt, 2001). This tendency to choose each decision with probability one half is not surprising given the simple intuition that one must not be predictable in this game. Nevertheless, it is useful to review the representation of the Nash equilibrium as an intersection of players’ best response functions, as explained in Chapter 5. First consider the Row player, who will choose Top if Column is expected to choose Left. Since Row’s payoffs are 72 and 36 in the top part of Table 10.1 and are 36 and 72 in the bottom part, it is apparent that Row would choose Top as long as Left is thought to be more likely than Right. This best-response behavior is represented by the solid line in Figure 10.1 that begins in the top/left corner, continuing along the top until there is an abrupt drop to the bottom when the probability of Right is 0.5. (Please ignore the curved line for now.)



**Figure 10.1. Best Response Functions for a Symmetric Matching Pennies Game:
The Effect of Noisy Behavior**

The column player's best response line is derived analogously and is shown by the thick dashed line that starts in the lower/left corner, rises to 0.5 and then crosses over horizontally to the right side, because Column would want to play Right whenever the probability of Top is greater than 0.5. The intersection of these two lines is in the center of the graph, at the point where each probability is 0.5, which is the Nash equilibrium in mixed strategies.

II. Noisy Best Responses

The previous analysis is not altered if we allow a little “noise” in the players' responses to their beliefs. Noisy behavior of this type was noticed by psychologists who would show subjects two lights and ask which was brighter, or let them hear two sounds and ask which was louder. When the signals (lights or sounds) were not close in intensity, almost everybody would indicate the correct answer, with any errors being caused by mistakes in recording decisions. As the two signals became close in intensity, then some people would guess incorrectly, e.g. because of bad hearing, random variations in ambient noise, or distraction and boredom. As the intensity of the signals approached equality, the proportions of guesses for each signal would approach 1/2 in the absence of measurement bias. In other words, there was not a sharp break where the stronger signal was selected with certainty, but rather a “smooth” tendency to guess the stronger signal more as its intensity increased. The probabilistic nature of this behavior is captured in the phrase: “probabilistic choice” or “noisy best response.”

This probabilistic-choice perspective can be applied to the matching pennies game. The intuitive idea is that Row will choose Top with high probability when the expected payoff for Top is a lot larger than the expected payoff for Bottom, but that some randomness will start to become apparent when the expected payoffs for the two decisions are not that far apart. To proceed with this argument, we begin by calculating these expected payoffs. Let p denote Row's beliefs about the probability of Right, so $1-p$ is the probability of Left. From the top row of Table 10.1, we see that if Row chooses Top, then Row earns 72 with probability $1-p$ and 36 with probability p , so the expected payoff is:

$$\text{Row's Expected Payoff for Top} = 72(1-p) + 36(p) = 72 - 36p.$$

Similarly, by playing Bottom, Row earns 36 with probability $1-p$ and 72 with probability p , so the expected payoff is:

$$\text{Row's Expected Payoff for Bottom} = 36(1-p) + 72(p) = 36 + 36p.$$

It follows that the difference in these expected payoffs is: $(72 - 36p) - (36 + 36p)$, or equivalently:

$$\text{Row's Expected Payoff Difference (for Top Minus Bottom)} = 36 - 72p.$$

When p is near zero, this difference is 36, but as p approaches $1/2$ this difference goes to zero, in which case Row is indifferent and would be willing to choose either decision or to flip a coin.

If Row were perfectly rational (and responded to arbitrarily small expected payoff differences), then Top would be played whenever p is even a little less than $1/2$. The curved line in Figure 10.1 shows some departure from this rationality. Notice that the probability that Row plays Top is close to 1 but not quite there when p is small, and that the curved line deviates more from the best-response line as p approaches $1/2$.

The curved “noisy best-response line” for Row intersects Column’s dashed best-response line in the center of Figure 10.1, so a relaxation of the perfect rationality assumption for Row will not affect the equilibrium prediction. Similarly, suppose that we allow some noise in Column’s best response, which will “smooth off” the sharp corners for the thin dashed line that represents Column’s best responses, as shown by the upward sloping dashed line on the right side of Figure 10.2. This starts in the lower-left corner, rises with a smooth arc as it levels off at 0.5 in the center of the graph before curving upward along the upper-right boundary of the graph. Since this line will intersect Row’s downward sloping noisy response line in the center of the graph, we see that the “fifty-fifty” prediction is not affected if we let each player’s decision be somewhat noisy.




Figure 10.2. Best Responses (Left Side) and Noisy Best Responses (Right Side)

Finally, you should think about what would happen if the amount of randomness in behavior were somehow reduced. This would correspond to a situation where each person takes full advantage of even a slight tendency for the other to choose one decision even slightly more often than the other. These sharp responses to small probability differences would cause the curved lines on the right to become more like the straight-line best response functions, i.e. the “corners” on in the curved lines would become sharper. In any case, the intersection would remain in the center, so adding noise has no effect on predictions in the symmetric matching pennies game.

III. The Effects of Payoff Imbalances

As one would expect, the fact that the noisy best response lines always intersect in the center of Figure 10.2 is due to the balanced nature of the payoffs for this game (Table 10.1). An unbalanced payoff structure is shown in Table 10.2, where the Row player’s payoff of 72 in the Top/Left box has been increased to 360. Recall that the game was balanced before this change, and the choice proportions should be 1/2 for each player. The increase in Row’s Top/Left payoff from 72 to 360 would make Top a more attractive choice for a wide range of beliefs, i.e. Row will choose Top unless Column is almost sure to choose Right. Intuitively, one would expect that this change would move Row’s choice proportion up from the 1/2 level that is observed in the balanced game. This intuition is apparent in the choice data for an experiment done with Veconlab software, where the payoffs were in pennies, and with random matching. Each of the three sessions involved 10-12 players, with 25 periods of random matching. The proportion of Top choices is 67% for the “360 treatment” game in Table 10.2.

Table 10.2. An Asymmetric Matching Pennies Game
(Row’s payoff, Column’s payoff)

Row Player:		Column Player:	
		Left	Right
Top		360, 36 \Rightarrow	\Downarrow 36, 72
Bottom		36, 72 \Uparrow	\Leftarrow 72, 36


This intuitive “own payoff effect” of increasing Row’s Top/Left payoff is consistent with the Nash equilibrium prediction. First, notice that there is no equilibrium in non-random strategies, as can be seen from the arrows in Table 10.2, which go in a clockwise circle. To derive the mixed equilibrium prediction,

let p denote Row's beliefs about the probability of Right, so $1-p$ is the probability of Left. Thus Row's expected payoffs are:

$$\text{Row's Expected Payoff for Top} = 360(1-p) + 36(p) = 360 - 324p.$$

$$\text{Row's Expected Payoff for Bottom} = 36(1-p) + 72(p) = 36 + 36p.$$


It follows that the difference in these expected payoffs is: $(360 - 324p) - (36 + 36p)$, or equivalently, $324 - 360p$. Row is indifferent if the expected payoff difference is zero, i.e. if $p = 324/360 = 0.9$. Therefore, Row's best response line stays at the top of the left panel in Figure 10.3 as long as the probability of Right is less than 0.9. The striking thing about this figure is that Column's best response line has not changed from the symmetric case; it rises from the Bottom/Left corner and crosses over when the probability of Top is $1/2$. This is because Column's payoffs are exactly reflected (36 and 72 on the Left side, 72 and 36 on the right side). In other words, the only way that Column would be willing to randomize is if Row chooses Top and Bottom with equal probability.



**Figure 10.3. The 360 Treatment:
Best Responses (Left Side) and Noisy Best Responses (Right Side)**

The previous treatment change made Top more attractive for Row by raising Row's Top/Left payoff from 72 to 360. In order to produce an imbalance that makes Top *less* attractive, we reduce Row's Top/Left payoff in Table 10.1 from 72 to 40. Thus Row's expected payoff for Top is: $40(1-p) + 36(p) = 40 - 4p$, and the expected payoff for Bottom is: $36(1-p) + 72(p) = 36 + 36p$. These are equal when $p = 4/40$, or 0.1. What has happened in the left panel of Figure

10.3 is that the reduction in Row's Top/Left payoff has pushed Row's best response line to the bottom of the figure, unless Column is expected to play Right with a probability that is less than 0.1. This downward shift in Row's best response line is shown on the left side of Figure 10.4. The result is that the best response lines intersect where the probability of Top is 1/2, which was the same prediction obtained from the left sides of Figures 10.2 and 10.3. Thus a change in Row's Top/Left payoff from 36 to 72 to 360 does not change the Nash equilibrium prediction for the probability of Top. The mathematical reason for the result is that Column's payoffs do not change, and Row must choose each decision with equal probability or Column would not want to choose randomly.



**Figure 10.4. The 40 Treatment:
Best Responses (Left Side) and Noisy Best Responses (Right Side)**

These invariance predictions are not borne out in the data from the Veconlab experiments reported here. Each of the three sessions involved 25 periods for each treatment, with the order of treatments being alternated. The percentage of Top choices increased from 24% to 74% when the Row's Top/Left payoff was increased from 40 to 360. The Column players reacted to this change by choosing Right only 36% of the time in the 40 treatment, and 67% of the time in the 360 treatment.

The qualitative effect of Row's own-payoff effect is captured by a noisy-response model where the sharp-cornered best response functions on the left sides of Figures 10.3 and 10.4 are rounded off, as shown on the right sides of the figures. Notice that the intersection of the curved lines implies that Row's proportion of Top choices will be below 1/2 in the 40 treatment and above 1/2 in

the 360 treatment, and the predicted change in the proportion of Right decisions will not be as extreme as the movement from 0.1 to 0.9 implied by the Nash prediction. The actual data averages are shown by the black dots on the right sides of Figures 10.3 and 10.4. These data averages exhibit the strong own-payoff effects that are predicted by the curved best-response lines, and the prediction is fairly accurate, especially for the 40 treatment.

The quantitative accuracy of these predictions is affected by the amount of curvature that is put into the noisy response functions, and is ultimately a matter of estimation. Standard estimation techniques are based on writing down a mathematical function with a “noise parameter” that determines the amount of curvature and then choosing the parameter that provides the best fit. The nature of such a function is discussed next.

IV. Probabilistic Choice (Optional)

Anyone who has taken an introductory psychology course will remember the little stimulus-response diagrams. Stochastic or noisy response models were developed after researchers noticed that responses could not always be predicted with certainty. The work of a mathematical psychologist, Duncan Luce (1959), suggested a way to model noisy choices, i.e. by assuming that response probabilities are increasing functions of the strength of the stimulus. For example, suppose that a person must judge which of two very faint sounds is loudest. The probability associated with one sound should be an increasing function of the decibel level of that sound. Since probabilities must sum to 1, the choice probability associated with one sound should also be a *decreasing* function of the intensity of the *other* sound.

In economics, the stimulus intensity associated with a given response (decision) might be thought of as the expected payoff of that decision. Suppose that there are two decisions, D_1 and D_2 , with expected payoffs that we will represent by E_1 and E_2 . For example, the decisions could be Row’s choice of Top or Bottom, and the expected payoffs would be calculated using the equations given earlier, e.g. $E_1 = 360(1-p) + 36p$, where p represents the probability which the Row player thinks Column will choose Right. Instead of using the mathematical expressions for these expected payoffs, we will just use the shorthand notation, π_1 and π_2 , which is both simpler and more general since it applies to any two decisions. (Think of π , the Greek letter pronounced like “pie,” as shorthand for payoff.)

Using Luce’s suggestion, the next step is to find an increasing function, i.e. one with a graph that “goes uphill.” (Mathematically speaking, a function, f , is increasing if a $f(\pi_1) > f(\pi_2)$ whenever $\pi_1 > \pi_2$, but the intuitive idea is that the slope in a graph is from lower-left to upper-right.) Several examples of increasing functions are the linear function: $f(\pi_1) = \pi_1$, or the exponential function, $f(\pi_1) =$

$\exp(\pi_1)$. The linear function obviously has an uphill slope; its graph is just the 45-degree line. The exponential function has a curved shape, like a hill that keeps getting steeper, like a snow-capped Mt. Fuji in a Japanese woodblock print.

Once we have found an increasing function, we might be tempted to assume that the probability of the decision is determined by the function itself: $\Pr(D_1) = f(\pi_1)$ and $\Pr(D_2) = f(\pi_2)$. The problem with this approach is that it does not ensure that the two probabilities sum to 1. This is easily fixed by a simple trick with a fancy name: “normalization.” Just divide each function by the sum of the functions:

$$(10.1) \quad \Pr(D_1) = \frac{f(\mathbf{p}_1)}{f(\mathbf{p}_1) + f(\mathbf{p}_2)} \quad \text{and} \quad \Pr(D_2) = \frac{f(\mathbf{p}_2)}{f(\mathbf{p}_1) + f(\mathbf{p}_2)}.$$

If you are feeling a little unsure, try adding the two ratios in (10.1) to show that $\Pr(D_1) + \Pr(D_2) = 1$. Now let’s see what this gives us for the linear case:

$$(10.2) \quad \Pr(D_1) = \frac{\mathbf{p}_1}{\mathbf{p}_1 + \mathbf{p}_2} \quad \text{and} \quad \Pr(D_2) = \frac{\mathbf{p}_2}{\mathbf{p}_1 + \mathbf{p}_2}.$$

Suppose $\pi_1 = \pi_2 = 1$. Then each of the probabilities in (10.1) will equal $1/(1+1) = 1/2$. This result holds as long as the expected payoffs are equal, even if they are not both equal to 1. This makes sense; if each decision is equally profitable, then there is no reason to prefer one over the other, and the choice probabilities should be $1/2$ each. Next notice that if $\pi_1 = 2$ and $\pi_2 = 1$, the probability of choosing decision D_1 is higher ($2/3$), and this will be the case whenever $\pi_1 > \pi_2$.

One potential limitation to the usefulness of the payoff ratios in (10.2) is that expected payoffs may be negative if losses are possible. This problem can be avoided if we choose a function in (10.1) that cannot have negative values, i.e. $f(\pi_1) > 0$ even if $\pi_1 < 0$. This non-negativity is characteristic of the exponential function, which can be used in (10.1) to obtain:

$$(10.3) \quad \Pr(D_1) = \frac{\exp(\mathbf{p}_1)}{\exp(\mathbf{p}_1) + \exp(\mathbf{p}_2)} \quad \text{and} \quad \Pr(D_2) = \frac{\exp(\mathbf{p}_2)}{\exp(\mathbf{p}_1) + \exp(\mathbf{p}_2)}.$$

This avoids the possibility of negative probabilities, and all of the other useful properties of (10.2) are preserved. The probabilities in (10.3) will sum to one, they will be equal when the expected payoffs are equal, and the decision with the higher expected payoff will have a higher choice probability. The probabilistic choice model that is based on exponential functions is known as the *logit model*.

The curved lines in the right parts of the figures in this chapter were constructed using logit choice functions, but not quite the ones in equation (10.3). It is true that (10.3) applied to the expected payoffs for the matching pennies games will produce curved response functions, but the lines will not have as much curvature as those in the figures in this chapter. The lines drawn with equation (10.1) have corners that are too sharp to explain the “own payoff effects” that we are seeing in the matching pennies games. Just as we could make it harder to distinguish between the width of two pins by making them each half as thick, we can add more noise or randomness into the choice probabilities in (10.3) by reducing all expected payoffs by a half or more. Intuitively speaking, dividing all expected payoffs by 100 may inject more randomness, since dollars become pennies, and non-monetary factors (boredom, indifference, playfulness) may have more influence. The right panels of the figures in this chapter were obtained by using the logit model with all payoffs being divided by 10:

$$(10.4) \quad \Pr(D_1) = \frac{\exp(\mathbf{p}_1/10)}{\exp(\mathbf{p}_1/10) + \exp(\mathbf{p}_2/10)}, \quad \Pr(D_2) = \frac{\exp(\mathbf{p}_2/10)}{\exp(\mathbf{p}_1/10) + \exp(\mathbf{p}_2/10)}.$$

At this point, you are probably wondering, why 10, why not 100? The degree to which payoffs are “diluted” by dividing by larger and larger numbers will determine the degree of curvature in the noisy response functions. Thus we can think of the number in the denominator of the expected payoff expressions as being an error parameter that determines the amount of randomness in the predicted behavior. The *logit error parameter* will be called μ , and it is used in the logit choice probabilities:

$$(10.5) \quad \Pr(D_1) = \frac{\exp(\mathbf{p}_1/\mu)}{\exp(\mathbf{p}_1/\mu) + \exp(\mathbf{p}_2/\mu)}, \quad \Pr(D_2) = \frac{\exp(\mathbf{p}_2/\mu)}{\exp(\mathbf{p}_1/\mu) + \exp(\mathbf{p}_2/\mu)}.$$

The logit form in (10.5) is flexible, since the degree of curvature is captured by a parameter that can be estimated. It also has the intuitive property that an increase in the payoffs will reduce the “noise,” i.e. will reduce the curvature of the noisy best response curves in the figures (see question 4).

Extensions and Further Reading

The use of probabilistic choice functions in the analysis of games was pioneered by McKelvey and Palfrey (1995), and the intersections of noisy best response lines in Figures 10.2-10.4 correspond to the predictions of their *quantal response equilibrium*. (A “quantal response” is essentially the same thing as a

noisy best response.) The approach taken in this chapter is based on Goeree, Holt, and Palfrey (2004), who introduce the idea of a regular quantal response equilibrium, and who discuss its empirical content. Goeree, Holt, and Palfrey (2000) use this approach to explain behavior patterns in a number of matching pennies games. Their analysis also includes the effects of risk aversion, which can explain the over-prediction of own-payoff effects for high payoffs that is seen on the right side of Figure 10.3. Risk aversion introduces diminishing marginal utility that reduces the attractiveness of the large 360 payoff for the row player, and this shifts that person's best response line down so that the intersection is closer to the data average point.

All of the games considered in this chapter involved two decisions, but the same principles can be applied to games with more decisions (see Goeree and Holt, 1999; Capra, Goeree, Gomez, and Holt, 1999, 2001, which will be discussed in the next two chapters). Of particular interest is Composti (2003), who had University of Virginia soccer players attempt penalty kicks into a net that was divided into three sectors by hanging strips of cloth. Thus the kicker had three decisions corresponding to the sector of the kick, say L , M , and R . Similarly, the goalie had three decisions, corresponding to the direction of the defensive dive. In a symmetric treatment, the kicker received \$3 for each point scored, regardless of sector, and the goalie received \$3 each time a kick was unsuccessful for any reason. The penalty kicks were approximately evenly divided between the two sides, 13 attempts to L and 11 attempts to R , with only 6 to the middle. In an asymmetric treatment, the kickers received \$6 for each kick scored on the R side, and \$3 for each kick that scored to the M or L sides. The goalie's incentives did not change (\$3 for each kick that did not score, regardless of direction). This asymmetry in the kicker's payoffs caused a strong own-payoff effect, with 25 attempts to R , 5 to L , and none to the middle. This study exhibits the advantages of an experiment, where asymmetries can be introduced in a manner that is not possible from an analysis of data from naturally occurring soccer games.

The focus in this chapter has been on games, but probabilistic choice functions can be applied to simple decision problems. For example, recall that a risk neutral person would make exactly 4 safe choices in the lottery choice menu presented in Chapter 4, and a person with "square root utility" would make 6 safe choices. These predictions produce lines with sharp corners, e.g. the dashed line in Figure 4.2, which looks a lot like the curved lines in the figures in this chapter. The actual data averages shown in that figure have smoothed corners that would result from a probabilistic choice function, and Holt and Laury (2001) did find that such a function provided a good explanation of the actual pattern of choice proportions.

Questions

1. In a matching pennies game played with pennies, one person loses a penny and the other wins a penny, so the payoffs are 1 and -1 . Show that there is a simple way to transform the game in Table 10.1 (with payoffs of 40 and 80) into this form. (Hint: first divide all payoffs by 20, and then subtract a constant from all payoffs. This would be equivalent to paying in 20-cent coins if they existed, and charging an entry fee to play.)
2. Consider a soccer penalty kick situation where the kicker is equally skillful at kicking to either side, but the goalie is better diving to one side. In particular, the kicker will always score if the goalie dives away from the kick. If the goalie dives to the side of the kick, the kick is always blocked on the goalie's right but is only blocked with probability one half on the goalie's left. Represent this as a simple game, in which the goalie earns a payoff of $+1$ for each blocked kick and -1 for each goal, and the kicker earns -1 for each blocked kick and $+1$ for each goal. A 0.5 chance of either outcome results in an expected payoff of 0. Determine the equilibrium probabilities used in the Nash equilibrium.
3. Show that the two choice probabilities in (10.5) sum to 1.
4. Show that doubling all payoffs, i.e. both π_1 and π_2 , has the same effect as reducing the noise parameter μ in (10.5) by half.
5. Show that multiplying all payoffs in (10.2) by 2 will not affect choice probabilities.
6. Show that adding a constant amount, say x , to all payoffs in (10.5) will have no effect on the choice probabilities. *Hint:* Use the fact that an exponential function of the sum is the product of the exponential functions of the two components of the product: $\exp(\pi+x) = \exp(\pi)\exp(x)$.

Chapter 11. The Traveler's Dilemma

Each player's best decision in a (single play) prisoner's dilemma is to defect, regardless of what the other person is expected to do. The dilemma is that both could be better off if they resist these incentives and cooperate. The "traveler's dilemma" is a game with a richer set of decisions; each person must claim a money amount, and the payment is the minimum of the claims, plus a (possibly small) payment incentive to be the low claimant. Both would be better off making identical high claims, but each has an incentive to "undercut" the other to obtain the reward for having the lower claim. This game is similar to a prisoner's dilemma in that the unique equilibrium involves payoffs that are lower than can be achieved with cooperation. But here there is a more interesting dimension, since the optimal decision in the traveler's dilemma *does* depend on what the other person is expected to do. Thus the game is more sensitive to interactions of imprecise beliefs and small variations in decisions. These interactions can cause data patterns to be quite far from Nash predictions, as will be illustrated with a set of iterated spreadsheet calculations. This is one of my favorite games! It can be played with the instructions in the Appendix or with the Vecon program TG. In addition, it is possible for students to play an online demonstration version of the Traveler's Dilemma, where the "other decision" is taken from a database of decisions made by University of Virginia law students who were taking a class in behavioral game theory. This online simulation can be accessed from: <http://veconlab.econ.virginia.edu/tddemo.htm>

I. A Vacation with an Unhappy Ending?

Once upon a time, two travelers were returning from a tropical vacation where they had purchased identical antiques that were packed in identical suitcases. When both bags were lost in transit, the passengers were asked to produce receipts, an impossible task since the antiques were purchased with cash. The airline representative calmly informed the travelers that they should go into separate rooms and fill out claim sheets for the contents of the suitcases. He assured them that both claims would be honored if a) they are equal, and b) they are no greater than the liability limit of \$200.00 per bag in the absence of proof of purchase. There was also a minimum allowed claim of \$80.00 to cover the cost of the luggage and inconvenience. One of the travelers expressed some frustration, since the value of the suitcase and antique combined was somewhat above \$200.00. The other traveler, who was even more pessimistic, asked what would happen if the claims were to differ. To this, the reply was: "If the claims are unequal, then we will assume that the higher claim is unjustly inflated, and we will reimburse you both at an amount that equals the *minimum* of the two claims.

In addition, there will be a \$5.00 penalty assessed against the higher claimant and a \$5.00 reward added to the compensation of the lower claimant.” The dilemma was that each could obtain reasonable compensation if they made matching high claims of \$200, and resisted the temptation to come down to \$199 to capture the \$5 reward for being low. If one person expects the other one to ask for \$199, the best response is \$198, but what if the other person is thinking in similar thoughts? This line of reasoning leads to an unfortunate possibility that a cycle of anticipated moves and counter-moves may cause each traveler to put in very low claims. Unfortunately, there was no way for the two travelers to communicate between rooms, and the stress of the trip led each to believe that there would be no sharing of any unequal compensation amounts received. The unhappy ending has both making low claims, and the alternative happy ending involves two high claims. The actual outcome is an empirical question that we will consider below.

This game was introduced by Basu (1994), who viewed it as a dilemma for the game theorist. With a very low penalty rate, there is little risk in making a high claim, and yet each person has an incentive to “undercut” any common claim level. For example, suppose that both are considering claims of \$200; perhaps they whispered this to each other on the way to the separate rooms. Once alone, each might reason that a deviation to \$199 would reduce the minimum by \$1, which is more than made up for by the \$5 reward for being low. In fact, there is no belief about the other’s claim that would make one want to claim the upper limit of \$200. If one expects any lower claim, then it is better to undercut that lower claim as well. Reasoning in this manner, we can rule all common claims as candidates for equilibrium, except for the lowest possible claim of \$80. Similar reasoning shows that no configuration with unequal claims can be an equilibrium (Question 1).

The unique Nash equilibrium for the traveler’s dilemma has another interesting property: it can be derived from an assumption that each person knows that the other is perfectly rational. Recall that there is no belief about the other’s claim that would justify a claim of \$200. Since they each know that the other will never claim \$200, then the upper bound has shifted to \$199, and there is no belief that justifies a claim of \$199. To see this, suppose that claims must be in integer dollar amounts and note that \$199 is a best response to the other’s claim of \$200 but it is not a best response to any lower claim. Since \$200 will not be claimed by any rational person, it follows that \$199 is not a best response to any belief about the other’s decision. Reasoning in this manner, we can rule out all successively lower claims except for the very lowest feasible claim. (This argument based on common knowledge of rationality is called “rationalizability,” and the minimum possible claim in this game is the unique *rationalizable* equilibrium. Economists often place juicy, persuasive adjectives in front of the word equilibrium, sometimes to no avail. This will turn out to be one of those times.)

The dilemma for the theorist is that the Nash equilibrium is not sensitive to the size of the penalty/reward level, as long as this level is larger than the smallest possible amount by which a claim can be reduced. For example, the unilateral incentive to deviate from any common claim in the traveler's dilemma is not affected by if the penalty/reward rate is changed from \$5 to \$4. If both were planning to choose a claim of \$200, then one person's deviation to \$199 would reduce the minimum to \$199, but would result in a reward of \$4 for the deviator. Thus the person deviating would earn $\$199 + \4 instead of the \$200 that would be obtained if they both claim \$200. The same argument can be used to show that there is an incentive to undercut any common claim, whether the penalty/reward rate is \$2 or \$200. Thus the Nash equilibrium is not sensitive to this penalty/reward rate as long as it exceeds \$1, whereas one might expect observed claims to be responsive to large changes in this payoff parameter.

II. Data

The traveler's dilemma was analyzed by Capra et al. (1999). This was an exciting experiment. We expected the penalty/reward parameter to have a strong effect on actual claim choices, even though the unique Nash prediction would be independent of changes in this parameter. The design involved two parts, A and B, with different penalty/reward parameters. Each session involved about 10-12 subjects, who were randomly paired at the start of each round, for a series of 10 rounds. The part A data averages for four of the treatments are plotted in Figure 11.1, where the horizontal axis shows the round number and the penalty/reward parameter is denoted by R .

With a high penalty/reward parameter of \$0.80, the claims average about \$1.20 in round 1 and fall to levels approaching \$0.80 in the final four rounds, as shown by the thick solid line at the bottom of the figure. The data for the \$0.50 treatment start somewhat higher but also approach the Nash prediction in the final rounds. In contrast, the round 1 averages for the \$0.10 and \$0.05 treatments, plotted as dashed lines, start at about a dollar above the Nash prediction and actually rise slightly, moving *away* from the Nash prediction. The data for the intermediate treatments (\$0.20 and \$0.25) are not shown, but they stayed in the middle range (\$1.00 to \$0.50) below the dashed lines and above the solid lines, with some more variation and crossing over.




Figure 11.1. Data for the Traveler's Dilemma (Capra et al. 1999)

III. Learning and Experience

Notice that the most salient feature of the traveler's dilemma data, the strong effect of the penalty/reward parameter, is not predicted by the Nash equilibrium. Some might dismiss this game on the grounds that it is somewhat artificial. There is some truth to this, although some standard economic games do involve payoffs that depend on the minimum price, as is the case with "Bertrand" price competition (Chapter 15) or the next section's minimum-effort coordination game. An alternative perspective is that this game involves an intentionally abstract setting, which serves as a paradigm for some types of strategic interactions. The traveler's dilemma is no more about lost luggage than the prisoner's dilemma is about actual prisoners. If standard game theory cannot predict well in such simple situations, then some rethinking is needed. At a minimum, it would be nice to have some idea of when the Nash equilibrium will be useful and when it will not. Even better would be a theoretical apparatus that explains both the convergence to Nash predictions in some treatments and the divergence in others. The rest of this chapter pertains to some possible approaches to this problem. We begin with an intuitive discussion of learning.

Behavior in an experiment with repeated random pairings may evolve from round to round as people learn what to expect. For example, consider the data in Table 11.1 from a classroom experiment conducted at the University of

Virginia. Claims were required to be between 80 and 200 cents, and the penalty/reward parameter was 10 cents. There were 20 participants, who were divided into 10 pairs. Each 2-person team had a handheld wireless PDA with a touch-sensitive color screen that showed the HTML displays from the Veconlab software. It was a nice spring day, and this class was conducted outside on the “lawn,” with students being seated informally in a semi-circle. The table shows the decisions for five of the teams during the first four rounds. The round is listed on the left, and the average of all 10 claims is shown in the second column. The remaining columns show the team’s own decisions, which are listed above the decision of the other team for that round (shown in parentheses).

First consider the round 1 decision for “Stacy/Naomi,” who claimed 150. The other team was lower, at 135, so the earnings for Stacy/Naomi were the minimum, 135, minus the penalty of 10 cents, or 125. “SuzSio” began lower, at 100, and encountered a claim of 133. This team then cut their claim to 95, and encountered an even higher claim of 191 in round 2. This caused them to raise their claim to 125, and they finished round 10 (not shown) with a claim of 160.

Table 11.1. Traveler’s Dilemma Data for a Classroom Experiment with $R = 10$
Key: Own Claim (Other’s Claim)

Round	Average (12 claims)	SuzSio	K Squared	Kurt/Bruce	JessEd	Stacy/Naomi
1	137	100 (133)	080 (195)	139 (140)	133 (100)	150 (135)
2	131	095 (191)	098 (117)	135 (140)	80 (130)	127 (200)
3	144	125 (135)	096 (117)	135 (100)	199 (199)	134 (200)
4	142	115 (125)	130 (100)	125 (115)	198 (115)	150 (1.34)

The team “K Squared” (Katie and Kari) began round 1 with a decision of 80 cents, which is the Nash equilibrium. They had the lower claim and earned the minimum, 80, plus the 10 cents for being low. After observing the other’s claim of 195 for that round, they raised their claim to 98 in round 2, and eventually to 120 in round 10. The point of this example is that the best decision in a game is not necessarily the equilibrium decision; it makes sense to respond to and anticipate the way that the others are actually playing. Here the Nash equilibrium decision of 80 played in the first four periods would have earned 90 in each round, for a total of 360. “K Squared” actually earned 409 by adapting to

observed behavior and taking more risk. A claim of 200 in all rounds would have been even more profitable (question 4).

Teams who were low relative to the other's claim tended to raise their claims, and teams that were high tended to lower them. This qualitative adjustment rule, however, does not explain why "SuzSio's" claim was lowered even after it had the lower claim in round 1. This reduction seems to be in anticipation of the possibility that other claims might fall. In the final round, the claims ranged from 110 to 200, with an average of 146 and a lot of dispersion.

One way to describe the outcome is that people have different beliefs based on different experiences, but by round 10 most expect claims to be about 150 on average, with a lot of variability. If claims had converged to a narrow band, say with all at 150, then the "undercutting" logic of game theory might have caused them to decline as all seek to get below 150. But the variability in claims did not go away, perhaps because people had different experiences and different reactions to those experiences. This variability made it harder to figure out the best response to others' claims. In fact, claims did not diminish over time; if anything there was a slight upward trend. The highest earnings were obtained by the "JessEd" team, which had a relatively high average claim of 177.

The part A treatment for this session was followed by 10 more rounds with a higher payoff parameter (50 cents). This created a strong incentive to be the low claimant, and claims did decline to the Nash equilibrium level of 80 after the first several rounds. Thus we see that convergence to a Nash equilibrium seems to depend on the *magnitudes* of incentives, not just on whether one decision is slightly better than another. An analysis that is sensitive to the magnitudes of payoff differences will be considered next. The goal is to develop a theoretical explanation of the observed convergence to the Nash prediction in the high- R treatments and the divergence in the low- R treatments.

IV. A Computational Model of Noisy Behavior (Optional)

Recall from Chapter 11 that stochastic response functions can be used to capture the idea that the magnitudes of payoff differences matter. The form of these functions is based on the intuitive idea that a person is much more likely to be able to determine which of two sounds is louder when the decibel levels are not close together. In these games, the expected payoff is the stimulus analogous to the decibel level, so let's begin by calculating expected payoffs for each decision. To keep the calculations from becoming tedious, suppose that there are only 13 possible decisions in even ten-cent amounts: 80, 90, 100, ... 200. A person's beliefs will be represented by 13 probabilities that sum to 1.

These belief probabilities can be used to calculate expected payoffs for each decision, and the logit functions introduced in the previous chapter can then be used to determine the choice probabilities that result from the initial beliefs.

The intuitive idea is that a decision is more likely if its expected payoff is higher. In particular, choice probabilities are assumed to be increasing (exponential) functions of expected payoffs, normalized so that all probabilities sum to 1.

Some notation will be useful for the calculation of expected payoffs. Let P_i be the probability associated with claim i , so beliefs are characterized by $P_{80}, P_{90}, \dots, P_{200}$. First consider a claim of 80, which will tie with probability P_{80} and will be low with probability $1 - P_{80}$. The expected payoff is:

$$\text{expected payoff for 80} = 80P_{80} + (80 + R)(1 - P_{80}).$$

Notice that the first term covers the case of a tie and the second covers the case of having the low claim. For a claim of 90, we have to consider the chance of having the higher claim and incurring the penalty of R , so the expected payoff is:

$$\text{expected payoff for 90} = 90P_{90} + (90 + R)(1 - P_{80} - P_{90}) + (80 - R)P_{80}.$$

As before, the first term is for the possibility of a tie at 90, and the second is for the possibility that the other claim is above 90, which occurs with probability $1 - P_{80} - P_{90}$. The third term now reflects the possibility of having the higher claim. The payoff structure can be clarified by considering a higher claim, say 150.

$$\begin{aligned} \text{expected payoff for 150} = & 150P_{150} \\ (11.1) \quad & + (150 + R)[1 - P_{80} - P_{90} - \dots - P_{150}] \\ & + (80 - R)P_{80} + (90 - R)P_{90} + \dots + (140 - R)P_{140}. \end{aligned}$$

In order from top to bottom, the parts on the right side of (11.1) correspond to the cases of a tie, of having the lower claim, and of having the higher claim.

This section shows you how to do these calculations in an Excel spreadsheet that makes it possible to repeat the calculations iteratively in order to obtain predictions for how people might actually behave in a traveler's dilemma game when we do not assume perfect rationality. The logic of this spreadsheet is to begin with a column of probabilities for each claim (80, 90, ... 200). Given these probabilities, the spreadsheet will calculate a column of expected payoffs, one for each claim. Let the expected payoff for claim i be denoted by \mathbf{p}_i^e , where $i = 80, 90, 100$, etc. Then there is a column of exponential functions of expected payoffs, $\exp(\mathbf{p}^e/\mu)$, where the expected payoffs have been divided by an error parameter μ . We want choice probabilities to be increasing in expected payoffs, but these exponential functions cannot be used as probabilities unless they are normalized to ensure that they sum to 1. Thus the column of exponential

functions is summed to get the normalizing element in the denominator of the logit probability formula: $\exp(\mathbf{p}^e/\mu)/\sum_i(\exp(\mathbf{p}_i^e/\mu))$, which corresponds to equation (10.5) from the previous chapter.

The best way to read this section is to open up a spreadsheet. The instructions will be provided for Excel, but analogous instructions would work in other spreadsheet programs. Table 11.2 is laid out like a spreadsheet, with columns labeled A, B, ..., and rows labeled 1, 2,

Table 11.2. Excel Spreadsheet for Traveler's Dilemma Logit Responses

	A	B	C	D	E	F	G	H	I	J	K
1	$\mu =$	10									
2	$R =$	10									
3											
4											
5											
6	X	X-R	X+R	P	PX	F(X)	P(X-R)		\mathbf{p}^e	Exp(\mathbf{p}^e/m)	P
7	70					0		0			
8	80	70	90	1	80	1	70	70	80	2980	0.184
9	90	80	100	0	0	1	0	70	70	1096	0.067
10	100	90	110	0	0	1	0	70	70	1096	0.067
11	110	100	120	0	0	1	0	70	70	1096	0.067
12	120	110	130	0	0	1	0	70	70	1096	0.067
13	130	120	140	0	0	1	0	70	70	1096	0.067
14	140	130	150	0	0	1	0	70	70	1096	0.067
15	150	140	160	0	0	1	0	70	70	1096	0.067
16	160	150	170	0	0	1	0	70	70	1096	0.067
17	170	160	180	0	0	1	0	70	70	1096	0.067
18	180	170	190	0	0	1	0	70	70	1096	0.067
19	190	180	200	0	0	1	0	70	70	1096	0.067
20	200	190	210	0	0	1	0	70	70	1096	0.067
21				1	80					16140	1

Table 11.3. Column Key for Spreadsheet in Table 11.2
(formulas for row 8 should be copied down to row 20)

Column	Variable	Notation	Formula for Row 8
A	Claim	X	= A7 + 10
B	Payoff if Lower	X - R	= A8 - \$B\$2
C	Payoff if Higher	X + R	= A8 + \$B\$2
D	Probability	P	1

E	Product	PX	$= D8*\$A8$
F	Cumulative P	$F(X)$	$= F7+D8$
G	Product 2	$P(X-R)$	$= D8*\$B8$
H	Cumulative Product 2		$= H7 + G8$
I	Expected Payoff	p^e	$= E8 + (1-F8)*\$C8 + H7$
J	Exponential of payoff	$\exp(p^e/m)$	$= \exp(I8/\$B\$1)$
Cell J21	Sum of exponentials	$\sum_i (\exp(p_i^e/m))$	$= \text{sum}(J8..J20)$
K	Probability	$\exp(p^e/m) / \sum_i (\exp(p_i^e/m))$	$= J8/J\$21$

Step 1: Parameter Specification. It is convenient to have the error parameter, μ , be at a focal location in the upper left corner, along with the penalty/reward parameter. Put a value of **10** for the error parameter in cell **B1**, and put **10** for the penalty/reward parameter in cell **B2**. These numbers can later be changed to experiment with different amounts of noise for different treatments.

Step 2: Claim Values. Put a **70** in cell **A7**, so that the formula in **A8** can be: $=\text{sum}(A7+10)$, which will yield a value of 80. Then this formula can be copied to cells A9 to A20 by clicking on the lower right corner of the A8 box and dragging it down. (You could later go back and expand the table to allow for all penny amounts from 80 to 200, but the smaller table will do for now.)

Step 3: Initial Values. Before beginning to fill in the other cells, put values of 0 in cells F7 and H7; these zeros are used to start cumulative sums, in a manner that will be explained below.

Step 4: Formulas. Notice that the formula that you used in step 2 is shown at the top of the far right column of Table 11.3. The formulas for row 8 cells of other columns are also shown in the right column. Enter these in cells **B8**, **C8**, ... **K8**, making sure to place the \$ symbols in the places indicated. The dollar symbol forces the reference to stay fixed even when the formula is copied to another location. For example, the references to the error parameter will be $\$B\1 , with two \$ signs, since both the row and location of the reference to this parameter must stay fixed. A reference to cell $\$B8$, however, would keep the column fixed at B and allow one to copy the formula from row 8 to other rows. Please use the \$ symbols only where indicated, and not elsewhere. Finally, copy all of these formulas down from row 8 to rows 9-20.

Step 5: Sums. The numbers in column K will not make sense until you put the formula for the sum of exponentials into cell **J21**; this formula is given in the

next-to-last row of Table 11.3. Add a formula to sum the probabilities and payoffs by copying the formula in cell **J21** to cells **D21**, **E21**, and **K21**.

At this point, the numbers in your spreadsheet should match those in Table 11.2. The initial probability column reflected a belief that the other person would choose 80 with certainty. Therefore the expected payoff in column I is 80 if one matches this claim. Given these beliefs, any higher claim will result in a 10-cent penalty, and the payoff will be other's claim of 80, which is the minimum, minus 10, or 70 as shown in column I. This column will provide expected payoffs for each decision as the initial belief column D is changed, so let's look at the structure of the formula in column I:

$$\begin{aligned} \text{expected payoff for claim of 80} &= \text{E8} \\ (11.2) \qquad \qquad \qquad &+ (1 - \text{F8}) * \text{C8} \\ &+ \text{H7}. \end{aligned}$$

This formula has three parts that correspond to the three parts of equation (11.1), i.e. depending on whether the other's claim is equal to, higher, or lower than one's own claim. We will discuss these three elements in turn.

Case of a Tie: In equation (11.2), the first term, **E8**, was calculated as the probability in **D8** that the other chooses 80 times the claim itself in **A8**. Thus this first term covers the case of a tie.

Case of Having the Lower Claim: The second term in (11.2) involves **C8**, which is the claim plus the reward, R , so this term pertains to the case where one's claim is the lower one. The $(1 - \text{F8})$ term is the probability that the other's claim is higher, so **F8** is the probability that the other's claim is less than or equal to one's own claim of 80. Since there is no chance that the other is less than 80, **F8** is calculated as **F7**, which has been set to 0, plus the probability that the other's claim is exactly 80, i.e. **C8**. As this formula is copied down the column, we get a sum of probabilities, which can be thought of as the cumulative "less-than-or-equal-to" probabilities. Retracing our steps, one minus the less-than-or-equal-to probabilities in column F will be the probability that the other's claim is higher, so the second term in (11.1) has the $(1 - \text{F8})$ being multiplied by the claim plus the reward.

Case of Having the Higher Claim: Here we have to consider each of the claims that are lower than one's own claim. Column B has claims with the penalty subtracted, and these payoffs are multiplied by the probability of that claim to

yield the numbers in column G. Then column H calculates a cumulative “less-than-or-equal-to” sum of the elements in G. This sum is necessary since, for any given claim, there may be many lower claims that the other can make, with associated payoffs that must be multiplied by probabilities and summed to get an expected payoff.

Now that the expected payoffs are calculated in column I, they are put into exponential functions in column J after dividing by the error parameter in cell **B2**. These are then summed in cell **J21**, and the exponentials are divided by the sum to get the choice probabilities in column K. For example, in Table 11.2 the probability associated with a claim of 80 is about 0.184, as shown in cell **K8**. The other probabilities are about a third as high. These other probabilities are all equal, since the other person is expected to choose a claim of 80 for sure, and therefore all higher claims lead to earnings of 70, as compared with the 80 that could be earned by matching the other’s expected claim. This reduction in expected payoff is not so severe as to prevent deviations from happening, at least for the error parameter of 10 in cell **B1** that is being used here. Try a lower error parameter, say 5, to be sure that it results in a higher probability for the claim of 80, with less chance of an “error” in terms of making a claim that is higher than 80.

Next consider cell **E21**, which shows the expected value of the claim for the initial beliefs. If you still have a probability of 1 in cell **D8** for the lowest claim, then the number in **E21** should be 80. As you change the probabilities in column D (making sure that they sum to 1), the expected claim in **E21** will change.

Finally, we can use the choice probabilities in column K to calculate the resulting expected claims. To make this calculation, we essentially need to get the formula in **E21** over to the right of the K column. At this point, you should save your spreadsheet (in case something goes wrong) and then perform the final two steps:

Step 6. Highlight the shaded block of cells in Table 11.2, i.e. from **E6** to **K21**. Then copy and paste this block with the cursor in cell **L6**, which will essentially replicate columns E to K, putting them in columns L to R. Now you should see that the average claim in cell **L21** is about 133, way above the average claim on 80 based on the initial beliefs.

Step 7. The next step is to perform another iteration. You could mark the shaded block in Table 11.2 again (or skip this step if it is still on the clipboard) and then place the cursor in cell **S6** and copy the block. The new average claim in cell **S21**

should be 166. In two more iterations this average will converge to 178 (in cell **AG21**), and this average will not change in subsequent iterations.

At this point, there will be a vector of probabilities in column AG with the following interesting property: if these represent initial beliefs, then the stochastic best responses to these beliefs will yield essentially the same probabilities, i.e. the beliefs will be confirmed. This is the notion of equilibrium that was introduced in chapter 10. Even more interesting is the fact that the level of convergence, 178, is approximately equal to the average claim for the $R = 10$ treatment in Figure 11.1. Now try changing the payoff parameter from 10 to 50 to be sure that the average claims converge to 80, and check to be sure that these convergence levels are not too sensitive to changes in initial beliefs, e.g. if you put a 1 in the bottom (**A20**) element of column D instead of in the **A8** element. In this sense, the model of logit stochastic responses can explain why behavior converges to the Nash prediction in some treatments (50 and 80), and why data go the other side of the set of feasible decisions in other treatments (5 and 10).

In equilibrium, beliefs are confirmed, and the result is called a logit equilibrium, which was used by Capra et al. (1999) to evaluate data from the experiment. Our spreadsheet calculations in this section have been based on an error rate of 10, which is close to the value of 8.3 that was estimated from the actual data (using maximum likelihood techniques), with a standard error of 0.5.

V. Extensions and Further Reading

The experiment just discussed involved repeated random matching, so that people can learn from experience, even though people are not interacting with the same person in each round. Goeree and Holt (2001) provide experimental evidence for traveler's dilemma games played only once. There is no opportunity to learn from past experience in a one-shot game; each person must base their claim decision on *introspection* about what the other is likely to do, about what the other thinks they will do, etc. The average claims are strongly influenced by the size of the penalty/reward, even though changes in this parameter have no effect on the unique Nash equilibrium. Of course, it is not reasonable to expect data to conform to a Nash equilibrium in a one-shot game with no past experience, since this equilibrium (in pure strategies) implies that the other's claim is somehow known. The development of models of introspection is an important and relatively poorly understood topic in game theory. See Goeree and Holt (2003) for a formal model of introspection.

There are several other games with a similar payoff that depends on the minimum of all decisions. One such game, discussed in the next chapter, has the "weakest-link" property that the output is determined by the minimum of the individual effort levels. Similarly, the shopping behavior of informed consumers

may make firms' profits sensitive to whether or not their price is the minimum in the market. Capra et al. (2001) provide experimental data for a price competition game with "meet-or-release clauses," which release the buyer from the contract if a lower price offer is found and the original seller refuses to match. If all consumers are informed and all sellers are producing the same product, then each would rather match the other instead of losing all business as informed consumers switch. A unilateral price cut, may, however pick up some business for a group of uninformed consumers. If one firm sets a lower price than the other, then it will obtain the larger market share, and the high-price firm will have to match the other's price to get any sales at all. This is like the traveler's dilemma in that earnings are determined by the minimum price, with a penalty for having a higher price. In particular, each firm earns an amount that equals the minimum price times their sales quantity, but the firm that had the lower price initially will have the larger market share by virtue of picking up sales to the informed shoppers.

The Capra et al. (2001) price-competition game also has a unique Nash equilibrium price at marginal cost, since at any higher price there is a unilateral incentive to cut price by a very small amount and pick up the informed shoppers. The Nash equilibrium is independent of the number of informed buyers who respond to even small price differences. Despite this independence property, it is intuitively plausible that a large fraction of buyers who are uninformed about price differences would provide sellers with some power to raise prices. This intuition was confirmed by the results of the experiments reported by the authors. As with the traveler's dilemma, data averages were strongly affected by a parameter that determines the payoff differential for being low, even though this parameter does not affect the unique Nash equilibrium at the lowest possible decision.

Questions

1. Show that unequal claims cannot constitute a Nash equilibrium in a traveler's dilemma.
2. What is the Nash equilibrium for the traveler's dilemma game where there is a \$5 penalty and a \$5 reward, and claims must be between -\$50 and \$50? (A negative claim means that the traveler pays the airline, not the reverse.)
3. Consider a traveler's dilemma with N players who each lose identical items, and the airline requests claim forms to be filled out with the understanding that claims will be between \$80 and \$200. If all claims are not equal, there is a \$5 penalty if one's claim is not the lowest and a \$5 reward rate for the person with the lowest claim. Speculate on what the

effect on average claims of increasing N from 2 to 4 players in a setup with repeated random matchings.

4. Calculate what the earnings would have been for the “K squared” team if they had chosen a claim of 200 in each of the first four rounds of the game summarized in table 11.1.
5. Use the spreadsheet constructed in section IV to show that the choices converge to the Nash level of 80 when the error rate is reduced to 1 for the $R = 10$ treatment.

Chapter 12. Coordination Games

A major role of management is to coordinate decisions so that better outcomes can be achieved. The game considered in this chapter highlights the need for such external management, since otherwise players may get stuck in a situation where nobody exerts much effort because others are not expected to work hard either. In these games, the productivity of each person's effort depends on that of others, and there can be multiple Nash equilibria, each at a different common effort level. Behavior is sensitive to factors like changes in the effort cost or the size of the group, even though these have no effect on the set of Nash equilibria. The experiment can be run with the Veconlab game CG. Alternatively, the matrix game with 7 possible effort decisions discussed below is implemented by the default parameters for the large matrix game program, MGN. The instructions in the appendix are for the matrix-game version.

I. "The Minimum Effort Game? That's One I Can Play!"

Most productive processes involve specialized activities, where distinct individuals or teams assemble separate components that are later combined into a final product. If this product requires one of each of the components, then the number of units finally sold is sensitive to bottlenecks in production. For example, a marine products company that produces 100 hulls and 80 engines will only be able to market 80 boats. Thus there is a bottleneck caused by the division with the lowest output. Students (and professors too) have little trouble understanding the incentives of this type of "minimum-effort game," as one student's comment in the section title indicates.

The minimum-effort game was originally discussed by Rousseau in the context of a stag hunt, where a group of hunters form a large circle and wait for the stag to try to escape. The chances of killing the stag depend on the watchfulness and effort exerted by the encircling hunters. If the stag observes a hunter to be napping or hunting for smaller game instead, then the stag will attempt an escape through that sector. If the stag is able to judge the weakest link in the circle, then the chances of escape depend on the minimum of the hunters' efforts. The other aspect of the payoffs is that effort is individually costly, e.g., in terms of giving up the chance of bagging a hare or taking a rest.

The game in Table 12.1 is a 2x2 version of a minimum effort game. First, consider the lower-left corner where each person has a low effort, and payoffs are 70 each. If the Row player increases effort, while the Column player maintains low efforts, then the relevant Row payoff is reduced to 60, as can be seen from the top-left box. Think of it this way: the unilateral increase in effort will not increase

the minimum, but the extra cost reduces Row's payoff from 70 to 60, so the cost of this extra "unit" of effort is 10.

Table 12.1. A Minimum Effort Game (Row's payoff, Column's payoff)

		Column Player:	
		Low Effort (1)	High Effort (2)
Row Player:			
High Effort (2)		60, 70 ↓	⇒ 80, 80
Low Effort (1)		70, 70 ⇐	↑ 70, 60

Now suppose that the initial situation is a high effort for Row and a low effort for Column, as in the upper-left box. An increase in Column's effort will raise the minimum, which raises Row payoff from 60 to 80. Thus a "unit" increase in the minimum effort will raise payoffs by 20, holding one's own effort constant. These observations can be used to devise a mathematical formula for payoffs that will be useful in devising more complex games. Let the low effort be 1 and the high effort be 2, as indicated by the numbers in parentheses next to the row and column labels in the payoff table. Then the payoffs in this table are determined by the formula: $60 + (20 \text{ times the minimum effort}) - (10 \text{ times one's own effort})$. Let M denote the minimum effort and E denote one's own effort, so that this formula is:

$$(13.1) \quad \text{own payoff} = 60 + 20M - 10E.$$

For example, the upper right payoffs are determined by noting that both efforts are 2, so the minimum (M) is 2, and the formula yields: $60 + (20) \cdot 2 - (10) \cdot 2 = 80$.

The formula in (12.1) was used to construct the payoffs in Table 12.2, for the case of efforts that can range from 1 to 7. Notice that the four numbers in the bottom-left corner of the table correspond to the row payoffs in Table 12.1 above. With a larger number of possible effort levels, we see the dramatic nature of the potential gains from coordination on high-effort outcomes. At a common effort of 7, each person earns 130, which is almost twice the amount earned at the lowest effort level. Movements along the diagonal from the lower-left to the upper-right corner show the benefits from coordinated increases in effort; a one unit increase in the minimum effort raises payoff by 20, minus the cost of 10 for the increased effort, so each diagonal payoff is 10 larger than the one lower on the diagonal.

Besides the gains from coordination, there is an additional feature of Table 12.2 that is related to risk. When the row player chooses the lowest effort (in the bottom row), the payoff is 70 for sure, but the highest effort may yield payoffs that range from 10 to 130, depending on the column player's choice. This is the strategic dilemma in this coordination game: there is a large incentive to coordinate on high efforts, but the higher effort decisions are risky.

Table 12.2. The Van Huyck et al. (1990) Minimum Effort Game
With Row Player's Payoffs Determined by the Minimum of Others' Efforts

		Column's Effort (or Minimum of Other Efforts)						
		1	2	3	4	5	6	7
Row's Effort	7	10	30	50	70	90	110	130
	6	20	40	60	80	100	120	120
	5	30	50	70	90	110	110	110
	4	40	60	80	100	100	100	100
	3	50	70	90	90	90	90	90
	2	60	80	80	80	80	80	80
	1	70	70	70	70	70	70	70

An additional element of risk is introduced when more than two people are involved. Suppose that payoffs are still determined by the formula in (12.1), where M is the minimum of all players' efforts. The result is still the payoff table 12.2, where the payoff numbers pertain to the row player as before, but where the columns correspond to the minimum of the other player's efforts. For example, the payoffs are all 70 in the bottom row because Row's effort of 1 is the minimum regardless of which column is determined by the minimum of the others' efforts. Notice that the incorporation of larger numbers of players into this minimum-effort game does not alter the essential strategic dilemma, i.e. that high efforts involve high potential gains but more risk. At a deeper level, however, there is more risk with more players, since the minimum of a large number of independently selected efforts is likely to be small when there is some variation from one person to another. This is analogous to having a large number of hunters spread out in a large circle, which gives the stag more of a chance to find a sector where one of the hunters is absent or napping.

II. Nash Equilibria, Numbers Effects, and Experimental Evidence

These intuitive considerations (the gains from coordination and the risks of un-matched high efforts) are not factors in the structure of the Nash equilibria

in this game. Consider Table 12.1, for example. If the other person is going to exert a low effort, the best response is a low effort that saves on effort cost. Thus the lower-left outcome, low efforts for each, is a Nash equilibrium. But if the other player is expected to choose a high effort, then the best response is a high effort, since the gain of 20 for the increased minimum exceeds the cost of 10 for the additional unit of effort. Thus the high-effort outcome in the upper-right corner of Table 12.1 is also a Nash equilibrium. This is an important feature of the coordination game: there are multiple equilibria, with one that is preferred to the other(s).

The presence of multiple equilibria is a feature that differentiates this game from a prisoner's dilemma, where all players may prefer the high-payoff outcome that results from cooperative behavior. This high-payoff outcome is not a Nash equilibrium in a prisoner's dilemma since each person has a unilateral incentive to "defect." The equilibrium structure for the game in Table 12.1 is not affected by adding additional players, each choosing between efforts of 1 and 2, and with the row player's payoffs determined by the column that corresponds to the minimum of the others' efforts. In this case, adding more players does not alter the fact that there are two equilibria (in non-random strategies): all choose low efforts or all choose high efforts. Just restricting attention to the Nash equilibria would mean ignoring the intuition that adding more players would seem to make the choice of a high effort riskier, since it is more likely that one of the others will choose low effort and pull the minimum down.

The problem of multiple equilibria is more dramatic with more possible efforts, since any common effort is a Nash equilibrium in Table 12.2. To see this, pick any column and notice that the row player's payoffs are highest on the shaded diagonal. As before, this structure is independent of changes in the number of players, since such changes do not alter the payoff table. These considerations were the basis of an experiment conducted by Van Huyck et al. (1990), who used the payoffs from Table 12.2 for small groups (size 2) and large groups (size 14-16). The large groups played the same game ten consecutive times with the same group, and the lowest effort was announced after each round to enable all to calculate their payoffs (in pennies). Even though a majority of individuals selected high efforts of 6 or 7 in the first round, the minimum effort was no higher than 4 in the first round for any group. With a minimum of 4, higher efforts were "wasted," and effort reductions followed in subsequent rounds. The minimum fell to the lowest level of 1 in all groups, and almost all decisions in the final round were at the lowest level.

This experiment is important since previously it had been a common practice in theoretical analysis to *assume* that individuals could coordinate on the best Nash equilibrium when there was general agreement about which one is best, as is the case for Table 12.2. In contrast, the subjects in the experiment managed

to end up in the equilibrium that is *worst* for all concerned. With groups of size 2, individuals were able to coordinate on the highest effort, except when pairings were randomly reconfigured in each round. With random matching, the outcomes were variable, with average efforts in the middle range. In any case, it is clear that group size had a large impact on the outcomes, even though changes in the numbers of players had no effect on the set of equilibria.

The coordination failures for large groups captured the attention of macroeconomists, who had long speculated about the possibility that whole economies could become mired in low-productivity states, where people do not engage in high levels of market activity because no one else does. The macroeconomic implications of coordination games are discussed in Bryant (1983), Cooper and John (1998), and Romer (1996), for example.

III. Effort-Cost Effects

Next consider what happens when the cost of effort is altered. For example, suppose that the effort cost of 10 used to construct Table 12.1 is raised to 19, so that the payoffs (for an effort of E and a minimum effort of M) will be determined by: $60 + 20M - 19E$. In this case, a one unit increase in each person's effort raises payoffs by 20 minus the cost of 19, so the payoffs in the upper-right box of the table are only 1 cent higher than the payoffs in the lower-left box. Notice that this increase in effort cost did not change the fact that there are two Nash equilibria, and that both players prefer the high-effort equilibrium. However, simple intuition suggests that effort levels in this game may be affected by effort costs. From the row player's perspective, the top row offers a possible gain of only one cent and a possible loss of 19 cents, as compared with the bottom row.

Table 12.3. A Minimum-Effort Game With High Effort Cost
(Row's payoff, Column's payoff)

		Column Player:	
		Low Effort (1)	High Effort (2)
Row Player:			
High Effort (2)		42, 61	62, 62
Low Effort (1)		61, 61	61, 42

Goeree and Holt (1999, 2001) report experiments in which the cost of effort is varied between treatments, using the payoff formula:

$$(13.2) \quad \text{payoff} = M - cE,$$

where M is the minimum of the efforts, E is one's own effort, and c is a cost parameter that is varied between treatments. Decisions were restricted to be any amount between (and including) \$1.10 and \$1.70. As before, any common effort is a Nash equilibrium in this game, as long as the cost parameter, c , is between 0 and 1. This is because a unilateral decrease in effort by one unit will reduce the minimum by 1, but it will only reduce the cost by an amount that is less than 1. Therefore, a unit decrease in effort will reduce one's payoff by $1 - c$. Conversely, a unilateral increase in effort by one unit above some common level will not raise the minimum, but the payoff will fall by c . Even though deviations from a common effort are unprofitable when c is greater than 0 and less than 1, the magnitude of c determines the relative cost of "errors" in either direction. A large value of c , say 0.9, makes increases in effort more costly, and a small value of c makes decreases more costly.

Figure 12.1 shows the results for sessions that consisted of 10-12 subjects who were randomly paired for a series of 10 rounds. There were three sessions with a low effort cost parameter of 0.25, and the averages by round for these sessions are plotted as thin dashed lines. The thick dashed line is the average over all three treatments. Similarly, the thin solid lines plot round-by-round averages for the three sessions with a high effort cost parameter of 0.75, and the thick solid line shows the average for this treatment.

Efforts in the first round average in the range from \$1.35 to \$1.50, with no separation between treatments. Such separation arises after several rounds, and average efforts in the final round are \$1.60 for the low-cost treatment, versus \$1.25 for the high-cost treatment. Thus we see a strong cost effect, even though any common effort is a Nash equilibrium.

One session in each treatment seemed to approach the boundary, which raises the issue of whether behavior will "lock" on one of the extremes. This pattern was observed in a Veconlab classroom experiment in which efforts went to \$1.70 by the 10th period. This kind of extreme behavior is not universal, however. Goeree and Holt (1999) report a pair of sessions that were run for 20 rounds. With an effort cost of 25 cents, the decisions converged to about \$1.55, and with an effort cost of 75 cents the decisions leveled off at about \$1.38. Both of these outcomes seem to fit the pattern seen in Figure 12.1.




Figure 12.1. Average Efforts for the Goeree and Holt (1999) Coordination Experiment:
Thick Lines are Averages Over Three Sessions for Each Cost Treatment

IV. An Analysis of Noisy Behavior (Optional)

The traveler's dilemma and the coordination game have similar payoff structures, so it is straightforward to modify the spreadsheet in Table 11.2 so that it applies to the coordination game. Here are the steps:

Step 1. Save the traveler's dilemma spreadsheet under a different name, e.g. cg.xls, before deleting all information to in column L and farther to the right.

Step 2. The coordination experiment has fewer possible decisions, when restricted to 10-cent increments: 110, 120, ... 170. Therefore, delete all material in rows 7-10, and in rows 18-20, leaving row 21 as it is. You will have to enter the possible effort levels in column A: **110** in **A11**, **120** in **A12**, etc.

Step 3. There is no penalty/reward parameter in the coordination game, so enter a **0** in cell **B2**. Next put "C =" in cell **A3**, and enter a value of **0.75** in cell **B3**.

Step 4. The starting values for the cumulative column sums will have to be moved, so put values of **0** in cells **F10** and **H10**.

Step 5. To set the initial probabilities, change the entry in cell **D11** to **1**, and leave the other entries in that column to be 0.

Step 6. Next we have to subtract the effort cost from the expected payoff formula in cell **I11**. The cost per unit effort in cell B3 must be multiplied by the effort in column A to calculate the total effort cost. Thus you should *append* the term **-\$B\$3*\$A11** to the end of the formula that is already in cell **I11**. Then this modified formula should be copied into rows 12 to 17 in this column.

Table 12.4. Excel Spreadsheet for Coordination Game Logit Responses

	A	B	C	D	E	F	G	H	I	J	K
1	$\mu =$	10									
2	$R =$	0									
3	$C =$	0.75									
4											
5											
6	X	X	X	P	PX	F(X)	PX		p^e	$Exp(p^e/m)$	P
7	70										
8	80										
9	90										
10	100					0		0			
11	110	110	110	1	110	1	110	110	27.5	15.64	0.530
12	120	120	120	0	0	1	0	110	20	7.38	0.250
13	130	130	130	0	0	1	0	110	12.5	3.49	0.128
14	140	140	140	0	0	1	0	110	5	1.64	0.056
15	150	150	150	0	0	1	0	110	-2.5	0.77	0.026
16	160	160	160	0	0	1	0	110	-10	0.36	0.012
17	170	170	170	0	0	1	0	110	-17.5	0.17	0.006
18											
19											
20											
21				1	110				16140		1

Step 7. At this point, the numbers should match those in Table 12.4, with a new vector of probabilities in column K. To perform the first iteration, mark the block from **E6** to **K21**, copy, and place the cursor in cell L6 before pasting. This same procedure can be repeated 10 times, so that the average effort after 10 iterations will be found on the far right side of the spreadsheet, in cell **AG21**.

The expected effort after one iteration, found in L21, is about 119. This rises to 125 by the fifth iteration. When the initial probabilities in column D are changed to put a probability of 1 for the highest effort and 0 for all other efforts, the iterations converge to 126 by the 10th iteration. When the cost of effort is reduced from 0.75 to 0.25, the average claims converge to the 155-156 range after several iterations. These average efforts are quite close to the levels for the two treatment averages in Figure 12.1 in round 10. As noted in the previous chapter, the convergence of iterated stochastic responses indicates an equilibrium in which a given set of beliefs about others' effort levels will produce a matching distribution of effort choice probabilities. The resulting effort averages are quite good predictors of behavior in the minimum-effort experiment when the error rate is set to 10. Even though any common effort is a Nash equilibrium, intuition suggests that lower effort costs may produce higher efforts. These spreadsheet calculations of the logit equilibrium are consistent with this intuition; they explain why efforts are inversely correlated with effort costs. A lower error rate will push effort levels to one or another of the extremes, i.e. to 110 or to 170 (question 4).

V. Extensions

Goeree and Holt (1999) also considered the effects of raising the number of players per group from 2 to 3, holding effort cost constant, which resulted in a sharp reduction of effort levels. Effort-cost effects were observed as well in games in which payoffs were determined by the median of the three efforts. Some coordination games may be played only once; Goeree and Holt (2001) report strong effort-cost effects in such games as well. A theoretical analysis of these effects (in the context of a logit equilibrium) can be found in Anderson, Goeree, and Holt (2001).

There is an interesting literature on factors that facilitate coordination on good outcomes in matrix games, e.g. Sefton (1999), Straub (1995), and Ochs (1995). In particular, notions of "risk dominance" and "potential" provide good predictions about behavior when there are multiple equilibria. Anderson, Goeree, and Holt (2001) introduce the more general notion of *stochastic potential* and use it to explain results of coordination game experiments. Models of learning and evolution have also been widely used in the study of coordination games (see references in Ochs, 1995).

Questions

1. Consider the game in the table below, and find all Nash equilibria in pure strategies. Is this a coordination game?

Row Player:	Column Player:			
	Left (26%)	Middle (8%)	Non-Nash (68%)	Right (0%)
Top (68%)	200, 50	0, 45	10, 30	20, -250
Bottom (32%)	0, -250	10, -100	30, 30	50, 40

2. The numbers under each decision label in the above table show the percentage of people who chose that strategy when the game was played only once (Goeree and Holt, 2001). Why is the column player's most commonly used strategy labeled as "Non-Nash" in the table? (Note: it was not given this label in the experiment.) Conjecture why the Non-Nash decision is selected so frequently by column players?
3. Use the spreadsheet for the minimum-effort coordination game to fill in average efforts for the following Table. Use initial belief probabilities of $1/8$ (or 0.125) for each of the elements in column D between rows 11 to 17. Then discuss the effects of the error parameter on these predictions.

	Average Effort Predictions			
	$\mu = 1$	$\mu = 5$	$\mu = 10$	$\mu = 20$
High Cost ($c = 75$)				
Low Cost ($c = 25$)				

4. Consider the game shown below, which gives the Column Player a safe decision, S. Row's payoffs depend on a parameter X that changes from one treatment to the other. Goeree and Holt (2001) report that the magnitude of X affects the frequency with which the Row Player chooses Top, so the issue is whether this effect is predicted in theory. Find all Nash equilibria in pure and mixed strategies for $X = 0$, and for $X = 400$, under the assumption that each person is risk neutral.

		Column Player:		
		L	R	S
Row Player:	Top	90, 90	0, 0	X, 40
	Bottom	0, 0	180, 180	0, 40

5. (Advanced) Find the mixed-strategy equilibrium for the game shown in question 1. (Hint: In this equilibrium, column randomizes between Left and Middle. The probabilities associated with Non-Nash and Right are zero, so consider the truncated 2x2 game involving only Left and Middle for column and Top and Bottom for row. Finally, check to be sure that column is not tempted to deviate to either of the two strategies not used.)

Part IV. Market Experiments

The games in this part represent many of the standard market models that are used in economic theory. In a *monopoly*, there is a single seller, whose production decision determines the price at which the output can be sold. This model can be generalized by letting firms choose quantities independently, where the aggregate quantity determines the market price. This *Cournot* setup is widely used in economic theory, and it has the intuitive property that the equilibrium price outcomes are decreasing in the total number of sellers. The experiment discussed in Chapter 13 implements a setup with linear demand and constant cost, which permits an analysis of the effects of changing the number of competitors, including monopoly. The Cournot model may be appropriate when firms pre-commit to production decisions, but in many situations it may be more realistic to model firms as choosing prices independently.

Many important aspects of a market economy pertain to markets of intermediate goods along a supply chain. Chapter 14 considers experiments based on a very simple case of a monopoly manufacturer selling to a monopoly retailer. Frictions and distortions (“the bullwhip effect”) that may arise along longer supply chains are also discussed.

Chapters 15 and 16 pertain to price competition, imperfections, and the effects of alternative trading institutions when buyers are not simulated. Buyer and seller activities are essentially symmetric in the *double auction*, except that buyers tend to bid the price up, and sellers tend to undercut each other’s prices. All bids and asks are public information, as is the transaction price that result when someone accepts the price proposed by someone on the other side of the market. This strategic symmetry is not present in the *posted-offer auction*, where sellers post prices independently, and buyers are given a chance to purchase at the posted prices (further bargaining and discounting is not permitted). The posted offer market resembles a retail market, with sellers producing “to order” and selling at catalogue or “list” prices. In contrast, the double auction more closely resembles a competitive “open outcry” market. Collusion and the exercise of market power cause larger distortions in markets with posted prices, although these effects may be diminished if sellers are able to offer secret discounts.

The effects of asymmetric information about product quality are considered in Chapter 17. When sellers can select a quality grade and a price, the outcomes may be quite efficient under full information condition, but qualities fall to low levels when buyers are unable to observe the quality grade prior to purchase.

Many markets are for commodities or assets that provide a stream of benefits over time. For example, a share of stock may pay dividends over time.

Markets for assets are complicated by the fact that ownership provides two potential sources of value, i.e. the benefits obtained each period (services, dividends, etc.) and the capital gain (or loss) in the value of the asset. Asset values may be affected by market fundamentals like dividends and the opportunity cost of money (the interest that could be earned in a safe account). In addition, values can be affected by expectations about future value, and such expectations-driven values can cause trading prices to deviate from levels determined by market fundamentals. The final chapter in this section summarizes some experimental work on such asset markets where price bubbles and crashes are often observed.

Chapter 13. Monopoly and Cournot Markets

By restricting production quantity, a monopolist can typically obtain a higher price, and profit maximization involves finding a balance between the desire to charge a high price and the desire to maintain a reasonable sales quantity. Subjects in laboratory experiments with simulated buyers are able to adjust quantity so that prices are at or near the profit-maximizing monopoly level. This type of experiment can be done with the Veconlab Cournot program (CR) when the number of sellers is set to 1. This seller selects a quantity that determines the resulting price. The setup allows price to be subject to random shocks, which adds interest and realism.

The analysis of monopoly provides a natural bridge to the most widely used model of that oligopoly interaction of several sellers. The key behavioral assumption of this model is that each seller takes the others' quantity choices as given and fixed, and then behaves as a monopolist maximizing profits for the resulting residual demand. The popularity of the Cournot model is, in part, due to the intuitive prediction that the market price will decrease from monopoly levels toward competitive levels as the number of sellers is increased, and this tendency can be verified with laboratory experiments (e.g., using the CR game), or with the hand-run version that is given in the appendix.

I. Monopoly

A monopolist is defined as being a sole seller in a market, but the general model of monopoly is central to the analysis of antitrust issues because it can be applied more widely. For example, suppose that all sellers in a market are somehow able to collude and set a price that maximizes the total profit, which is then divided among them. In this case, the monopoly model would be relevant, either for providing a prediction of price and quantity, or as a benchmark from which to measure the success of the cartel.

The monopoly model is also applied to the case of one large firm and a number of small "fringe" firms that behave competitively (expanding output as long as the price that they can get is above the cost of each additional unit of output). The behavior of the firms in the competitive fringe can be represented by a supply function, which shows the quantity provided by these firms in total as a function of the price. Whether or not a group of small firms can be counted upon to behave competitively is a behavioral issue that might be investigated with an experiment. Let this fringe supply function be represented by $S(P)$, which is increasing in price if marginal costs are increasing for the fringe firms. Then the *residual demand* is defined to be the market demand, $D(P)$, minus the fringe supply, so the residual demand is $R(P) = D(P) - S(P)$. This residual demand

function indicates a relationship between price and the sales quantity that is not taken by the fringe firms; this is the quantity that can be sold by the large dominant firm. In this situation, it may be appropriate to treat the dominant firm as a monopolist facing a demand function $Q = R(P)$.

From the monopolist's point of view, the demand (or residual demand) function reveals the amount that can be sold for each possible price, with high prices generally resulting in lower sales quantities. It is useful to invert this demand relationship and think of price as a function of quantity, i.e. selling a larger quantity will reduce price. For example, a linear inverse demand function would have the form: $P = A - BQ$, where A is the vertical intercept of demand in a graph with price on the vertical axis, and $-B$ is the slope, with $B > 0$. The experiments to be discussed in this chapter all have a linear inverse demand, which for simplicity, will be referred to as the demand function.

The left side of Figure 13.1 shows the results of a laboratory experiment in which each person had the role of a monopoly seller in a market with a constant cost of \$1 per unit and a linear demand curve: $P = 13 - Q$, where P is price and Q is the quantity selected by the monopolist. Since the slope is minus one, each additional unit of output raises the cost by \$1 and reduces the price by \$1. The vertical axis in the figure is the average of the quantity choices, and it is apparent that the participants quickly settle on a quantity of about 6. The presumption is that this is the profit maximizing choice, as will be verified next.



Figure 13.1. Monopoly Versus Duopoly

The demand curve used in the experiment is also shown in the top two rows of Table 13.1. Notice that as price in the second row decreases from 12 to 11 to 10, the sales quantity increases from 1 to 2 to 3. The total revenue, PQ , is shown in the third row, and the total cost, which equals quantity, is given in the fourth row. Please take a minute to fill in the missing elements in these rows, and to subtract cost from revenue to obtain the profit numbers that should be entered in the fifth row. Doing this, you should be able to verify that profit is maximized with a quantity of 6 and a price of 7.

Table 13.1 Monopoly with Linear Demand and Constant Cost

Quantity	1	2	3	4	5	6	7	8	9	10	11	12
Price	12	11	10	9	8	7	6	5	4	3	2	1
TR	12	22	30									
TC	1	2	3									
Profit	11	20	27									
MR	12	10	8									
MC	1	1	1									

Even though the profit calculations are straightforward, it is instructive to consider the monopolist's decision as quantity is increased from 1 to 2, and then to 3, while keeping an eye on the effects of these increases on revenue and cost *at the margin*. The first unit of output produced yields a revenue of 12, so the marginal revenue of this unit is 12, as shown in the MR row at the left. An increase from $Q = 1$ to $Q = 2$ raises total revenue from 12 to 22, which is an increase of 10, as in the MR row. These additional revenues for the first and second units are greater than the cost increases of \$1 per unit, so the increases were justified. Now consider an increase to an output of 2. This raises revenue from 22 to 30, an increase of 8, and this marginal revenue is again greater than the marginal cost of 1. Notice that profit is going up as long as the marginal revenue is greater than the marginal cost, a process that continues until the output reaches the optimal level of 6, as you can verify.

One thing to notice about the MR row of table 13.1 is that each unit increase in quantity, which reduces price by 1, will reduce marginal revenue by 2, since marginal revenue goes from 10 to 8 to 6 etc. This fact is illustrated in Figure 13.2, where the demand line is the outer thick line, and the marginal revenue line is the thick dashed line. The marginal revenue line has a slope that is twice as negative as the demand line. The MR line intersects the horizontal marginal cost line at a quantity of 6. Thus the graph illustrates what you will see when you fill out the table, i.e. that marginal revenue is greater than marginal cost

for each additional unit, the 1st, 2nd, 3rd, 4th, 5th, and 6th, but the marginal revenue of the 7th unit is below marginal cost, so that unit should not be added. (The marginal revenues in the table will not match the numbers on the graph exactly, since the marginal revenues in the table are for going from one unit up to the next. For example, the marginal revenue in the 6 column of the table will be 2, which is the increase in revenue from going from 5 to 6 units. Think of this as 5.5, and the marginal revenue for 5.5 in the figure is exactly 2.)

In addition to the table and the graph, it is useful to redo the same derivation of the monopoly quantity using simple calculus. (A brief review of the needed calculus formulas is provided in the appendix to this chapter.) Since demand is: $P = 13 - Q$, it follows that total revenue, PQ , is $(13 - Q)Q$, which is a quadratic function of output: $13Q - Q^2$. Marginal revenue is the derivative of total revenue, which is $13 - 2Q$, which is a line that starts at a value of 13 when $Q = 0$ and declines by 2 for each unit increase in quantity, as shown in Figure 13.2. Since marginal cost is 1, it follows that marginal revenue equals marginal cost when $1 = 13 - 2Q$, ie. when $Q = 6$, which is the monopoly output for this market.




Figure 13.2. Monopoly Profit Maximization

The case against monopoly can be illustrated in Figure 13.2, where the monopoly price is \$7 and the marginal cost is only \$1. Thus buyers with valuations between \$1 and \$7 would be willing to pay amounts that covers cost, but the monopoly does not satisfy this demand, in order to keep price and earnings high. As shown in Chapter 2, the loss in buyers' value can be measured

as the area under the demand curve, and the net loss is obtained by looking at the area below demand, to the right of 6 and above the marginal cost line in the figure. This triangular area is a measure of the welfare cost of monopoly, as compared with the competitive outcome where price is equal to marginal cost (\$1) and the quantity is 12.

It is important to mention that the convergence of the quantity to the monopoly prediction on the left side of Figure 13.1 was observed in a context where the demand side of the market was simulated. This kind of experiment is probably appropriate if one is thinking of a market with a very large number of consumers, none of whom have any significant size or power to bargain for reductions from the monopoly price levels.

II. Cournot Duopoly

Next consider what would happen if a second firm were to enter the market that was considered in the previous section. In particular, suppose that both firms have constant marginal costs of \$1, and that they each select an output quantity, with the price then being determined by the sum of their quantities from the top two rows of Table 13.1. This duopoly structure was the basis for the second part of the experiment summarized in Figure 13.1. Notice that the quantity per seller starts out at the monopoly level of 6 in round 9, which forces price down to \$1. It is not surprising that the quantity would fall when profits, which had been high, are driven to 0 in this round. The incentive to cut output can be seen from the graph in Figure 13.2. Suppose that one seller (the entrant) knew the other would produce a quantity of 6. If the entrant were to produce 0, the price would stay at the monopoly level of \$7. If the entrant were to produce 1, the price would fall to \$6, etc. These price/quantity points for the entrant are shown as the large dots labeled “Residual Demand” in figure 13.2. The marginal revenue for this residual demand curve has a slope that is twice as steep, as shown by the thin dotted line that crosses MC at a quantity of 9. This crossing translates determines an output of 3 for the entrant, since the incumbent seller is producing 6. In summary, when one firm produces 6, the best response of the other is to choose a quantity of 3. This suggests why the quantities, which start at an average of 6 for each firm in period 9, begin to decline in subsequent periods, as shown on the right side of Figure 13.1. The outputs fall to an average of 4 for each seller, which suggests that this is the equilibrium, in the sense that if one seller is choosing 4, the best response of the other is to choose 4 also.

In order to show that the Cournot equilibrium is in fact 4 units per seller, we need to run through some other best response calculations, which are shown in Table 13.3. This table shows a firm’s profit for the example under consideration for each if its own output decisions (increasing from bottom to top) and for each output decision of the other firm (increasing from left to right). First consider the

column labeled “0” on the left side, i.e. the column that is relevant when the other firm’s output is 0. If the other firm produces nothing, then this is the monopoly case, and the profits in this column are just copied from the monopoly profit row of Table 13.1: a profit of 11 for an output of 1, 20 for an output of 2, etc. Note that the highest profit in this column is 36 in the upper left corner, at the monopoly output of 6. This profit has been highlighted by putting a dark border around the box. To be sure that you understand the table, you should fill in the two missing numbers in the right column. In particular, if both have outputs of 6, then the total output is 12, the price is ____, the total revenue for the firm is ____, the total cost is ____, and then the profit is 0 (please verify). Some of the other best response payoffs are also indicated by boxes with dark borders. Recall that the best response to another firm’s output of 6 is to produce 3 (as seen in Figure 13.2), and this is why the box with a payoff of 9 in the far-right column has a dark border.

Table 13.3 A Seller’s Own Profit Matrix and Best Responses
Key: Columns for Other Seller’s Outputs, and Rows for Own Outputs

	0	1	2	3	4	5	6
6	36	30	24	18	12	6	
5	35	30	25	20	15	10	
4	32	28	24	20	16**	12	8
3	27	24	21	18*	15	12	9
2	20	18	16	14	12	10	8
1	11	10	9	8	7	6	5

A Nash equilibrium in this duopoly market is a pair of outputs, such that each seller’s output is the best response to that of the other. Even though 3 is a best response to 6, the pair (6 for one, 3 for the other) is not a Nash equilibrium (problem 1). The payoff of 16, marked with a double asterisk in the table, is the location of a Nash equilibrium, since it indicates that an output of 4 is a best response to an output of 4, so if each firm were to produce at this level, there would be no incentive for either to change unilaterally. Of course, if they could coordinate on joint output reductions to 3 each, the total output would be 6 and the industry profit would be maximized at the monopoly level of 26, or 18 each. This joint maximum is indicated by the single asterisk at the payoff of 18 in the table. This joint maximum is not a Nash equilibrium, since each seller has an incentive to expand output (problem 2).

The fact that the Nash equilibrium does not maximize joint profit raises an interesting behavioral question, i.e. why couldn’t the subjects in the experiment

somehow coordinate on quantity restrictions to raise their joint earnings? The answer is given in the legend on the right side of Figure 13.1, which indicates that the matchings were random for the duopoly phase. Thus each seller was matched with a randomly selected other seller in each round, and this switching would make it difficult to coordinate quantity restrictions. In experiments with fixed matchings, such coordination is often observed, particularly with only two sellers. Holt (1985) reported patterns where sellers sometimes “walked” the quantity down in unison, e.g. both duopolists reducing quantity from 7 to 6 in one period, and then to 5 in the next, etc. This kind of “tacit collusion” occurred even though sellers could not communicate explicitly. There were also cases where one seller produced a very large quantity, driving price to 0, followed by a large quantity reduction in an effort to send a threat and then a conciliatory message and thereby induce the other seller to cooperate. Such tacit collusion is less common with more than 2 sellers. Part of the problem in these quantity-choice models is that when one seller cuts output, the other’s has a unilateral incentive to expand output, so when one shows restraint, the other has greater temptation.

The Nash equilibrium just identified is also called a Cournot equilibrium, after the French mathematician who provided an equilibrium analysis of duopoly and oligopoly models in 1838. Although the Cournot equilibrium is symmetric in this case, there can be asymmetric equilibria as well. A further analysis Table 13.3 indicates that there are other Nash equilibria, but all with a total quantity of 8, and hence an average of 4. Can you find these (problem 3)?




Figure 13.3. Cournot Duopoly

As was the case for monopoly, it is useful to illustrate the duopoly equilibrium with a graph. If one firm is producing an output of 4, then the other can produce an output of 1 (total quantity = 5) and obtain a price of 8, as shown by one of the residual demand dots in Figure 13.3. Think of the vertical axis as having shifted to the right at the other firm's quantity of 4, as indicated by the vertical dotted line, and the residual demand dots yield a demand curve with a slope of minus 1. As before, marginal revenue will have a slope that is twice as negative, as indicated by the heavy dashed line in the figure. This line crosses marginal cost just above the quantity of 8, which represents a quantity of 4 for this firm, since the other is already producing 4. Thus the quantity of 4 is a best response to the other's quantity of 4.

III. Cournot Oligopoly

The classroom experiment shown in Figure 13.4 was done with fixed matchings, for duopoly markets (left side) and then for three-firm markets (right side). Despite the fixed nature of the matchings, participants were not able to coordinate on output reductions below the duopoly prediction of 4 per seller. In the triopoly treatment, the outputs converge to 3 on average. Notice that outputs of 3 for each of three firms translates into an industry output of 9, as compared with 8 for the duopoly case and 6 for the monopoly case. Thus we see that increases in the number of sellers raise the total quantity and reduce price toward competitive levels.




Figure 13.4. A Classroom Oligopoly Experiment with Fixed Matchings

It is possible to redraw Figure 13.3 to show that if two firms each produce 3 units each, then the remaining firm would also want to produce 3 units (problem 4). Instead, we will use a simple calculus derivation. Suppose that there are two firms which produce X units each, so their total is $2X$. The remaining firm must choose a quantity Q to maximize its profit. The residual demand for the remaining firm is: $A - (2X)B - BQ$, where $A = 13$ and $B = 1$ in our example. Multiplying the residual demand by Q , we obtain the total revenue function for the remaining firm: $AQ - (2X)BQ - BQ^2$, which has two terms that are linear in Q and one term that is quadratic. The derivative of this total revenue expression with respect to Q is: $A - 2XB - 2BQ$. (The derivatives of the linear terms are just the coefficients of those terms and the derivative of the quadratic term is taken by moving the 2 in the exponent down with the B .) This derivative is the expression for the marginal revenue that is associated with the residual demand function (the thick dashed line in Figure 13.3). By equating marginal revenue to marginal cost, one obtains a single equation, $A - 2XB - 2BQ = 1$. This equation could be solved for Q , which would show the best response output Q for any given level of the other firms' common outputs, X units each. Just as there was a symmetric equilibrium in the duopoly case, there will be one here, with $X = Q$, and with this substitution, the marginal revenue equals marginal cost condition becomes: $A - 2QB - 2BQ = 1$, which can be solved for $Q = (A-1)/4B$. With $A = 13$ and $B = 1$, this reduces to $12/4 = 3$, which is the Cournot equilibrium for the triopoly case.

IV. Extensions

The Cournot model is perhaps the most widely used model in theoretical work in Industrial Organization. Its popularity is based, to some extent, on the prediction that the equilibrium price will be a decreasing function of the number of sellers in a symmetric model. This prediction is supported by evidence from laboratory experiments that implement the Cournot assumption that firms select quantities independently (e.g. Holt, 1985).

The obvious shortcoming of the Cournot model is the specific way in which price is determined. The implicit assumption is that firms make quantity decisions independently, and then price is cut so that all production can be sold. For example, you might think a situation in which the quantities have been produced, so the short run supply curve is vertical at the total quantity, and price is determined by the intersection of this vertical supply function with the market demand function. In other words, the price-competition phase is extremely competitive. As we will see in later chapters, there is some game-theoretic and experimental evidence to support this view. Even so, a Cournot equilibrium would not be appropriate if it is price that firms set independently, and then produce to fill the orders that arrive. Independent price choice may be an appropriate assumption if firms mail out catalogues or post "buy now" prices on

the internet, with the ability to produce quickly to fill orders. Some of the richer models of price competition with discounts will be discussed in the coming chapters.

Appendix: Optional Quick Calculus Review

The derivative of a linear function is just its slope. So for the demand function in Figure 13.2, the slope is -1 (each unit increase in quantity decreases price by \$1). To find the slope using calculus, we need a formula: the derivative of BQ with respect to Q is just B , for any value of the slope parameter B . Thus the derivative of $(-1)Q$ is -1 . The slope of the demand curve is the derivative of $13 - Q$ and we know that the derivative of the second part is -1 , which is the correct answer, so the derivative of 13 must be 0 . In fact, the derivative of any constant is zero. To see this, note that the derivative is the slope of a function, and if you graph a function with a constant height, then the function will have a slope of 0 in the same manner that a table top has no slope. For example, consider the derivative of a more general linear demand function: $A - BQ$. The intercept, A , is just a constant (it does not depend on Q , which is variable, but rather it stays the same). So the derivative of A is 0 , and the derivative of $-BQ$ is $-B$, and therefore the derivative of $A - BQ$ is $0 - B = -B$. Here we have used the fact that the derivative of the sum of two functions is the sum of the derivatives. To summarize (ignore rules 4 and 5 for the moment):

1) **Constant Function:** $dA/dQ = 0$.

The derivative of a constant like A is just 0 .

2) **Linear Function:** $d(KQ)/dQ = K$.

The derivative of a constant times a variable, which has a constant slope, is just the constant slope parameter, i.e. the derivative of KQ with respect of Q is just K .

3) **Sum of Functions:** The derivative of the sum of two functions is the sum of the derivatives.

4) **Quadratic Function:** $d(KQ^2)/dQ = 2KQ$.

The derivative of a quadratic function is obtained by moving the 2 in the exponent down, so the derivative of KQ^2 is just $2KQ$.

5) **Power Function:** $d(KQ^x)/dQ = xKQ^{x-1}$.

The derivative of a variable raised to the power x is obtained by moving the x down and reducing the power by 1 , so the derivative of KQ^3 is just $3KQ^2$, the derivative of KQ^4 is $4KQ^3$, and in general, the derivative of KQ^x with respect to Q is xKQ^{x-1} .

For example, the monopolist being discussed has a constant marginal cost of \$1 per unit, and since there is no fixed cost, the total cost of producing Q units is just Q dollars. Think of this total cost function as being the product of 1 and Q , so the derivative of $1Q$ is just 1 using rule 2 above. If there had been a fixed cost of F , then the total cost would be $F + Q$. Note that F is just a constant, so its derivative is 0 (rule 1), and the derivative of this total cost function is the derivative of the first part (0) plus the derivative of the second part (1), so marginal cost is again equal to 1.

Next consider a case where demand is $P = A - BQ$, which has a vertical intercept of A and a slope of $-B$. The total revenue function is obtained by multiplying by Q to get: the total revenue function: $AQ - BQ^2$, which has a linear term with slope A and a quadratic term with a coefficient of $-B$. We know that the derivative of the linear part is just A (rule 2). The fourth rule indicates how to take the derivative of $-BQ^2$, you just move the 2 in the exponent down, so the derivative of this part is $-2BQ$. We add these two derivatives together to determine that the derivative of the total revenue function is: $A - 2BQ$, so marginal revenue has the same vertical intercept as demand, but the slope is twice as negative. This is consistent with the calculations in Table 13.1, where each unit increase in quantity reduces price by \$1 and reduces marginal revenue by \$2.

Questions

1. Use Table 13.3 to show that outputs of 6 and 3 do not constitute a Nash equilibrium to the duopoly model that is the basis for that table.
2. Use Table 13.3 to show that outputs of 3 for each firm do not constitute a Cournot Nash equilibrium.
3. Find a Nash equilibrium in Table 13.3 with the property the total quantity is 8 but one seller produces more than the other. Therefore, you must specify what the two outputs are, and you must show that neither seller has a unilateral incentive to deviate.
4. Redraw Figure 13.3 for the triopoly case, putting the vertical dotted line at a quantity of 6 (three for each of two other firms) and show that the residual marginal revenue for the remaining firm would cross marginal cost in a manner that make the output of 3 a best response for that firm.
5. (requires calculus) Suppose that demand is linear: $P = A - BQ$, and marginal cost is constant at C . Use the approach taken at the end of Section III to find the Cournot triopoly equilibrium in terms of the parameters A , B , and C . (To check your answer, set $C = 1$ and compare with the calculations in the reading.)

6. (requires calculus) Find the Cournot duopoly equilibrium for the model given in problem 5.
7. (requires calculus) Find a symmetric Cournot equilibrium for the model in question 5 with the number of firms set to N instead of 3. To check your answer, set $N = 3$ and compare.
8. (requires calculus) Show that the price is a decreasing function of the number of firms, N , in the Cournot equilibrium model for question 7.
9. (for the Appendix) Find the derivatives with respect to Q of: a) Q , b) $2Q$, c) $5 + 10Q$, d) 5 , and e) $A - BQ$ where A and B are positive parameters.
10. (for the Appendix) Find the derivatives with respect to y of : a) $10y$, b) $10y + 3$, c) $2y^2$, d) $y^2 + 3y$, and e) $13 - 4y + y^2$.

Chapter 14. Vertical Market Relationships

A complex economy is characterized by considerable specialization along the supply chain; with connections between manufacturers of intermediate products, manufacturers of final products, wholesalers, and retailers. There is some theoretical and empirical evidence that frictions and market imperfections may be induced by these “vertical” supply relationships. This chapter begins by considering the very specific case of a vertical monopoly, i.e. interaction between an upstream monopoly wholesaler selling to a firm that has a local monopoly in a downstream retail market. In theory, the vertical alignment of two monopolists can generate retail prices that are even higher than those that would result with a single integrated firm, i.e. a merger of the upstream and downstream firms. These effects can be investigated with the Veconlab Vertical Monopoly program (VM) or with the instructions for a hand-run version that can be found in Badasyan et al. (2004), which is the basis for most of the discussion in this chapter. The addition of more layers in the supply chain raises the question of the extent to which demand shocks at the retail level may cause even larger swings in orders and inventories at higher levels, a phenomenon that is known as the “bullwhip effect.”

I. Double Marginalization

Since Adam Smith, economists have been known for their opposition to monopoly. Monopolization (carefully defined) is a crime in U.S. antitrust law, and horizontal mergers are commonly challenged by the antitrust authorities if the effect is to create a monopoly. In contrast, there is much more leniency shown toward vertical mergers, i.e. between firms that do business with each other along a supply chain. One motivation for this relative tolerance of vertical mergers is that two monopolies may be worse than one, at least when these two monopolies are arrayed vertically. The intuition is that each monopolist restricts output to the point where marginal revenue equals marginal cost, a process that reduces production below competitive levels and results in higher prices. When this marginal restriction is made by one firm in the supply chain, it raises the price to a downstream firm, which in turn raises price again at the retail level, and the resulting “double marginalization” may be worse than the output restriction that would result if both firms were merged into a single, vertically-integrated monopolist.

The effects of double marginalization can be illustrated with a simple laboratory experiment based on a linear demand function: $P = 12 - Q$, which is the demand at the retail level. There is an upstream manufacturer with a constant marginal cost of production of \$2. The retailer purchases units and then

packages them, thereby incurring an additional cost of \$2 per unit. Thus there is a total cost (manufacturing plus retail) of \$4 per unit. This would be the marginal cost of a single vertically integrated firm that manufactures and sells the good in a retail monopoly market.

First consider the monopoly problem for this integrated firm, who will find the output that equates the marginal cost of 4 with the marginal revenue, as explained in the previous chapter. With a demand of $P = 12 - Q$, and an associated total revenue of $12Q - Q^2$, the marginal revenue will be $12 - 2Q$, since the marginal revenue is twice as steep as the demand curve. Equating this to the marginal cost of 4 and solving for Q yields a monopoly output of 4 and an associated price of 8. These calculations are illustrated by the red demand, MR, and MC lines on the left side of Figure 14.1, where the intersection at a quantity of 4 leads to a price of \$8.




Figure 14.1. A Classroom Vertical Monopoly Experiment

Next, consider the case in which the two firms are not vertically integrated. Each firm will be equating marginal revenue and marginal cost. It turns out that in these kinds of problems, it is easiest to start at the end of the supply chain (retail) and work backwards to consider the manufacturing level last. For any wholesale price, W , charged by the manufacturer, the marginal cost to the retailer will be $W + 2$, since each unit sold at retail must be purchased at a wholesale price W and then processed at a cost of 2. Thus the downstream retailer is a monopolist who will equate this marginal cost, $W + 2$, to the marginal

revenue, $12 - 2Q$, to obtain a single equation: $W + 2 = 12 - 2Q$ that can be solved for the retail order quantity Q as a function of the wholesale price: $Q = 5 - W/2$.

Next consider the upstream firm's optimal decision. The retail order is the demand function faced by the wholesaler, i.e. for each wholesale price it determines how much is sold to the retailer. Just as we draw the inverse of retail demand with price on the vertical axis, it is useful to graph the inverse of the wholesale demand with the wholesale price on the vertical axis. To do this, invert the wholesale demand function, $Q = 5 - W/2$, to obtain the wholesale price as a function of Q : $W = 10 - 2Q$. This inverse wholesale demand function is graphed on the left side of Figure 14.1 as the thin blue line with a vertical intercept of 10 and a horizontal intercept of 5. The associated marginal revenue line has the same vertical intercept but a slope that is twice as steep, as shown by the blue line that has a horizontal intercept of 2.5. The wholesaler will want to increase output until this marginal revenue is equal to the wholesale marginal cost of 2, as shown in the figure, where the intersection occurs at a quantity of 2, which determines a wholesale price of $W = 10 - 2Q = 10 - 4 = 6$. When the wholesale price is 6, the marginal cost to the retailer is $6 + 2 = 8$, as shown by the horizontal blue line at \$8 in the figure. This marginal cost line intersects the red retail marginal revenue line at a quantity of 2, which determines a retail price of \$10. To summarize, for the market in Figure 14.1 with a monopoly output of 4 and a price of 8, the presence of two vertically stacked monopolists reduces the output to 2 and raises the retail price to \$10. Thus the output restriction is greater for two monopolists than would be the case for one vertically integrated firm.

II. Some Experimental Evidence

The first 5 rounds of the experiment shown in Figure 14.1 were done with pairs of students, with one retailer and one wholesaler in each pair. The wholesale price prediction (\$4) and the retail price prediction (\$10) are indicated by the horizontal lines, and the average prices are indicated by the dots converging to those lines. Notice that wholesale prices are a little low relative to theoretical predictions; this could be due to the fact that the buyer, the retailer, is a participant in the experiment, not a simulated buyer as was the case for the retail demand in this experiment and in the monopoly experiments discussed in the previous chapter.

The center part of Figure 14.1 shows 5 rounds in which each of the 12 participants from the first part was put into the market as a vertically integrated monopolist facing the retail demand determined by $P = 12 - Q$. The average prices converge to the monopoly level of \$8 that is indicated by the horizontal line for periods 6-10. This reduction in retail price, from about 10 to about 8, resulted in an increase in sales quantity, and less of a monopoly output restriction, as predicted by theory. In addition, the profits of the integrated firm are larger, since

by definition, profit is maximized at the monopoly level. This increase in profitability associated with vertical integration can be verified by calculating the profits for each firm separately (problem 1).

Vertical integration may not always be feasible or desirable, or even cost efficient in all cases. An alternative, which works in theory and is observed in practice, is for the upstream firm to require the retailer to pay a fixed franchise fee in order to be able to sell the product at all. In particular, the wholesale firm chooses a wholesale price and a franchise fee. The retailer then may reject the arrangement, in which case both earn 0, or accept and place an order for a specified number of units. The idea behind the franchise fee is to lower the wholesale price from the \$6 charged when the firms are not integrated, to a level of \$2 that reflects the true wholesale marginal cost. The retailer's marginal cost would then be the new low wholesale price of \$2 plus the retail marginal cost of \$2, which equals \$4, the true social cost of producing and retailing each unit. When the retailer faces this marginal cost of \$4, it will behave as an integrated monopolist, choosing an output of 4, charging the monopoly price of \$8, and earning the monopoly profit of $8 \cdot 4 - 4 \cdot 4 = 16$. If this were the whole story, the wholesaler would have zero profits, since the units were sold at a price that equals marginal cost. The wholesaler can recover some of this monopoly profit by charging a franchise fee. A fee of \$8 would split the monopoly profit, leaving \$8 for each. In theory, the wholesaler could demand \$15.99 of the profit, leaving one penny for the retailer, under the assumption that the retailer would prefer a penny to a payoff of zero that results from rejecting the franchise contract. Intuition and other experimental evidence, however, suggest that such aggressive franchise fees would be rejected (see the Chapter 23 on Bargaining). In effect, retailers will be likely to reject contract offers that are viewed as being unfair, and such rejections may even have the effect of inducing the wholesaler to make a more moderate demand in a subsequent period.

The right side of Figure 14.1 shows average prices and franchise fees for periods 11-15 of the experiment, in which the participants were again divided into wholesaler/retailer pairs. The average franchise fees are the light dots that lie between \$6 and \$8. The wholesalers were not able to demand a high share of the monopoly profit in this experiment. Instead, it is apparent that they did not lower the price all of the way to their wholesale cost level of \$2. Average prices do fall below the wholesale price levels in the first 5 rounds, but they only fall to a level of about \$3. Thus fairness considerations seem to be preventing the franchise fee treatment from solving the vertical monopoly problem.

II. Extensions: The Bullwhip Effect

Durham (2000) uses posted-offer institution and a monopolistic upstream firm to compare treatments with one monopolistic downstream firm to three

competitive downstream firms. She finds that the presence of competitive downstream firms leads to outcomes similar to those under vertical integration. In effect, the competition at the downstream level takes out one of the sources of monopoly marginalization, so the price and quantity outcomes approach those that would result from a single monopolist. To summarize, the double marginalization problem associated with two monopolies may be alleviated to some extent by a vertical merger or by the introduction of competition downstream. Without these kinds of corrections, this analysis of the previous section would apply to an even greater extent with a longer supply chain, e.g. with a monopoly manufacturer selling to a wholesaler, who sells to a distributor, who sells to a retailer, where the firms at each stage are monopolists. In this context, the effects of successive “marginalizations” get compounded, unless these effects were somehow diluted by competition between sellers at each level.

A second source of inefficiency in a long supply chain is the possibility that orders and inventories are not well coordinated, and that information is not transmitted efficiently from one level to the next. For example, Procter & Gamble found that diaper orders by distributors too variable relative to consumer demand, and Hewlett-Packard found that printer orders made by retail sellers are much more variable than consumer demand itself. These and other examples are discussed in Lee et al. (1997a, 1997b).

There is a long tradition in business schools of putting M.B.A. students into a supply chain simulation, which is known as the “Beer Game.” There are four vertical levels: manufacturing, wholesale, distribution, and retail (Forrester, 1961). The participants in these classroom games must typically fill orders from the level below from inventory, and then place new orders to the level above. There is a cost of carrying unsold inventory, and there is also a cost associated with not being able to fill an order from below, i.e. the lost profit per unit on sales. The setup typically involves having a stable retail demand for several rounds before it is subjected to an unexpected and unannounced increase that persists in later periods. The effect of this demand increase is to cause larger and larger fluctuations in orders placed upstream, which is known as the “bullwhip effect.”

In classroom experiments with the Beer Game, the upstream sellers tend to attribute these large fluctuations with exogenous demand shifts, despite the fact that most of the fluctuation is due to the reactions of those lower in the supply chain (Sterman, 1989). See Lee, Padmanabhan, and Whang (1997a) for a theoretical analysis of quasi-rational behavior that may cause orders to fluctuate more than sales, and that may cause this distortion to be greater as one moves up the supply chain. For more recent experiments, see Holt and Moore (2004a).

Questions

1. Find the predicted profit for a monopoly seller in the market described in section I.
2. Find the profits for wholesaler and retailer for the market described in section I when the wholesale price is 6 and the retail price is 10, as predicted. Show that profits for these two firms are, in total, less than that of a vertically integrated monopolist, as calculated in the previous question.

Chapter 15. Market Institutions and Power

Traders in a double auction can see all transactions prices and the current bid/ask spread, as is the case with trading on the New York Stock Exchange. The double auction is an extremely competitive institution, given the temptation for traders to improve their offers over time in order to make trades at the margin. In contrast, markets with posted prices (Bertrand or Posted-Offer) allow sellers to pre-commit to fixed, take-it-or-leave-it prices that cannot be adjusted during a trading period. The focus of this chapter is price and efficiency outcomes of markets with posted prices, in particular when sellers possess *market power*, which is, roughly speaking, the ability of a firm to raise price profitably above competitive levels. In most cases, the market efficiency is higher in the double auction, since the price flexibility built into that institution tends to bring outcomes closer to a competitive (supply equals demand) outcome. Efficiency is also enhanced by the incorporation of more price flexibility into markets with posted prices, e.g., by allowing sellers to offer secret discounts to specific buyers. These various market institutions can be implemented “by hand” with the instructions in the appendix to Davis and Holt (1993), or with the Veconlab double-auction (DA), posted-offer (PO), or Bertrand (/BR) programs, which provide flexible setup options and automatic data and graphical summaries.

I. Introduction

The common perception that laboratory markets yield efficient competitive outcomes is surprising given the emphasis on market imperfections that pervades theoretical work in industrial organization. This apparent contradiction is resolved by considering the effects of trading institutions: As discussed in Chapter 2, competitive outcomes are typical in “double auction” markets, with rules similar to those used in many centralized financial exchanges. But most markets of interest to industrial organization economists have different institutional characteristics; sellers post prices and buyers must either buy at those prices or engage in costly search and negotiation to obtain discounts. Unlike more competitive double auctions, the performance of markets with posted prices can be seriously impeded by the presence of market power, price-fixing conspiracies, and cyclical demand shocks.

In many markets, it is common for prices to be set by the traders on the thin side of the market. Sellers, for example, typically post prices in retail markets. For this reason, theoretical (Bertrand) models are often structured around an assumption that prices are listed simultaneously at the beginning of each “period.” Using laboratory experiments, it is possible to make controlled comparisons between markets with posted prices and more symmetric institutions such as the “double

auction," where both buyers and sellers post bids and asks in an interactive setting that resembles a centralized stock market. Laboratory double auctions yield efficient, competitive outcomes in a surprisingly wide variety of settings, sometimes even in a monopoly (Smith, 1962, 1981; Davis and Holt, 1993). In contrast, prices in markets with posted prices are often above competitive levels (Plott, 1989; Davis and Holt, 1993).

Figures 15.1 and 15.2 show a matched pair of markets, one double auction and one posted-offer auction, which were run under research conditions with payments equaling 1/5 of earnings. The Veconlab setup provides options for specifying the numbers of buyers and sellers, the of units available to each buyer and each seller, and the values and costs associated with these units. In both markets, there were 6 buyers and sellers, each with three units. The values and costs were arrayed in a manner that predicted earnings per person were approximately equal. In each case, the buyer values were reduced after period 5 to shift demand down by \$4, thereby lowering the competitive price prediction by \$2, as shown on the left sides of each figure. For example, the two buyers with the highest value units, at \$15 in period 6, had these values reduced to \$11 in period 6, whereas all costs stayed the same in all periods.

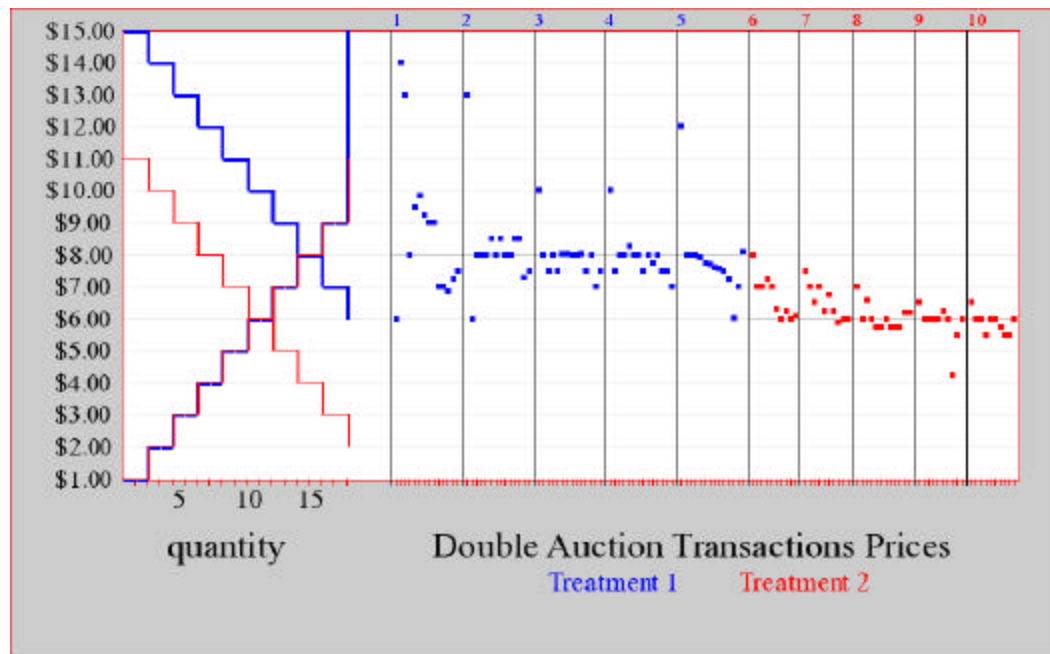


Figure 15.1. Price Sequence for a Double Auction

First, consider the double auction, in Figure 15.1, where the dots between the first two vertical black lines represent transactions prices in the order in which they

were observed. Recall that each trader in a double auction may make a price offer (bids for buyers and asks for sellers) or may accept the best offer on the other side of the market. Thus traders typically see a sequence of declining ask prices and/or increasing bid prices, until the bid ask spread narrows and one person accepts another's proposal. The computerized market maker ensured that a person making a successful bid or ask would obtain the best available terms. For example, suppose that the highest bid were \$5 and there were asking prices of \$10 and \$8, and that a buyer decided to accept by bidding \$8. If another ask were to come in at \$7 prior to the time at which the buyer confirmed the acceptance bid of \$8, then the buyer would obtain the unit at the price of \$7.

Despite the considerable price variation in the first period, the double auction market achieved 99 per cent of the maximum possible surplus, and this efficiency measure averaged 98 per cent in the first 5 rounds, with quantities of 14 that equal the competitive prediction. All traders were told that some payoff parameters may have changed in round 6, but sellers who observed their own costs had not changed had no idea whether buyer values or others' costs had gone up or down. Transactions prices began round 6 at the old competitive level of \$8, but fell to \$6 by the end of the period. This is a typical pattern, where the high-value and low-cost units trade early in a period at levels close to those in the previous periods, but late trades for units near the margin are forced to be closer to the supply-demand intersection price. The transaction quantity was at the competitive prediction of 10 in all rounds of the second treatment. As early-period prices fell toward the competitive level after round 6, the price averages were about equal to the new prediction of \$6, and efficiencies were at about 96 percent.

Recall from the discussion in Chapter 2 that buyers do not post bids in a posted-offer market. Sellers post prices independently at the start of each period, along with the maximum number of units that are offered for sale. Then buyers are selected in a random order to make desired purchases. The period ends when all buyers have finished shopping or when all sellers are out of stock. As shown in Figure 15.2, prices in the posted offer market also began in the first period with considerable variation, although all prices were far from the competitive prediction. Prices in periods 2-4 seemed to converge to a level at about \$1 above the competitive prediction. Both the average quantity (12) and the efficiency (86 percent) were well below the competitive levels observed in the matched double auction. Prices fell slowly after the demand shift in period 6, but they never quite reached the new competitive prediction, and efficiencies averaged 86 percent again in the last 5 periods. The lowest efficiencies were observed in the first two periods of each treatment (73 and 71 percent in periods 1 and 2, and 67 and 77 percent in periods 3 and 4), illustrating the more sluggish adjustment of the posted offer market. What is happening here is that sellers are not able to see buyer bids or to learn from them, and instead, they must look at sales data to make inferences about

how price should be altered. By the final period of each treatment, efficiencies had climbed above 90 percent, and the transactions quantity had reached the relevant competitive prediction.

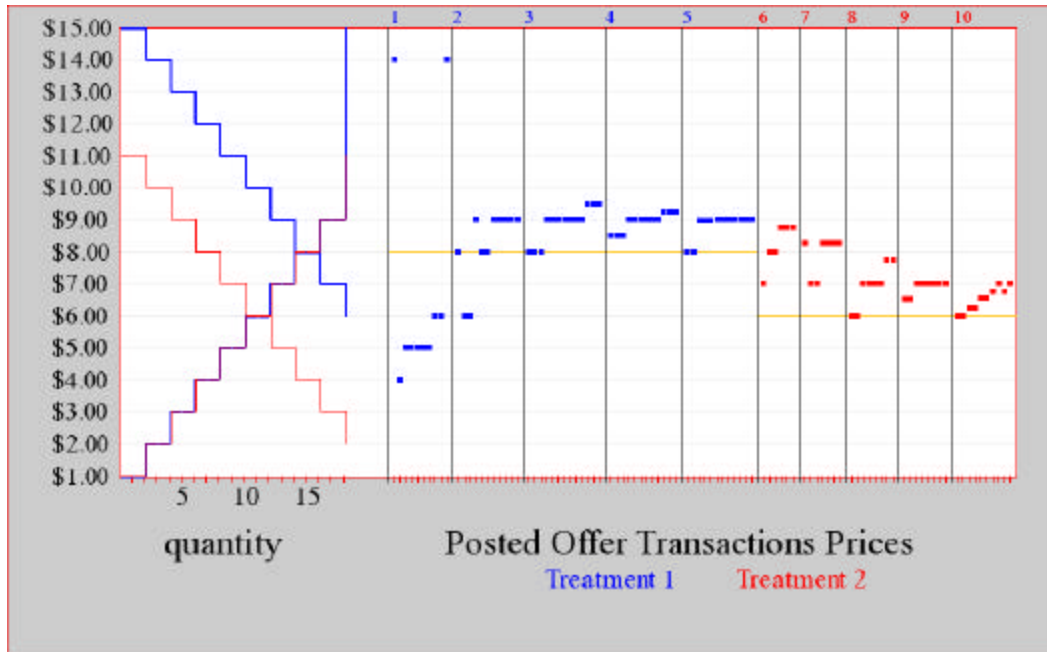


Figure 15.2. Price Sequence for a Posted-Offer Auction

The results shown in Figures 15.1 and 15.2 are fairly typical. Compared to double auctions, laboratory posted-offer markets converge to competitive predictions more slowly (Ketcham, Smith and Williams, 1984) and less completely (Plott, 1986, 1989). Even in non-monopolized designs with stationary supply and demand functions, traders in a posted-offer market generally forego about 10% of the possible gains from trade, whereas traders in a double auction routinely obtain 95-98% of the total surplus in such designs (Davis and Holt, 1993, chapters 3 and 4).

The sluggish price responses shown in Figure 15.2 are even more apparent in markets with sequences of demand shifts that create a boom and bust cycle. *** first consider a case where the supply curve stayed stationary as demand shifted up again and again in a sequence of periods, creating an upward momentum in expectations. This boom was followed by a sequence of downward demand shifts that reduced the competitive price predictions incrementally until it returned to the initial level. Prices determined by double auction trading track the predicted price increases and decreases fairly accurately, with high efficiencies. This is because the demand shifts are conveyed by the intensity of buyer bidding behavior during each

period, so that sellers could learn about new market conditions as they started making sales. In contrast, prices in posted price markets are selected before any shopping begins, so sellers cannot spot changes in market conditions, but rather, must try to make inferences from sales quantities. When posted-offer markets were subjected to the same sequence of demand increases and decreases mentioned above, the actual trading prices lagged behind competitive predictions in the upswing, and prices continued to rise even after demand started shifting downward. Then prices fell too slowly, relative to the declining competitive (supply-equals-demand) predictions. The result was that prices stayed too high on the downswing part of the cycle, and these high prices caused transactions quantities to fall dramatically, essentially drying up the market for several periods.

Davis and Holt (1996) replicate these sluggish adjustment patterns on posted-price markets subjected to a series of demand increases, followed by a series of demand decreases. Even in markets with posted prices, sellers do receive some useful information from buyer behavior. For example, sellers might be able to observe excess demand on the upswing, as buyers seek to make purchases but are not accommodated by the production quantities. Davis and Holt ran a second treatment with the boom/bust cycle of demand shifts, with an added feature that let sellers observe excess demand by frustrated buyers at the prices that they set. This excess demand information improved price responsiveness and raised market efficiency, but not to the high levels observed in double auctions. Similar increases in efficiency were observed when sellers were allowed to offer a single price reduction (“clearance sale”) after the posted prices had been submitted and some sales had been made at those prices.

II. The Exercise of Seller Market Power without Explicit Collusion

One of the major factors considered in the antitrust analysis of mergers between firms in the same market is the possibility that a merged firm may be able to raise price, to the detriment of buyers. Of course, any seller may raise a price unilaterally, and so the real issue is the extent to which price can be raised profitably. Such a price increase is more likely to raise profits if others in the market are not in a position to absorb increases in sales at lower prices, so the capacities of other sellers may constrain a firm’s market power, and a merger that reduces others’ capacities may create market power.

Even if a firm loses some sales as a result of a price increase, this may not be very detrimental to the firm’s profit if these “marginal” units are low-profit units, i.e. if the costs of these units are close to the initial price. For example, consider the supply and demand structure on the left side of Figure 15.3. The supply function has a flat spot at a price of \$2.60, where it crosses demand, and this competitive price is indicated by the horizontal line below the dots that represent transactions prices for the first 4 rounds of trading, which was done with posted-offer rules. The failure of

price to converge to the competitive prediction was due to the fact that two of the sellers had a number of units with costs at \$2.60, so they had little incentive to cut price and sell these units, especially at prices just above this level. The narrow (1 unit) gap between demand and supply at prices just above the competitive level means that a seller who refuses to sell 2 of these low-profit marginal units would shift supply to the left by 2 units and raise the supply/demand intersection to the next highest supply step, at \$2.80.



Figure 15.3 A Posted-Price Market with Seller Market Power

To summarize, market power to raise price above competitive levels can exist when competitors' capacities are limited, and when a firm's marginal units are selling at low per-unit profits. In addition, market power can be influenced by the nature of the trading institution. Vernon Smith (1981) investigated an extreme case of market power, with a single seller, and found that even monopolists in a double auction are sometimes not able to raise prices above competitive levels, whereas posted-offer monopolists are typically able to find and enforce monopoly price levels. The difference is that a seller in a posted offer auction sets a single, take-it-or-leave-it price, so sellers are not tempted to cut price late in the period in order to sell marginal units. This temptation is present in a double auction. In particular, a monopolist who cuts price late in a trading period to sell marginal units will have a harder time selling units at near-

monopoly levels at the start of the next period, and this buyer resistance may cause prices in a double auction monopoly to be lower than the monopoly prediction.

This effect of the trading institution is also apparent from experiments conducted by Holt, Langan, and Villamil (1986) with the design from Figure 15.3, using 5 sellers and 5 buyers. As noted above, two of the sellers had higher capacities that included a number of low-profit units with costs at the competitive level of \$2.60. In contrast, the large vertical difference between demand and the competitive price for all “included” units means that buyers had a strong incentive to buy these high-value units, even if price were increased above competitive levels. Thus the design was intended to create an asymmetry between buyers who were eager to buy, and sellers with little incentive to sell units at the margin, and hence with a large incentive to try to raise prices above competitive levels. Despite this asymmetry, prices converged to competitive levels in about half of the sessions. Some modest price increases above competitive levels were observed in the other sessions, indicating that sellers could sometimes exercise market power even in a double auction.

Davis and Williams (1991) replicated the Holt, Langan, and Villamil results for double auctions with the market power design in Figure 15.3, and the resulting prices were slightly above competitive levels. In addition, they ran a new series of sessions using posted-offer trading, which generated somewhat higher price sequences like those observed in the figure.

Supra-competitive prices in posted-offer markets are not surprising, since experimental economists have long known that prices in posted-offer auctions tend to converge to competitive levels from above, if at all. This raises the question of whether these high prices can be explained by game-theoretic calculations. It is straightforward to specify a game-theoretic definition of market power, based on the incentive of one seller to raise price above a common competitive level (Holt, 1989). In other words, market power is said to exist when the competitive equilibrium is not a Nash equilibrium.

Although it is typically easy to check for the profitability of a unilateral deviation from a competitive outcome, it may be more difficult to identify the Nash equilibrium for a market with posted prices. The easiest case is where firms do not have constraints on what they can produce, a case commonly referred to as Bertrand competition. For example, if each firm has a common, constant marginal cost of C , then no common price above C would be a Nash equilibrium, since each firm would have an incentive to cut price slightly and capture all market sales. The Bertrand prediction for a price competition game *played once* is for a very harsh type of competition that drives price to marginal cost levels, even with only two or three firms.

Even with a repeated series of market periods, the one-shot Nash predictions may be relevant if random matchings are used to make the market interactions have a one-shot nature, where nobody can punish or reward other's pricing decisions in subsequent periods. Most market interactions are repeated, but if the number of market periods is fixed and known, then the one-shot Nash prediction applies in the final period, and a process of backward induction can be used to argue that prices in all periods would equal this one-shot Nash prediction. In most markets, however, there is no well-defined final period, and in this case, there is a possibility that a kind of tacit cooperation might develop. In particular, a seller's price restraint in one period might send a message that causes others to follow suite in subsequent periods, and sellers' price cuts might be deterred by the threat of retaliatory price cuts by others. There are many ways that such tacit cooperation might develop, as supported by punishment and reward strategies, and the multiplicity of possible arrangements typically makes a theoretical analysis impossible without strong simplifying assumptions. Here is where experiments can help.

Figure 15.4 shows the results of a session with posted prices in which participants were matched in groups of 3, with no capacity constraints. Demand was simulated, with aggregate quantity being determined by the function, $Q = 12$

P . In the first 10 rounds, all sellers had a constant marginal cost of \$1, so the Bertrand prediction is \$1. The 3-person groups were fixed, and this is probably a factor in the ability of sellers in one of the six triopoly groups to maintain prices at about \$9, well above the Bertrand prediction. The price dots in the figure represent averages over all 3-person groups for each period, and these averages mask the fact that the 5 other triopoly groups were pricing much nearer to the Bertrand prediction of \$1. The marginal cost was reduced to \$0 for the final 10 periods, and matchings were reconfigured randomly after each period for this second treatment. The random matching is probably a factor in the somewhat sharper convergence of average prices to the Bertrand prediction in the final periods.

To summarize, we have identified several factors that may facilitate pricing at supra-competitive levels: 1) the price fixity and asymmetry of the posted-offer institution, relative to the double auction, 2) the extent to which market interactions are repeated, and 3) the structure of the seller costs and capacities, which may enable one seller to raise price without bsing significant sales to others. The next section expands on the third factor, by providing a more detailed analysis of price competition with capacity constraints.




Figure 15.4. A Posted-Price Market with No Capacity Constraints and No Market Power

III. Market Power and Equilibrium Price Distributions

When sellers all offer the same homogeneous product, buyers with good price information will flock to the seller with the lowest price. This can create a price instability, which may lead to randomized price choices. First consider a specific example in which demand is inelastic at 3 units for all prices up to a limit price of V , i.e. the demand curve is flat at a height of V for less than 3 units and become vertical at a quantity of 3 for prices below V . There are two sellers, each with a capacity to produce 2 units at zero cost. Thus the market supply is vertical at a quantity of 4 units, and supply and demand intersect at a price of \$0 and a quantity of 3 (question 2). The sellers are identical, so each can only expect to sell to half of the market, 1.5 units on average, if they offer the same price. If the prices are different, the seller with the lower price sells 2 units and the other only sells 1 unit (as long as price is no greater than V). There is no Nash equilibrium at any common price between 0 and V , since each seller could increase expected sales from 1.5 to 2 by decreasing price slightly. Nor is a common price of \$0 a Nash equilibrium, since earnings are zero and either seller would have an incentive to raise price and earn a positive amount on the one-unit residual demand (question 3). Thus there is no equilibrium in pure (non-random) strategies, and therefore, one would not expect to see stable price patterns.

An Edgeworth Cycle

One possibility, first considered by Edgeworth, is the idea that prices would cycle, with each firm undercutting the other's price in a downward spiral. At some point prices go so low that one seller may raise price. In the example discussed above, suppose that one seller has a price of that is a penny above a level p between 0 and V . The other seller could cut price to p and sell 2 units, thereby earning $2p$, or raise price to V , sell the single residual demand unit, and earn V . Thus cutting price would be better if $2p > V$, or if $p > V/2$, and raising price would be better if $p < V/2$. This reasoning suggests that prices might fall in the range from V down to $V/2$, at which point one seller would raise price to V and the cycle would begin again. The problem with this argument is that if one knows that the other will cut price by one cent, then the best reaction is to cut by 2 cents, but then the other would want to cut by 3 cents, etc. Thus the declining phase of prices is likely to be sporadic and somewhat unpredictable. This raises the possibility that prices will be random, i.e. that there will be a Nash equilibrium in mixed strategies of the type considered in Chapter 5. In this section, we consider such an equilibrium, in which prices can vary continuously on some interval. But first we must introduce some notation for a distribution of prices.

Price Distributions

With a continuous distribution of prices, the probability that another seller's price is less than or equal to any given amount p will be an increasing function of p , and we will denote this probability by $F(p)$. For example, suppose that price is equally likely to be any dollar amount between 0 and \$10.00. This price distribution could be generated by the throw of a 10-sided die, with sides marked 1, 2, ... 10, and using the outcome to determine price. Then the probabilities would be given by the numbers in the second row of Table 15.1. For each value of p in the top row, the associated number in the second row is the probability that the other seller's price is less than or equal to p . For a price of $p = \$10$, all prices that might be chosen by the other seller are less than or equal to \$10 by assumption, so the value of $F(p)$ in the second row of the far-right column is 1. Since all prices are equally likely, half of them will be less than or equal to the midpoint of \$5, so $F(p) = 0.5$ when $p = \$5$. The other numbers in the table are calculated in a similar manner. A mathematical formula for this uniform distribution of prices would be: $F(p) = p/10$, as can be verified by direct calculation (question 1).

Table 15.1 A Uniform Distribution of Prices

p	0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10
$F(p)$	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1

Uniform price distributions, such as the one shown in Table 15.1, were encountered in the earlier chapter on sequential search and will come up frequently in some of the subsequent chapters on auctions.

The function, $p/10$, shown in Table 15.1 is increasing in p , but it increases at a uniform rate, since all prices in the range from 0 to 10 are equally likely. Other distributions may not be uniform, e.g. some central ranges of prices may be more likely than extreme high or low prices. In such cases, $F(p)$ would still be an increasing function, but the rate of increase would be faster in a range where prices are more likely to be selected. In a mixed-strategy equilibrium, the distribution of prices, $F(p)$, must be determined in a manner to make each seller indifferent over the range of prices, since no seller would be willing to choose price randomly unless all prices in the range of randomization yield the same expected payoff. As was the case in Chapter 5, we will calculate the equilibrium under the assumption that players are risk neutral, so that indifference means equal expected payoffs (with risk averse players, we would need to force indifference in terms of expected utilities).

Mixed Strategies for the Duopoly Example

Again, the discussion pertains to a duopoly example with zero costs and capacities of 2 units for each seller. This supply is vertical at a quantity of 4. As before, assume that demand is inelastic at 3 units for all prices below V . Thus the competitive price is \$0, which is not a Nash equilibrium, as explained above. Begin by considering a price, p , in the range from 0 to V . The probability that the other seller's price is less than or equal to p is $F(p)$, which is assumed to be increasing in p . When the other seller's price is lower, a firm sells a single unit, with earnings equal to p . When the other's price is higher, one's own sales are 2 units, with earnings of $2p$. Since prices are continuously distributed, we will ignore the possibility of a tie, and hence, $F(p)$ is the probability associated with having the high price and earning the payoff of p , and $1 - F(p)$ is the probability associated with having the low price and earning the payoff of $2p$. Thus the expected earnings are: $F(p)p + [1 - F(p)]2p$. This expected payoff must be constant for each price in the range over which a firm randomizes, to ensure that the firm is indifferent between prices in this range. To determine this constant, note that setting the highest price, V , will result in sales of only 1 unit, since the other seller will have a lower price for sure. Thus we set the expected payoff expression equal to the constant V :

$$(15.1) \quad F(p)p + [1 - F(p)]2p = V.$$

To summarize, equation (15.1) ensures that the expected payoff for all prices is a constant and equals the earnings level for charging the buyer price limit of V . This equation can be solved for the equilibrium price distribution $F(p)$:

$$(15.2) \quad F(p) = \frac{2p - V}{p}$$

for $V/2 \leq p \leq V$. Notice that equation (15.2) implies that $F(p) = 1$ when $p = V$. In words, the probability that the other seller's price is less than or equal to V is 1. The lower limit of the price distribution, $p = V/2$, is the value of p for which the right side of (15.2) is 0. Note that this is also the lower bound of the Edgeworth cycle discussed above. The fact that the lower bound is half of the buyer reservation value V is due to the assumption that the high-price firm only sells half of what the low-price firm sells (question 4).

IV. An Experiment on the Effects of Market Power

There may be several reasons for observing prices that are above competitive levels in a design like that considered in the previous section. For example, with only two sellers, a type of tacit collusion may be possible, especially if the sellers interact repeatedly. Another possible reason is that demand is inelastic and that the excess demand is only one unit at prices above the competitive level of \$0 in this example. A final reason is that earnings would be zero at the competitive outcome, which might produce erratic behavior. These types of arguments led Davis and Holt (1994) to consider a design with two treatments, each with the same aggregate supply and demand functions, but with a reallocation of units that creates market power. In particular, a reallocation of capacity from one seller to others changed the Nash equilibrium price from the competitive price (Bertrand result) to higher (randomized) prices over the range spanned by the "Edgeworth cycle."

Consider Design 1 on the left side of Figure 15.5. As indicated by the seller numbers below each unit, sellers S1-S3 each have 3 units, and S4 and S5 each have a single, low-cost unit. The demand curve has a vertical intercept of r and intersects supply at a range of prices from the highest competitive price, p_c , to the level of the highest cost step, c . The demand is simulated with the high-value units being purchased first. This demand process ensures that a unilateral price increase above a common price p_c would leave the 8 high-value units to be purchased by the other sellers, whose capacity totals to 8 units. Thus a unilateral increase from a common competitive price will result in no sales, and hence, will be unprofitable. It follows that no seller has market power in this design.

Market power is created by giving seller S3's two high-cost units to S1 and S2, as shown in Design 2 on the right side of the figure. Now each of the large

sellers, S1 and S2, has 4 units. If one of these sellers were to raise price unilaterally to the demand intercept, r , one of these 4 units would sell since the other 4 sellers only have enough capacity to sell 7 of the 8 units that are demanded at prices above the competitive level. By making the demand intercept high relative to the high-cost step, we make such deviations profitable for the two large sellers, thereby creating market power. In this case, it is possible to calculate the price distributions in the mixed-strategy equilibrium, by equating sellers' expected payoffs to a constant (since a seller would only be willing to randomize if expected payoffs are independent of price on some range). These calculations parallel those in the previous section, but the analysis is more tedious, given the asymmetries in sellers' cost structures (see Davis and Holt, 1994, for details). For Design 2, the range of randomization is shown as the darkened region on the vertical axis. Note that this design change holds constant the number of sellers and the aggregate supply and demand arrays, so that price differences can be attributed to the creation of market power and not to other factors such as a small number of sellers or a low excess supply at supra-competitive prices.




Figure 15.5. Capacity Allocations for a No-Power/Power Design (Source: Davis and Holt, 1994)

Figure 15.6 shows the results of a session in which 30 periods with the Power design were followed by 30 periods of the No-Power design. The period numbers and treatments are shown at the bottom of the figure. The supply and demand functions are reproduced on the left side (on a different scale from figure 15.5) to indicate the relative positions of the key price prediction: the competitive price is 309. Demand was determined by a passive, price-taking simulated buyer, and sellers were told the number of periods and all aspects of the demand and supply

structure. In each period, *S1*'s posted price is plotted as a box, *S2*'s price is plotted as a cross, and the other three small sellers' prices are plotted as dots.

In all but one of the first nine periods of this session, seller *S2* (cross) stays at the demand intercept price, 589, selling the residual demand of 1 unit. This action lured the other sellers up, and seller *S1* (box) reached the demand intercept price in period 9. Then *S1*'s (cross) price cut in period 10 started a general price decline, which was stopped as cross goes back up to 539 for four periods. The final 12 periods of this treatment contain two relatively tight price cycles.



Figure 15.6. Price Series for a Power/No-Power Sequence (Source: Davis and Holt, 1994)

Key: *S1*'s prices are indicated by boxes, *S2*'s prices by crosses, and the other prices by dots.

Units are reallocated after period 30 to take away market power (a movement from Design 2 to Design 1 from Figure 15.5). The initial high price by *S2* (cross) results in no sales, and the aggressive competition that follows drives prices to the competitive (and now Nash) levels. This market efficiently exploits all gains from trade in the absence of market power. Market power also resulted in large price increases in five other sessions, holding constant the number of sellers and the aggregate supply and demand structure. In fact, the price-increasing effect of market power was considerably greater than the difference between the Nash/Bertrand price in the No-Power treatment and the mean of the mixed price distribution in the Power treatment. Since explicit communications between sellers were not permitted, it is appropriate to use the term “tacit collusion” to describe this

ability of sellers to raise prices above the levels determined by a Nash equilibrium. To summarize, market power has a double impact: it raises the predicted mean price prediction, and it facilitates tacit collusion that raises prices above the Nash prediction.

V. Extensions

Much of the more recent effects of market power can be understood by reconsidering the narrowness of the vertical gap between demand and supply to the left of the supply/demand intersection (see Figure 15.5). This narrow gap implies that a seller who refuses to sell units with these relatively high costs will not forego much in the way of earnings. This seller might profit from such withholding if the price increase on low cost units sold would more than compensate for the lost earnings on these “marginal” units. Suppose that the demand/supply gap is large during a boom period, which makes marginal units more profitable and limits the exercise of market power. In contrast, the gap might fall during a contraction, thereby enabling sellers to raise prices profitably. Thus the effects of market power in some markets may be counter-cyclical (Reynolds and Wilson, 1997 and Wilson, 1997). A similar consideration may arise in the analysis of mergers that may create synergies which lower costs. If these cost reductions occur on units near the margin, i.e. near the supply/demand intersection, then the cost reduction on the marginal unit may make the exercise of market power less profitable. Davis and Wilson (1998) discuss this case, and also some contrary cases where cost-reducing synergies of a merger may create power where none exists previously.

A second aspect of market power is that unilateral price increases are more likely to be successful if competitors are not able to move in and expand their production. Goodby (1997) explores the exercise of market power in situations where producers must acquire pollution permits to cover the byproducts of certain production activities. Then the acquisition of other sellers’ permits may limit their capacities, and hence to enable a firm to raise price profitably.

Questions

1. Show that the formula, $F(p) = p/10$, produces the numbers shown in the second row of table 16.1. How would the formula have to be changed if prices were uniform on the interval from \$0 to \$20, e.g. with half of the prices below \$10, 1/4 below \$5, etc.?
2. Use the information given in the duopoly example from Section III to construct the supply and demand graph, and verify that they intersect at a quantity of 3 and a price of \$0.

3. For the duopoly example in section III, show that prices of \$0 for both sellers do not constitute a Nash equilibrium, i.e. show that a unilateral price increase by either seller will raise earnings.
4. Consider a modification of the duopoly example in section III, where demand is vertical at 6 units for all prices below V , and each seller has a capacity of 5 units at zero cost. What is the range of prices over which an Edgeworth cycle would occur? Find the mixed-strategy Nash equilibrium distribution of prices, and compare the range of randomization with the range of the Edgeworth cycle.
5. Answer question 4 for the case in which each seller's costs are constant at \$1 per unit (with $V > 1$).
6. Construct an example with at least 3 sellers, in which a merger of two of these sellers reduces the costs of marginal units and makes the exercise of market power more profitable. Illustrate your answer with a graph that shows each seller's ID, and the way in which costs are reduced.
7. Construct an example (again with at least 3 sellers) of a merger that destroys market power by reducing costs on marginal units. Draw a graph that shows demand and seller costs and IDs for each unit.

Chapter 16. Collusion and Price Competition

Ever since Adam Smith, economists have believed that sellers often conspire to raise price. Such collusion involves trust and coordination, and therefore, the plan may fall apart if some sellers defect. In fact, Smith's oft-quoted warning about the likelihood of price fixing is immediately qualified: "In a free trade an effectual combination cannot be established but by the unanimous consent of every single trader, and it cannot last longer than every single trader continues of the same mind. The majority of a corporation can enact a bye-law with proper penalties, which will limit the competition more effectually and more durably than any voluntary combination whatever." (Smith, 1776, p. 144). Price-fixing is illegal in the U.S. and most other developed economies, and hence it is difficult to study. Moreover, conspirators will try to keep their activities secret from those who have to pay the high prices. Without good data on participants and their costs, it is difficult to evaluate the nature and success of collusion, and the causes of breakdowns in pricing discipline. Laboratory experiments are not hampered by these data problems, since controlled opportunities for price fixing can be allowed, holding constant other structural and institutional elements that may facilitate supra-competitive pricing.

I. Early Experiments on Collusion

Isaac and Plott (1981) allowed sellers in double auctions to go to a corner of the room and discuss prices between trading periods. These conspiracies were not very effective in actually raising transactions prices in the fast-paced competition of a double auction, where sellers are faced with the temptation to cut prices in order to sell marginal units late in the trading period. Isaac, Ramey, and Williams (1984) subsequently showed that price-fixing activities can be more effective in posted-offer auctions, since sellers are not permitted to lower prices once they are posted. Most markets of interest to industrial organization economists cannot be classified as continuous double auctions (where all price activity is public) or as posted offer markets (which do not permit discounts and sales). This raises the issue of how effective price collusion would be in markets with a richer array of pricing strategies and information conditions.

II. Collusion with Secret Discounts

Davis and Holt (1998) report an experiment in which sellers are able to offer secret discounts to individual buyers after list prices are posted publicly. There were three buyers and three sellers in each session. At the beginning of each period, the buyers were taken from the room under the guise of assigning different redemption values to them. While buyers were out of the room, sellers were

allowed to push their chairs back from their visually isolated cubicles so that they could see each other and discuss price. They were not permitted to discuss their own production costs or to divide up earnings. Then they returned to their computers as the buyers came back into the room. At that time, sellers would enter their posted prices independently, without further discussion. The buyers were not aware of the seller price discussions.

The structure of the values and costs produced the supply and demand functions shown on the left side of Figure 16.1. The intersection of these functions determines the competitive price, which is shown as a horizontal thin line. If all three sellers could set a price that maximized total earnings for the three of them, then they would each sell a single unit at a price indicated by the thick horizontal line in the graph.




Figure 16.1. Prices for a Session with No Collusion (Lower Sequence of Squares) and for a Second Session with Collusion (Upper Sequence of Circles)
Source: Davis and Holt (1996)

First consider a session in which buyers were taken from the room for value assignments, as in all sessions, but sellers were *not* permitted to fix prices. The prices for this session are shown as the lower sequence in Figure 16.1. The circles are the list prices, and the units sold are shown as small dots which create a short horizontal line that begins in the circle and may extend to the right if more than two units are sold at this price. The prices for each period are separated by vertical lines, and the prices for the three sellers are shown in order from left to

middle to right, for sellers S1, S2, and S3 respectively. Notice that the low-price firm sells the most units in the first two rounds, and that the other prices fall quickly. Prices are roughly centered around the competitive price in the later periods. The competitive nature of the market (without collusion) was an intentional design feature.

The upper sequence in Figure 16.1 shows a pattern of prices, marked with squares, that developed when sellers were able to collude while buyers were out of the room. Attempts to fix a price resulted in high but variable prices in the early periods. Sellers agreed on a common price in the 5th round, but all buyers made purchases for seller S1, which created an earnings disparity. Seller S1 then suggested that they take turns having the low price, and that he, S1, be allowed to go first! This agreement was adopted, and S1 made all sales in period 6. The low price position was rotated to from seller to seller in subsequent rounds, much as the famous "phases of the moon" price fixing conspiracy involving electrical equipment in the sixties. Notice that there was some experimentation with prices above and below the joint-profit-maximizing collusive level, but that prices stayed at approximately this level in most periods. This rotation scheme was quite inefficient, since each seller had a low-cost unit that would not sell when it was not their turn to have the low price. In other sessions with collusion, earnings were enhanced by agreeing to choose a common price and to limit sales to one unit each, thereby avoiding the inefficiencies caused by the rotation.

Prices for a second session with collusion are shown in Figure 16.2. Again there is some price variability until sellers agree on a common price, this time in the fourth period. The failure of seller S2 to make a sale at this price forces them to deal with the allocation issue, which is solved more efficiently by an agreement that each will limit sales to 1 unit. This agreement breaks down several periods later as S2, whose price is always listed in the center, lists a price below the others. A high price is reestablished in the final 6 periods, but prices remained slightly below the joint-profit-maximizing monopoly level. In two of these final periods, the sellers agreed to raise price slightly and hold sales to one unit each, but in each case S2 sold two units, leaving S3 with nothing. These defections were covered up by S2, who did not admit the extra sale, but claimed in the subsequent meeting that "this is economics," and that there is less sold at a higher price. On both occasions, the others went along with this explanation and agreed to lower price slightly.

Markets for producer goods or major consumer purchases typically differ from the typical posted-offer institution in that sellers can offer private discounts from the "list" prices. The effectiveness of conspiracies in such markets is important for antitrust policy, since many of the famous price-fixing cases, like the electrical equipment bidding conspiracy of the late 1950's, involve producer goods. Moreover, sales contracts and business practices that may deter discounts have been the target of antitrust litigation, as in the Federal Trade Commission's *Ethyl* case,

where the FTC alleged that certain best-price policies deterred sellers from offering selective discounts. The anti-competitive nature of these best-price practices is supported by results of some experiments run by Grether and Plott (1984), who used a market structure that was styled after the main characteristics of the market for lead-based anti-knock gasoline additives that was litigated in the Ethyl case. Holt and Scheffman (1987) provide a theoretical analysis of best-price policies and of the experimental data reported by Grether and Plott.

In order to evaluate the effects of discount opportunities for the market structure considered in Figures 16.1 and 16.2, Davis and Holt (1996) ran a third series of sessions, in which sellers could collude as before, but when buyers returned and saw the sellers' posted prices, they could request discounts. The way this worked was that a buyer whose turn it was to shop could either press a buy button for a particular seller or a request discount button. The seller would then type in a price, which could be equal to the list price (no discount) or lower. The seller response was not observed by other sellers, so the discounts were given secretly. Sellers were free to discount selective to some buyers and not others, and to hold discounts until later in a period.




Figure 16.2. Prices for a Second Posted Offer Session with Collusion
Source: Davis and Holt (1996)

Prices were much lower in the collusion sessions with opportunities for discounting than in the collusion sessions with no such opportunities. In one session, one of the sellers got so mad at the others that she refused to speak with them in the discussion period between periods. Even in groups that maintained an active price-fixing discussion, the results were often surprising. Consider, for example, the price sequence in Figure 16.3. As before, the small squares indicate list prices, and the small dots indicate actual sales, often at levels well below posted list price. In period 3, for example, all sellers offer the same list price, but S1 sold two units at a deep discount. In periods 7 and 8, seller 2 began secret discounting, as indicated by the dots below the middle price square. These discounts caused S1 to have no sales in these periods, and S1 responded with a sharp discount in period 9 and a lower list price in period 10. After this point, discounts were pervasive, and the outcome was relatively competitive. This competitive outcome is similar to the results of several other sessions with discounting.




Figure 16.3. Prices for a Session with Collusion and Secret Discounts
Source: Davis and Holt (1996)

One factor that hampered the ability of sellers to maintain high prices in the face of secret discounts is the inability to identify who is cheating on the agreement. Antitrust investigations have found that many successful price-fixing conspiracies, especially those with many participants, involve industries where a trade association reports reliable sales information for each seller (Hay and Kelly, 1974). This observation motivated the Davis and Holt (1998) treatment where sellers were given reports of each person's sales quantity at the end of each round. All other procedures were identical, and secret discounts were permitted. Even with discounting, this *ex post* sales information permitted sellers to raise prices about half way between the competitive price and the collusive price (the two horizontal lines in the figures).

III. Extensions

To summarize, the effects of market power and explicit collusion are much more severe in markets where sellers post prices on a take-it-or-leave-it basis. The opportunities to offer price reductions during the trading period makes collusion difficult to coordinate collusive price increases and to exercise power in the absence of collusion. These results are consistent with the antitrust hostility to industry practices that are seen as limiting sellers' options to offer selective discounts.

There have been a number of follow-up studies on the effects of price collusion. Isaac and Walker (1985) found that the effects of collusion in sealed-bid are similar to those in posted-price auctions. This is not surprising, since a sealed bid auction is similar to a posted-offer auction, except that only a single unit or prize is typically involved. Collusion in sealed bid auctions is an important topic, since many price-fixing conspiracies have occurred in such auctions, and since many companies are relying more and more on auctions to procure supplies. The design of the Irrigation Reduction Auction discussed in Chapter 22 was based to a large extent on efforts to thwart collusion, and for that reason, the auction allowed bidders to adjust their bids in a series of rounds, with no fixed final round.

Another interesting issue is the extent to which the pattern of bids might be used to infer collusion (Davis and Wilson, 1998). The idea here is that there may be less correlation between losing bids and costs when the losing bidders have agreed in advance to bid high and lose (Porter and Zona, 1993). Cason (1997) investigates the possibility that the NASDAQ dealers' convention of relying on "even eighths" may have facilitated collusion.

There have been allegations that competitors who post prices on computer networks may be able to signal threats and cooperative intentions. For example, some bidders in early FCC bandwidth auctions used decimal places to attach zip codes to bids in an attempt to deter rivals by threatening to bid on licenses that rivals were trying to obtain. Similarly, some airlines attached aggressive letter combinations (e.g. FU) to ticket prices posted on the Airline Tariff Publishing (ATP)

computerized price system. Cason (1995) reports that such nonbinding (“cheap talk”) communications can raise prices, but only temporarily. See Holt and Davis (1990) for a similar result for a market where sellers could post non-binding intended prices before posting actual prices. These non-binding price announcements would correspond to the posting of intended future prices on the ATP, which are visible to competitors before such prices are actually available for consumers. Cason and Davis (1995) find that such price signaling opportunities had more of an effect in a multi-market setting, but even here the effects of purely nonbinding price announcements was quite limited.

Chapter 17. Product Quality, Asymmetric Information, and Market Failure

When buyers can observe both price and product quality prior to purchase, there is pressure on sellers to provide good qualities at reasonable prices. But when quality is not observed, a seller may be tempted to cut quality, and the result will be disappointed buyers who are hesitant to pay the high prices needed to cover the costs to the seller of providing a high quality. When buyers expect to encounter low-quality items, or “lemons,” the resulting downward pressure on price may force sellers to cut both price and quality. The “lemons market” terminology is due to George Akerlof (1970), who explained how the pressure of competition may cause quality to deteriorate to such low levels that the market may fail to exist. The market failure that is predicted in this case can be studied with the help of laboratory experiments. These experiments can be run by hand, using the instructions provided, or they can be run with the Veconlab program LE.

I. Product Quality

The markets considered in this chapter have the property that sellers can choose the quality of the product that they sell. A high quality good is more costly to produce, but it is worth more to buyers. These costs and benefits raise the issue of whether or not there is an optimal quality. Many, or even most, markets have the property that buyers are diverse in their willingness to pay for quality increases, so there will typically be a variety of different quality levels being sold at different prices. Even in this case, sellers who offer high-cost, high-quality items may face a temptation to cut quality slightly, especially if buyers cannot observe quality prior to purchase. Quality might be maintained in these markets, even with asymmetric information about quality, if sellers can acquire and maintain reputations, reported or signaled by warranties and return policies. Before considering the effects of such policies, it is useful to examine how markets might fail if quality cannot be observed in advance by buyers.

For simplicity, consider a case where all buyers demand at most one unit each and have the same preferences for quality, which is represented by a numerical grade, g . The maximum willingness to pay for a unit of the commodity will be an increasing function of the grade, which will be denoted by $V(g)$. Similarly, let the cost per unit be an increasing function, $C(g)$, of the grade. The net value for each grade g is the difference: $V(g) - C(g)$. The optimal grade maximizes this difference. The key thing to notice is that the optimal grade may not be the maximum feasible grade. For example, it is often prohibitively expensive to remove all impurities or reduce the risk of product failure to zero. The important behavioral

issue in this type of market is the extent to which competition in the market place will force quality to near-optimal levels.

II. A Classroom Experiment

Holt and Sherman (1990) use 6 quality grades in a research experiment, but some of the main conclusions can be illustrated with a simpler classroom experiment with 3 grades (Holt and Sherman, 1999). Each buyer only demands one unit of the commodity, and buyers have identical valuations. The value of the commodity depends on the quality grade: \$4.00 for grade 1, \$8.80 for grade 2, and \$13.60 for grade 3. Thus, with four buyers and a grade of 1, for example, the market demand would be vertical at a quantity of 4 units for any price below \$4, as shown by the demand curve in the lower part of Figure 17.1. The demands for higher grades are similar, with cutoff prices of \$8.80 and \$13.60 for grades 2 and 3.




Figure 17.1. Demand and Supply Arrays by Grade

Market supply is determined by the sellers' costs given to sellers. Each seller has a capacity to produce two units, with the cost of the second unit being \$1 higher than the cost of the first unit. For a grade of 1, the costs for the first and

second units produced are \$1.40 and \$2.40 respectively, so each individual seller's supply curve would have two steps, before becoming perfectly inelastic at two units for prices above \$2.40. With three identical sellers, the market supply will also have two steps with three units on each step. The market supply for grade 1 is labeled S1 in the lower part of Figure 17.1, and it crosses the D1 curve at a price of \$2.40. The supply and demand curves for the other grades are shown above those for grade 1. It is apparent from Figure 17.1 that the sum of consumer and producer surplus is maximized at a grade of 2, which is the optimal quality.

The results of a classroom experiment are shown in Table 17.1. This was run by had with the instructions from the Appendix, with 4 buyers and 3 sellers. In the first three rounds, sellers chose price and grade and wrote them on record sheets. When all had finished, these prices and grades were written on the blackboard, and buyer were selected one by one in random order to make purchases. Two of the three quality choices in the first round were at the maximum level. Each of these high-quality sellers sold a single unit, but the seller who offered a grade of 2 sold 2 units and earned more. The buyers preferred the grade 2 good since it provided more surplus relative to the price charged. All three sellers had settled on the optimal grade of 2 by the third round, and the common price of \$5.60 was approximately equal to the competitive level determined by the intersection of supply and demand at this grade.

Table 17.1 Results from a Classroom Experiment
Source: Holt and Sherman(1999)

	<i>Seller 1</i>	<i>Seller 2</i>	<i>Seller 3</i>
Period 1 (full information)	\$11.50 grade 3 sold 1	\$6.00 grade 2 sold 2	\$12.00 grade 3 sold 1
Period 2 (full information)	\$5.75 grade2 sold 2	\$5.50 grade 2 sold 1	\$1.90 grade 1 sold 1
Period 3 (full information)	\$5.65 grade2 sold1	\$5.60 grade 2 sold 2	\$5.60 grade 2 sold 1
Period 4 (only price information)	\$2.40 grade1 sold 1	\$5.60 grade2 sold 1	\$2.40 grade 1 sold 2
Period 5 (only price information)	\$2.40 grade 1 sold 1	\$1.65 grade 1 sold 1	\$5.50 grade1 sold 2

The first 3 “full information” rounds were followed by two rounds in which sellers selected price and grade as before, but only the price was written on the board for buyers to see while shopping. Two of the sellers immediately cut grade to 1 in round 4, although they also cut the price, so buyers would be tipped off if they interpret a low price as a signal of a low grade. In the final round, seller 3 offers a low grade of 1 at a price of \$5.50, which had been the going price for a grade of 2. The buyer who purchased from this seller must have anticipated a grade of 2, and this buyer lost money in the round.

Figure 17.2 shows the results of an experiment with the same structure, but which was done with the Veconlab software, which made it possible to run more rounds. There were 7 buyers and 5 sellers, which determined the supply and demand functions for the three grades as shown on the left side of Figure 17.2. Again, most of the sales are at an optimal grade of 2 in the full-information rounds (1-8), although sellers managed to keep prices above competitive levels in these rounds, especially for the grade 2 items. When only price information was allowed to be posted in round 9, the grade levels immediately fell to 1, although most of the sales were at price levels that would result in losses for buyers.




Figure 17.2. Data from a Veconlab Classroom Experiment

Extensions and Further Reading

There have been a number of experimental studies of markets with asymmetric quality information. Lynch et al. (1986) document lemons outcomes in double auction markets, and they show that performance can be improved with certain types of warranties, requirements for truthful advertising, etc. The lemons market experiment summarized in this chapter is based on the setup in Holt and Sherman (1990), who used laboratory experiments to evaluate factors that affect the degree to which quality deteriorates when it is not observed by buyers. DeJong, Forsythe, and Lundholm (1985) allowed sellers to make price and quality representations, but the quality representation did not have to be accurate, and the buyer had imperfect information about quality even after using the product. Miller and Plott (1985) report experiments in which sellers can make costly decisions that "signal" high quality, which might prevent quality deterioration. This literature is surveyed in Davis and Holt (1993, chapter 7) and Holt (1995).

Chapter 18. A Limit Order Asset Market

Many electronic markets for assets are run as “call markets” in which traders submit limit orders to buy or sell, and these orders are used to generate a single, market-clearing price when the market is “called.” Thus all asset shares trade at the same price; the shares are purchased by those who submitted buy orders with a maximum willingness to pay limit at or above this price, and the shares are sold by those who submitted sell orders with a minimum willingness to accept limit at or below this price. The experiment discussed in this chapter involves trading of shares that pay a randomly determined dividend. Traders are endowed with asset shares and cash that can either be used to purchase shares or earn interest in a safe account. The interest rate determines the value of the asset on the basis of fundamentals: dividends and associated probabilities, and the redemption value of each share after a pre-announced final round of trade. This fundamental value serves as a benchmark from which price bubbles can be measured. This setup is easy to implement using the Veconlab Limit Order Market (LOM) program. Large bubbles and crashes are typical for inexperienced traders, and the resulting discussion can be used to teach lessons about present value, backward induction, etc.

I. Bubbles and Crashes

Despite the widespread belief that prices in equity markets rise steadily over the long term, these markets exhibit strong swings in price that do not seem to be justified by changes in the underlying economic fundamentals. Keynes’ explanation for these swings was that many (or even most) investors are less concerned with the fundamentals that determine long-term future profitability of a company than with what the stock might sell for in several weeks or months. Such investors will try to identify stocks that they think other investors will flock to, and this herding may create its own self-confirming upward pressure on price.

The psychology behind this process is described by Charles Mackay’s (1841) account of the Dutch “tulipmania” in the 17th century:

Nobles, citizens, farmers, mechanics, seamen, footmen, maid-servants, even chimney-sweeps and old clotheswomen, dabbled in tulips. Houses and lands were offered for sale at ruinously low prices, or assigned in payment of bargains made at the tulip-mart. Foreigners became smitten with the same frenzy, and money poured into Holland from all directions.

Of course, tulips are not scarce like diamonds. They can be produced, and it was only a matter of time until the correction, which in Mackay's words happened:

At last, however, the more prudent began to see that this folly could not last for ever. Rich people no longer bought the flowers to keep them in their gardens, but to sell them again at cent per cent profit. It was seen that somebody must lose fearfully in the end. As this conviction spread, prices fell, and never rose again. Confidence was destroyed, and a universal panic seized upon the dealers.

Even if traders realize that prices are out of line with production costs and profit opportunities in the long run, a kind of overconfidence may lead traders to believe that they will be able to sell at a high level, or that they will be able to sell quickly enough in a crash to preserve most of the capital gains that they have accumulated. The problem is that there are no buyers at anything like previous price levels in the event of a crash, and often all it takes is slight decline, perhaps from an exogenous shock, to spook all buyers and stimulate the subsequent free fall in price.

Price bubbles and crashes can be recreated in computer simulations by introducing a mix of trend-based and fundamentals-based traders. Then price surges can be stimulated by positive exogenous shocks, and such simulations can even produce negative bubbles that follow negative shocks (Steiglitz and Shapiro, 1998). In a negative bubble, the trend traders are selling because they expect prior price decreases to continue and this sell pressure draws prices down if the fundamentals-based traders do not have the resources to correct the situation.

Computer simulations are suggestive, especially if they mimic the price and trading volume patterns that are observed in stock market booms and crashes. The obvious question, however, is how long human traders would stick to mechanical trading rules as conditions change. One problem with running a laboratory experiment with human subjects, however, is to decide how owners of unsold shares are compensated when the experiment ends.

Smith, Suchanek, and Williams (1988) addressed the endpoint problem by pre-specifying a redemption value for the asset after 20 periods. Traders were endowed with asset shares and cash, and they could buy and sell assets at the start of each period (via double auction trading). Assets owned at the end of a period paid a dividend that was typically randomly determined from a known distribution. Cash did not earn any interest, and all cash held at the end was converted to earnings at a pre-announced rate.

In most of their sessions, the final redemption value was set to zero. To a risk-neutral person, each share in this market then provides value that equals the

expected value of the dividend times the number of periods remaining. For example, consider an asset that pays \$0.50 or \$1.50, each with probability 1/2, so the expected dividend is \$1.00. This asset would only be worth a number of dollars that equals the number of periods remaining. Thus with 20 rounds, the expected value of the asset in round t that is $20 - t$, i.e. the fundamental value of the asset declines linearly over time.




Figure 18.1. Average Transactions Prices in a 15 Round Double Auction with a Declining Fundamental Value

With this declining-value setup, price bubbles were observed in many (but not all) sessions. Figure 18.1 shows the results of a typical price bubble for a session where the final redemption value was \$0.00 and the expected dividend was \$0.16. Since there were 15 rounds, the fundamental value line starts at \$2.40 and declines by \$0.16 per period, reaching \$0.16 in the final round. The figure shows the average transactions prices for each round, as determined by double auction trading. The prices were below the fundamental value line in early rounds. Then prices increase steadily and rise above the declining value line. People who buy in one round and see the price rise, would often try to buy again, and others wanting to join in on the gains would begin to buy as well. The price trajectory flattened out and then fell as the final round approached, it would become obvious that nobody would pay much more than the expected dividend in the final round. Such bubbles were observed with undergraduates, graduate students, business students, and even a group of commodity traders. Despite the

lack of realism of having a fixed endpoint, the fact was that speculative bubbles could form, *even if* people knew that there was a fixed end point. These price patterns suggest that there might be more frenzy in markets with no final period.

The observed bubbles are surprising given the obvious nature of the zero-value in the final round. Smith, Suchanek, and Williams did run some sessions with a non-zero final redemption value, which was set to be equal to the sum of the realized dividend payments for all rounds. Thus the final redemption value would not be known in advance, but rather, would be determined by the random dividend realizations, and would generally be decreasing from period to period (question 1).

Since most financial assets do not have predictable, declining fundamental values, it is instructive to set up experiments with constant or increasing values. Ball and Holt (1998) achieved this in a classroom experiment by using a fixed probability of breakage to induce a preference for the present (when the asset is not broken) over the future. In their setup, there was a $1/6$ chance that each share would be destroyed, and they threw a six-sided die separately for each asset share at the end of each trading period (after trades were made via double auction and after dividends of \$1 per share were paid). Any shares that remained after at the end of a known final round could be redeemed for at a pre-announced rate of \$6 per share.

One (incorrect) way to value this asset would be to take the final redemption value of \$6 and add the total dividends of \$1 per round, with some adjustment for the fact that an asset only has a $5/6$ chance of surviving for another period, e.g. $\$6 + (5/6) + (5/6)(5/6) + \dots$. The problem with this approach is that the \$6 final payment is unlikely to be received. Even an asset that has survived up to the beginning of the final period has a $5/6$ chance of producing the \$6 redemption value. So at the start of the final round, the value of the asset is the dividend of \$1 plus the $5/6$ chance of getting \$6, which for a risk-neutral person would be $1 + (5/6)6 = 1 + 5 = 6$. Thus an asset that is worth \$6 at the end of the final round is also worth \$6 at the beginning of the round. Working backwards, the next step is to think about the value of the asset at the beginning of the next-to-last round. This asset will pay \$1 in that round and have a $5/6$ chance of surviving to be worth \$6 at the beginning of the last round, which for a risk-neutral person becomes: $\$1 + (5/6)^6 = 6$. Thus the asset is valued at \$6 in the last two rounds. In this manner, it can be shown that the fundamental value of the asset is \$6 in all rounds, assuming risk neutrality.

This setup, with a constant fundamental value of \$6, produced price bubbles in some, but not all, trading sequences. In one sequence, prices started at a little below \$6 in round 1 and rose to near \$9 in round 5, before prices fell near the final round (9). An obvious question is whether bubbles would become more severe with more rounds, since the final round would seem more distant in early

rounds, although the asset destruction process would choke off trade as the number of shares fell. Another problem is that the double auction trading used in this experiment is relatively time-consuming, even if trading is limited to 3 minute periods. Section III presents an alternative approach based on limit order trading with a single market-clearing price when the market is called. But first, it will be useful to discuss how assets with time streams of future dividends should be valued in the present.

II. A Digression on Present Value

Suppose that money can be invested in a safe account that earns interest at a rate of r per period. The simplest case to consider is that of an asset only pays a dividend of D dollars once, at the end of the first period, and then the asset loses all value. Valuing the asset means deciding what to pay now for a dividend of D that arrives one period later. This asset is worth less than D now, since you could take D dollars, invest at an interest rate r , and earn $(1+r)D$, by the end of the period, which is greater than D . Thus it is better to have D dollars now than to have an asset that pays D dollars one period from now. This is the essence of the preference for present over future payments, which causes one to “discount” such future payments. To determine how much to discount, let's consider an amount that is less than D , and in particular consider $D/(1+r)$ dollars now and invest at rate r , the amount obtained at the end of the period is the investment of $D/(1+r)$ times $1+r$, which equals D . To summarize, the *present value* of getting D dollars one period from now is $D/(1+r)$. Similarly, the present value of any amount F to be received one period from now is $F/(1+r)$. Conversely, an amount V invested today yields $V(1+r)$ one period from now. These observations can be summarized:

$$\text{present value of future payment } F: \quad V = F/(1+r)$$

$$\text{future value of present investment } V: \quad F = V(1+r).$$

Similarly, to find the future value of an amount V that is invested for one period and then reinvested at the same interest rate, we multiply the initial investment amount by $(1+r)$ and then by $(1+r)$ again. Thus the future value of V dollars invested for two periods is $V(1+r)^2$. Conversely, the present value of getting F dollars two periods from now is $F/(1+r)^2$.

$$\text{present value of future payment } F \text{ (2 periods later):} \quad V = F/(1+r)^2$$

$$\text{future value (2 periods later) of present investment } V: \quad F = V(1+r)^2.$$

In general, the present value of an amount F received t periods into the future is $F/(1+r)^t$. With this formula, one can value a series of dividend payments for any finite value of t . For example, the present value of an asset that pays a dividend D for two periods and then is redeemed for $\$R$ at the end of period 2 would be: $D/(1+r) + D/(1+r)^2 + R/(1+r)^2$, where the final term is the present value of the redemption value.

The final issue to be addressed is how to value a series of dividend payments that has no terminal point. The main result is:

$$(18.1) \quad V = \frac{D}{1+r} + \frac{D}{(1+r)^2} + \frac{D}{(1+r)^3} + \dots + \frac{D}{(1+r)^t} + \dots = \frac{D}{r}.$$

As indicated by the right-hand equality, this present value turns out to be D/r . One way to verify this is to use a mathematical formula for finding the sum of an infinite series:

$$(18.2) \quad 1 + x + x^2 + x^3 + \dots = \frac{1}{1-x} \quad \text{if } x < 1.$$

To apply this formula, we need to regroup the terms in equation (18.1) so that they begin with a 1:

$$(18.3) \quad V = \frac{D}{1+r} \left[1 + \frac{1}{(1+r)} + \frac{1}{(1+r)^2} + \dots \right]$$

Now we can apply the formula in (18.2) to evaluate the sum of terms in the square brackets in (18.3) by letting $x = 1/(1+r)$, which is less than 1 as required. Since $1-x = 1-1/(1+r) = r/(1+r)$, it follows that $1/(1-x)$ is $(1+r)/r$, which equals the sum of the terms in square brackets on the right side of (18.3). By making this substitution into (18.3) and letting the $(1+r)$ terms cancel from the numerator and denominator, we obtain the resulting present value for the infinite series of dividends of D dollars per period is D/r .

It will be instructive to develop an alternative derivation of the result that the present value of an infinite sum of dividends, D per period, is D/r . First, since the future is infinite and the dividends do not change, it follows that the future always looks the same, even as time passes. Thus the present value of the future stream of dividends will be the same in each period. Let this present value be V . We can think of V as being the present value of getting a dividend at the end of the round, plus the present value of getting the asset (always worth V) back at the end of the round. Thus $V = D/(1+r) + V/(1+r)$, where the first term is the effect of the end-of-period dividend, and the second term represents what the asset could

be sold for at the end of the round after getting the dividend. This equation can be solved to obtain: $V = D/r$.

III. The Limit Order Market Experiment

In this section, we consider a simple asset pricing model with a safe asset that pays an interest rate of r per period (e.g. an insured savings account) and a risky asset that either pays a high dividend, H , or a low amount, L , each with probability $1/2$. Let the expected value of the dividend be denoted by $D = (H+L)/2$. As shown in the previous section, the value of the asset would be D/r if the number of periods were infinite. In the experiment to be discussed, $H = \$1.00$, $L = \$0.40$, so $D = \$0.70$. With an interest rate of 10 per cent, $r = 0.1$, and hence the present value of an infinite-period asset would be $\$0.70/0.1 = \7.00 . This is the present value of future dividends, and it does not change as time passes, since the future is always the same. The trick in setting up a *finite-horizon* experiment with an asset value that is constant over time is to have the final-period redemption value be equal to D/r , which would be the present value of the dividends if they were to continue forever. In this manner, the infinite future is incorporated into the redemption value per share in the final period, and the resulting asset value will be flat over time, even as the final round is approaching. Thus the redemption value was set to be $\$7.00$.

As mentioned in the introduction, trades were arranged by letting people submit limit orders to buy and/or sell, and then using these order to determine a single market-clearing price at which all trades were executed when the market is closed or “called.” This setup is typically called a Call Market, since the process ends and trades are determined at a pre-specified time, e.g. at the end of the business day, or an hour before the financial markets open in the morning. The experiment used the Veconlab LOM program, and the call occurred when all traders had submitted orders, or when the experimenter pressed the “Stop” button on the Admin Results page.

A sell order involves specifying the maximum number of shares to be sold, and a *minimum* price that will be accepted. A buy order also stipulates a maximum number of shares, but the difference is that the price is the *maximum* amount that the person is willing to pay. Buy orders are ranked from high to low in a pseudo demand array, and the sell orders are ranked from low to high in a pseudo supply array. If more than one trader submits the same limit order price, then a randomly determined priority number is used to decide who gets listed first, i.e. who is highest in the bid queue or who is lowest in the offer queue. Then these supply and demand arrays are crossed to determine the market-clearing price and quantity. For example, if one buyer bids $\$5$ for 2 shares and if another bids $\$4$ for 2 shares, then the demand array would have steps at $\$5$ and $\$4$. Suppose that the only sell order involved 3 shares at a limit price of $\$3$. Then 3

shares would trade at a price of \$4, since there are 4 shares demanded at any price below \$4 and only 2 shares demanded at any higher price (problem 3).

The first experiment to be discussed involved 12 traders, each with an endowment of \$50 in cash and 6 shares. The other parameters were as discussed above: an interest rate of 10 per cent, dividends of either \$0.40 or \$1.00 per share, and a final-period redemption payment of \$7 per share at the end period, 20. Traders in this and all subsequent markets to be discussed were paid an amount that was 1/100 of total earnings.




Figure 18.2. A 20 Round Call Market with a Constant Fundamental Value of \$7.00

The trading activity is shown in figure 18.2, where the horizontal red line at \$7.00 indicates the fundamental (present) expected value of the asset. The first round results are to the left of the first vertical black line on the left side of the figure. The clearing price was \$10 in that round, as indicated by the blue line, which has a width that is proportional to the number of shares traded. The bids and asks for traded shares are shown in yellow and orange, respectively. There were more transactions in round 2, but the price stayed the same, before falling in round 3. The second price drop in round 6 was followed by a sharp rise and continued high prices near \$12, which persisted until a steady decline begin in round 13 and continued until prices fell to near the \$7.00 level in the final round.

The second market, shown in Figure 18.3, also involved 12 traders and 20 rounds, with the same parameters which yield a present value of \$7 per share. In this case, the trading started at \$12, which had been the high in the previous case. Prices rise steadily to a level of about \$28 that is four times the level of the present value as calculated from the market fundamentals. Note that people who bought at prices above this value may have done so profitably if they were able to sell at higher prices. Once the price began to drop dramatically in round 18, the volume of trade fell sharply, since few buyers were not willing to pay prices that were anything close to the \$28 level that had prevailed prior to round 17.




Figure 18.3. A 20 Round Call Market with a Constant Fundamental Value of \$7.00

The sharp drops that sometimes occurred just prior to the final rounds of the 20 round markets was the motivation for conducting some with a 40 round horizon, but with all other parameters unchanged. One of these markets, again involving 12 traders, is shown in Figure 18.4. Here prices begin at exactly \$7 in period 1, and then begin a steady rise in subsequent periods. Price increases become more dramatic after period 25. Price reaches a high of \$257 in round 31, which is over 30 times the fundamental value! The crash, when it comes, starts slowly and then accelerates, with a drop from \$200 to \$100 in round 34, and continued drops, on somewhat higher trading volume after that. Individual earnings were highly variable, with one person earning over \$70 (after dividing

final cash by 100), and with several people ending up with less than \$5. The experiment lasted about an hour.




Figure 18.4. A 40 Round Call Market with a Constant Fundamental Value of \$7.00

IV. Other Research on the Call Market Institution

There is a large and growing literature in which experiments are used to study behavior in financial markets. One important topic is the extent to which traders with inside information about the value of an asset are able to exploit this information and earn more than uninformed traders. In some cases, market prices reveal inside information and the trading prices converge to levels that would be expected if this information were public (add references).

There is a second, related strand of the literature on behavior in call markets where buyers and sellers submit limit prices that determine the common, market-clearing price. These are called “call markets” because the final market-clearing price is calculated when the market is called, usually at a pre-announced time. Call market trading is commonly used for electronic trading of stocks on exchanges where there is not sufficient trading volume to support the continuous trading of a double auction process like that used for the New York Stock

Exchange. And a call market mechanism is used to determine the opening prices on the New York Stock Exchange (Cason and Friedman, 1997).

A classic study of computerized call markets is that of Smith, et al. (1982), who provided buyers with redemption values for each unit and sellers with costs for each unit. (This kind of call market can be run with the Veconlab software using the CM program on the Markets Menu.) Smith compared the outcomes of call markets with those of a double auction, in a stationary environment where costs and values stay the same from period to period. Price convergence to competitive levels was generally faster and more reliable in the double auction, but one variant of the call market trading rules yielded comparable efficiency. Friedman (1993) reports almost price formation results for call markets that are almost as reliable as those of comparable double auctions.

An example of a call market run with this software is shown in Figure 18.5. Values and costs were determined randomly at the start of the first round and remained the same for all 5 rounds. These values and costs determine the supply and demand functions, shown in blue on the left side of the figure. The clearing price converges to the competitive prediction, and the bid and ask arrays, shown in yellow and orange respectively on the left side of the figure, flatten out at approximately this level.




Figure 18.5. A Five-Round Classroom Call Market

Traders with multiple units in a call market may attempt to exercise market power. For example, a seller with a cost that is near the competitive price may wish to hold it back in an effort to shift the revealed supply curve up and thereby raise the price received on the seller's other (low-cost) units. Similarly, buyers may seek to hold back bids for marginal units in an effort to lower price. The McCabe, Rassenti, and Smith (1993) results revealed that traders on each side of the market sometimes tried to exercise market power in call markets, as discussed in Chapter 16. Despite these incentives, the markets that were run tended to be quite efficient, at levels comparable to those of double auctions.

Kagel and Vogt (1993) report results of two-sided call markets where buyer values and seller costs were determined randomly at the start of each trading period. The resulting prices are compared with predictions determined by a Nash equilibrium in which each person's bid or ask is a function of their realized value or cost (see the next chapter on private value auctions for this type of calculation). Again, the Kagel and Vogt result is that the call market is slightly more efficient than the double auction. This result must be considered in light of the fact that call markets are easier to administer and can be run "after hours," with traders who connect from remote locations and who do not need to monitor trading activity as carefully, since all included trades will be at the same price.

Cason and Friedman (1997) also compare call market prices to the Nash predictions. They find some persistent differences, and they use a model of learning to explain these differences.

Questions

1. Consider the setup discussed in section I, with no interest payments. Suppose that the experiment lasts for 15 periods, with dividends that are either \$0.12 or \$0.24, each with probability $1/2$, and with a final redemption value that equals the sum of the dividends realized in the 15 periods. There is no interest paid on cash balances held in each round. Thus the actual final redemption value will depend on the random dividend realizations. Calculate the expected value of the asset at the start of the first period, before any dividends have been determined. On average, how fast will the expected value of the asset decline in each round? What is the highest amount the final redemption value could be, and what is the lowest amount?
2. Consider an asset that pays a dividend of \$2 per period and has a $1/10$ chance of being destroyed (after the dividend payment) in each round.

The asset can be redeemed for \$20 at the end of round 5. What is the fundamental value of the asset in each round?

3. Graph the supply and demand arrays used in the example in section III (buyers bid \$5 for 2 shares and \$4 for 2 shares, and only sell order is for 3 shares at a minimum \$3).

Part V. Auctions

Auctions can be very useful in the sale of perishable commodities like fish and flowers. The public nature of most auctions is also a desirable feature when equal treatment and “above-board” negotiations are important, as in the public procurement of milk, highway construction, etc. These days, internet auctions can be particularly useful for a seller of a rare or specialty item who needs to connect with geographically dispersed buyers. And auctions are being increasingly used to sell bandwidth licenses, as an alternative to administrative or “beauty contest” allocations, which can generate wasteful lobbying efforts. (Such lobbying is considered in the chapter on rent seeking in Part VII.)

Participants of all ages typically find auction games to be exciting, given the win/lose nature of the competition. The two main classes of models are those where bidders know their own private values, and those where the underlying common value of the prize is not known. Private values may differ from person to person, even when the characteristics of the prize are perfectly known. For example, two prospective house buyers may have different family sizes or numbers of vehicles, and therefore, the same square footage may be of much less useful to one than to another, depending on how it is configured into common areas, bedrooms, and parking. An example of a common value auction is the bidding for an oil lease, where each bidder makes an independent geological study of the likely recovery rates for the tract of land being leased. Winning in such an auction can be stressful if it turns out that the bidder overestimated the prize value, a situation known as the “winner’s curse.”

Private value auctions are introduced in Chapter 19, where the focus is on comparisons of bidding strategies with Nash equilibrium predictions for different numbers of bidders. There is also a discussion of several interesting variations: 1) an “all-pay” provision that forces all bidders to pay their own bids, whether or not they win, 2) a “Santa Clause” provision that returns a fraction of auction revenues to the bidders (distributed equally), and 3) a “second-price” rule that only requires the high bidder to pay the second highest bid price.

Common value auctions are considered in Chapter 21, where the focus is on the effects of the numbers of bidders on the winner’s curse. The buyer’s curse, discussed in Chapter 20, is a similar phenomenon that arises with a bidder seeking to purchase a firm of unknown profitability from the current owner. In each case, the curse effect arises because bidders may not realize that having a successful bid is an event that conveys useful information about others’ valuations or estimates of value.

The internet has opened up many exciting possibilities for introducing new auction designs, and laboratory experiments can be used to “testbed” possible

procedures prior to the final selection of the auction rules. The Georgia Irrigation Auction described in Chapter 22 was largely designed and run by experimental economists (Cummings, Holt, and Laury, 2004). It permitted the Georgia Environmental Protection Department to use about five million dollars to compensate some farmers for not using their irrigation permits during a severe draught year. The Veconlab game captures some (but not all) features of the auction that was run in March 2001.

The irrigation reduction auction is an example of a multi-unit auction, which raises a number of interesting possibilities. One possibility is to let winning bidders be paid their per-acre bid amounts, which means that some farmers will be receiving compensation rates that are higher than those received by others. This is called a “discriminatory” auction. An alternative would have been to will pay all at the per-acre bid rate of the marginal bidder, i.e. the lowest bid that was not accepted. This is called a “uniform price auction.” Some of these and other variations are also discussed in Chapter 22.

Chapter 19. Private Value Auctions

Internet newsgroups and online trading sites offer a fascinating glimpse into the various ways that collectables can be sold at auction. Some people just post a price, and others announce a time period in which bids will be entertained, with the sale going to the person with the highest bid at the time of the close. These bids can be collected by email and held as “sealed bids” until the close, or the highest standing bid at any given moment can be announced in an ascending bid auction. Trade is motivated by differences in individual values, e.g. as some people wish to complete a collection and others wish to get rid of duplicates. This chapter pertains to the case where individual valuations differ randomly, with each person knowing only their own “private” value. The standard model of a private-value auction is one where bidders’ values are independent draws of a uniform distribution. For example, one could throw a 10-sided die twice, with the first throw determining the “tens” digit and the second throw determining the “ones” digit. The chapter begins with the simplest case, where the bidder receives a value and must bid against a simulated opponent, whose bids are in turn determined by a draw from a uniform distribution. Next, we consider the case where the other bidder is another participant. In each case, the slope of the bid/value relationship is compared with the Nash prediction. These auctions can be done by hand with dice (see the Appendix), or with the Veconlab game PV. The webgame allows options that include an “all-pay” rule, a “second price auction,” and a “Santa Clause” provision that returns a fraction of auction revenue to the bidders.

I. Introduction

The rapid development of *e-commerce* has opened up opportunities for creating new markets that coordinate buyers and sellers at diverse locations. The vast increase in the numbers of potential traders online also makes it possible for relatively thick markets to develop, even for highly specialized commodities and collectables. Many of these markets are structured as auctions, since an auction permits bids and asks to be collected 24 hours a day over an extended period. Gains from trade in such markets arise because different individuals have different values for a particular commodity. There are, of course, many ways to run an auction, and different sets of trading rules may have different performance properties. Sellers, naturally, would be concerned with selecting the type of auction that will enhance sales revenue. From an economist’s perspective, there is interest in finding auctions that promote the efficient allocation of items to those who value them the most. If an auction fails to find the highest-value buyer,

this may be corrected by trading in an “after market,” but such trading itself entails transactions costs, which may be five per cent of the value of the item, and even more for low-value items where shipping and handling costs are significant.

Economists have typically relied on theoretical analysis to evaluate efficiency properties of alternative sets of auction rules. The seminal work on auction theory can be found in a 1962 *Journal of Finance* paper by William Vickrey, who later received a Nobel Prize in economics. Prior to Vickrey, an analysis of auctions would likely be based on a Bertrand-type model in which prices are driven up to be essentially equal to the resale value of the item. It was Vickrey’s insight that different people are likely to have different values for the same item, and he devised a mathematical model of competition in this context. The model is one where there is a probability distribution of values in the population, and the bidders are drawn at random from this population. For example, suppose that a buyer’s value for a car with a large passenger area and low gas mileage is inversely related to how long they have to drive to work. There is a distribution of commuting distances in the population, so there will be a distribution of individual valuations for the car. Each person knows their own needs, and hence their own valuation, but they do not know for sure what other bidders’ values are.

The mathematical model in its simplest form is one where the population distribution of values is uniform on an interval. The uniform distribution, which was so useful in simplifying the analysis of sequential search in chapter 8, is also convenient here. Before discussing the Vickrey model, it is useful to begin with an overview of different types of auctions.

II. Auctions: Up, Down, and the “Little Magical Elf”

Suppose that you are a collector of cards from *Magic: The Gathering*. These cards are associated with a popular game in which the contestants play the role of wizards who duel with spells from cards in a deck. The cards are sold in random assortments, so it is natural for collectors to find themselves with some redundant cards. In addition, rare cards are more valuable, and some out-of-print cards sell for hundreds of dollars. Let’s consider the thought process that may occur to you as you contemplate a selling strategy. Suppose that you log into the newsgroup site and offer a bundle of cards in exchange for a bundle that may be proposed by someone else. You do receive some responses, but the bundles offered in exchange contain cards that you do not desire, so you decide to post a price for each of your cards. Several people seek to buy at your posted price, but suppose that someone makes a price offer that is a little above your initial price in anticipation of excess demand. This causes you to suspect that others are willing to pay more as well, so you post another note inviting bids on the cards, with a one-week deadline. The first bid you receive on a particularly nice card is \$40,

and then a second bid comes in at \$45. You wonder whether the first bidder would be willing to go up a bit, say to \$55. If you go ahead and post the highest current bid on the card each time a new high bid is received, then you are essentially conducting an *English auction* with ascending bids. On the other hand, if you do not announce the bids and sell to the high bidder at closing time, then you will have conducted a *first-price sealed-bid auction*. David Lucking-Reiley (1999) reports that the most commonly used auction method for *Magic* cards is the English auction, although some first-price auctions are observed.

Lucking-Reiley also found a case where these cards were sold in a descending-bid *Dutch auction*. Here the price begins high and is lowered until someone agrees to the current proposed price. This type of descending-bid auction is used in Holland to sell flowers. Each auction room contains several clocks, marked with prices instead of hours. Carts of flowers are rolled into the auction room on tracks in rapid succession. When a particular cart is “on deck,” the quality grade and grower information are flashed on an electronic screen. Then the hand of clock falls over the price scale, from higher to lower prices, until the first bidder presses a button to indicate acceptance. This process proceeds quickly, which results in a high sales volume that is largely complete by late morning. The auction house is located next to the Amsterdam airport, so that flowers can be shipped by air to distant locations like New York and Tokyo.

If you know your own private value for the commodity being auctioned, then the Dutch auction is like the sealed-bid auction. This argument is based on the fact that you learn nothing of relevance as the clock hand falls in a Dutch auction (Vickrey, 1962). Consequently, you might as well just choose your stopping point in advance, just as you would select a bid to submit in a sealed-bid auction. These two auction methods also share the property that the winning bidders pay a price that equals their own bid. In each case, a higher bid will raise the chances of winning, but the higher bid lowers the value of winning. In both auctions, it is never optimal to bid an amount that equals the full amount that you are willing to pay, since in this case you would be indifferent between winning and not winning. To summarize, the descending-bid Dutch auction is strategically equivalent to the first-price, sealed-bid auction in the case of known private prize values.

This equivalence raises the issue of whether there is a type of sealed-bid auction that is equivalent to the ascending-bid English auction. With a known prize value, the best strategy in an English auction is to stay in the bidding until the price just reaches your own value. For example, suppose that one person’s value is \$50, another’s is \$40, and a third person’s value is \$30. At a price of \$20, all three are interested. When the auctioneer raises the price to \$31, the third bidder drops out, but the first two continue to nod as the auctioneer raises the bid amount. At \$41, however, the second bidder declines to nod. The first bidder,

who is willing to pay more, will agree to \$41 but should feel no pressure to express an interest at a higher price since nobody else will speak up. After the usual “going once, going twice...” warning, the prize will be sold to the first bidder at a price of \$41. Notice that the person with the highest value purchases the item, but only pays an amount that is approximately equal to the second highest value.

The observation that the bidding in an English auction stops at the second-highest value led Vickrey to devise a *second-price sealed-bid auction*. As in any sealed bid auction, the seller collects sealed bids and sells the item to the person with the highest bid. The winning bidder, however, only has to pay the second-highest bid. Vickrey noted that the optimal strategy in this auction is to bid an amount that just equals one’s own value. To see why this is optimal, suppose that your value is \$10. If you bid \$10 in a second-price auction, then you will only win when all other bids are lower than your own. If you decide instead to raise your bid to \$12, you increase your chances of winning, but the increase is *only* in those cases where the second highest bid is above \$10, causing you to lose money on the “win.” For example, if you bid \$12 and the next bid is \$11, you pay \$11 for an item that is only worth \$10 to you. Thus it is never optimal to bid above value in this type of auction. Next consider a reduction to a bid, say to \$8. If the second highest bid were below \$8, then you would win anyway and pay the second with or without the bid reduction. But if the second bid were above \$8, say at \$9, then your bid reduction would cause you to lose the auction in a case where you would have won profitably. In summary, the best bid is your own value in a second-price auction. If everybody does, in fact, bid at value, then the high-value person will win, and will pay an amount that equals the second highest value. But this is exactly what happens in an English auction where the bidding stops at the second-highest value. Thus the second-price auction is, in theory, equivalent to the English auction.

Lucking-Reiley (2000) points out that stamp collectors have long used Vickrey-like auctions as a way of including bidders who cannot attend an auction. For example, if two distant bidders mail in bids of \$30 and \$40, then the bidding would start at \$31. If nobody entered a higher bid, then the person with the higher bid would purchase at \$31. If somebody else agrees to that price, the auctioneer would raise the price to \$32. In fact, the auctioneer would continue to “go one up” on any bidder present until the higher mail-in bid of \$40 is reached. This mixture of an English auction and a sealed-bid second price auction was achieved by allowing “proxy bidding,” since it is the auctioneer who is entering bids based on the limit prices submitted by mail. It is a natural extension to entertain only mail-in bids and to simulate the English auction by awarding the prize to the high bidder at the second bid. Lucking-Reiley (2000) found records of a pure second-price stamp auction held in Northampton, Massachusetts in 1893. He also notes

that the most popular online auction, eBay, allows proxy bidding, which is explained: “Instead of having everyone sit at their computers for days on end waiting for an auction to end, we do it a little differently. Everyone has a little magical elf (a.k.a. proxy) to bid for them and all you need to do is tell your elf the most that you want to spend, and he’ll sit there and outbid the others for you, until his limit is reached.”

III. Bidding Against a Uniform Distribution

This section describes an experiment conducted by Holt and Sherman (2000) in which bidders received private values, and others’ bids were simulated by the throw of ten-sided dice. This experiment lets one begin to study the tradeoffs involved in optimal bidding without having to do a full game-theoretic analysis of how others’ bids are actually determined in a market with real (non-simulated) bidders. At the start of each round, the experimenter would go to each person’s desk and throw a ten-sided die three times to determine a random number between \$0.00 and \$9.99. This would be the person’s private value for the prize being auctioned. Since each penny amount in this interval is equally likely, the population distribution of values in this setup is uniform. After finding out their value, each person would select a bid knowing that the “other person’s” bid would be randomly determined by three throws of the ten-sided die. If the randomly determined “other person’s bid” turns out to be lower than the bidder’s own bid, then the bidder would earn the difference between the value and the bid selected. If the other’s bid were higher, then the bidder would earn nothing.

Suppose that the first three throws of the die determine a value that will be denoted by v , where v is now some known dollar amount between \$0.00 and \$9.99. The only way to win money is to bid below this value, but how much lower? The strategic dilemma in an auction of this type is that a higher bid will increase the chances of winning, but the value of winning with a higher bid is diminished because of the higher price that must be paid. Optimal bidding involves finding the right point in this tradeoff, given one’s willingness to tolerate risk, which can be considerable since the low bidder in the auction earns nothing.

This strategic tradeoff can be understood better by considering the bidder’s expected payoff under a simplifying assumption of risk neutrality. This expected payoff consists of two parts: the probability of winning and the payoff conditional on winning. A person with a value of v who wins with a bid of b will have to pay that bid amount, and hence will earn $v - b$. Thus the expected payoff is the product of the winner’s earnings and the probability of winning:

$$(19.1) \quad \text{Expected Payoff} = (v - b)\Pr(\text{winning with } b).$$

The probability of winning with a bid of b is just the probability that this bid is above the simulated other bid, i.e. above the result of the throws of the 10-sided die. The other bid is equally likely to be any penny amount: \$0.00, \$0.01, ... \$9.99. For simplicity, we will ignore ties and assume that the other person will win in the event of a tie. Then a bid of 0 would win with probability 0, a bid of \$10.00 would win with probability 1. This suggests that the probability of winning is $b/10$, which is 0 for a bid of 0 and 1 for a bid of 10. For a bid of \$5, the probability of winning is exactly 1/2 according to this formula, which is correct since there are 500 ways that the other bid is below \$5 (\$0.00, \$0.01, ... \$4.99), and there are 500 ways that the bid is greater than or equal to \$5 (\$5.00, \$5.01, ... \$9.99). Using this formula ($b/10$) for the probability of winning, the expression for the bidder's expected payoff in (19.1) can be expressed:

$$(19.2) \quad \text{Expected Payoff} = (v-b)(b/10) = vb/10 - b^2/10.$$

This expected payoff exhibits the strategic dilemma discussed earlier. The payoff conditional on winning, $v - b$, is decreasing in the bid amount, but the probability of winning, $b/10$, is increasing in b . The optimal bid involves finding the right balance between these two good things, high payoff and high probability of winning.

The bidder knows the value, v , at the time of bidding, so the function on the right side of (19.2) can be graphed to find the highest point, as in Figure 19.1 for the case of $v = \$8$. The bid, b , is on the horizontal axis, the function starts with a height of 0 when $b = 0$ since a 0 bid has no chance of winning. This at the other end, a bid that equals the value v will also yield a 0 expected payoff, since the payoff for bidding the full value of the prize is 0 regardless of whether one wins or loses. In between these two points, the expected payoff function shows a hill-shaped graph, which rises and then falls as one moves to higher bids (from left to right).




Figure 19.1. Expected Payoffs with a Private Value of \$8.00

One way to find the best bid is to use equation (19.2) to set up a spreadsheet to calculate the expected payoff for each possible bid and find the best one (question 1). For example, if $v = \$4$, then the expected payoffs are: \$0 for a bid of \$0, \$0.30 for a bid of \$1, \$0.40 for a bid of \$2, \$0.30 for a bid of \$3, and \$0.00 for a bid of \$4. Filling in the payoffs for all possible bids in penny increments would confirm that the best bid is \$2 when one's value is \$4. Similarly, it is straightforward to show that the optimal bid is \$2.50 when one's value is \$5. These calculations suggest that the best strategy (for a risk-neutral person) is to bid one half of one's value.

The graphical intuition behind bidding half of value is shown in Figure 19.1. When the value is \$8, the expected payoff function starts at the origin of the graph, rises, and falls back to \$0 when one's bid is equal to the value of \$8. The expected payoff function is quadratic, and it forms a hill that is symmetric around the highest point. The symmetry is consistent with the fact that the maximum is located at \$4, halfway between \$0 and the value of \$8.

At the point where the function is flat, the slope of a dashed tangent line is zero, and the tangency point is directly above a bid of \$4 on the horizontal axis. This point could be found graphically for any specific private value. Alternatively, we can use calculus to derive a formula that applies to all possible values of v . (The rest of this paragraph can be skipped by those who are already familiar with calculus, and those who need even more review should see the appendix to Chapter 13.) A reader who is not familiar with calculus should go ahead and read the discussion that follows; the only thing you will need besides a little intuition is a couple of rules for calculating derivatives (finding the slopes of tangent lines). First consider a linear function of b , say $4b$. This is a straight line with a slope of 4, so the derivative is 4. This rule generalizes to any linear function with a constant slope of k : the derivative of kb with respect to b is just k .

The second rule that will be used is that the derivative of a quadratic function like b^2 is linear: $2b$. The intuition is that the slopes of tangent lines to a graph of the function b^2 become steeper as b increases, which the slope, $2b$, is an increasing function of b (the 2 is due to the number 2 in the exponent of b^2). This formula is easily modified to allow for multiplicative constants, e.g. the derivative of $3b^2$ is $3(2b)$, or the derivative of $-b^2$ is $-2b$.

The expected payoff in equation (19.2) consists of two terms. The first one, can be written as $(v/10)b$, which is a linear function of b with a slope of $v/10$. Therefore the derivative of the expected payoff will have a $v/10$ term in it, as can be seen on the right side of (19.3):

$$(19.3) \quad \text{Derivative of Expected Payoff} = \frac{v}{10} - \frac{2b}{10}.$$

The second term in the expected payoff expression (19.2) is $-b^2/10$, and the derivative of this term is $-2b/10$ because the derivative of b^2 is $2b$. To summarize, the derivative on the right side of (19.3) is the sum of two terms, each of which is the derivative of the corresponding term in the expected payoff in (19.2).

The optimal bid is the value of b for which the slope of a tangent line is 0, so the next step is to equate the derivative in (19.3) to 0:

$$(19.4) \quad \frac{v}{10} - \frac{2b}{10} = 0.$$

This equation is linear in b , and can be solved to obtain the optimal bidding strategy:

$$(19.5) \quad b^* = \frac{v}{2} \quad (\text{optimal bid for risk neutral person}).$$

The calculus method is general in the sense that it yields the optimal bid for all possible values of v , whereas the graphical and numerical methods had to be done separately for each value of v being considered.

To summarize, the predicted bid is a linear function of value, with a slope of 0.5. The actual bid data in the Holt and Sherman experiment formed a scatter plot with an approximately linear shape, but most bids were above the half-value prediction. A linear regression yielded the estimate:

$$(19.6) \quad b = 0.14 + 0.667v \quad (R^2 = 0.91),$$

where the intercept of 14 cents, with a standard error of 0.61, was not significantly different from 0. The slope, with a standard error of 0.017, was significantly different from 1/2.

By bidding above one half of value, bidders are obtaining a higher chance of winning, but a lower payoff if they win. A willingness to take less money in order to reduce the risk of losing and getting a zero payoff may be due to risk aversion, as discussed in the next section. There could, of course, be other explanations for the over-bidding, but the setup with a simulated other bidder does permit us to rule out some possibilities. Since the other bidder was just a roll of the dice, the over-bidding cannot be due to issues of equity, fairness, or rivalistic desires to win or reduce the other's earnings.

The experiment described in this section with simulated "other bids" is essentially an individual decision problem. An analogous game can be set up by providing each of two bidders with randomly determined private values drawn from a distribution that is uniform from \$0 to \$10. As before, a high bid results in earnings of the difference between the person's private value and the person's own bid. A low bid results in earnings of 0. This is known as a *first-price auction* since the high bidder has to pay the highest (first) price.

Under risk neutrality, the Nash equilibrium for this game is to bid one half of value. The proof of this claim is essentially an application of the analysis given above, where the distribution of the other's bids is uniform on a range from 0 to \$10. Now suppose that one person is bidding half of value. Since the values are uniformly distributed, the bid of one-half of value will be uniformly distributed from 0 to a level of \$5, which is one half of the maximum value. If this person's bids are uniformly distributed from 0 to \$5, then a bid of 0 will not win, a bid of \$5 will win with a probability that is essentially 1, and a bid of \$2.50 will win with probability 1/2. So the probability of winning with a bid of b is $b/5$. Thus the expected-payoff function is given in equation (19.2) if the 10 in each denominator is replaced by a 5. Equations (19.3) and (19.4) are changed similarly. Then multiplying both sides of the revised (19.4) by 5 yields the $v/2$ bidding rule in (19.5).

IV. Bidding against a Uniform Distribution with Risk Aversion (optional)

The analysis in this next section shows that a simple model of risk aversion can explain the general pattern of overbidding that is implied by the regression equation (19.6). The analysis uses simple calculus, i.e. the derivative of a "power function" like kx^s , where k is a constant and the variable x is raised to the power s . The derivative with respect to x of this "power function" is a new function where the power has been reduced by 1 and the original power enters multiplicatively. Thus the derivative of kx^s is ksx^{s-1} , which is called the

“power-function rule” of differentiation. A second rule that will be used is that the derivative of a product of two functions is the first function times the derivative of the second, plus the derivative of the first function times the second. This is analogous to calculating the change in room size (the product of length and width) as the length times the change in width plus the change in length times the width. This “product rule” is accurate for “small” changes. Using the power-function and the product rules, those not familiar with calculus should be able to follow the arguments given below, but if you have trouble, skip to the discussion that is just below the generalized bidding rule in (19.11).

The most convenient way to model risk aversion in an auction is to assume that utility is a nonlinear function, i.e. that the utility of a money amount $v-b$ is a power function $(v-b)^{1-r}$ for $0 \leq r < 1$. When $r = 0$, this function is linear, which corresponds to the case of risk neutrality. If $r = 1/2$, then the power $1-r$ is also $1/2$ so the utility function is the square root function. A higher value of r corresponds to more risk aversion. Now the expected payoff function in (19.2) must be replaced with an expected utility function, which is the utility of the payoff times the probability of winning:

$$(19.7) \quad \text{Expected Utility} = (v-b)^{1-r} \left(\frac{b}{10} \right).$$

As before, the optimal bid is found by equating the derivative of this function to zero. The expected utility on the right side of (19.7) is a product of two functions of b , so we use the power-function rule to obtain the derivative (first function times the derivative of the second plus the derivative of the first times the second function). The derivative of $b/10$ is $1/10$, so the first function times the derivative of the second is the first term in (19.8). Next, note that the power function rule implies that the derivative of $(v-b)^{1-r}$ is $-(1-r)(v-b)^{-r}$, which appears in the second term in (19.8).

$$(19.8) \quad (v-b)^{1-r} \left(\frac{1}{10} \right) - (1-r)(v-b)^{-r} \left(\frac{b}{10} \right).$$

This derivative can be rewritten by putting parts common to each term in the parentheses, as shown on the left side of (19.9):

$$(19.9) \quad (v-b) \left(\frac{(v-b)^{-r}}{10} \right) - (1-r)b \left(\frac{(v-b)^{-r}}{10} \right) = 0.$$

Multiplying both sides by $(v-b)^{-r}/10$, one obtains:

$$(19.10) \quad v - b - (1-r)b = 0.$$

This equation is linear in b , and can be solved to obtain the optimal bidding strategy:

$$(19.11) \quad b^* = \frac{v}{2-r} \quad (\text{optimal bid with risk aversion}).$$

This bidding rule reduces to the optimal bidding rule for risk neutrality (bidding half of value) when $r = 0$. Increases in r will raise the bids. When $r = 1/2$, the bids will be two-thirds of value, which is consistent with the results of the regression equation (19.6).

V. Bidding Behavior in a Two-Person, First-Price Auction

Recall that the predicted bidding strategy is linear, with a slope of 1/2 when bidders are risk neutral. Bids exceeded the $v/2$ line in experiment with simulated other bids, and the same pattern emerges when the other bidder is a subject in the experiment. Figure 19.2 shows the bid/value combinations for 10 rounds of a classroom experiment done with the *Veconlab* software. There were 10 bidding teams, each composed of 1-2 students on a networked PC. The teams were randomly matched in a series of rounds. The bidding pattern is approximately linear, except for a leveling off with high values (a pattern that has also been noted by Dorsey and Razzolini (2002)). Only a small fraction of the bids are at or below the Nash prediction of $v/2$, which is graphed as the solid line. In fact, the majority of bids are above the dashed line with slope of .667 obtained from regression for the case discussed in the previous section; the bid-to-value ratio averaged over all bids in all rounds is about 0.7. This pattern of overbidding relative to the Nash prediction is typical, and the most commonly mentioned explanation is risk aversion.




Figure 19.2. Observed Bidding Strategies for Ten Subjects in a Classroom Experiment

An analysis of risk aversion for this game parallels the previous section's analysis of bidding against a simulated bid that is uniform on the interval from 0 to 10 dollars. Suppose that the equilibrium bidding strategy with risk aversion is linear: $b = \beta v$, where $0 < \beta < 1$. Then the lowest bid will be 0, corresponding to a value of 0, and the highest bid will be 10β , corresponding to a value of 10. The distribution of bids is represented in Figure 19.3. For any particular value of β , the bids will be uniformly distributed from 0 to 10β , as indicated by the dashed line with a constant height representing a constant probability for each possible bid. The figure is drawn for the case of $\beta = 0.6$, so bids are uniform from \$0 to \$6. A bid of 0 will never win, a bid of $10\beta = 6$ will win with probability 1, and an intermediate bid will win with probability of $b/6$. In the general case, a bid of b will win with probability $b/10\beta$.



Figure 19.3. A Uniform Distribution Bids on [0, 6]

In a two-person auction when the other's bid is uniform from 0 to 10β , the probability of winning in the expected payoff function will be $b/10\beta$ instead of the ratio $b/10$ that was used in section III. The rest of the analysis of optimal bidding in that section is unchanged, with the occurrences of 10 replaced by 10β , which cancels out of the denominator of equation (19.9) just as the 10 did. The resulting equilibrium bid is given in equation (19.11) as before. This bidding strategy is again linear, with a slope that is greater than one half when there is risk aversion ($r > 0$). Recall that a risk aversion coefficient of $r = 1/2$ will yield a bid line with slope $2/3$, so the bids in Figure 19.2 are roughly consistent with risk aversion of at least $1/2$.

VI. Extensions and Further Reading

Some of the earliest experimental tests of the Vickrey model are reported by Coppinger, Smith, and Titus (1980). In particular, they observed overbidding relative to Nash predictions in a private-value, first-price auction. The effects of risk aversion on individual behavior was explored in a series of papers by Cox, Roberson, and Smith (1982), and Cox, Smith, and Walker (1985, 1988). See Kagel (1995) for a survey of auction experiments. Goeree, Holt, and Palfrey (2002) provide an analysis of risk aversion and noisy behavior in first-price auctions.

Up to now, we have only considered two-person auctions. Auctions with more bidders are more competitive, which will cause bids to be closer to values. The formula in (19.5) can be generalized to the case of N risk-neutral bidders drawing from the same uniform private value distribution:

$$(19.12) \quad b^* = \frac{(N-1)v}{N} \quad (\text{optimal bid with risk neutrality}),$$

with a further upward adjustment for the case of risk aversion. Notice that (19.12) specifies bidding half of value when $N = 2$, and that bids rise as a proportion of value for larger group sizes. As the number of bidders becomes very large, the ratio, $(N-1)/N$, becomes closer and closer to 1, and bids converge to values. This increase in competition causes expected payoffs to converge to zero as bid/value differences shrink.

The Vecon software permits a number of other interesting variations. One option is to require all bidders to pay their own bids, with the prize still going to the highest bid. For example, if a bidder with a value of \$4 bids \$1 and if a second bidder with a value of \$9 bids \$5, then the first earns $-\$1$ and the second earns $\$9 - \$5 = \$4$. In this strategic setting, bidders with low values (and low chances of winning) will reduce their bids towards 0, to avoid having to pay in the

event of a loss. Those with a relatively high value will risk a higher bid, so that the bidding strategy will have a curved shape, only rising significantly above 0 at very high values. This curved pattern can be seen in the bids in Figure 19.4, which were observed in a research experiment conducted at the University of Virginia (with all payments made in cash).




Figure 19.4. Observed Bids for 10 Subjects in an All-Pay Auction Experiment

This “all-pay” auction is sometimes used to model lobbying competition or a patent race, where the contender with the highest effort will win, but even the losers incur the costs associated with their efforts.

A second option provides for returning a fraction of the auction revenue to all bidders, divided equally among the bidders. This is analogous to using the revenue from the sale of prize to reduce taxes for all bidders. Such revenue rebates tend to raise the bids.

Questions

1. This question lets you set up a simple spreadsheet to calculate the optimal bid. Begin by putting the text “V =” in cell A1 and any numerical value, e.g. 8, in cell B1. Then put the possible bids in the A column in fifty-cent increments: 0 in A3, 0.5 in A4, 1 in A5, etc. (You can use penny amounts if you prefer.) The expected payoff formula in column B, beginning in cell B3, should contain a formula with terms involving \$B\$1 (the prize value) and \$A3 (the bid). Use the expected payoff formula in equation (19.2) to complete this formula, which is then copied down to the lower cells in column B. Verify the expected payoff numbers for a value of \$4 that were reported in section II. By looking at the expected payoffs in column B, one can determine the bid with the highest expected payoff. If the value of 8 is in cell B1, then the optimal bid should be 4. Then change the value in cell B1 to 5 and show that the optimal bid is \$2.50.
2. Recall that equations (19.1)-(19.5) pertain to the case of a bidder who is bidding against a simulated other bidder. Write out the revised versions of equations (19.2), (19.3), (19.4), and (19.5) for the two-person auction, assuming risk neutrality.
3. Write out the revised versions of equations (19.2), (19.3), (19.4), and (19.5) for the two-person auction under risk aversion with a power-function utility and parameter r .

Chapter 20. The Takeover Game

This chapter pertains to a situation in which a buyer cannot directly observe the underlying value of some object. A buyer's bid will only be accepted if it is higher than the value to the current owner. The danger is that a purchase is more likely to be made precisely when the owner's value is low, which may lead to a loss for the buyer. The tendency to purchase at a loss in such situations is called the "buyer's curse." The Veconlab game "BC" sets up this situation, or alternatively, the instructions in the appendix can be used with 10-sided dice.

I. *Wall Street* (the Film)

In the 1987 Hollywood film *Wall Street*, Michael Douglas plays the role of a corporate raider who acquires TELDAR Enterprises, with the intention of increasing its profitability by firing the union employees. After the acquisition, the new owners become aware of some previously hidden business problems (a defect in an aircraft model under development) that drastically reduce the firm's profit potential. As the stock is falling, Douglas turns and gives the order "Dump it." This event is representative of a wave of aggressive acquisitions that swept through Wall Street in the mid-1980's. Many of these takeovers were motivated by the belief that companies could be transformed by infusions of new capital and better management techniques. The mood later turned less optimistic as acquired companies did not achieve profit goals.

With the advantage of hindsight, some economists have attributed these failed mergers to a selection bias: it is the *less profitable* companies that are more likely to be sold by owners with inside information about problem areas that raiders may not detect. In a bidding process with a number of competitors, the bidder with over-optimistic expectations is likely to end up making the highest tender offer, and the result may be an acquisition price that is not justified by subsequent profit potential. The tendency for the winning bidder to lose money on a purchase is known as the "winner's curse." Even with only a single potential buyer, a bid is more likely to be accepted if it is too high, and the resulting potential for losses is known as the "buyer's curse."

In a sense, winning in a bidding war can be an informative event. A bid that is accepted, at a minimum, indicates that the bid exceeds the current owner's valuation. Thus there would be no incentive for trade if the value of the company were the same for the owner and the bidder. But even when the bidder has access to superior capital and management services, the bidder may end up paying too much if the intrinsic valuation cannot be determined in advance.

II. The Buyer's Curse Experiment

Some elements that affect a firm's intrinsic profitability are revealed by accounting data and can easily be observed by both current and prospective owners. Other aspects of a firm's operations are private, and internal problems are not likely to be revealed to outsiders. The model presented in this section is highly stylized in the sense that all profitability information is private and known only by the current owner. The prospective buyer is unsure about the exact profitability and has only probabilistic information on which to make a bidding decision. The prospective buyer, however, has a productivity advantage that will be explained below.

The first scenario to be considered is one where the value of the firm to the current owner is equally likely to be any amount between 0 and 100. This current owner knows the exact value, V , and the prospective buyer only knows the range of possible equally likely values. In an experiment, one could throw a ten-sided die twice at the owner's desk to determine the value, with the throw being unobserved by the buyer. To provide a motivation for trade, the buyer has better management skills. In particular, the value to the buyer is 1.5 times the value to the current owner. For example, if a bidder offers 60 for a company that is only worth 50 to the current owner, then the sale will go through and the bidder will earn 1.5×50 minus the bid of 60, for a total of 15.

Table 20.1 Round 1 Results from a Classroom Experiment
(University of Virginia, Econ 482, Spring 2002)

	Proposer 1	Proposer 2	Proposer 3	Proposer 4	Proposer 5	Proposer 6
Buyer Bid	60	49	50	36	50	0
Owner Value	21	23	31	6	43	57
Owner Response	accept	accept	accept	accept	accept	reject
Buyer Value	32	35	46	10	65	86
Buyer Earnings	-28	-14	-4	-26	15	0

Table 20.1 shows the first round results from a classroom experiment. The proposer names (and their comments) have been removed. Notice that proposers in this round were particularly unfortunate. Except for proposer 6 (who knew "too much" as we will see below), the proposer bids averaged 49. All of these bids were accepted, and the owner values (21, 23, 31, 6, and 43 for Proposers 1 to 5 respectively) averaged 25. The first four proposers lost money,

as shown by the bottom row of the table. The fifth proposer earned 1.5 times 43 on an accepted bid of 50, for earnings of 15 cents. This tendency for buyers to make losses is not surprising *ex post*, since a bid of about 50 will only be accepted if the seller value is lower than 50. The seller values averaged only 25, and 1.5 times this amount is 37.5, which is still below 50.




Figure 20.1. Expected Outcomes for a Bid of 60

The tendency for buyers to lose money is illustrated in Figure 20.1, where payoffs are in pennies. Owner values are uniformly distributed from 0 to 99, with a probability of 1/100 of each value, as indicated by the dashed line with a height of 0.01. A bid of 60 will only be accepted if the owner's value is lower, and since all lower bids are equally likely, the expected owner value is 30 for an accepted bid of 60. This average owner value of 30 translates into an average value of 1.5 times as high, as indicated by the vertical line at 45. So an accepted bid of 60 will only yield expected earnings of 45. This analysis is easily generalized. A bid of b will be accepted if the owner value is below b . Since owner values for accepted bids are equally likely to be anywhere between 0 and b , the average owner value for an accepted bid is $b/2$. This translates into a value of $1.5(b/2)$, or $(3/4)b$, which is less than the accepted bid, b . Thus any positive bid of b will generate losses on average in this setup, and the optimal bid is zero.

The feedback received by bidders is somewhat variable, so it is difficult to learn this through experience. Only three of the 13 accepted bids resulted in positive earnings, and losses resulted in lower bids in the subsequent round in all cases. But accepted bids with positive earnings were followed by bid increases, and rejections were followed by bid increases in about half of the cases (excluding the proposer who bid 0 and was rejected every time). The net effect of these reactions caused average bids to stay at about 30 cents for rounds 2-5. The only person who did not end round 5 with a cumulative loss was the person who bid 0 in all rounds. This person, a second-year physics and economics major, had

figured out the optimal bidding strategy on his own. The results of this experiment are typical of what is observed in research experiments summarized next.

III. Further Reading

The first buyer's curse experiment was reported in Bazerman and Samuelson (1983), who used a discrete set of equally likely owner values (0, 20, 40, 60, 80, and 100). Ball, Bazerman, and Carroll (1990) used a grid from 0 to 100 with M.B.A. subjects, and the bids did not decline to zero. The average bid stayed above 30, and the modal bid was around 50.

In all of the setups described thusfar, the optimal bid is zero. This is no longer the case when the lowest possible owner value is positive. When seller values range from 50 to 100 and buyer values are 1.5 times the seller value, for example, the optimal bid is at the upper end of the range, i.e. at 100 (see questions 2-4). Holt and Sherman (1994) ran experiments with this setup and found bids to be well below the optimal level. Similarly, this high-owner-value setup was used in periods 6-10 of the classroom experiment described above, and the average bids were 75, 76, 78, 83, and 83 in these rounds, well below the optimal level of 100. Holt and Sherman attributed the failure to raise bids to another type of error, i.e. the failure to realize that an increase in the bid will pick up relatively high-value objects at the margin. For example, a bid of 70 will pick up objects with seller values from 50 to 70, with an average of 60. But raising the bid from 70 to 71 will pick up a purchase if the seller value is at the upper end of this range, at 70. Failure to recognize this factor may lead to bids that are too low, which they termed the "loser's curse."

Questions

1. How would the analysis of optimal bidding in section II change if the value to the bidder were a constant K times the value to the owner, where $0 < K < 2$ and the owner's value is equally likely to be any amount between 0 and 100? (Thus an accepted bid yields a payoff for the bidder that is K times the owner's value, minus the amount of an accepted bid.) In particular, does the optimal bid depend on K ?
2. Suppose that the owner values are equally likely to be any amount from 50 to 99, so that a bid of 100 will always be accepted. Assume that owners will not sell when they are indifferent, so a bid of 50 will be rejected and will produce zero earnings. The value to the bidder is 1.5 times the value to the owner. Show that the lowest bid of 50 is not optimal. (Hint: Compare the expected earnings for a bid of 50 with the expected earnings for a bid of 100.)

3. For the setup in question 2, a bid of 90 will be accepted about four-fifths of the time. If 90 is accepted, the expected value to the owner is between 50 and 90, with an average of 70. Use this information to calculate the expected payoff for a bid of 90, and compare your answer with the expected payoff for a bid of 100 obtained earlier.
4. The analysis for your answers to questions 2 and 3 (with owner values between 50 and 100) suggests that the best bid in this setup is 100. Let's model the probability of a bid b being accepted as $(b - 50)/50$, which is 0 for a bid of 50 and 1 for a bid of 100. From the bidder's point of view, the expected value to the owner for an accepted bid of b is 50 plus half of the distance from 50 to b . The bidder value is 1.5 times the owner value, but an accepted bid must be paid. Use this information in a spreadsheet to calculate the expected payoff for all bids from 50 to 100, and thereby to find the optimal bid. The five columns of the spreadsheet should be labeled: (1) Bid b , (2) Acceptance Probability for a Bid b , (3) Expected Value to Owner (conditional on b being accepted), (4) Expected Value to Bidder (conditional on b being accepted), (5) Expected Payoff for Bidder of Making a Bid b (which is the product of columns 2 and 4). Alternatively, you may use this information to write the bidder's expected payoff as a quadratic function of b , and then use calculus by setting the derivative of this function to zero and solving for b .

Note: A classroom experiment for the setup in questions 2-4 resulted in average bids of about 80. This information will not help you find the optimal bid.

Chapter 21. The Winner's Curse

In many bidding situations, the value of the prize would be approximately the same for all bidders, although none of them can assess exactly what this “common value” will turn out to be. In such cases, each bidder may obtain an estimate of the prize value, and those with higher estimates are more likely to make higher bids. As a result, the winning bidder may overestimate the value of the prize and end up paying more than it is worth. This “winner’s curse” is analogous to the buyer’s curse discussed in the previous chapter. An auction where each bidder has partial information about an unknown prize value can be run using 10-sided dice to determine common value elements, as explained in the Appendix. Alternatively, it can be run with the Veconlab game “CV,” which has default settings that permit an evaluation of increases in the number of bidders, which tends to produce a more severe winner’s curse. An alternative formulation of a common value auction can be implemented with the program MV that permits a mixture of common and private values for a number of alternative types of auctions, including ascending bid (English) and descending bid (Dutch).

I. “I Won the Auction But I Wish I Hadn’t”

Once the author asked several flooring companies to make bids on replacing some kitchen floor tile. Each bidder would estimate the amounts of materials needed and the time required to remove the old tile and install the floor around the various odd-shaped doorways and pantry area. One of the bids was somewhat lower than the others. I expected him to be happy when he was informed that the job was his. Instead, he exhibited considerable anxiety about whether he had miscalculated the cost, although in the end he did not withdraw the bid. This story illustrates the fact that winning an auction can be an informative event, or equivalently, the fact that you win produces new information about the unknown value of the prize. A rational bidder should anticipate this information in making the bid originally. This is a subtle strategic consideration that is almost surely learned by (unhappy) experience.

The possibility of paying too much for an object of unknown value is particularly dangerous for one-time auctions in which bidders are not able to learn from experience. Suppose that two partners in an insurance business work out of separate offices and sell separate types of insurance, e.g. business insurance from one office and life from the other. Each partner can observe the other’s “bottom-line” earnings, but cannot determine whether the other one is really working. One of the partners is a single mother who works long hours, with good results despite a low earnings-per-hour ratio. The other one, who enjoys somewhat of a lucky

market niche, is able to obtain good earnings levels while spending a large fraction of each day “socializing” on the internet. After several years of happy partnership, each decides to try to buy the whole business from the stockholders, who are current and former partners. When they make their bids, the one with the profitable niche market is likely to think the other office is as profitable as the person’s own office, and hence to overestimate the value of the other office. As a result, this partner could end up acquiring the business for a price that exceeds its value.

The tendency for winners’ value estimates to be biased has long been known in the oil industry, where drilling companies must make rather imperfect estimates of the likely amounts of oil that can be recovered on a given tract of land being leased. For example, Robert Wilson, a professor in the Stanford Business School, was once consulting with an oil company that was considering bidding at less than half of the estimated lease value. When he inquired about the possibility of a higher bid, he was told that firms that bid this aggressively on such a lease are no longer in this business (source: personal communication at a conference on auction theory in the early 1980’s).

The intuition underlying the unprofitability of aggressive bidding is that the firm with the highest bid in an auction is likely to have overestimated the lease value. The resulting possibility of winning “at a loss” is more extreme when there are many bidders, since the highest estimate out of a large number of estimates is more likely to be biased upward, even though each individual estimate is *ex ante* unbiased. For example, suppose that a single value estimate is unbiased. Then the higher of two unbiased estimates will be biased upward. Analogously, the highest of three unbiased estimates will be even more biased, and the highest of 100 estimates may show an extreme bias toward the largest possible estimation error in the upward direction. Knowing this, a bidder in an auction with many bidders should treat their own estimates as being inflated, which will likely turn out to be true if they win, and the resulting bids may end up being low relative to the estimates. This *numbers effect* can be particularly sinister, since the normal strategic reaction to increased numbers of bidders is to bid higher, as in a private value auction.

Wilson (1967, 1968) specified a model with this common-value structure and showed that the Nash equilibrium with fully rational bidders involved some downward adjustment of bids in anticipation of a winner’s-curse effect. Each bidder should realize that their bid is only relevant for payoffs if it is the highest bid, which means that they have the highest value estimate. So prior to making a bid, the bidder should consider what they would want to bid knowing that all other estimates are lower than theirs in the event that they win the auction. The *a priori* correction of this over-estimation produces bids that will earn positive payoffs on average.

II. A Simple Shoe-Box Model

Consider a simple situation where each of two bidders can essentially observe half of the value of the prize, as would be the case for two bidders who can drill test holes on different halves of a tract of land being leased for an oil well. Another example would be the case for the two insurance partners discussed in the previous section. In particular, let the observed value components or “signals” for bidders 1 and 2 be denoted by v_1 and v_2 respectively. The prize value is the average of the two signals:

$$(21.1) \quad \text{Prize Value} = \frac{v_1 + v_2}{2}.$$

Both bidders know their own estimates, but they only know that the other person’s estimate is the realization of a random variable. In the experiment to be discussed, each person’s value estimate is drawn from a distribution that is uniform on the interval from \$0 to \$10. Bidder 1, for example, knows v_1 and that v_2 is equally likely to be any amount between \$0 and \$10.

Consider an example of a bidder in the first round of a common-value experiment with this setup. The bidder had a relatively high value signal of \$8.69, and submitted a bid of \$5.03 (presumably the three-cent increase above \$5 was intended to outguess anyone who might bid an even \$5). The other person’s signal was \$0.60, so the prize was only worth the average of \$8.69 and \$0.60, which is \$4.64. The bidder won this prize, but paid a price of \$5.03, which resulted in a loss. Three of the five winning bidders in that round ended up losing money, and about one out of five winners lost money in each of the remaining rounds.

The prize value function in (21.1) can be generalized for the case of a larger number of bidders with independent signals, by taking an average, i.e. dividing the sum of all signals by the number of bidders. In a separate classroom experiment, conducted with 12 bidders, the sole winning bidder ended up losing money in three out of five rounds. As a result, aggregate earnings were zero or negative for most bidders. A typical case was that of bidder 9, who submitted the high bid of \$6.10 on a signal of \$9.64 in the fifth and final round. The average of all 12 signals was \$5.45, so this person lost 65 cents for the round. These results illustrate why the winner’s-curse effect is often more severe with large numbers of bidders.

Once an economist at the University of Virginia was asked to consult for a major telecommunications company that was planning to bid in the first major U.S. auction for radio wave bandwidth to be used for personal communications services. This company was a major player, but incredibly, it decided not to bid

at the last minute. The company representative mentioned the danger of overpayment for the licenses at auction. This story does illustrate a point: that players who have an option of earning zero with non-participation will never bid in a manner that yields negative expected earnings. The technical implication is that expected earnings in a Nash equilibrium for a game with an exit option cannot be negative. This raises the issue of how bidders rationally adjust their behavior to avoid losses in a Nash equilibrium, which is the topic of the next section.

III. The Nash Equilibrium

As was the case in the chapter on private-value auctions, the equilibrium bids will end up being linear functions of the signals, of the form:

$$(21.2) \quad b_i = \mathbf{b} v_i, \quad \text{where } 0 < \mathbf{b} < 1 \text{ and } i = 1, 2.$$

Suppose that bidder 2 is using a special case of this linear bid function by bidding exactly one half of the signal v_2 . In the next several pages, we will now use this assumed behavior to find the expected value of bidder 1's earnings for any given bid, and then we will show that this expected payoff is maximized when bidder 1 also bids one half of the signal v_1 , i.e. $\beta = 0.5$. Similarly, when bidder 1 is bidding one half of their signal, the best response for the other bidder is to bid one half of their signal too. Thus the Nash equilibrium for two risk-neutral bidders is to bid half of one's signal.

Since the arguments that follow are more mathematical than most parts of the book, it is useful to break them down into a series of steps. We will assume for simplicity that bidders are risk neutral, so we will need to find bidder 1's expected payoff function before it can be maximized. (It turns out to be the case that risk aversion has no effect on the Nash equilibrium bidding strategy in this case, as noted later in the chapter.) In an auction where the payoff is 0 in the event of a loss, the expected payoff is the probability of winning times the expected payoff conditional on winning. So the first step is finding the probability of winning given that the other bidder is bidding one half of their signal value. The second step is to find the expected payoff, *conditional on winning with a particular bid*. The third step is to multiply the probability of winning times the expected payoff conditional on winning, to obtain an expected payoff function, which will be maximized using simple calculus in the final step. The result will show that a bidder's best response is to bid half of the signal when the other one is bidding in the same manner, so that this strategy is a Nash equilibrium.

Before going through these steps, it may be useful to review how we will go about maximizing a function. Think of the graph of a function as a hill, which is increasing on the left and decreasing on the right, as for Figure 19.2 in the chapter on private-value auctions. At the top of the hill, a tangent line will be horizontal (think of a graduation cap balanced on the top of someone's head). So to maximize the function, we need to find the point where the slope of a tangent line is 0. The slope of a tangent line can be found by drawing the graph carefully, drawing the tangent line, and measuring its slope, but this method is not general since specific numbers are needed to draw the lines. In general, the slope of the tangent line to a function can be found by taking the derivative of the function. Both the probability of winning and the conditional expected payoff will turn out to be linear functions of the person's bid b , so the expected-payoff product to be maximized will be quadratic, with terms involving b and b^2 . A linear term, like $b/2$, is a straight line through the origin with a slope of $1/2$, so the derivative is the slope, $1/2$. The quadratic term, b^2 , also starts at the origin, and it increases to 1 when $b = 1$, to 4 when $b = 2$, to 9 when $b = 3$, and to 16 when $b = 4$. Notice that this type of function is increasing more rapidly as b increases, i.e. the slope is increasing in b . Here all you need to know is that the derivative of a quadratic expression like b^2 is $2b$, which is the slope of a straight line that is tangent to the curved function at any point. This slope is, naturally, increasing in b . Armed with this information, we are ready to find the expected payoff function and maximize it.

Step 1. Finding Bidder 1's Probability of Winning for a Given Bid

Suppose that the other person (bidder 2) is known to be bidding half of their signal. Since the signal is equally likely to be any value from \$0 to \$10, the second bidder's bid is equally likely to be any of the 1000 penny amounts between \$0 to \$10, as shown by the dashed line with height of 0.001 in Figure 21.1. When $b_2 = v_2/2$, then bidder 1 will win if $b_1 > v_2/2$, or equivalently, when the other's value is sufficiently low: $v_2 < 2b_1$. The probability of winning with a bid of b_1 is the probability that $v_2 < 2b_1$. This probability can be assessed with the help of Figure 21.1. Suppose that bidder 1 makes a bid of \$2, for example, as shown by the short vertical bar. We have just shown that this bid will win if the other's value is less than $2b_1$, i.e. less than \$4 in this example. Notice that four-tenths of the area under the dashed line is to the left of \$4. Thus the probability that the other's value is less than \$4 is 0.4, calculated as $4/10 = 2b_1/10$. A bid of 0 will never win, and a bid of \$5 will always win, and in general, we have the result:

$$(21.3) \quad \text{Probability of Winning (with a bid of } b_1) = \frac{2b_1}{10}.$$




Figure 21.1. A Uniform Distribution on the Interval [0, 10]

Step 2. Finding the Expected Payoff Conditional on Winning

Suppose that bidder 1 bids b_1 and wins. This happens when $v_2 < 2b_1$. For example, when the bid is \$2 as in Figure 21.1, winning would indicate that $v_2 < \$4$, i.e. to the left of the vertical bar in the figure. Since v_2 is uniformly distributed, it is equally likely to be any penny amount less than \$4, so the expected value of v_2 would be \$2 once we find out that the bid of \$2 won. This generalizes easily; the expected value of v_2 conditional on winning with a bid of b_1 is just b_1 (question 1). Bidder 1 knows the signal v_1 and expects the other signal to be equal to the bid b_1 if it wins, so the expected value of the prize is the average of v_1 and b_1 :

$$(21.4) \text{ Conditional Expected Prize Value (winning with a Bid of } b_1) = \frac{b_1 + v_1}{2}.$$

Step 3. Finding the Expected Payoff Function

The expected payoff for a bid of b_1 is the product of the probability of winning in (21.3) and the difference between the conditional expected prize value in (21.4) and the bid:

$$(21.5) \quad \text{Expected Payoff} = \frac{2b_1}{10} \left(\frac{b_1 + v_1}{2} - b_1 \right) = \frac{b_1 v_1}{10} - \frac{b_1^2}{10}.$$

Step 4. Maximizing the Expected Payoff Function

In order to maximize this expected payoff, we will set its derivative equal to 0. The expression on the far right side of (21.5) contains two terms, one that is quadratic in b_1 and one that is linear. Recall that the derivative of $(b_1)^2$ is $2b_1$ and the derivative of a linear function is its slope coefficient, so:

$$(21.6) \quad \text{Expected Payoff Derivative} = \frac{v_1}{10} - \frac{2b_1}{10}.$$

Setting this derivative equal to 0 and multiplying by 10 yields: $v_1 - 2b_1 = 0$, or equivalently, $b_1 = v_1/2$. To summarize, if bidder 2 is bidding half of value as assumed originally in (21.2), then bidder 1's best response is to bid half of value as well, so the Nash equilibrium bidding strategy is for each person to behave in this manner.

$$(21.7) \quad b_i = \frac{v_i}{2} \quad i=1,2. \text{ (equilibrium bid).}$$

Normally in a first-price auction, one should "bid below value," and it can be seen that the bid in (21.7) is less than the conditional expected value in (21.4). Finally, recall that this analysis began with an assumption that bidder 2 was using the strategy in (21.2) with $\beta = 1/2$. This may seem like an arbitrary assumption, but it can be shown that the only linear bidding strategy for this auction must have a slope of 1/2 (question 3).

IV. The Winner's Curse

When both bidders are bidding half of the value estimate, as in (21.7), then the one who wins will be the one with the higher estimate, i.e. the one who overestimates the value of the prize. Another way to see this is to think about what a naïve calculation of the prize value would entail. One might reason that since the other's value estimate is equally likely to be any amount between \$0 and \$10, the expected value of the other's estimate is \$5. Then knowing one's own estimate, say v_1 , the expected prize value is $(v_1 + 5)/2$. This expected prize value, however, is not conditioned on winning the auction. Notice that this unconditional expected value is greater than the conditional expected value in (21.4) whenever the person's bid is less than \$5, which is always the case when bids are half of the value estimate. Except for this boundary case where the bid is exactly \$5, the naïve value calculation will result in an overestimate of value, which can lead to an excessively high bid and negative earnings.

Holt and Sherman (2000) conducted a number of common-value auctions with a prize value that was the average of the signals. Subjects began with a cash balance of \$15 to cover any losses. Figure 21.2 shows bids and signals for the final 5 rounds of a single session. Two of the bidders can be distinguished from the others by the diamond and square marks. Notice that both of these are bidding in an approximately linear manner, but that they and the others are bidding above the $v/2$ line, which represents the Nash equilibrium. The dashed line shows the

regression line that was fit to all bids over all sessions, so we see that this session was a little high but not atypical in the nature of the bid/value relationship.




Figure 21.2. Bids for Eight Subjects in the Final Five Rounds of a Common-Value Auction
Key: Solid Line: Nash Equilibrium, Dashed Line: Regression for All Sessions
(Holt and Sherman, 2000)

The experiment session shown in Figure 21.2 was a research experiment, with students recruited from a variety of economics classes at the University of Virginia. Figure 21.3 shows analogous data for a classroom experiment with one person paid partial earnings at random. These were students in an experimental economics class, who had already read about private-value auctions and were relatively well-versed in game theory. Notice that the bids are lower, but still generally above the Nash prediction. Some people, like the person with bids marked as diamonds, were bidding about half of their signals, however.




Figure 21.3. Bids for Ten Subjects in a Classroom Common-Value Auction (final 5 rounds)

V. Extensions and Further Reading

The winner's curse was first discussed in Wilson (1969), and applications to oil lease drilling were described in Capen, Clapp, and Campbell (1971). Kagel and Levine (1986) provided experimental evidence that it could be reproduced in the laboratory, even after subjects obtain some experience. The setup used by Kagel and Levine can be implemented in class with the Veconlab program MV (under auctions). The literature on common-value auctions is surveyed in Kagel (1995).

Questions

1. Sketch a version of Figure 21.1 for the case where bidder 1's bid is \$1. Show the probability distribution for the other person's value, v_2 , conditional on the event that the first person's bid is the high bid. Then calculate the expected value of v_2 , conditional on bidder 1 winning with a bid of \$1. Then explain in words why the expected value of v_2 conditional on the winning bid b_1 is equal to b_1 .

2. (This question is somewhat mechanical, but it lets you test your knowledge of the bid derivations in the reading.) Suppose that the prize value in equation (21.1) is altered to be the *sum* of the signals instead of the average. The signals are still uniformly distributed on the range from \$0 to \$10. a) What is the probability of winning with a bid of b_1 if the other person bids an amount that equals their signal, i.e. if $b_2 = v_2$. b) What is the expected prize value conditional on winning with a bid of b_1 ? c) What is the expected payoff function for bidder 1? d) (Calculus required.) Show that the bidder 1's expected payoff is maximized with a bid that is equal to the bidder's own signal.
3. The goal of this question is to show that the only linear strategy of the form in equation (21.2) is one with a slope of $1/2$ when the prize value is the average of the signals. a) What is the probability of winning with a bid of b_1 if the other person bids: if $b_2 = \beta v_2$. (Hint: your answer should imply (21.3) when $\beta = 1/2$.) b) What is the expected prize value conditional on bidding with a bid of b_1 ? c) What is the expected payoff function for bidder 1? d) (Calculus required.) Show that the bidder 1's expected payoff is maximized with a bid that equals one half of the bidder's own signal.

Chapter 22. Multi-Unit Auctions: Irrigation Reduction and FCC Bandwidth (in progress)

This chapter pertains to situations where a single auction is used to arrange multiple transactions at the same time. The first part of the chapter describes an auction that was used to determine which tracts of land in drought-stricken Southwest Georgia would not be irrigated during the 2001 growing season. The auction design was based on a series of laboratory experiments in which participants were put in the role of being farmers with one or more tracts of land, each consisting of a specified number of acres. In each round of the auction, the bids are ranked from low to high, with low bids being provisionally accepted (up to a total expenditure set by the experimenter). Bidders do not know which round will be the final round that actually determines who must forego irrigation for that growing season and be compensated. The game can be used to motivate a discussion of the advantages of auctions, and of the many possible alternative ways that such an auction can be run. The actual auction was conducted with a web-based network of computers at eight different locations, and a similar structure is used in the Veconlab “WA” program. Alternatively, a discussion of whether or not to use a uniform price can be motivated by the hand-run instructions in the appendix. These instructions require that participants first be endowed with some small items, like ball point pens, which can be sold to the experimenter. The second part of the chapter contains a discussion of the widely discussed Federal Communications Commission (FCC) auctions for communications bandwidth. In both of the multi-unit auctions considered in this chapter, laboratory experiments have been used extensively to help design and test the procedures that were used.

I. Dry 2K

In early 2000, just after the publicity associated with the new century and the Y2K computer bugs, a severe draught plagued much of the Southeastern United States. Some of the hardest hit localities were in South Georgia, and one of the Atlanta newspapers ran a regular “Dry 2K” update on conditions and conservation measures. Of particular concern were the record low levels for the Flint River, which threatened wildlife and fish in the river and the oyster fishery in Florida. As a result, pressure was mounted to release water from a lake north of Atlanta into a parallel river to protect the oyster fishery, and the threat to Atlanta drinking water was a factor in the final decision to take drastic action.

In April 2000, the Georgia legislature passed the Flint River Drought Protection Act, mandating the use of an “auction-like process” to restrict

agricultural irrigation in certain areas if the Director of the Georgia Environmental Protection Department called a drought emergency. The unspecified nature of the mandated auction made an ideal situation for laboratory testing of alternative auction mechanisms that might be recommended. This chapter summarizes the experiments and subsequent auction results, which are reported in Cummings, Holt, and Laury (2003).

The Flint River, which begins from a drainage pipe near the Atlanta Hartsfield Airport, grows to a size that supports some barge traffic by the time it reaches the Florida state line and later empties into the Gulf. About 70 percent of the water usage in this river basin is agricultural irrigation. Farmers have permits, which were obtained without charge, for particular circular irrigation systems that typically cover areas from 50 to 300 acres. Water is not metered, and therefore, it is liberally dumped into the fields during dry periods, creating the green circles visible from the air, which makes restrictions on irrigation easy to monitor. The idea behind the law was to use the economic incentives of a bidding process to select relatively low-use-value land to retire from irrigation. Farmers would be compensated to reduce any negative political impacts. The state legislature set aside 10 million dollars, taken from its share of the multi-state Tobacco industry settlement, to pay farmers not to use one or more permits for irrigation. Therefore, what was being purchased by the state was the right to irrigate, not the use of the land itself. In reality, it is unlikely that land would be planted in a drought year in the absence of irrigation permit. Besides being non-coercive and sensitive to economic use value, the auction format had the advantage of being fair and easy to implement relative to administrative processes. Speed was also an important factor, given the limited time between the March 1 deadline for the declaration of a drought emergency and the optimal time for planting crops several weeks later. Finally, the diverse geographic locations of the farmers required that an auction collect bids from diverse locations, which suggested the use of web-based communications between officials at a number of bidding sites.

Any auction would involve a single buyer, the state Environmental Protection Department (EPD), and many sellers, the farmers with permits. (Technically speaking, the EPD officials were careful not to use purchase or sale terminology, which would have ownership implications, but this terminology will be used here for simplicity.) Permits varied in size in terms of the numbers of acres covered, and in terms of whether the water would come from the surface (river or creeks) or from wells. This is a multi-unit auction, since the state could “purchase” many permits, i.e. compensate the permit holders for not irrigating the covered areas for the specified growing season.

The goals of the people running the auction were that the auction not be viewed as being arbitrary or unfair, and that the auction take out as much irrigation as possible (measured in acres covered by repurchased permits) given

the budget available to be spent. Of course, economists would also be concerned with economic efficiency, i.e. that the auction would take less productive land out of irrigation.

A number of different types of auctions could be considered, and all involved bids being made on a per-acre basis so that bids for different sized tracks could be compared and ranked, with the low bids being accepted. One method, a discriminative auction, would have people submitting sealed bids, with the winning low bidders each receiving the amounts that they bid. For example, if the bids were \$100, \$200, \$300, and \$400 per acre for 4 permits, and if the two lowest were accepted, then the low bidder would receive \$100 per acre and the second low bidder would receive \$200 per acre. This auction is “discriminative” since different people receive different amounts for approximately the same amount of irrigation reduction per acre. In contrast, a uniform-price auction would establish a cutoff price and pay all bidders at or below this level an amount that equals the cutoff price. If the cutoff price were \$200 in the above example, then the two low bidders would each receive \$200, despite the fact that one bidder was willing to accept a compensation of only \$100 per acre. This uniform-price auction would have been a multi-unit, low-bidder-wins version of the “second-price auction” discussed in Chapter 25, whereas the discriminative auction is analogous to the first-price auction discussed in that chapter. Just as bidding behavior will differ between first and second-price auctions, bidding behavior will differ between discriminative and uniform price multi-unit auctions, so it is not obvious which one will provide the greatest irrigation reduction for a given expenditure. This is where the experiments are used. In addition, it was widely believed that collusion between farmers might become problematic, given that many of them are friends and relatives, and that many shared a distrust of Atlanta-based officials. Hence it was desirable to evaluate alternative auction procedures under conditions where the bidders could discuss bidding strategy and results freely.

The initial experiments were run in May of 2000, almost a year before the first actual auction, so we will discuss the experiments first. Early discussions with state officials indicated that discriminative auctions were preferred, to avoid the apparent “waste” of paying someone more than they bid, which would happen in a uniform-price auction in which all bidders would the same amount per acre. After some initial experiments, it became clear that a multi-round auction would remove a lot of the uncertainty that bidders would face in such a new situation. In a multi-round auction, bids would be collected, ranked, and provisional winners would be posted, but the results would not be effected if the officials running the auction decided to accept revised bids in a subsequent round. Bids that were unchanged between rounds would be carried over, but bidders would have the option of lowering or raising their bids, based on the provisional results. This

process was perceived as allowing farmers to find out approximately what the going price would be and then to compete at the margin to be included in the irrigation reduction.

Some of the early experiments were run about a year before the actual auction and were watched by state EPD officials. A conscious decision was made to provide an amount of context and realism that is somewhat unusual for laboratory experiments. Participants were recruited in groups ranging in size from 8 to over 42, and were told that they would have the role of farmers. Some of the participants were students from Atlanta, and others were farmers and locals from the South Georgia area. One final test involved over 50 local participants bidding simultaneously at 3 different locations near Albany, to test the software and communication with officials in Atlanta. The final auction, conducted in March 2001, involved about 200 farmers in 8 different locations.

Since most farmers had multiple tracts of differing sizes and productivities, each participant in the experiment was given three tracts of land. A tract was described by a specific number of acres and a use value, which is the amount of money that would be earned if that tract would be irrigated and farmed for the current growing season. Participants were told that the land would not be farmed if the irrigation permit were sold.

For example, consider a person with three permits, with acreage and use values per acre as shown in the first two columns of Table 22.1. The first permit covers 100 acres, and has a use value of 100 per acre, so the bidder would require an amount above 100 per acre as compensation for not irrigating and using the land. In this example, the bid was 120 per acre on permit 1, and this bid was accepted, so the earnings are 120 (per acre) times 100 (acres) so the total shown in the right-hand column is 12,000. If this bid had been rejected, the farmer would have farmed the land and would have earned the use value of 100 per acre times the number of acres, for a total that would have been 10,000. To be sure that you understand these calculations, please calculate the earnings for permit 2, with an accepted bid, and for permit 3, with a rejected bid. These calculations can be entered in the far-right column of the table.

Table 22.1 Sample Earnings Calculations for the Irrigation Reduction Auction

	<i>Total Acres</i>	<i>Use Value (per acre)</i>	<i>Bid (per acre)</i>	<i>Auction Outcome</i>	<i>Earnings</i>
Permit 1	100	100	120	accepted	12,000
Permit 2	50	200	250	accepted	_____
Permit 3	100	300	400	rejected	_____

After some experimentation with alternative sets of procedures and some consultation with EPD officials, we ended up focusing on two alternative setups. One proposed auction design involved a single-round, sealed-bid discriminative auction, and another involved a multi-round version of the same auction. In either case, all bids would be collected on bid sheets and ranked from low to high. Then starting with the low bids, the total expenditures would be calculated for adding permits with higher bids. A cutoff for inclusion was determined when the total expenditures reached the amount allotted by the auctioneer. For example, if the amount to be spent had been announced to be 50,000, and if the lowest 20 bids yielded a total expenditure of 49,500, and a 21st bid would take the expenditure above this limit, then only 20 bids would be accepted. This cutoff would determine earnings in the single-round sessions, with earnings on permits with rejected bids being determined by their use values. In the sessions with a multi-round setup, the cutoff would be calculated, and the permits with bids below the cutoff would be announced as being “provisionally accepted.” Then new bids would be accepted, which would then be ranked, with a new announcement of which tracts had provisionally accepted bids. This process would continue until the experimenter decided to stop the auction, at which time the accepted bids would be used to determine earnings.

A typical session began with a “trainer” auction for colored writing pens (see the instructions in the appendix to this chapter). This practice auction was intended to convey the main features of submitting different bids for different tracts, which would be ranked, with some low bids being accepted and others not. Participants were allowed to collude on bids in any manner, except that they were not permitted to keep people away from the bid submission area. Some people discussed price in small groups, and others made public suggestions.

Figure 22.1 shows the results of two sessions run with identical cost structures, except that the larger one had been scaled up to accommodate more participants. The line labeled “land values” shows the permit use values per acre, with each step having a width equal to the number of acres for the permit with a use value at that step. These opportunity costs, arrayed in this manner from low to high, constitute a supply function. In each of these sessions, we announced a fixed total budget to be used to purchase permits. So a low total acreage can be purchased at a high price per acre, and a high total acreage can be purchased at a low price per acre. This fixed budget then generates a curve that is analogous to a demand function. If B is the total budget available to purchase Q total acres at a price per acre of P , then all money is spent if $PQ = B$, or if $P = B/Q$, which generates the negative relationship between P and Q that is graphed on the left side of Figure 22.1. The connected lines on the right side of the figure show the average price per acre of the provisionally accepted bids, by round, for each of the sessions. Notice that these price averages converge to the competitive

predictions, despite the fact that participants could freely discuss the bidding process, announce suggested bids, etc. Participants were not told in advance how many rounds there would be.




Figure 22.1 Results of Two Sessions with Multiple Rounds
Source: Cummings, Holt, and Laury (2003)

One advantage of laboratory experiments is that procedures can be test-bedded and unanticipated problems can be discovered and fixed before the real auction. In one of the sessions with student subjects, the a participant asked what would happen if there happened to be a number of bidders tied at the cutoff bid that exhausted the announced budget, and if there was not enough money in the budget to cover all bids at the tie level. This possibility was not covered clearly in the instructions, and the experimenter in charge of the session announced that all of those tied would be included as provisional winners, or as final winners if this were the final round. In this session, a tie arose at a price about 5 percent above the competitive level, and all bids at that level were provisionally included. In the subsequent round, more bids came in at the tie level from the previous round, and this accumulation of bids at the focal tie point continued. The resulting payment to subjects needed to include all tied bids ended up being twice the budgeted amount. This would have been analogous to spending 20 million in the actual auction instead of the 10 million budgeted by the legislature!




Figure 22.2 Results of Two Sessions with Different Tie Breaking Rules
Source: Cummings, Holt, and Laury (2003)

A second session was run later that day with different people. The procedures were identical, except that it was announced that the bids to be included would be randomly determined in the event of a tie. The prices converged to the competitive level, as indicated in Figure 22.2.

There were a number of other procedural changes that were implemented as a result of the experiments. For example, we noticed that low bids tended to increase when we announced the cutoff bid, i.e. the highest provisionally accepted bid, at the end of each round. This pattern of increasing low bids is intuitive, since the low bidders faced less risk with bid increases if they knew about how high they could have gone in the previous round. This induced us to run some sessions where we only announced the permit numbers (but not the actual bids) for those permits that were provisional winners at the end of each round. This reduction in information resulted in less of an upward creep of low bids in successive rounds, and as before, the high bids tended to fall as bidders scrambled to become included. These modifications in tie breaking-rules and the post-round announcement procedures were incorporated into the auction rules used by the EPD in the subsequent auction.

The final auction was conducted in April 2001, with assistance from a number of experimental economists and graduate students and staff from Georgia State University, where the experiments had been done and where the results were collected and displayed to top EPD officials, who met in the Experimental Economics Laboratory at Georgia State. Almost 200 farmers turned up at 8 locations at 8 a.m. on the designated Saturday morning, along with numerous television reporters and spectators. The procedures were quite close to those that had been implemented in the experiments, except that there were no redemption

values, and all acreage amounts were those registered with the EPD. Each bid was a signed contract, with one copy going to the bidder and another to the bid officials on site. All bids were entered in computers by auction officials, as had been done in the experiments, and the resulting bids from the eight locations were ranked and projected in Atlanta, where officials then discussed whether to stop the auction, and if not, how much money to release to determine provisional winners for that round. This non-fixity of the budget differed from the procedures that had been used in the experiments, but it did not contradict any of the published auction rules. The changes in the provisionally released budget caused the cutoff bid to stay roughly the same in the first four rounds, ending up at \$125 per acre in round 4. The director of the EPD then decided to release more money and raise the cutoff bid to \$200 in round 5, when the experiment was terminated.

Figure 22.3 shows the distributions of bids in the relevant range, from \$130-\$210, a range that covered the cutoffs between bids that were provisionally accepted and those that were not. As can be seen from the figure, the bids in this range fell from one round to the next, enabling more acres to be acquired for cutoffs in this range. If a fixed budget had been used, as in the experiments, then the cutoff bid would have fallen from round to round, and the economists involved as advisors were concerned that the increasing budgets used from round to round may have discouraged some farmers from making bid reductions at the margin. In addition, the dramatic final-round increase in the budget expenditure and in the cutoff bid might have serious consequences for bidding in a future auction that used the same procedures.

The 2001 Irrigation Auction was considered to be a success. In particular, bids were received on about 60 percent of the acres that were eligible to be retired from irrigation. In total, about 33,000 acres were taken out of irrigation, at an average price of about \$135 per acre.

In 2002, the state officials, in consultation with experimental economists (Laury and Cummings) decided to run the auction as a single-round discriminative auction, with sealed bids being accepted by mail. This method was less expensive to administer, and attendance at the auction site was less important, since farmers were already familiar with the compensation and low-bids-win features of the auction that had been implemented in the previous year. In addition, the mail-in procedures may have enabled more people to participate, since attendance at the auction site was not required. A reserve price (maximum bid) of \$150 was imposed, so that bids above this amount would not be considered. In this second auction, 41,000 acres were removed from irrigation, at an average cost of \$143 per acre (McDowell, 2002).

//




Figure 22.3 Results from the Final Irrigation Reduction Auction
Source: Cummings, Holt, and Laury (2003)

II. FCC Bandwidth Auctions (unfinished)

III. Demand Reduction in Multiunit Auctions with Uniform Prices (unfinished)

List and Lucking-Reiley (2000)

IV. Extensions (unfinished)

For a nice discussion of the considerations involved in designing multi-unit auctions, see Klemperer (2002).

Part VI. Bargaining and Fairness

Bilateral bargaining is pervasive, even in a developed economy, especially for large purchases like automobiles or specialty items like housing. Bargaining is also central in many legal and political disputes, and it is implicit in family relations, care of the elderly, etc. The highly fluid give and take of face-to-face negotiations makes it difficult to specify convincing structured models that permit the calculation of Nash equilibria. Experimental research has responded in two ways. First, it is possible to look at behavior in unstructured bargaining situations to spot interesting patterns of behavior, like the well known “deadline effect,” the tendency to delay agreements until the last moment. The second approach is to limit the timing and sequence of decisions, in order to learn something about fairness and equity considerations in a simplified setting. The ultimatum bargaining game discussed in Chapter 23 is an example of this latter approach. In an ultimatum game, one person makes a proposal about how to split a fixed amount of money, and the other either accepts or rejects, in which case both earn zero. Seemingly irrational rejections in such games have fascinated economists, and more recently, anthropologists. A generalization of the ultimatum game is one where the responder can either accept the original proposal or make a counter proposal, which is then either accepted or rejected by the original proposer, and there is a discussion of behavioral patterns in such multi-stage games.

The scenarios discussed in Chapter 24 also have alternating decision structures, but the focus is on manipulations that highlight issues of fairness, trust, and reciprocity. The “trust game” begins when one person is endowed with some cash, say \$10, of which part or all may be passed to the other person. The money passed is augmented, e.g. tripled, and any part of the resulting amount may be passed back to the original person. A high level of trust would be indicated if the first person passed the full \$10, expecting the second person to reciprocate and pass back at least a third of the resulting \$30. Of course, the proposer might prefer to pass nothing if no pass-back is anticipated. The second setup that is considered, the “reciprocity game,” has more of a market context, but the underlying behavioral factors are similar. Here, employer announces a wage, and seeing this, the worker chooses an effort level, which is costly for the worker but which benefits the employer. The issue is whether fairness considerations, which are clearly present in bilateral negotiations, will have an effect on market clearing in more impersonal market settings. The final part of the chapter also pertains to a two-stage game in which the “principal” (employer) selects a contract, which is either accepted or rejected by the “agent” (worker), who then chooses an effort. The menu of contracts allows incentives (targets and probabilistic punishments) and/or non-binding bonuses that are provided *ex post*.

Chapter 23. Ultimatum Bargaining

This game provides one person with the ability to propose a final offer that must be accepted or rejected, just as a monopolist may post a price on a take-it-or-leave-it basis. The difference between the usual monopoly situation and ultimatum bargaining is that there is only one buyer of a single item in ultimatum bargaining, so a rejection results in zero earnings for both buyer and seller. Although rejections of positive amounts of money, however small, may surprise some, they will not surprise anyone who has participated in this type of bargaining, e.g. with the game “UG.” The webgame also has a setup option that allows the responder a range of intermediate responses that involve an equally proportional reduction (“squish”) of the proposed earnings for each person. This squish would be analogous to a partial purchase in the bilateral monopoly situation, at least under constant cost and value conditions.

I. “This is My Final Offer, Take it or Leave It”

Many economic situations involve a final offer from one person to another, where rejection means zero earnings on the transaction for both. For example, suppose that a local monopolist can produce a unit of a commodity for \$5. The sole buyer needs the product and is willing to pay any amount up to \$15 but not a penny more, since \$15 is what it would cost to buy the product elsewhere and pay to have it shipped into the local market. Then the “surplus” to be divided is \$10, since this is the difference between the value and the cost. The buyer knows the seller’s cost, and hence knows that a price of \$10 would split the surplus. A higher price corresponds to offering a lower amount to the buyer. The seller has a strategic advantage if the seller can make a take-it-or-leave-it offer, which we assume to be the case. This setup is called an “ultimatum game” since the proposer (seller) makes a single offer to the responder (buyer), who must either accept the proposed split of the surplus (as determined by the price) or reject, which results in zero earnings for both.

The ultimatum game was introduced by Guth et al. (1982) because it highlights an extreme conflict between the dictates of selfish, strategic behavior and notions of fairness. If each person only cares about his or her own earnings, and if more money is preferred to less, then the proposer should be able to get away with offering a very small amount to the responder. The monopolist in the previous paragraph’s example could offer a price of \$14.99, knowing that the buyer would rather pay this price than a price of \$15.00 (including shipping) from a seller in another market. The \$14.99 price is unfair in the sense that \$9.99 of the available surplus goes to the seller, and only a penny goes to the buyer. Even so,

the buyer who only cares about getting the lowest possible price should accept any price offer below \$15.00, and hence the seller should offer \$14.99.

It is easy to think of situations in which a seller might hesitate to exploit a strong strategic advantage. Getting a reservation for dinner after graduation in a college town is the kind of thing one tries to do six months in advance. Restaurants seem to shy away from allocating the scarce table space on the basis of price, which is usually not raised on graduation day. Some moderate price increases may be hidden in the form of requiring the purchase of a “special” graduation meal, but this kind of price premium is nowhere near what would be needed to remove excess demand, as evidenced by the long lead time in accepting reservations. A possible explanation is that an exorbitant price might be widely discussed and reported, with a backlash that might harm future business. The higher the price, the lower the cost of rejecting the deal. In the monopoly example discussed above, a price of \$14.99 would be rejected if the buyer is willing to incur a one-cent cost in order to punish the seller for charging such a price.

Several variants of the ultimatum game have been widely used in laboratory experiments because this game maximizes the tension between fairness considerations and the other extreme where people only care about their own earnings. Moreover, in the lab it is often possible to set up a “one-shot” situation with enough anonymity to eliminate any considerations of reputation, reward, and punishment. The next section describes an ultimatum experiment where laboratory control was a particularly difficult problem.

II. Bargaining in the Bush

Jean Ensminger (2001) conducted an ultimatum experiment in a number of small villages in East Africa. All participants were members of the Orma clan. The Orma offer an interesting case where there is considerable variation in the extent of integration into a market economy, which may affect attitudes towards fairness. The more nomadic families raise cattle and live largely off of the milk and other products, with very little being bought or sold in any market. Although some nomadic families have high wealth, which is kept in the herd, they typically have low incomes in terms of wage payments. Other Orma, in contrast, have chosen a more sedentary lifestyle for a variety of reasons, including the encroachment of grazing lands. Those who live sedentary lives in villages typically purchase food with money income obtained as wages or crop revenues. Thus exposure to a market economy is quite variable and is well measured as direct money income. Such income is not highly correlated with wealth, since some of the wealthiest are self-sufficient nomadic families with large herds. Many bargaining situations in a market context end up with individuals agreeing

to “split the difference,” so a natural conjecture is that behavior in an ultimatum game will be related to the degree of exposure to a market economy.

A survey was completed for each household, and at least one adult was recruited from each household to play “fun games for real money.” The experiments were conducted in grass houses, which enabled Ensminger to isolate groups of people during the course of the experiment. A “grand master,” who was known to all villagers, read the instructions. This person would turn away to avoid seeing decisions as they were made. The amount of money to be divided between the proposer and responder in each pair was set to be approximately equal to a typical day’s wages (100 Kenyan shillings). The game was only played once. Each proposer would make an offer by moving some of the shillings to the other side of the table, before leaving the room while the responder was allowed to accept or reject this offer. Ensminger reports that the people enjoyed the games, despite some amusement at the “insanity” and “foolishness” of Western ways. She tried to move from one village to another before word of results arrived.




Figure 23.1. Distribution of Offers in an Ultimatum Game with the Orma (56 Pairs, One-Day Wage Stake; Source: Ensminger, 2001)

The data for 56 bargaining pairs are shown in Figure 23.1. The average offer was 44 percent of the stake, with a clear mode at 50 percent. The lowest offer was 30 percent, and even these unequal splits were rarely rejected (as indicated by the light part at the bottom of the bar). If people had foreseen that the 40 percent would all be accepted, then they might have lowered their offers.

It is clear that the modal offer is not optimal in the sense of expected-payoff maximization against the actual *ex post* rejection pattern. People who made these generous offers almost always mentioned fairness as the justification in follow-up interviews. Ensminger was suspicious, however, and approached some reliable informants who revealed a different picture of the “talk of the village.” The proposers were apparently *obsessed* with the possibility that low offers would be rejected, even though rejections were thought to be unlikely. Such obsessions suggest that expected-payoff maximization may not be an appropriate assumption for large amounts of money, e.g., a day’s income. This conjecture would be consistent with the payoff-scale effects on risk aversion discussed previously in Chapter 4.

Nevertheless, it cannot be the case that individuals making relatively high offers were doing so *solely* out of fairness considerations, since offers were much lower in a second experiment in which the proposer’s split of the same amount of money was automatically implemented. This game, without any possibility of rejection, is called a “dictator game.” The average offer fell from 44 percent in the ultimatum game to 31 percent when there was no possibility of rejection. Even though there seems to be some strategic reaction by proposers to their advantage in the dictator game, the modal offer was still at the “fair” or “fifty-fifty” division, and less than a tenth of the proposers kept all of the money.

Ensminger used a multiple regression to evaluate proposer offers in both the ultimatum and dictator games. In both cases, the presence of wage income is significantly related to offers; those with such market interactions tend to make higher offers. Her conjecture is that people exposed to face-to-face market transactions may be more used to the notion of “splitting the difference.” Variables such as age, gender, education, and wealth (in cattle equivalents) are not significant. This intra-cultural finding is also evident in a cross-cultural study involving 15 small-scale societies on five continents (Henrich et al., 2001). This involved one-shot ultimatum games with comparable procedures that were conducted by anthropologists and economists. There was considerable variation, with the average offer ranging from 0.26 to 0.58 of the stake. The societies were ranked in two dimensions: the extent of economic cooperation and the extent of market integration. Both variables were highly significant in a regression, explaining about 61 percent of the variation, whereas individual variables like age, sex, and relative wealth were not. The lowest offers were observed in societies where very little production occurred outside of family units (e.g., the Machiguenga of Peru). The highest offers were observed in a society where production involved joint effort; offers above one half for the Lamelara of Indonesia, who are whale hunters in large sea-going canoes.

A one-shot ultimatum game, played for money, is a strange new experience for these people, and behavior seemed to be influenced by parallels

with social institutions in some cases. The Ache of Paraguay, for example, made generous offers, above 0.5 on average. The proposed sharing in the ultimatum game has some parallels with a practice whereby Ache hunters with large kills will leave them at the edge of camp for others to find and divide.

III. Bargaining in the Lab

Ultimatum games have been conducted in more standard, student subject pools in many developed countries, and with stakes that are usually about \$10. The mean offer, as a fraction of the stake, is typically about 0.4, and there seems to be less variation in the mean offer than was observed in the small-scale societies. Roth et al. (1991) report ultimatum game experiments that were run in four universities, each in a different country. The modal offer was 0.5 in the U.S. and Slovenia, as was the case for the Orma and for other studies in the U.S. (e.g., Forsythe et al., 1988). The modal offer was somewhat lower, 0.4, in Israel and Japan. Rejection rates in Israel and Japan were no higher than in the other two countries, despite the lower mean offers, which led the authors to conjecture that the differences in behavior across countries were due to different cultural norms. These differences in behavior reported by Roth et al. (1991) were, however, lower than the differences observed in the 15 small-scale societies, which are less homogeneous in the nature of their economic production activities.

Recall that the Orma made no offers below 0.3. In contrast, some student subjects in developed countries make offers of 0.2 or lower, and these low offers are rejected about half of the time. For example, consider the data in Figure 23.2, which is for a one-shot ultimatum game that was done in class, but with full money payments and a \$10 stake. The mean offer was about 0.4 (as compared with 0.44 for the Orma), but about a quarter of the offers were 0.2 or below, and these were rejected a third of the time.

Ultimatum bargaining behavior is relatively sensitive to various procedural details, and for that reason extreme care was taken in the Henrich et al. cross-cultural study. Hoffman et al. (1994) report that the median offer of 5 in an ultimatum game was reduced to 4 by putting the game into a market context. The market terminology had the proposer play the role of a seller who chose a take-it-or-leave-it price for a single unit. Interactions with posted prices are more anonymous than face-to-face negotiations, and market price terminology in the laboratory may stimulate less generous offers for this reason. The median offer fell again, to 3, when high scores on a trivia quiz were used to decide which person in each pair would play the role of the seller in the market. Presumably, this role-assignment effect is due to the fact that people may be more willing to accept an aggressive (low) offer from a person who earned the right make the offer.




Figure 23.2. Distribution of Offers in a One-Shot Ultimatum Game (11 Pairs, \$10 Stake, Cash Payments, University of Virginia, Spring 2001)

It is somewhat unusual for a seller to offer only one unit for sale to a buyer, and a multi-unit setup provides the buyer with an option for partial rejection. For example, suppose that there are 10 units for sale. Each costs \$0 to produce, and each unit is worth \$1 to the buyer. The seller posts a price for the 10 units as a group, but the buyer can decide to purchase a smaller number, which reduces each party's earnings proportionately. For example, suppose that the seller posts a price of \$6, which would provide earnings of \$6 for the seller and \$4 for the buyer. By purchasing only half of the units, the buyer reduces these earnings to \$3 for the seller and \$2 for the buyer. This partial rejection option is implemented in the *Veconlab* software as a "squish" option (Andreoni, Castillo, and Petrie, 2003). If this option is permitted, the responder can choose a fraction that indicates the extent of acceptance: with 0 being full rejection and 1 being full acceptance. This option was used in the classroom experiment shown in Figures 23.3 (first round) and 23.4 (fifth and final round). The extreme offer of 0.2 in the first round was squished by a half. Offers in the fifth round were more clustered about 0.5.




Figure 23.3. Distribution of Offers in Round 1 of a Classroom Ultimatum Game (6 Pairs, University of Virginia, Spring 2002)




Figure 23.4. Distribution of Offers in Round 5 of a Classroom Ultimatum Game (6 Pairs, University of Virginia, Spring 2002)

IV. Multi-Stage Bargaining

Face-to-face negotiations are typically characterized by a series of offers and counter-offers. The ultimatum game can be transformed into a game with many stages by letting the players take turns making proposals about how to split an amount of money. If there is an agreement at any stage, then the agreed split is implemented. A penalty for delay can be inserted by letting the size of the stake shrink from one stage to the next. For example, the amount of money that could be divided in the first stage may be \$5, but a failure to reach an agreement may reduce this stake to \$2 in the second stage. If there is a pre-announced final stage and the responder in that stage does not agree, then both earn zero. Thus the final stage is like an ultimatum game, and can be analyzed as such.

Now consider a game with only two stages, with a money stake, $\$Y$, in the first stage and a lower amount, $\$X$, in the final stage. For simplicity, suppose that the initial stake is \$15, which is reduced to \$10 if no agreement is reached in the first round. *Warning: the analysis in this paragraph is based on the (very unreasonable) assumption that people are perfectly rational and care only about their own payoffs.* Put yourself in the position of being the proposer in the first stage of this game, with the knowledge that both of you would always prefer the action with the highest payoff, even if that action only increases one's payoff by a penny. If you offer the other player too little in the first stage, then this offer will be rejected, since the other player becomes the proposer in the final (ultimatum game) stage. So you must figure out how much the other person would expect to earn in the final stage when they make a proposal to split the \$10 that remains. In theory, the other person could make a second-stage offer of a penny, which you would accept under the assumption that you prefer more money (a penny) to less (zero). Hence the other person would expect to earn \$9.99 in the final stage, assuming perfect rationality and no concerns for fairness. If you offer them less than \$9.99 in the first stage, they will reject, and if you offer them more they will accept. The least you could get away with offering in the first stage is, therefore, just a little more than \$9.99, so you offer \$10, which is accepted.

In the previous example, the theoretical prediction is that the first-stage offer will equal the amount of the stake that would remain if bargaining were to proceed to the second stage. This result can be generalized. If the size of the money stake in the final stage is $\$X$, then the person making the offer in that stage “should” offer the other a penny, which will be accepted. The person making the proposal in the final stage can obtain essentially the whole stake, i.e. $\$X - \0.01 . Thus the person making an offer in the first stage can get away with offering a slightly higher amount, i.e. $\$X$, which is accepted. In this two-stage game, the initial proposer earns the amount by which the pie shrinks, $\$Y - \X , and the other person earns the amount remaining in the second stage, $\$X$.

An experiment with this two-stage structure is reported in Goeree and Holt (2001). The size of the stake in the initial stage was \$5 in both cases. In one treatment, the pie was reduced to \$2.00, so the first-stage offer should be \$2.00, as shown in the middle column of Table 21.1. The average first-stage offer was \$2.17, quite close to this prediction. In a second treatment, shown on the right side of the table, the pie shrunk to \$0.50, so the prediction is for the first-stage offer to be very inequitable (\$0.50). The average offer was somewhat higher, at \$1.62, as shown in the right column of the table. The reduction in the average offer was much less than predicted by the theory, and rejections were quite common in this second treatment. Notice that the cost of rejecting a low offer is low, and knowing this, the initial proposers were reluctant to exploit their advantage fully in this second treatment.

Table 21.1. A Two-Stage Bargaining Game Played Once
(Source: Goeree and Holt, 2001)

	Treatment 1	Treatment 2
Size of Pie in First Stage	\$5.00	\$5.00
Size of Pie in Second Stage	\$2.00	\$0.50
Selfish Nash First-Stage Offer	\$2.00	\$0.50
Average First-Stage Offer	\$2.17	\$1.62

The effects of payoff inequities are even more dramatic in a second two-stage experiment reported by Goeree and Holt (2000). The initial proposers in this design made seven choices corresponding to seven different bargaining situations, with the understanding that only one of the situations would be selected afterwards, at random, before the proposal for that situation was communicated to the other player. The initial pie size was \$2.40 in all seven cases, but the second-stage pie size varied from \$0.00 to \$2.40, with five intermediate cases.

The two extreme treatments for this experiment are shown in Table 21.2. In the middle column, the pie shrinks from \$2.40 to \$0.00. This gives the initial proposer a large strategic advantage, so the initial offer would be a penny in a game between two selfish, rational players. This outcome would produce a sharp asymmetry in the payoffs in favor of the proposer. Adding insult to injury, the instructions for this case indicated that the initial proposer would receive a fixed payment of \$2.65 in addition to the earnings from the bargaining, whereas the

initial responder would only receive \$0.25. Only if the initial proposer were to offer the whole pie of \$2.40 in the first stage would final earnings be equalized. This is listed as the “egalitarian first-stage offer” of \$2.00 in the middle column. The average of the actual offers (\$1.59) shown in the bottom row, is well above the Nash prediction, and is closer to the egalitarian offer of \$2.40.

The other extreme treatment is shown in the right-hand column of Table 21.2. Here the pie does not shrink at all in the second stage, so the theoretical prediction is that the first-stage offer will be \$2.40. The initial responder now has the strategic advantage, since the pie remains high in the final (ultimatum) stage when this person has the turn to make the final offer. To make matters even more asymmetric, this strategic advantage is complemented with a high fixed payment to the initial responder. The only way for the disadvantaged initial proposer to obtain equal earnings would be to offer nothing in the first stage, despite the fact that the theoretical prediction (assuming selfish behavior) is \$2.40. The average of the observed offers, \$0.63, is much closer to the egalitarian offer.

Table 21.2. A Two-Stage Bargaining Game with Asymmetric Fixed Payments
(Source: Goeree and Holt, 2000)

Size of Pie in First Stage	\$2.40	\$2.40
Size of Pie in Second Stage	\$0.00	\$2.40
Proposer Fixed Payment	\$2.65	\$0.25
Responder Fixed Payment	\$0.25	\$2.65
Egalitarian First-Stage Offer	\$2.40	\$0.00
Selfish Nash First-Stage Offer	\$0.01	\$2.40
Average First-Stage Offer	\$1.59	\$0.63

To summarize, the first-stage offers in this experiment should be equal to the remaining pie size, but the asymmetric fixed payments were structured so that the egalitarian offers would be inversely related to the remaining pie size. This inverse relationship was generally present in the data for the seven treatments. The authors show that the data patterns are roughly consistent with an enriched model in which people care about relative earnings as well as their own earnings. For example, a person may be willing to give up some money to avoid having the other person earn more, which is an aversion to disadvantageous inequity. Roughly speaking, think of this as an “envy effect.” It is also possible that people might wish to avoid making significantly more than the other person, which

would be an aversion to advantageous inequity. Think of this as a kind of “guilt effect,” which is likely to be weaker than the envy effect. These two effects are captured by a formal model of inequity aversion proposed by Fehr and Schmidt (1999). Goeree and Holt (2000) estimate guilt and envy parameters for their data and conclude that the envy effect is more pronounced.

V. Extensions and Further Reading

The first ultimatum experiment was reported by Guth et al. (1982), and the results were replicated by Forsythe et al. (1988), who introduced the dictator game. Economists and others have been fascinated by behavior in these games where there is a high tension between notions of fairness and strategic, narrowly self-interested behavior. Ensminger (2001) reports that one of her African subjects jovially remarked: “I will be spending years trying to figure out what this all meant.”

Many economists were initially skeptical of the high degree of seemingly non-strategic play and costly rejections. One reaction was that “irrational” rejections would diminish when the stakes of the game are increased. Hoffman et al. (1996) increased the stakes from \$10 to \$100, which did not have much effect on initial proposals. The effects of high stakes have also been studied by Slonim and Roth (1998) and List and Cherry (2000).

Rejections are not, of course, irrational if individuals have preferences that depend on relative earnings. For example, a responder may prefer that both earn equal zero amounts to a situation with inequitable positive earnings. Bolton and Ockenfels (1998) proposed a model with preferences based on relative earnings, and Fehr and Schmidt (1999) developed a closely related model of inequity aversion that was mentioned above.

Finally, there have been many experiments in other, less-structured negotiations, where the order of proposal and response is not imposed by the experimenter. For example, Hoffman and Spitzer (1982, 1985) used an open, unstructured setting to evaluate the ability of bargainers to agree on efficient outcomes, irrespective of property rights (the Coase theorem). See Roth (1995) for a survey of the literature on bargaining experiments.

Chapter 24. Trust, Reciprocity, and Principal-Agent Games

Although increases in the size of the market may promote productive specialization and trade, the accompanying increase in anonymity raises the need to trust in trading relationships. The trust game sets up a stylized situation where one person can decide how much of an initial stake to keep and how much to pass to the other person. All money passed is augmented, and the responder then decides how much of this augmented amount to keep and how much to pass back to the initial decision maker. The Veconlab game, TG, puts participants into this situation so that they can experience the tension between private motives and the potential gains from trust and cooperation.

The “gift-exchange” view of wage setting is that employers set wages above market-clearing levels in an effort to elicit high effort responses, even though those responses are not explicitly rewarded *ex post* after wages have been set. The reciprocity game, RG, is one where each employer is matched with a worker who first sees the wage that is offered and then decides on the level of costly effort to supply. The issue is whether notions of trust and reciprocity may have noticeable effects in a market context.

The final part of the chapter pertains to another employer/employee relationship, where the employer begins by selecting the type of contract, from a menu that includes either *ex ante* incentives or less-structured *ex post* bonuses. This “principal-agent game” can be implemented with the Veconlab game, PA.

I. The Trust Game

The trust game was examined in an experiment by Berg, Dickhaut, and McCabe (1995), with the objective of studying trust and reciprocity in a controlled setting. A standard version begins when one person in each pair is given \$10. The first mover must decide how much (if any) of this money to pass to the other person, and how much to keep. Money that is passed is tripled before it is given to the second mover, who decides how much (if any) to return to the first person. The first mover earns the amount kept initially plus any money that is returned. The second mover earns the amount that is kept in the second stage. The game would typically be explained as an “investment game,” in order to avoid the suggestive “trust” terminology. If this game is only played once, as is often the case in experiments, then the subgame perfect Nash equilibrium for selfish players is for the second person to keep all that is passed, and hence for the first mover to pass nothing. The action of passing money initially would signal

that the first mover trusts the second mover to return a reasonable amount of money, perhaps due to a feeling of reciprocity generated by the initial action.

Berg, Dickhaut and McCabe ran this experiment with 32 pairs of participants in a single round interaction. Of these, almost all of the senders (30 of 32) passed a positive amount of money, and about a third of the responders returned more than was sent. Both of these behavior patterns are inconsistent with the prediction of a subgame-perfect Nash equilibrium for purely selfish players. The average amount passed was \$5.16 (out of \$10), and after being tripled, the average amount returned was 18 percent.

Table 24.1 shows the amounts passed and returned for 6 pairs of individuals in a demonstration experiment run at the University of Virginia with full payment and with the standard parameters (\$10 given to proposers, with the amount passed being tripled). Thus, in the aggregate, proposers are given \$60. Based on the Berg, Dickhaut, and McCabe results, one would expect about half (\$30) to be passed, which when tripled, would become \$90, and about 18% of that (\$16) to be returned. For the six pairs shown in the table, the aggregate amount passed was \$33.00, which was tripled to \$99.00, but only \$10.00 was returned. The responders in this experiment were clearly not living up to the expectations that proposers had, unless proposers somehow were just trying to increase responder earnings.

Table 24.1 A Demonstration Trust Game Experiment

	<i>ID1</i>	<i>ID2</i>	<i>ID3</i>	<i>ID4</i>	<i>ID5</i>	<i>ID6</i>
Amount Passed	\$0.00	\$10.00	\$2.00	\$10.00	\$1.00	\$10.00
Amount Returned	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$10.00
Proposer's Earnings	\$10.00	\$0.00	\$8.00	\$0.00	\$9.00	\$10.00
Responder's Earnings	\$0.00	\$30.00	\$6.00	\$30.00	\$3.00	\$20.00

The results of a classroom trust game are shown in Figure 24.1; the matchings were random in the first treatment and fixed in the second. This setting clearly induced a significant amount of money being passed, on average, and return rate that was just enough, on average, to reimburse the first movers.




Figure 24.1 A Classroom Trust Game

There has been a fair amount of discussion of possible reasons for deviations of behavior from the Nash predictions for selfish players. Cox (1999) replicated the Berg, Dickhaut, and McCabe results, and compared the amounts passed with behavior in a second treatment where the responder did not have the opportunity to reciprocate. The amounts passed were not significantly different for the two treatments, which is evidence that the pass behavior may not be due to anticipated reciprocity. Cox suggested that the positive amounts typically passed and tripled were due to altruism, i.e. a kindness or “other-regarding preference” for increasing another person’s earnings, even at some cost to oneself. Altruism is often mentioned as an explanation for voluntary contributions in public goods games discussed in a Chapter 26. If altruism is driving these contributions, then it cannot be the whole story, since the responder behavior is not very cooperative. One difference is that what is passed is tripled and what is returned is not altered, so the cost of altruism is lower for proposers. Cox argues that altruism is a factor in responder behavior, but he notes the reciprocity may also be involved. Such reciprocity could arise because participants may feel some social pressure to return a part of the gains from money passed, despite the anonymity of the experimental procedures. A similar explanation would be that concern for the other’s earnings goes up if the other person has shown some initial kindness.

Many economic interactions in market economies are repeated, and this repetition may enable trading partners to develop trust and reciprocal arrangements. For one thing, either person has the freedom to break out of a binary trading relationship in the event that the other’s performance is not

satisfactory. Even when a breakup is not possible, the repetition may increase levels of cooperation. For example, Cochar, Van Phu, and Willinger (2000) report results of a trust game that was run both as a single-shot interaction and as a repeated interaction for 7 periods, but with payoff parameters scaled down to account for the higher number of periods. There were some other minor procedural differences that will not be discussed here. The main result is that the amounts passed and returned were higher in the repeated-interaction treatment. In the final period, which was known in advance, the amount passed stayed high, but the amount returned was very low.

II. A Labor Market Game

This game implements a setup where participants are matched in pairs, with one setting a wage and the other choosing an effort level. The person in the employer role is free to choose any wage between specified limits. This wage is paid irrespective of the worker's subsequent effort decision, which must be between 0 and some upper limit. A higher level of worker effort is costly to the worker and beneficial to the employer. This is sometimes called a reciprocity game, since the employer may offer a high wage in the hope that the worker will reciprocate with high effort. In the experiment to be discussed, the wage was between \$0 and \$10, and the worker effort was required to be between 0 and 10. Each additional unit of effort reduced the worker's earnings by \$0.25, and added \$3.00 to the employer's earnings. Thus economic efficiency would require an effort of 10 in this context. Employers did not know workers' costs, and workers did not know the value of effort to the employer.

The results of a classroom reciprocity game with these parameters are shown in Figure 24.2. The first 10 periods were done with random matchings, which gives workers little incentive to provide acceptable efforts. The average efforts, shown by the small dots, leveled off at about 2 in rounds 3-8, followed by a sharp fall at the end. The final period efforts of 0 indicate that workers were aware that they had no incentive to be cooperative if the employers had no opportunity to reciprocate. The second treatment was done with fixed matchings, and the increased incentive to cooperate resulted in much higher effort levels (until the final period) and higher earnings. It is curious that the repeated nature of the matching process increased trust and reciprocity more in this game than in the trust game in Figure 24.1.




Figure 24.2. A Classroom Reciprocity Game

III. A Principal-Agent Game

The principal-agent model is one in which one person, the principal, hires the other, the agent, to perform some task. The game begins when the principal selects the nature of a contract payment, which can be contingent on observed and verifiable outcome measures. Then the agent may either reject the contract, and both earn default payments, or accept and perform the required task, with payment from the principal depending on outcomes as specified in the contract. As before, a high effort by the agent is costly to the agent and beneficial to the principal. This model has been widely applied in labor economics and contract law, and generalizations are used in many other two-person, sequential-decision settings where payments are determined by explicit or implicit contracts.

Figure 24.3 shows data from a classroom principal-agent experiment run in a Contract Law class at the University of Virginia Law School. Each of the 12 employers began by choosing a contract that specified a suggested effort and a fixed wage to be paid regardless of the worker's effort. Employers also chose between an "incentive contract" with a potential punishment and a "bonus contract" with a non-binding promise of an *ex post* bonus if a suggested level of effort was met. The wage had to cover the worker's effort cost at the suggested effort level, and upper limits on possible penalty and bonus payments were enforced. A penalty could only be imposed if low effort could be verified by a third party, which occurred with a pre-specified probability of 1/3. Workers could either reject the contract or accept and select an effort from 1 to 10. About three-fourths of the contracts were bonus contracts. Effort levels declined, and about

half of the efforts ended up at the minimum level of 1. Efforts were lower for penalty contracts.




Figure 24.3. A Classroom Principal-Agent Game

The particular setup used in this classroom exercise is based on the experiments reported in Fehr, Klein, and Schmidt (2001) "Fairness, Incentives, and Contractual Incompleteness," Working Paper 72, Institute of Empirical Research in Economics, University of Zurich, although their "hand-run" procedures are not exactly replicated by the Veconlab interactive program. The default payoffs are obtained by dividing the Fehr, et al. parameters by 10 (division by 8 would reflect the token to DM to U.S. dollar conversion rate more accurately). The results of their research experiments are roughly consistent with the classroom data patterns shown in Figure 24.3. In particular, bonus contracts were more widely used than incentive contracts.

IV. Extensions (unfinished)

For experimental evidence of reciprocity in a related context, see Fehr, Kirchsteiger, and Riedl (1993).

Part VII. Public Choice

Public expenditures are typically influenced by voting and other political processes, and the outcomes of such processes can be very sensitive to the rules that govern voting. Chapter 25 is an introduction to the results of voting experiments, and some related field experiments from the political science literature.

Many public programs and policies are designed to remedy situations where the actions taken by some people affect the well being of others. The classic example is that of the provision of a “public good” like national defense, which can be consumed freely by all without crowding and the possibility of exclusion. Chapter 26 introduces a model of voluntary contributions to a public good, where the private cost to each person is less than the social benefit. The setup of the game allows the independent alteration of (internal) private cost and (external) public benefit. Non-selfish motives for giving are discussed in the context of laboratory experiments.

A similar situation with “external” benefits arises when it only takes one “volunteer” to provide the good outcome preferred by all. The “volunteer’s dilemma” discussed in Chapter 27 is whether to incur the private cost to achieve this good outcome. This game is different from a standard public goods setup in that the private benefit to the volunteer exceeds the cost. For example, suppose that you would rather jump in the icy river and rescue someone, but you would prefer not to jump if one or more other bystanders decide to volunteer. Data from volunteer’s dilemma games are often used to evaluate Nash predictions of the effects of changing the number of potential volunteers.

External effects on others’ well being may arise in a negative sense as well, as is the case with pollution or overuse of a common resource. The “common pool resource” game in Chapter 28 is a setting where each person’s efforts to secure benefits from a shared resource tend to diminish the benefits that others derive from their efforts, as might happen with excessive harvests from a fishery. The resource is like a public good in the sense that the problem arises from non-exclusion, but the difference is the presence of congestion effects, i.e. it is not non-rival. The simplest common pool resource game is one in which people select efforts independently, and the benefits achieved are diminished as total effort increases. The focus of the discussion is on the extent of overuse.

The final chapter introduces a problem of wasteful competition that may arise in non-market allocation procedures. In “beauty contest” competitions for a broadcast license, for example, the contenders may spend considerable amounts of real resources in the process of lobbying. These expenditures are not recovered by unsuccessful contenders. Such lobbying expenses are examples of “rent

seeking,” and the total cost of such activities may even exceed the value of the prize (full dissipation of rents). Chapter 29 is based on Gordon Tullock’s model of rent seeking, where the probability of obtaining the prize is the proportion of total effort that one exerts. Factors that affect rent dissipation in theory are evaluated in the context of laboratory experiments.

Chapter 25. Voting (in progress)

Decentralized equilibrium outcomes are highly efficient in many economic markets, and as a consequence, economists sometimes fall into the trap of thinking of the outcomes of a political process as necessarily having some kind of intrinsic value. In contrast, political scientists are well aware that the outcome of a vote may depend critically on factors like the structure of the agenda, whether straw votes are allowed, and the extent to which people vote strategically or naively (political scientists use a less pejorative term, “sincerely”). Voting experiments allow one to evaluate alternative political institutions, and to study how people actually behave in controlled laboratory conditions. These studies are complemented by field experiments that seek to create contexts with more external validity, e.g. having voters view media ads with differing formats. These laboratory and field experiments are selectively surveyed in this chapter. Voting experiments can be run with the Veconlab program VT, or with the instructions provided in the Appendix. Parts of this chapter are based on Anderson and Holt (1997) and parts are based on a survey by Holt and Moore (2004b).

Chapter 26. Voluntary Contributions

This chapter is based on the standard voluntary contributions game, in which the private net benefit from making a contribution is negative unless others reciprocate later or unless the person receives satisfaction from the benefit provided to others. The setup makes it possible to evaluate independent variations of the private internal benefit and the public external benefit to others. These and other treatment manipulations in experiments are used to evaluate alternative explanations for observed patterns of contributions to a public good. The public goods game, PG, should be conducted prior to class discussion. Alternatively, a hand-run version using playing cards is easy to implement using the instructions provided in the appendix. In some contexts, a target level of contributions is required for a public good to be provided, and this setup can be implemented with the provision-point public goods program (PPG).

I. “Economists Free Ride, Does Anyone Else?”

The selfish caricature of *homo economicus* implies that individuals will “free ride” off of the public benefits provided by others’ activities. Such free riding may result in the under-provision of public goods. A pure public good, like national defense, has several key characteristics:

- 1) It is *jointly provided* and *non-excludable* in the sense that the production required to make the good available to one person will ensure that it is available to all others in the group. In particular, access cannot be controlled.
- 2) It is *non-rivalled* in the sense that one person’s consumption of the good is not affected by another’s, i.e. there is no congestion.

When a single individual provides a public good, like shoveling a sidewalk, the private provision cost may exceed the private benefit, even though the social benefit for all others’ combined may exceed the provision cost incurred by that person. The resulting misallocations have been recognized since Adam Smith’s (1776) discussion of the role of government in the provision of street lamps.

In many cases, there is not a bright-line distinction between private and public goods. For example, parks are generally considered to be public goods, although they may become crowded and require some method of exclusion. A good that is rivalled but non-excludable is sometimes called a “common-pool resource,” which is the topic of a subsequent chapter. With common-pool

resources like ground water resources, fisheries, or public grazing grounds, the problem is typically one of how to manage the resource to prevent overuse, since individuals may not take into account the negative effects that their own usage has on what is available for others. In contrast, most public goods are not provided by nature, and the major problems often pertain to provision of the appropriate amounts of the good. This chapter is mainly focused on the voluntary provision of such goods, although there is some discussion of voting and mechanisms that may be used to improve the situation.

Goods are often produced by those who receive the greatest benefit, but under-provision can remain a problem as long as there are some public benefits to others that are not fully valued by the provider. Education, for example, offers clear economic advantages to the student, but the public at large also benefits from having a well-educated citizenry. A public goods problem remains when not all of the benefits are enjoyed by the provider, and this is one of the rationales for the heavy public involvement in school systems. The mere presence of public benefits does not justify public provision of such goods, since public provision may involve inefficiencies and distortions due to the need to collect taxes. The political problems associated with public goods are complicated when the benefits are unequally distributed, e.g., public broadcasting of cultural materials.

In close-knit societies, public goods and common-pool resource problems may be mitigated by the presence of social norms that dictate or reward other-regarding behavior. For example, Hawkes (1993) reports that large game hunters in primitive societies are expected to share a kill with all households in a village, and sometimes even with those of neighboring villages. Such widespread sharing seems desirable given the difficulty of meat storage and the diminishing marginal value of excess consumption in a short period of time. These social norms transform a good that would normally be thought of as private into a good that is jointly provided and non-excludable. This transformation is desirable since the private return from large-game hunting is lower than the private return from gathering and scavenging. For example, the “!Kung” of Botswana and Namibia are a hunter-gatherer society where large prey (e.g. warthogs) are widely shared, whereas small animals and plant food are typically kept within the household. On the basis of published accounts, Hawkes (1993) estimates that, at one point, male large-game hunters acquired an average of 28,000 calories per day, with only about a tenth of that (2,500) going to the person’s own household. In contrast, the collection of plant foods yielded an estimated range around 5,000 calories per day, even after accounting for the extra processing required for the preparation of such food. Notice that large-game hunting by males was more productive for the village as a whole, but had a return for one’s household of about half the level that could be obtained from gathering activities. Nevertheless, many of the men continued to engage in hunting activities, perhaps due to the tendency for all to

share their kills in a type of reciprocal arrangement. Hawkes, however, questioned the reciprocity hypothesis since some men were consistently much better hunters than others, and yet all were involved in the sharing arrangements. She stressed the importance of more private, fitness-related incentives, and she conjectured that successful large-game hunters have more allies and better opportunities for mating.

Economists are sometimes ridiculed for ignoring social norms and simply assuming that individuals will typically free-ride on others' generosity. An early public goods experiment is reported by two sociologists, Marwell and Ames (1981). Their experiment involved groups of high school students who could either invest an initial endowment in a "private exchange" or a "public exchange." Investment in the public exchange produced a net loss to the individual, even though the benefits to others were substantially above the private cost to the individual subject. Nevertheless, the authors observed significant amounts of investment in the public exchange, with the major exception being when the experiment was done with a group of economics doctoral students. The resulting paper title began: "Economists Free Ride, Does Anyone Else?"

The Marwell and Ames paper initiated a large literature on the extent to which subjects in experiments incur private costs in activities that benefit others. A typical experiment involves dividing subjects into groups, and giving each one an endowment of "tokens" that can be invested in a private exchange or account, with earnings per token that exceed the earnings per token obtained from investment in the public account. For example, each token might produce 10 cents for the investor when invested in the private account and only 5 cents for the investor *and for each of the others* when invested in the public account. In this example, the social optimum would be to invest all tokens in the public account as long as the number of individuals in the group, N , is greater than 2, since the social benefit $5N$ would be greater than the private benefit of 10 for each token when $N > 2$. The 10-cent private return can be thought of as the opportunity cost of investment in the public account. The ratio of the per capita benefit to the opportunity cost is sometimes called the "marginal per capita return" or MPCR, which would be $5/10 = 0.5$ in this example. A higher MPCR reduces the net cost of making a contribution to the public account. For example, if the private account returns 10 cents and the per-capita return on the public account is raised to 9 cents, there is only a 1 cent private loss associated with investment. Many of the experiments involved changes in the MPCR.

A second treatment variable of interest is the number of people involved, since a higher group size increases the social benefit of an investment in the public exchange when the MPCR is held constant. For example, the social benefit in the example from the previous paragraph is $5N$, which is increasing in N . Alternatively, think about what you would do if you could give up ten dollars

in order to return a dollar to every member of the student body at your university, including yourself. Here the MPCR is only 0.1, but the public benefit is extremely large. The motives for contribution to a public good are amplified if others are expected to reciprocate in some future period. These considerations suggest that contributions might be sensitive to factors such as group size, the MPCR, and whether or not the public goods experiment involves repetition with the same group. In multi-round experiments, individuals are given new endowments of tokens at the start of each round, and groupings can either be fixed or randomly reconfigured in subsequent rounds.

The upshot is that the extent of voluntary contributions to public goods depends on a wide variety of procedural factors, although there is considerable debate about whether contributions are primarily due to kindness, to reciprocal reactions to others' kindness, or to confusion. For example, there is likely to be more confusion in a single-shot investment decision, and many people may initially divide their endowment of "tokens" equally between the two types of investment, public and private. This is analogous to the typical choice of dividing one's retirement fund contributions equally between stocks and bonds at the start of one's career. Many of the Marwell and Ames experiments involved a single decision, e.g. administered by a questionnaire that was mailed to high school students. Subsequent experiments by economists (that did *not* involve economics doctoral students) showed that repetition typically produced a declining pattern of contributions. Some people may stop contributing if others are observed to be free riding. One motive for making contributions may be to prevent others from behaving in this manner, in the hope that contributions will be reciprocated in subsequent rounds. This reciprocity motive is obviously weaker in the final rounds. Although contributions tend to decline in the final periods, some people do contribute even in the final period. We begin with a summary of an experiment with a one-shot setup. Multi-round experiments will be considered in later sections.

II. Single-Round Experiments

Ensminger (2001) reports the results of a public goods experiment involving young men of the Orma, a society from Kenya that is described in Chapter 23. Participants were divided into groups of 4, and each person was given 50 shillings that could be kept or invested in a "group project." Investments were made by placing tokens in envelopes, which were shuffled to preserve anonymity. Then the contents were emptied and publicly counted by a member of the group before being doubled by the experimenter and divided equally among the four participants. For example, if one person contributed two shillings, these would be doubled and the resulting four shillings would be distributed, one per

person. Thus a contribution of two shillings yields a private return of only 1, so the MPCR is only 0.5.

The data for this single-round game are shown in Figure 26.1. The modal contribution is four-tenths of the endowment, with a fair amount of variation. In fact, a quarter of the 24 participants contributed the entire 50 shillings. The average contribution is about 60%, which is at the high end of the 40-60% range observed in the first round of most public goods experiments done in the U.S. (Ledyard, 1995). Ensminger conjectures that contributions were enhanced by familiarity with the “Harambee” institution used to arrange for funding of public projects like schools. This practice involves the specification of income-based suggested voluntary contributions for each household, with social pressures for compliance. In fact, some of the participants commented on the similarity of the Harambee and the experimental setup, and a major Harambee solicitation was in progress at the time of the experiment. The absence of complete free riding by anyone is notable in Figure 26.1.




Figure 26.1. Contributions to a “Group Project” for Young Orma Males, with Group Size = 4 and an MPCR = 0.5. Source: Ensminger (2001)

MPCR Effects

The next topic is the extent to which contributions are affected by treatment variables like MPCR and group size. Goeree, Holt, and Laury (2002) report an experiment in which the participants had to make decisions for ten different treatments in which an endowment of 25 tokens was allocated to public and private uses. Subjects were paid \$6 and were told that only one of these treatments would be selected *ex post* to determine additional earnings. This is a

single-round experiment in the sense that there is no feedback obtained between decisions, and only one will count. A token kept was worth 5 cents in all ten cases. In one treatment, a token contributed would return 2 cents to each of the 4 participants. This yields an MPCR of 0.4 since each nickel foregone from the private return yields 2 cents from the public return, and $2/5 = 0.4$. In another treatment, a token contributed would return 4 cents to each of the 4 participants, so these treatments hold the group size constant and increase the MPCR from 0.4 to 0.8. This doubling of the MPCR essentially doubled average contributions, from 4.9 to 10.6 tokens. Of the 32 participants, 25 increased their contributions, 3 decreased, and 4 showed no change. The null hypothesis, that increases are just as likely as decreases, can be rejected using any standard statistical test. Although these are the only two treatments that allow a direct test of the MPCR effect, this effect has been observed in many other experiments, most of which have involved more than a single round (see the next section).

Internal and External Return Effects

In each of the treatments just discussed, the benefit to oneself for contributing is exactly equal to the benefit that every other person receives. In many public goods settings, however, the person making a voluntary contribution may enjoy a greater personal benefit than others. For example, benefactors typically give money for projects that they particularly value for some reason. It is quite common for gifts to medical research teams to be related to illnesses that are present in the donor's family. Even in Adam Smith's streetlight example, a person who erects a light over the street would pass by that spot more often, and hence would receive a greater benefit than any other randomly selected person in the town.

This difference between donor benefits and other public benefits can be examined by introducing the distinction between the "internal" to the person making the contribution and the "external" return that is enjoyed by each of the other people in the group. Recall that one of the Goeree, Holt, and Laury (2002) treatments involved a return of 2 cents for oneself and for each other person when a token worth 5 cents was contributed. Thus the internal return to oneself is $2/5 = 0.4$, and the external return to others is also $2/5 = 0.4$. There was another treatment in which the return to oneself was raised from 2 cents to 4 cents, whereas the return to each of the three others was held constant at 2 cents. Thus the internal return was increased from 0.4 to 0.8 ($=4/5$), with the external return constant at 0.4. This increase in the internal return essentially lowered the cost of contributing regardless of the motive for contributing (generosity, confusion, etc.) This doubling of the internal return essentially doubled the average observed contribution from 4.9 tokens to 10.7 tokens. Again, the effect was highly significant, with 25 people increasing their contributions, 3 decreasing, and 4

showing no change as the internal return increased holding the external return constant.

Next consider the effects of changing the external return, holding the internal return constant at 2 cents per token contributed (for an internal return of $R_I = 0.4$). The external return was raised from 2 cents to 6 cents, which raised the external return from $R_E = 2/5 = 0.4$ to $R_E = 6/5 = 1.2$. This tripling of the external return approximately doubled the average observed contribution, from 4.9 tokens to 10.5 tokens. (Contributions increased for 23 subjects, decreased for 2, and showed no change for the other 7). The ten treatments provide a number of other opportunities to observe external-return effects. Table 26.1 shows the average number of tokens contributed with group size fixed at 4, where the external return varies from 0.4 to 0.8 to 1.2. The top row is for a low internal return, and the bottom row is for a high internal return. With one exception, increases in the external return (moving from left to right) result in increases in the average contribution. A similar pattern (again with one exception) is observed in Table 26.2, which shows comparable averages for the treatments in which the group size was 2 instead of 4.

Table 26.1. Average Number of Tokens Contributed with Group Size = 4
(Source: Goeree, Holt, and Laury, 2002)

External Return	Low External $R_E = 0.04$	Medium External $R_E = 0.08$	High External $R_E = 1.2$	Very High External $R_E = 2.4$
Low Internal Return: $R_I = 0.4$	4.9	-	10.5	-
High Internal Return: $R_I = 0.8$	10.7	10.6	14.3	-

Table 26.2. Average Number of Tokens Contributed with Group Size = 2
(Source: Goeree, Holt, and Laury, 2002)

External Return	Low External $R_E = 0.04$	Medium External $R_E = 0.08$	High External $R_E = 1.2$	Very High External $R_E = 2.4$
Low Internal Return: $R_I = 0.4$	-	-	7.7	-
High Internal Return: $R_I = 0.8$	6.7	12.4	11.7	14.5

Finally, consider the vertical comparisons between two numbers in the same column of the same table, i.e. where the internal return is increased and the external return is held constant. The average contribution increases in all three

cases. Overall, the lowest contributions are for the case of low internal and external returns (upper-left corner of Table 26.1), and the highest contributions are for the case of a high internal return and a very high external return (bottom-right corner of Table 26.2).

Group-Size Effects

A comparison of the two tables affords some perspective on group-size effects. There are four cases where group size is changed, holding the returns constant, and the increase in group size (from Table 26.2 to Table 26.1) raises average contributions in all cases but one.

Economic Altruism

These comparisons (and some supporting statistical analysis) lead the authors to conclude that contributions respond positively to increases in internal return, external return, and group size, with the internal-return effect being strongest for the experiment being reported here.

Despite the individual differences and the presence of some unexplained variations in individual decisions, the overall data patterns in these single-round experiments are roughly consistent with a model in which people tend to help others by contributing some of their tokens. Such contributions are more common when 1) the private cost of contributing is reduced, 2) the benefit to each other person is increased, and 3) the number of others who benefit is increased. Even a perfectly selfish free-rider may decide to contribute if it is thought that this might induce others to contribute in the future, but there is no opportunity for reciprocity in these one-round experiments. These results suggest that many individuals are not perfectly selfish free-riders. Moreover, any altruistic tendencies are not exclusively a “warm-glow” feeling from the mere act of contributing, but rather, altruism is in part *economic* in the sense that it depends on the costs to oneself and the benefits to others.

III. Multi-Round Experiments

Economists began running multi-round public goods experiments to determine whether free-riding behavior would emerge as subjects gained experience and began to understand the incentives more clearly. The typical setup involves 10 rounds with a fixed group of individuals. Figure 26.2 shows the fractions of endowment contributed for a classroom experiment conducted with the Veconlab software. There is considerable variability across individuals. Player 1 (solid thin line) has relatively low contributions, with a spike upward in round 5. Player 2 contributes the full amount in all periods, and player 3 contributes nothing in any period. Average contributions for all 5 individuals (the thick solid line) start at about 0.5 and decline slowly, with an upward surge in

later periods caused largely by player 4's change from 0 to 25. Average contributions fall in the final periods, but not to zero. The overall pattern shown here is somewhat typical of other 10-round public goods experiments; contributions tend to start in the 0.4-0.6 range and decline over time, with erratic movements as individuals become frustrated and try to signal (with high contributions) or punish (with low contributions). Contributions generally fall in the final period, when an attempt to signal others to contribute in the future would be useless, but some people continue to contribute even in the final period.

As is the case with one-round public goods experiments, there is evidence that contributions in multi-round experiments respond to treatment variables that alter the benefits and costs of contributions. In a classic study, Isaac and Walker(1988b) used two different MPCR levels, 0.3 and 0.75, with the order being alternated in every other session. Six sessions were conducted with groups of size 4, and six sessions were conducted with groups of size 10. Group composition remained fixed for all




Figure 26.2 Data from a Five -Person Classroom Public Goods Game with an Internal Return of 0.6 and an External Return of 0.8 (UVA, Econ 482, Fall 2001)

10 rounds. The average fractions of the endowment contributed are shown in bold-faced font in the lower-left part of Table 26.3. The increase in MPCR

approximately doubles contributions for both group sizes. The increase in group size raises contributions for the low MPCR, but not for the high MPCR.

These patterns are reinforced by the Isaac, Walker, and Williams (1994) results in the far-right column of Table 26.3. This column shows average contributions for three groups of size 40, with an MPCR that goes from very low (0.03) to low (0.3) to high (0.75). Notice that both treatment variables tend to have no effect when contributions go above about 40 per cent of the endowment, as can be seen from the entries in the lower-right corner of the table.

Table 26.3. Average Fraction of Endowment Contributed for Ten Rounds
(Source: **Isaac and Walker, 1988b**; Isaac, Walker, and Williams, 1994)

	Group Size		
	$N = 4$	$N = 10$	$N = 40$
Very Low MPCR (0.03)	-	-	0.18
Low MPCR (0.3)	0.18	0.26	0.44
High MPCR (0.75)	0.43	0.44	0.39

The effects of varying the internal and external returns in multi-round public goods experiments are reported in Goeree, Holt, and Laury (2002). Subjects were paired in groups of size two, with new matchings in each round. Nobody was paired with the same person twice. With the internal return fixed at $R_I = 0.8$, the average number of tokens contributed (out of an endowment of 25) increased from 5 to 7.8 to 10 to 11.2 as the external return was increased from 0.4 to 0.8 to 1.2 to 2.4. Holding the external return fixed at $R_E = 1.2$, a reduction in the internal return from 0.8 to 0.4 caused average contributions to fall from 10 to 4.4. A (random-effects) regression based on individual data was used to evaluate the significance of these and other effects. The estimated regression (with standard errors shown in parentheses) is:

$$\text{Contribution} = -1.37 - 0.37 \cdot \text{Round} + 9.7 \cdot R_I + 3.2 \cdot R_E + 0.16 \cdot \text{Other}_{t-1}$$

(3.1) (0.1) (3.6) (1.2) (0.02)

The final “ Other_{t-1} ” variable is the contribution made by the partner in the previous period. Notice that contributions tend to decline over time, and that the internal-return effect is stronger than the external return effect. The positive effect of the previous partner’s contribution suggests that attitudes toward others may be influenced by their behavior. This is consistent with the results reported by van Dijk, Sonnemans, and van Winden (1997), who measured attitudes toward

others before and after a multi-round public goods experiment. Changes in attitudes were apparent and depended on the results of the public goods game in an intuitive manner.

IV. Further Reading

There is considerable interest in the study of factors that influence voluntary contributions. An analysis of these factors based only on economic theory may be insufficient, since many of the relevant issues are behavioral and must be addressed with experiments. For example, List and Lucking-Reiley (2001) report a controlled field experiment in which different mail solicitations for donations mentioned different amounts of “seed money” gifts, i.e. pre-existing gifts that were announced in the solicitation mailing. The presence of significant seed money had a very large effect on contribution levels. A related issue is the extent to which the public provision of a public good will “crowd out” private voluntary donations (Andreoni, 1993).

There have been many public goods experiments that focus on the effects of factors like culture, gender, and age on contributions in public goods games, with somewhat mixed results. For example, there is no clear effect of gender (see the survey in Ledyard, 1995). Goeree, Holt, and Laury (2002) found no gender differences on average, although men in their sample were more likely to make extremely low or high contributions than women.

A second strand of the literature consists of papers that examine changes in procedures, e.g. whether groups remain fixed or are randomly reconfigured each round (e.g., Croson, 1996), and whether or not the experimenter or the participants can observe individual contribution levels. Letting people talk about contributions between rounds tends to raise the level of cooperation, an effect that persists to some extent even if communication is subsequently stopped (Isaac and Walker, 1988a). Contributions are also increased in the presence of opportunities for individuals to levy punishments and rewards on those who did or did not contribute (Andreoni, Harbaugh, and Vesterlund, 2001).

A third strand of the literature pertains to variations in the payoff structure, e.g. having the public benefits be nonlinear functions of the total contributions by group members. In the linear public goods games described thusfar, it is optimal for a selfish person to contribute nothing when the internal return is less than one. In this case, the Nash equilibrium involves zero contributions by all, which is analogous to defection in a prisoner’s dilemma. Using nonlinear payoff functions makes it possible to design experiments in which the Nash equilibrium for selfish players is to contribute some, but not all, tokens to the public good. See Laury and Holt (2000) for a survey of experiments with interior Nash equilibria. Alternatively, the payoff structure can be altered by requiring that total contributions exceed a specified threshold before any benefits are derived

(Bagnoli and McKee, 1991, and Croson and Marks, 2000). Such “provision points” introduce a coordination problem, since it is not optimal to contribute unless one expects that other contributions will be sufficient to reach the required threshold. (Provision-point public goods games can be implemented with the Veconlab game PPG.)

Despite the large number of public goods experiments, there is still a lively debate about the primary motives for contribution. One approach is to model individuals’ utility functions as depending on both the individual’s own payoff and on payoffs received by others. Altruism, for example, can be introduced by having utility be an increasing function of the sum of others’ payoffs. Anderson, Goeree, and Holt (1998), for example estimate that individuals in the Isaac and Walker (1988b) experiments are willing to give up at most about ten cents to give others a dollar, and similar estimates were obtained by Goeree, Holt, and Laury (2002). On the other hand, some people may not like to see others’ earnings go above their own, which suggests that relative earnings may matter (Bolton and Ockenfels, 1998; Fehr and Schmidt, 1999). In multi-round experiments, individuals’ attitudes toward others may change in response to others’ behavior, and people may be willing to contribute as long as enough others do so. Many of these alternative theories are surveyed in Holt and Laury (1998).

All of the discussion thusfar has pertained to voluntary contributions, but there are many mechanisms that have been designed to induce improved levels of provision. A critical problem is that people may not know how others value a public good, and people have an incentive to understate their values if tax shares or required efforts depend on reported values. In theory, there are some mechanisms that do provide people with an incentive to report values truthfully, and these have been tested in the laboratory. For example, see Chen and Plott (1996) or the survey in Ledyard (1995). In practice, most public goods decisions are either directly or indirectly made on the basis of political considerations, e.g. lobbying and “log rolling” or vote trading. In particular, provision decisions may be sensitive to the preferences of the “median voter” with preferences that split others more or less equally. Moreover, the identity of the median voter may change as people move to locations that offer the mix of public goods that suits their own personal needs. See Hewett et al. (2002) for a classroom experiment in which the types and levels of public goods are determined by voting and by relocation or “voting with feet.” Anderson and Holt (2002b) survey experiments that focus on such public choice issues.

Questions

1. What was the approximate MPCR for the large game hunters in the !Kung example described in section I?

2. Suppose that Ensminger's experiment (discussed in section II) had used groups of size 8 instead of size 4, with all contributions being doubled and divided equally. What would the resulting MPCR have been?
3. How would it have been possible for Ensminger to double the group size and hold the MPCR constant? How could she have increased the MPCR from 0.5 to 0.75, keeping the group size fixed at 4?
4. Ensminger wrote a small code number on the inside of each person's envelope, so that she could record which people made each contribution. The others could not see the codes, so contributions were anonymous. What if she had wanted contributions to be "double anonymous" in the sense that nobody, not even the experimenter, could observe who made which contribution? Can you think of a feasible way to implement this treatment and still be able to pay people based on the outcome of the game?
5. Suppose that Ensminger had calculated payoffs differently, by paying each of the four people in a group an amount that was one-third of the doubled total contributions of the *other three* participants. If contributions were 1, 2, 3, and 4 shillings for people with IDs 1, 2, 3, and 4, calculate each person's return from the group project.
6. What would the internal and external returns be for the example in question 5?
7. The internal-return effects in Tables 26.1 and 26.2 are indicated by the vertical comparisons in the same table. The numbers effects are indicated by the comparisons from one table to the matched cell in the other. How many numbers-effect comparisons are there, and how many are in the predicted direction (more contributions with higher group size)?
8. The external-return comparisons in Tables 26.1 and 26.2 are found by looking at averages in the same row, with higher average contributions anticipated as one moves to the right. How many external-return comparisons are there in the two tables combined, and how many are in the predicted direction? (Hint: Do not restrict consideration to averages in adjacent columns.)
9. Which two entries in Table 26.1 illustrate the MPCR effect?
10. Calculate the MPCR for each of the two treatments described in the Instructions Appendix to Chapter 26.

Chapter 27. The Volunteer's Dilemma

In some situations, it only takes a single volunteer to provide a public benefit. A dilemma arises if the per-capita value of this benefit exceeds the private cost of volunteering, but each person would prefer that someone else incur this cost. For example, each major country on the UN Security Council may prefer that a proposal by a small country is vetoed, but each would prefer to have another country incur the political cost of a veto that is unpopular with many small member nations. This dilemma raises interesting questions, such as whether volunteering is more or less likely in cases with large numbers of potential volunteers. The volunteer's dilemma game provides data that can be compared with both intuition and theoretical predictions. This setup is implemented as the webgame VG, or it can be run by hand with playing cards and the instructions provided in the appendix.

I. Sometimes It Only Takes One Hero

After seeing the actress Teresa Saldana in the film *Raging Bull* in 1982, a crazed fan from Scotland came to Los Angeles and assaulted her as she was leaving home for an audition. Her screams for help attracted a group of people, but it was a man delivering bottled water who instinctively charged in and risked injury or worse by grabbing the assailant and holding him while the police and the ambulance arrived. Ms. Saldana survived her knife wounds; she has continued acting and has been active with victims' support groups. This is an example of a common situation where many people may want to see a situation corrected, but all that is needed is for one person to incur a cost to correct it. Another example is a case where several politicians may each prefer that someone else propose a legislative pay increase. But if nobody else is going to make such a proposal, then each would rather take the political heat and make it. These are all cases where it only takes one person's costly commitment to change an outcome for all of the others, e.g. with a veto or the provision of information that can be used by all. This type of situation is called a *volunteer's dilemma*, since people would prefer that others volunteer, but they prefer to volunteer if nobody else is going to do so.

II. Experimental Evidence

This dilemma was first studied by Diekmann (1985, 1986), where the focus was on the effects of increasing the numbers of potential volunteers. Let N be the number of potential volunteers. Each person must decide whether or not to incur a cost of C and volunteer. Each receives a high payoff value of V if at least one person volunteers, and a lower payoff L if nobody in the group volunteers.

Thus there are three possible payoff outcomes, either you volunteer and earn a certain return of $V - C$, or you do not volunteer and earn either V or L , depending on whether or not another person volunteers.

Consider a two-person setup, with $V - C > L$. This inequality ensures that it cannot be a Nash equilibrium for both people to refrain from volunteering, nor can it be an equilibrium for both to volunteer, since each can save C by free riding on the other's behavior. There are equilibria in which one person volunteers and the other does not, since each would prefer to volunteer if the other is expected to refrain (see question 6 for an example). Finally, there is also a symmetric equilibrium in which each person volunteers with a probability p^* , which is a topic that we will consider subsequently. Intuitively, one would expect that the equilibrium probability of volunteering would be lower when there is a larger number of potential volunteers, since there is less risk in relying on others' generosity in this case. This intuition is borne out in an experiment reported by Frazen (1995), with $V - L = 100$ and $C = 50$. The volunteer rates (middle column of Table 27.1) decline as group size is increased from 2 to 7, and these rates level off after that point. The third column indicates that the probability of obtaining no volunteer is essentially 0 for larger groups.

Table 27.1 Group Size Effects in Volunteer's Dilemma Experiment
(Source: Frazen, 1995)

Group Size	Individual Volunteer Rate	Rate of No-Volunteer Outcome
2	0.65	0.12
3	0.58	0.07
5	0.43	0.06
7	0.25	0.13
9	0.35	0.02
21	0.30	0.00
51	0.20	0.00
101	0.35	0.00

Next, consider the outcome for a classroom experiment conducted with payoffs: $V = \$25.00$, $L = \$0.00$, and $C = \$5.00$. The setup involved 12 participants who were randomly paired in groups of size 2 for each of 8 rounds, followed by 8 rounds of random matching in groups of size 4. (As with all classroom experiments reported in this book, participants were generally pairs of

students at a single computer, and one participant pair was selected *ex post* at random to be paid a small fraction of earnings.) The round-by-round averages for the 2-person treatment are generally in the 0.5 to 0.7 range, as shown on the left side of Figure 27.1. The increase in group size reduces the volunteer rate to about 0.4.




Figure 27.1. Volunteer's Dilemma With Group Size Change (UVA, Econ 482, Fall 2001)

III. The Mixed-Strategy Equilibrium

Most of the theoretical analysis of the volunteer's dilemma has been focused on the symmetric Nash equilibrium with random strategies. The dashed lines in Figure 27.1 indicate the Nash mixed-strategy equilibrium volunteer rates for the two treatments. In such an equilibrium, each person must be indifferent between volunteering and refraining, since otherwise it would be rational to choose the preferred action. In order to characterize this indifference, we need to calculate the expected payoffs for each decision. First, the expected payoff from volunteering is $V - C$ since this decision rules out the possibility of an L outcome. A decision not to volunteer has an expected payoff that depends on the others' volunteer rate, p . For simplicity, consider the case of only one other person (group size 2). In this case, the payoff for not volunteering is either V (with probability p) or L (with probability $1 - p$), so the expected payoff is:

$$(27.1) \quad \text{Expected Payoff (not volunteer)} = pV + (1-p)L \quad (\text{for } N = 2).$$

To characterize indifference, we must equate this expected payoff with the payoff from volunteering, $V - C$, to obtain an equation in p :

$$(27.2) \quad V - C = pV + (1-p)L \quad (\text{for } N = 2),$$

which can be solved for the equilibrium level of p :

$$(27.3) \quad p = 1 - \left(\frac{C}{V - L} \right) \quad (\text{equilibrium volunteer rate for } N = 2).$$

Recall that $V = \$25.00$, $L = \$0.00$, and $C = \$5.00$ for the parameters used in the first treatment in Figure 27.1. Thus, the Nash prediction for this treatment is $4/5$, as indicated by the horizontal dashed line with a height of 0.8. The actual volunteer rates in this experiment are not this extreme, but rather are closer to 0.6.

The formulas derived above must be modified for a larger group sizes, e.g. $N = 4$ for the second treatment. The payoff for a volunteer decision is independent of what others do, so this payoff is $V - C$ regardless of group size. If one does not volunteer, however, the payoff depends on what the $N - 1$ others do. The probability that one of them does not volunteer is $1 - p$, so the probability that none of the $N - 1$ others volunteer is $(1 - p)^{N-1}$. (This calculation is analogous to saying that if the probability of Tails is $1/2$, then the probability of observing two Tails outcomes is $(1/2)^2 = 1/4$, and the probability of having $N-1$ flips all turn out Tails is $(1/2)^{N-1}$.) To summarize, when one does not volunteer, the low payoff of L is obtained when none of the others volunteer, which occurs with probability $(1 - p)^{N-1}$. It follows that the high payoff V is obtained with a probability that is calculated: $1 - (1 - p)^{N-1}$. These observations permit us to express the expected payoff for not volunteering as the sum of terms on the right side of the following equation:

$$(27.4) \quad V - C = V[1 - (1 - p)^{N-1}] + L(1 - p)^{N-1}.$$

The left side is the payoff for volunteering, so (27.4) characterizes the indifference between expected payoffs for the two decisions that must hold if a person is willing to randomize. It is straightforward (question 7) to solve this equation for the probability, $1 - p$, of not volunteering:

$$(27.5) \quad 1-p = \left(\frac{C}{V-L} \right)^{\frac{1}{N-1}}$$

so the equilibrium volunteer rate is,

$$(27.6) \quad p = 1 - \left(\frac{C}{V-L} \right)^{\frac{1}{N-1}} \quad (\text{equilibrium volunteer rate}).$$

Notice that (27.6) reduces to (27.3) when $N = 2$. It is easily verified that the volunteer rate is increasing in the value V obtained when there is a volunteer and decreasing in the value L obtained when there is no volunteer. As would be expected, the volunteer rate is also decreasing in the volunteer cost C . (These claims are fairly obvious from the positions of C , V , and L in the ratio on the right side of (27.6), but they can be verified by changing the payoff parameters in the spreadsheet program that is created to answer question 8.)

Next, consider the effects of a change in the number of potential volunteers, N . For the setup in Figure 27.1, recall that $V = \$25.00$, $L = \$0.00$, $C = \$5.00$, and therefore, the volunteer rate is 0.8 when $N = 2$. When N is increased to 4, the formula in (27.6) yields $1 - (1/5)^{1/3}$, which is about 0.41 (question 8), as shown by the dashed horizontal line on the right side of Figure 27.1. This increase in N reduces the predicted volunteer rate, which is consistent with the qualitative pattern in the data (although the $N=2$ prediction is clearly too high relative to the actual data). It is straightforward to use algebra to show that the formula for the volunteer rate in (27.6) is a decreasing function of N . This formula can also be used to calculate the Nash volunteer rates for the eight different group-size treatments shown in Table 27.1 (question 8).

Finally, we will evaluate the chances that none of the N participants volunteer. The probability that any one person does not volunteer, $1 - p$, is given in (27.5). Since the N decisions are independent, like flips of a coin, the probability that all N participants decide not to volunteer is obtained by raising the right side of (27.5) to the power N :

$$(27.7) \quad \text{probability that nobody volunteers} = (1-p)^N = \left(\frac{C}{V-L} \right)^{\frac{N}{N-1}}$$

The right side of (27.7) is an increasing function of N (see question 9 for a guide to making spreadsheet calculations). As N goes to infinity, the exponent of the expression on the far right side of (27.7) goes to 1, and the probability of getting

no volunteer goes to $C/(V-L)$. The resulting prediction (1/5) is inconsistent with the data from Table 27.1, where the probability of finding no volunteer is zero for larger group sizes. What seems to be happening is that there is some randomness, which causes occasional volunteers to be observed in any large group.

One interesting feature of the equilibrium volunteer rate in (27.6) is that the payoffs only matter to the extent that they affect the ratio, $C/(V-L)$. Thus, a treatment change that shifts both V and L down by the same amount will not affect the difference in the denominator of this ratio. This invariance is the basis for a treatment change designed by a student who ran an experiment in class on this topic using the Veconlab software. The student decided to lower these parameters by \$15, from $V = \$25$ and $L = \$0$ to: $V = \$10$ and $L = -\$15$. The values of C and N were kept fixed at \$5 and 4. Table 27.2 summarizes these payoffs for alternative setups. All payoffs are positive in the “sure gain treatment,” but a large loss is possible in the “loss treatment.” This pair of treatments was designed with the goal of detecting whether the possibility of large losses would result in a higher volunteer rate in the second treatment. Such an effect could be due to “loss aversion,” which was discussed in chapter 8.

Table 27.2 Volunteer’s Dilemma Payoffs for a Gain/Loss Design

Own Decision	Number of Other Volunteers	Own Payoff	Sure Gain Treatment	Loss Treatment
Volunteer	any number	$V - C$	\$20.00	\$5.00
Not Volunteer	None	N	\$0.00	-\$15.00
Not Volunteer	At least one	V	\$25.00	\$10.00

The results for a single classroom experiment with these treatments are summarized in Figure 27.2. The observed volunteer rates closely approximate the common Nash prediction for both treatments, so loss aversion seems to have had no effect. This example was included to illustrate an invariance feature of the Nash prediction and to show how students can use theory to come up with clever experiment designs. The actual results are not definitive, however, because of the constraints imposed by the classroom setup: no replication, very limited incentives, and no credible way of making a participant pay actual losses.

Further Reading

Other aspects of the volunteer’s dilemma are discussed in Diekmann and Mitter (1986) and Diekmann (1993), who reports experiments for an asymmetric setup.




Figure 27. 2. Volunteer’s Dilemma with Gain/Loss Change (UVA, Econ 482, Fall 2001)

Goeree and Holt (1999c) discuss the volunteer’s dilemma and other closely related games where individuals are faced with a binary decision, e.g. to enter a market or not or to vote or not. As they note, the volunteer’s dilemma is a special case of a “threshold public good” in which each person has a binary decision of whether to contribute or not. The public benefit is not available unless the total number of contributors reaches a specified threshold, say M , and the volunteer’s dilemma is a special case where $M = 1$. Goeree and Holt analyze the equilibria for these types of binary-choice games when individual decisions are determined by probabilistic choice rules of the type discussed in chapters 11-13. Basically, the effect of introducing “noise” into behavior is to pull volunteer rates towards 0.5. This approach can explain why the probability of a no-volunteer outcome is *not* observed to be an increasing function of group size (as predicted in the Nash equilibrium). The intuition for why large numbers may decrease the chances of a no-volunteer outcome is that “noise” makes it more likely that at least one person in a large group will volunteer due to random causes. This intuition is consistent with the data from Frazen (1995) and is inconsistent with

the Nash prediction that the probability of a no-volunteer outcome is an increasing function of group size.

Questions

1. Consider a volunteer's dilemma in which the payoff is 25 if at least one person volunteers, minus the cost (if any) of volunteering, the payoff for no volunteer is 0, and the cost of volunteering is 1. Consider a symmetric situation in which the group size is 2 and each decides to volunteer with a probability p ? Find the equilibrium probability.
2. How does the answer to question 1 change if the group size is 3?
3. The Appendix to this chapter provides instructions for a volunteer's dilemma with group sizes of 2 and 4. Calculate the Nash equilibrium volunteer rate in each case.
4. For the setup in question 3, calculate the probability of obtaining no volunteer in a Nash equilibrium, and show that this probability is 16 times as large for group size 4 as it is for group size 2.
5. In deciding whether or not to volunteer, which decision involves more risk? Do you think a group of risk-averse people would volunteer more often in a Nash (mixed-strategy) equilibrium for this game than a group of risk-neutral people? Explain your intuition.
6. The setup in question 1 can be written as a two-by-two matrix game between a row and column player. Write the payoff table for this game, with the row payoff listed first in each cell. Explain why this volunteer's dilemma is not a prisoner's dilemma? Are there any Nash equilibria in pure (non-randomized) strategies for this game?
7. Derive equation (27.5) from (27.4), by showing all intermediate steps.
8. To calculate group-size effects for volunteer's dilemma experiments, set up a spreadsheet program. The first step is to make column headings for the variables that you will be using. Put the column headings, N , V , L , C , $C/(V-L)$, and P in cells A1, B1, C1, D1, E1, and F1 respectively. Then enter the values for these variables (here begin with the setup from the second treatment in Figure 27.1): 4 in cell A2, 25 in cell B2, 0 in cell C2, and 5 in cell D2. Then enter a formula: $=D2/(B2-C2)$ in cell E2. Finally enter the formula for the equilibrium probability from equation (27.6) into cell F2; the Excel code for this formula is: $=1-\text{power}(E2,1/(A2-1))$, which raises the ratio in cell E2 to the power $1/(N-1)$, where N is taken from cell A2. If you have done this correctly, you should obtain the equilibrium volunteer rate mentioned in the chapter, which rounds off to 0.41. Experiment with some changes in V , L , and C to determine the effects of changes in these parameters on the volunteer rate. In order to obtain predictions from the Frazen model, enter the values 2, 100, 0, 50 in

cells A3, B3, C3, and D3 respectively, and copy the formulas in cells E2 and F2 into cells E3 and F3. The resulting volunteer rate prediction in cell F3 should be 0.5. Then enter the other values for N vertically in the A column: these values are 3, 5, 7, 9, 21, 51, and 101. The predictions for these treatments can be obtained by copying the values and formulas for cells B3-F3 downward in the spreadsheet.

9. To obtain predictions for the probability of getting no volunteers at all, add a column heading in cell G1 (“No Vol.”) and enter the formula from equation (27.7) in excel code into cell G2: `=power(E2, A2/(A2-1))` and copy this formula down to the lower cells in the G column. What does this prediction converge to as the number of participants becomes large? Hint: as N goes to infinity, the exponent in equation (27.7) converges to 1.
10. Explain the volunteer’s dilemma to your roommates or friends and come up with at least one example of such a dilemma based on some common experience from class, dorm living, a recent film, etc.

Chapter 28. Externalities, Congestion, and Common Pool Resources

Many persistent urban and environmental problems result from excessive use or “exploitation” of a shared resource. For example, an increase in fishing activity may reduce the catch per hour for all fishermen, or a commuter’s decision to enter a tunnel may slow down the progress of other commuters. Each individual may tend to ignore the negative impact of their activity on other’s harvests or travel times. Indeed, people may not even be aware of the negative effects of their decisions if the effects are small and dispersed across many others. This negative externality is typically not priced in a market, and overuse can result. This chapter considers two somewhat stylized paradigms of overuse or congestion.

In the Common Pool Resource game, individual efforts to secure more benefits from the resource have the effect of reducing the benefits received by others. In technical terms, the average and marginal products of each person’s effort is decreasing in the total effort of all participants. The Veconlab game, CP, may be used to compare behavior in alternate treatments.

In Market Entry games, participants decide independently whether or not to enter a market or activity for which the earnings per entrant are a decreasing function of the number of entrants. The outside option earnings from not entering are fixed. There are typically many Nash equilibria, including one involving randomized strategies. Kahneman once remarked on the tendency for payoffs to be equalized for the two decisions: “To a psychologist, it looks like magic.” This impression should change for those who have participated in a classroom entry game experiment, which should be run prior to the class discussion. Although entry rates are typically distributed around the equilibrium prediction, the rates are variable and too high since each person ignores the negative effects of their decisions on other entrants. The Veconlab ME program provides for some policy options (“traffic report information” and tolls) that correct these problems.

I. Water Wars

Some of the most difficult resource management problems in both developing and advanced economies involve water. Consider a setup where there are upstream and downstream users, and in the absence of norms or rules, the downstream users are left with whatever is not taken upstream. This is an example of a “common pool resource,” where each person use may reduce the benefits obtained by others. Ostrom and Gardner (1993) describe such a situation

involving farmers in Nepal. If upstream users tend to take most of the water, then their usage may be inefficiently high from a social point of view, e.g. if the downstream land is more fertile bottom land. Here the inefficiency is due to the fact that the upstream farmers draw water down until the value of its marginal product is very low, even though more water downstream may be used more productively.

For example, the Thambesi system in Nepal is one where the headenders have established clear water rights over those downstream. The farmers located at the head of each rotation unit take all of the water they need before those lower in the system. In particular, the headenders grow water-intensive rice during the pre-monsoon season, and those lower in the system cannot grow irrigated crops at this time. If all farmers were to grow a less-water intensive crop (wheat), the area under cultivation during the pre-monsoon season could be expanded dramatically, nearly ten-fold (Yoder, 1986, p. 315). The irrigation decisions made by the headenders raise their incomes, but at a large cost in terms of lost crop value downstream. A commonly used solution to this problem is to limit usage, and such limits may be enforced by social norms or by explicit penalties. In either case, enforcement may be problematical. In areas where farmers own marketable shares of the water system, these farmers have an incentive to sell them so that water is diverted to its highest-value uses (Yoder, 1986).

Besides overuse, there is another potential source of inefficiency if activities to increase water flow have joint benefits to both types of users. For example, irrigation canals may have to be maintained or renewed annually, and this work can be performed shared by all users, regardless of location. In this case, the benefit of additional work is shared by all users, so each person who contributes work will only obtain a part of the benefits, even though they incur the full cost of their own effort, unless their contribution is somehow reimbursed or otherwise rewarded. Note that the level of work on the joint production of water flow should be increased as long as the additional benefit, in terms of harvest value, is greater than the cost, in terms of lost earnings on other uses of the farmers' time. Thus each person has an incentive to free ride on others' efforts, and this perverse incentive may be stronger for downstream users if their share of the remaining water tends to be low.

The possibility of underprovision can be illustrated with a numerical example. Suppose for simplicity that there is one upstream user and four downstream users, and that each has up to 4 units (e.g. days) of labor to allocate to the establishment of the shared irrigation system. The value of the total water produced is a function of the total effort is given in the second row of Table 28.1.

Table 28.1 Payoffs from Total Effort

<i>Total Effort</i>	1	2	3	4	5	6	7	8
Total Payoff	6	42	72	96	114	126*	132	132
Upstream	2	14	24	32	38	42	44	44
Downstream	1	7	12	16	19	21	22	22
Downstream	1	7	12	16	19	21	22	22
Downstream	1	7	12	16	19	21	22	22
Downstream	1	7	12	16	19	21	22	22

In the absence of any restriction, the upstream user is able to capture 1/3 of the value of the water, shown in the third row, and each of the four downstream users are able to capture only 1/6. If the opportunity cost of effort is 7, then the social optimum level of effort would be 6 (as indicated by the asterisk in the top row), since an additional (7th) unit of effort only increases the value by 6 (from 126 to 132), which is less than the cost of 7. If the upstream user were forced to act alone, that person would be willing to provide all 4 available units, but the total value of water produced, 96, would be far below the social optimum of 126. Yet none of the downstream users would have any incentive to join in, since with 4 units already provided, their share (16 each) would only increase to 19 if any one of them contributed a unit of effort, a gain of 3 that is less than the opportunity cost of 5. This is a Nash equilibrium in which the downstream users free ride, and the common resource is underprovided. Indeed there are no Nash equilibria where the downstream users contribute anything in this example (see problem 1).

To summarize, there are two main sources of inefficiency associated with common pool resources, i.e. overuse and underprovision. Overuse can occur if people do not consider the negative effects that their own use decisions have on the benefit of the resource that remains for others. Under-provision can occur if the benefits of providing the resource are shared by all users, but each person incurs the full cost of their own contribution to the group effort.

II. Traffic

Many of the top political issues in large urban areas involve commuting and traffic control. Large investments to provide fast freeways, tunnels, and bridges often are negated with a seemingly endless supply of eager commuters during certain peak periods. The result may be travel via the new freeways and tunnels is no faster than the older alternative routes via a maze of surface roads. To illustrate this problem, consider a stylized setup in which there are N commuters who must choose between a slow, reliable surface route and a potentially faster route (freeway or tunnel). However, if large numbers decide to enter the faster route, the resulting congestion will cause traffic to slow to a crawl, so entry can be a risky decision. This setup was implemented in a laboratory

experiment run by Anderson, Holt, and Reiley (2004). In each session, there were 12 participants who had to choose whether to enter a potentially congested activity or exit to a safe alternative. The payoffs for entrants were decreasing in the total number of entrants in a given round, as shown in Table 28.2.

Table 28.2 Individual Payoffs from Entry

<i>Entrants</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>10</i>	<i>11</i>	<i>12</i>
Individual Payoff	\$4.00	\$3.50	\$3.00	\$2.50	\$2.00	\$1.50	\$1.00	\$0.50	\$0.00	\$0.50	\$1.00	\$1.50

You should think of the numbers in the table as the net monetary benefits from entry, which go down by \$0.50 as each additional entrant increases the time cost associated with travel via the congested route. In contrast, the travel time associated with the alternative route (the non-entry decision) is less variable. For simplicity, the new benefit for each non-entrant is fixed at \$0.50, regardless of the number who select this route. Obviously, entry is the better decision as long as the number of entrants is not higher than 8, and the equilibrium number of entrants is, therefore, 8. In this case, each person earns \$0.50, and the total earnings (net benefit) for the 12 people together is \$6.00. In contrast, if only 4 enter, they each earn \$2.50 and the other 8 earn \$0.50, for total earnings of \$14. It follows that an entry restriction will more than double the social net benefit.

It is straightforward to calculate the social benefit associated with other entry levels (question 2), and the results are shown in Table 28.3.

Table 28.3 Social Benefit for all 12 Participants Together

<i>Entrants</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>10</i>	<i>11</i>	<i>12</i>
Total Payoff	\$9.50	\$12	\$13.50	\$14	\$13.50	\$12	\$9.50	\$6	\$1.50	\$4.00	\$10.50	\$18

If uncontrolled entry decisions lead to approximately 8 entrants, the number that equalizes earnings from the two alternative decisions, and the total benefit to all participants is inefficiently low. The reason for this is that each additional person who decides to enter lowers the net entry benefit by \$0.50 for themselves *and for all other entrants*, and it is this “external” effect that may be ignored by each entrant. For example if there are currently 4 entrants, the addition of a fifth would reduce their earnings by \$2 (4x\$0.50) and only increase the entrant’s earnings by \$1.50 (from the non-entry payoff of \$0.50 to the entry payoff of \$2), so total earnings decrease by \$0.50. A sixth entrant would reduce total earnings by even more, \$1.50 (a reduction of 5x\$0.50 that is only partly offset by the \$1 increase in

the entrant's earnings), since there are more people to suffer the consequences. As even more people enter, the earnings reduction caused by an additional entrant becomes higher, and therefore, the total payoff falls at an increasing rate. This is the intuition behind the very high costs associated with dramatic traffic delays during peak commuting periods. Any policy that pushes some of the traffic to non-peak periods will be welfare improving, since the cost of increased congestion in those periods is less than the benefit from reducing peak congestion.

In laboratory experiments, as on the open road, people ignore the negative external effects that their earnings have on others' earnings. The data from one such experiment, shown in Figure 28.1, illustrate this feature. The average entry rate in the first 10 rounds is 0.69, which is quite close to the two-thirds (8 out of 12) rate that would occur in equilibrium. There is, however, considerable variability, which in real traffic situations would be amplified by accidents and other unforeseen traffic events.




Figure 28. 1. Entry Game With Imposition of Entry Fee
Source: Anderson, Holt, and Reiley (2004)

After 10 rounds, all entrants were required to pay an entry fee of \$2, which reduces all numbers in the bottom row of Table 28.2 by \$2. In this case, the equilibrium number of entrants is 4, i.e. an entry rate of 0.33, as shown by the

dashed line on the right side of Figure 28.1. The average entry rate for these final 10 rounds was quite close to this level, at 0.38.

Even though the fee reduces entry to (nearly) socially optimal levels, the individuals do not enjoy any benefit. First for simplicity, let's ignore the variability and assume that the fee reduces entry to the 0.33 rate. In equilibrium entrants earn \$2.50 minus the \$2 fee, and non-entrants earn \$0.50, so everyone earns exactly the same amount, \$0.50, as if they did not enter at all. Where did the increase in social benefit go? The answer is in the collected fees of \$8 ($4 \times \2). This total fee revenue of \$8 is available for other uses, e.g. tax reductions and road construction. So if the commuters in this model are going to benefit from the reduction in congestion, then they must get a share of the benefits associated with the fee collections.

A second series of sessions was run with a redistribution of fee revenues, with each person receiving $1/12$ of collections. In particular, after 10 rounds with no fee, participants were allowed to vote on the level of the fee for rounds 11-20, and then for each 10-round interval after that. The voting was carried out by majority rule after a group discussion, with ties being decided by the chair, who was selected at random from the 12 participants. In one session, the participants first voted on a fee of \$1, which increased their total earnings, and hence was selected again after round 20 to hold in rounds 21-30. After round 30, they achieved a further reduction in entry and increase in total earnings by voting to raise the fee to \$2 for the final 10 rounds. In a second session with voting, the group reached the optimal fee of \$2 in round 20, but the chair of the meeting kept arguing for a lower fee, and these arguments resulted in reductions to \$1.75, and then to \$1.50, before it was raised to a near-optimal level of \$1.75 for the final 10 rounds. Interestingly, the discussions in these meetings did not focus on maximizing total earnings. Instead, many of the arguments were based on making the situation better for entrants or for non-entrants (different people had different points of view, based on what they tended to do). In each of these two sessions, the imposition in the entry fees resulted in an approximate doubling of earnings by the final 10 rounds. Even though the average payoffs for entry (paying the fee) and exit are each equal to \$0.50 in equilibrium, for any fee, the fee of \$2 maximized the amount collected, and hence maximized the total earnings when these include a share of the collections.

Regardless of average entry rate, there is considerable variation from round to round, as shown in Figure 28.1 and in the data graphs for the other sessions where fees were imposed, either by the experimenter or as the result of a vote. This variability is detrimental at any entry rate above one-third. For simplicity, suppose that there is no fee, so that entry would tend to bounce up and down around the equilibrium level of 8. It can be seen from Table 28.3 that an increase in entry from 8 to 9 reduces total earnings (including the shared fee) from

\$6 to \$1.50, whereas a reduction in entry from 8 to 7 only raises earnings from \$6 to \$9.50. In general, a bounce up in entry reduces earnings by more than a bounce down by an equal amount, so symmetric variation around the equilibrium level of 8 would reduce total earnings below the amount, \$6, that could be expected in equilibrium with no variation.

Everyone who commutes by car knows that travel times can be variable, especially on routes that can become very congested. In this case, anything that can reduce variability will tend to be an improvement in terms of average benefits. In addition, most people dislike variability, and the resulting risk aversion, which has been ignored up to now, would make a reduction in variability more desirable. Sometimes variability can be reduced by better information. In many urban areas, there is a news-oriented radio station that issues traffic reports every 10 minutes, e.g. “on the eights” at WTOP on Washington, DC. This information can help commuters avoid entry into already overcrowded roads and freeways. A stylized type of traffic information can be implemented in an experiment by letting each person find out the number who have already entered at any time up to the time that they decide to enter. In this case, the order of entry is determined by the participants themselves, with some people making a quick decision, and others waiting. The effect of this kind of information is to bring about a dramatic (but not total) reduction in variability, as can be seen for the experiment shown in Figure 28.2.

In a richer model, some commuters would have had higher time values than others, and hence would have placed a higher value on the reduction in commuting time associated with reduced congestion. The fee would allow commuters to sort themselves out in the value-of-time dimension, with high-value commuters being willing to pay the fee. The resulting social benefit would be further enhanced by the effects of using price (the entry fee) to allocate the good (entry) to those who value it the most. This observation is similar to a result reported in Plott (1983), where the resale of licenses to trade, from those with low values to those with high values, raised total earnings and market efficiency.

III. Common Pool Resource Experiments

Each person in the market entry game has two decisions, enter or not. In contrast, most common pool resource situations are characterized by a range of possible decisions, corresponding to the intensity of use. For example, a lobster fisherman not only decides whether to fish on any given day, but also (if permitted by regulations and weather) how many days a year to go out and how many hours to stay out each day. Let the fishing effort for each person be denoted by x_i , where the i subscript indicates the particular person. The total effort of all N fishermen is denoted by X , which is just the sum of the individual x_i for $i = 1, \dots, N$. For simplicity, assume that no person is more skilled than any other, so it is

natural to assume that a person's fraction of the total harvest is equal to their fraction of the total effort, x_i/X .




Figure 28. 2. Entry Game with Information About the Number of Prior Entrants
Source: Anderson, Holt, and Reiley (2004)

In the experiment to be discussed, each person selected an effort level simultaneously, and the total harvest revenue, Y , was a quadratic function of the total effort:

$$(28.1) \quad Y = AX - BX^2,$$

where $A > 0$, $B > 0$, and $X < A/B$ (to ensure that the harvest is not negative). This function has the property that the average revenue product, Y/X , is a decreasing function of the total effort. This average revenue product is calculated as the total revenue from the harvest, divided by the total effort: $(AX - BX^2)/X = A - BX$, which is decreasing in X since B is assumed to be positive. The fact that average revenue is decreasing in X means that an increase in one person's effort will reduce the average revenue product of effort for everyone else. This is the negative externality for this common pool resource problem.

Effort is costly, and this cost is assumed to be proportional to effort, so the cost for individual i would be Cx_i , where $C > 0$ represents the opportunity cost of

the fisherman's time. Assuming that the harvest in equation (28.1) is measured in dollar value, then the earnings for person i would be their share, x_i/X , of the total harvest in equation (28.1) minus the cost of Cx_i that is associated with their effort choice:

$$(28.2) \quad \text{earnings} = (AX - BX^2)(x_i / X) - Cx_i = (A - BX)x_i - Cx_i,$$

where right side of equation (28.2) is the average product of effort times the individual's effort, minus the cost of the person's effort decision. Note that this earnings expression on the right side of equation (28.2) is analogous to that of a firm in a Cournot market, facing a demand (average revenue) curve of $A - BX$ and a constant cost of C per unit.

A rational, but selfish, person in this context would take into account the fact that an increase in their own effort reduces their own average revenue product, but they would not consider the negative effect of an increase in effort on the others' average revenue products. This is the intuitive reason that unregulated exploitation of the resource would result in too much production relative to what is best for the group in the aggregate (just as uncoordinated production choices by firms in a Cournot market results in a total output that is too high relative to the level that would maximize joint profits).

In particular, the Nash/Cournot equilibrium for this setup would involve each person choosing the activity level, x_i , that maximizes the earnings in equation (28.2) taking the others' activities as given. Suppose that there are a total of N participants, and that the $N - 1$ others choose a common level of y each. Thus the total of the others' decisions is $(N - 1)y$, and the total usage is: $X = x_i + (N - 1)y$. With this substitution, the earnings in equation (28.2) can be expressed:

$$(28.3) \quad \text{earnings} = Ax_i - B(N - 1)yx_i - Bx_i^2 - Cx_i,$$

which can be maximized by equating marginal revenue with marginal cost C . The marginal revenue is just the derivative of total revenue (as explained in Chapter 13 and its appendix), and the marginal cost is the derivative of total cost, Cx_i . Therefore, the earnings in equation (28.3) can be maximized by setting the derivative of this expression with respect to x_i equal to zero:

$$(28.4) \quad A - B(N - 1)y - 2Bx_i - C = 0.$$

As discussed in Chapter 13, this equation can be used to determine a common decision, x , by replacing both y and x_i with x and solving to obtain:

$$(28.5) \quad x = \frac{A - C}{(N + 1)B}.$$

The nature of this common pool resource problem can be illustrated by in terms of the specific setup used in a research experiment run with the Veconlab software by Guzik (2004), who at the time was an undergraduate at Middlebury College. He ran experiments with groups of 5 participants with $A = 25$, $B = 1/2$, and $C = 1$. The cost was implemented by giving each person 12 tokens in a round. These could either be kept, with earnings of 1 each, or used to extract the common pool resource. Thus the opportunity cost, C , is 1 for each unit increase in x_i . With these parameters and $N = 5$, equation (28.5) yields a common equilibrium decision of $x = 8$. It can be shown that the socially optimal level of appropriation is 4.8 (question 6).

The default setting for the Veconlab software presents the setup with a mild environmental interpretation, i.e. in terms of a harvest from a fishery. Adjustments to the software and the use of supplemental instructions permitted Guzik to present the setup using either environmental terminology, workplace terminology or a neutral (“magic marbles”) terminology. In each of these three treatments (environmental, workplace, and marbles) there were 8 sessions, each with 5 participants. The each session was run for 10 rounds with fixed matchings. The purpose of the experiment was to see if the mode of presentation or “frame” would affect the degree to which people overused the common resource. There were no clear framing effects: the average decision was essentially the same for each treatment: 7.81 (environmental), 8.04 (marbles), and 7.8 (workplace). These small differences were not statistically significant, although there was a significant increase in variance for the non-neutral frames (workplace and environmental). Although the author was able to detect some small framing effects after controlling for demographic variables, the main implication for experimental work is that if you are going to look at economic issues, then the frame may not matter much in these experiments, although one should still hold the frame constant and change the economic variables of interest.

The results reported by Guzik (2004) are fairly typical of those for common pool resource experiments; the average decisions are usually close to the Nash prediction, although there may be considerable variation across people and from one round to the next. However, if participants are allowed to communicate freely prior to making their decisions, then the decisions are reduced toward socially optimal levels (Ostrom and Walker, 1991).

IV. Extensions

Common pool resource problems with fisheries are discussed in Hardin (1968). For a wide-ranging coverage of issues related to the tragedy of the commons, see Ostrom, E., R. Gardner, and J. K. Walker (1994) and the special section of the Fall 1993 *Journal of Economic Perspectives* devoted to “Management of the Local Commons.” The first common pool resource experiments are reported in Garnder, Ostrom, and Walker (1990), who evaluate the effects of subject experience, and by Walker, Gardner, and Ostrom (1990), who consider the effects of changing the endowment of tokens.

For results of a common pool resource experiment done in the field, see Cardenas, Stranlund, and Willis (2000), who conducted their experiment in rural villages in Columbia. The main finding is that the application of rules and regulations that are outside of informal community institutions tends to increase the effect of individualistic motives and Nash-like results. They conclude “... modestly enforced government controls of local environmental quality and natural resource use may perform poorly, especially as compared to informal local management” (p. 1721). Other recent work is summarized in Guzik (2004).

Questions

1. Consider the numerical example in section I, and suppose that one of the downstream users provides an effort of 1. Explain what the best response of the upstream user would be. If that best response were carried out, would the downstream user still be willing to supply a unit of effort?
2. Calculate the social benefits associated with 4, 6, 8, 10, and 12 entrants for the setup in Tables 28.2 and 28.3 (show your work).
3. Use the numbers in Table 28.3 to explain why unpredictability in the number of entrants is bad, e.g. why a “50/50” of either 6 entrants or 10 entrants is not as good as the expected value (8 entrants for sure) in terms of expected earnings.
4. Explain in your own words why each increase in the number of entrants above 4 results in ever greater decreases in the social benefit in Table 28.3.
5. What would the equilibrium entry rate be for the game described in section II if the entry fee were raised to \$3 per entrant? How would the total amount collected in fees change by an increase from \$2 to \$3? How would the total amount in fees collected change as the result of a fee decrease from \$2 to \$1?
6. (requires calculus) Derive a formula (in terms of A , B , C , and N) for the socially optimal level of appropriation of the common pool resource model in section III. What is the socially optimal level of appropriation per person for the parameters: $A = 25$, $B = 1/2$, $C = 1$, and $N = 5$?

Chapter 29. Rent Seeking

Administrators and government officials often find themselves in a position of having to distribute a limited number of prized items (locations, licenses, etc.). Contenders for these prizes may engage in lobbying or other costly activities that increase their chances for success. No single person would spend more on lobbying than the prize is worth, but with a large number of contenders, expenditures on lobbying activities may be considerable. This raises the disturbing possibility that the total cost of lobbying by all contenders may “dissipate” a substantial fraction of the prize value. In economics jargon, the prize is an “economic rent,” and the lobbying activity is referred to as *rent seeking*. Rent seeking is thought to be more prevalent in developing countries, where some estimates are that such non-market competitions consume a significant fraction of national income (Krueger, 1974). One only has to serve as a Department Chair in a U.S. university, however, to experience the frustrations of rent seeking. The game considered in this chapter is one in which the probability of obtaining the prize is equal to one’s share of the total lobbying expenditures of all contenders. The game illustrates the potential costs of administrative (non-market) allocation procedures. It can be implemented with playing cards, as indicated in the appendix, or it can be run on the Veconlab software (game RS).

I. Government with “a Smokestack on Its Back”

One of the most spectacular successes of recent government policy has been the auctioning off of bandwidth used to feed the exploding growth in the use of cell phones and pagers. The U.S. Federal Communications Commission (FCC) has allocated major licenses with a series of auctions that raised billions of dollars without adverse consequences. Some European countries followed with similarly successful auctions, which also raised many more billions than expected in some cases. The use of auctions collects large numbers of potential competitors, and the commodities can be allocated quickly to those with the highest valuations. In the U.S., this bandwidth was originally reserved for the armed forces, but it was underused at the end of the Cold War, and the transfer created a large increase in economic wealth, without significant administrative costs. It could have easily been otherwise. Radio and television broadcasting licenses were traditionally allocated via administrative proceedings, which are sometimes called “beauty contests.” The successful contender would have to convince the regulatory authority that service would be of high quality and that social values would be protected. This often required establishing a technical expertise and an effective lobbying presence. The lure of extremely high

potential profits was strong enough to induce large expenditures by aspiring providers. Those who had acquired licenses in this manner were strongly opposed to market-based allocations that forced the recipients to pay for the licenses. Similar economic pressures may explain why beauty contest allocations continue to be prevalent in many other countries.

The first crack in the door appeared in the late 1980's, when the FCC decided to skip the administrative proceedings in the allocation of hundreds of regional cell phone licenses. The forces opposed to the pricing of licenses did manage to block an auction, and the licenses were allocated by lottery instead. There were about 320,000 applications for 643 licenses. Each application involved significant paper work (legal and accounting services), and firms specializing in providing completed applications began offering this service for about \$600 per application. The resources used to provide this service have opportunity costs, and Hazlett and Michaels (1993) estimated the total cost of all submitted applications to be about \$400,000. The lottery winners often sold their licenses to more efficient providers, and resales were used to estimate that the total market value of the licenses at that time was about a billion dollars. Each individual lottery winner earned very large profits on the difference between the license value and the application cost, but the costs incurred by others were lost, and the total cost of the transfer of this property was estimated to be about *forty percent* of the market value of the licences. There are, of course, the indirect costs of subsequent transfers of licenses to more efficient providers, a process of consolidation that may take many years. In the meantime, inefficient provision of the cellular services may have created ripples of inefficiency in the economy. It is episodes like this that once caused Milton and Rose Friedman (1989), in a discussion of the unintended side effects of government policies, to remark: "Every government measure bears, as it were, a smokestack on its back."

In a classic paper, Gordon Tullock (1967) pointed out that the real costs associated with competitions for government granted "rents" may destroy or "dissipate" much of the value of those rents in the aggregate, even though the winners in such contests may earn large profits. This destruction of value is often invisible to those responsible, since the contestants participate willingly, and the administrators may enjoy the process of being lobbied. Moreover, some of the costs of activities like waiting in line and personal lobbying are not directly priced in the market. These costs can be quite apparent in a laboratory experiment of the type to be described next.

II. Rent Seeking in the Classroom Laboratory

In many administrative (non-market) allocation processes, the probability of obtaining a prize or monopoly rent is an increasing function of the amount spent in the competition. In the FCC auctions, the chances of winning a license

were approximately equal to the applicant's efforts as a proportion of the total efforts of the other contestants. This provides a rationale for a standard mathematical model of "rentseeking" with N contestants. The effort for person i is denoted by x_i for $i = 1, \dots, N$. The total cost of each effort is a cost c times the person's own effort, i.e. cx_i . In the simplest symmetric model, the value of the rent, V , is the same for all, and each person's probability of winning the prize is equal to their effort as a fraction of the total effort of all contestants:

$$(29.1) \quad \text{expected payoff} = \frac{x_i}{\sum_{j=1, \dots, N} x_j} V - cx_i.$$

Here the expected payoff is the probability of success, which is the fraction of total effort, times the prize value V , minus the cost of effort. This latter cost is not multiplied by any probability since it must be paid whether or not the prize is obtained.

Goeree and Holt (1999) used the payoff function in (1) in a classroom experiment with $V = \$16,000$, $c = \$3,000$, and $N = 4$. Each of the four competitors consisted of a team of 2-3 students. Efforts were required to be integer amounts. This requirement was enforced by giving 13 playing cards of the same suit to each team. Rent-seeking effort was determined by the number of cards that the person placed in an envelope, as described in the instructions to this chapter. Each team incurred a cost of $\$3,000$ for each card played, regardless of whether or not they obtained the prize. The cards played were collected, shuffled, and one was drawn to determine which contender would win the $\$16,000$ prize. The number of cards played varied from team to team, but on average there were slightly more than three cards played by each team. Thus a typical team incurred $3 \times \$3,000$ in expenses, and the total lobbying cost for all four teams was over $\$36,000$, all for a prize worth only $\$16,000$.

Similar results were obtained in a separate classroom experiment using the game RS from the Veconlab setup. The parameters were the same (four competitors, a $\$16,000$ value, and a $\$3,000$ cost per unit of effort), and the 12 teams were randomly put into groups of four competitors in a series of rounds. The average number of lobbying effort units was 3 in the first round, which was reduced to a little over 2 in rounds 4 and 5. Even with two units expended by each team, the total cost would be 2 (number of effort units) times 4 (teams) times the $\$3,000$ cost of effort, for a total cost of $\$24,000$. This again resulted in over-dissipation of the rent, which was only $\$16,000$.

III. The Nash Equilibrium

A Nash equilibrium for this experiment is a lobbying effort for each competitor such that nobody would want to alter their expenditure given that of the other competitors. First, consider the case of four contenders, a \$16,000 value, and a \$3,000 effort cost. Note that efforts of 0 cannot constitute an equilibrium, since each person could deviate to an effort of 1 and obtain the \$16,000 prize for sure at a cost of only \$3,000. Next suppose that each person is planning to choose an effort of 2 at a cost of \$6,000. Since the total effort is 8, the chances of winning are $2/8 = 1/4$ so the expected payoff is $16,000/4 - 6,000$, which is *minus* \$2,000. This cannot be a Nash equilibrium, since each person would have an incentive to deviate to 0 and earn nothing instead of losing \$2,000. Next, consider the case where each person's strategy is to exert an effort of 1, which produces an expected payoff of $16,000/4 - 3,000 = 1,000$. To verify that this is an equilibrium, we have to check to be sure that deviations are not profitable. A reduction to an effort of 0 with a payoff of 0 is obviously bad. A unilateral increase to an effort of 2, when the others maintain efforts of 1, will result in a $2/5$ chance of winning, for an expected payoff of $16,000(2/5) - 6,000 = 400$, which is also worse than the payoff of 1,000 obtained with an effort of 1.

In the symmetric equilibrium, the effort of 1 for each of the four contestants results in a total cost of $4(3,000) = 12,000$, which dissipates three-fourths of the 16,000 rent. This raises the issue of whether a reduction in the cost of rent seeking efforts might reduce the extent of wasted resources. In the context of the FCC lotteries, for example, this could correspond to a requirement of less paper work for each separate application. Suppose that the resource cost of each application is reduced from \$3,000 to \$1,000. This cost reduction causes the Nash equilibrium level to rise from 1 to 3 (see question 1).

IV. A Mathematical Derivation of the Equilibrium (Optional)

The Nash equilibrium calculations done thusfar were based on considering a particular level of rent-seeking activity, common to each person, and showing that a deviation by one person alone would not increase that person's expected payoff. This is a straightforward, but tedious, approach, and it has the additional disadvantage of not explaining how the candidate for a Nash equilibrium was found in the first place. A calculus derivation is relatively simple, and is offered here as an option for those who are familiar with basic rules for taking derivatives (some of these rules are introduced in the appendix to chapter 13). First, consider the expected payoff function in (29.1), modified to let the decisions of the $N-1$ others be equal to a common level, x^* .

$$(29.2) \quad \text{expected payoff} = \frac{x_i}{x_i + (N-1)x^*} V - cx_i.$$

This will be a “hill-shaped” (concave) function of the person’s own rent-seeking activity, x_i , and the function is maximized “at the top of the hill” where the slope of the function is zero. To find this point, the first step is to take the derivative and set it equal to zero. The resulting equation will determine player i ’s best response when the others are choosing x^* . In a symmetric equilibrium with equal rent-seeking activities, it must be the case that $x_i = x^*$, which will yield an equation that determines the equilibrium level of x^* . This analysis will consist of two steps: finding the derivative of player i ’s expected payoff and then using the symmetry condition.

Step 1. In order to use the rule for taking the derivative of a power function, it is convenient to express (29.2) as a power function:

$$(29.3) \quad \text{expected payoff} = x_i (x_i + (N-1)x^*)^{-1} V - cx_i.$$

The final term on the right side of (29.3) is linear in x_i , so its derivative is $-c$. The first term on the right side of (29.3) is the product of x_i and a power function that contains x_i , so the derivative is found with the product rule (derivative of the first function times the second, plus the first function times the derivative of the second), which yields the first two terms on the left side of (29.4):

$$(29.4) \quad (x_i + (N-1)x^*)^{-1} V - x_i (x_i + (N-1)x^*)^{-2} V - c = 0.$$

Step 2. Next we use the fact that $x_i = x^*$ in a symmetric equilibrium. Making this substitution into (29.4) yields a single equation in the common equilibrium level, x^* :

$$(29.5) \quad (x^* + (N-1)x^*)^{-1} V - x^* (x^* + (N-1)x^*)^{-2} V - c = 0,$$

which can be simplified to:

$$(29.6) \quad \frac{1}{Nx^*} V - \frac{1}{N^2 x^{*2}} V - c = 0.$$

Finally, we can multiply both sides of (29.6) by $N^2 x^*$, which yields an equation that is linear in x^* that reduces to:

$$(29.7) \quad x^* = \frac{(N-1)}{N^2} \frac{V}{c}.$$

It is straightforward to verify that $x^* = 1$ when $V = \$16,000$, $c = \$3,000$, and $N = 4$. When the cost is reduced to $\$1,000$, the equilibrium level of rent-seeking activity is raised to 3. Notice that the predicted effect of a cost reduction is that the total amount spent on rent-seeking activity is unchanged.

Finally, consider the total cost of all rent-seeking activity, which will be measured by the product of the number of contenders, the effort per contender, and the cost per unit of effort: Nx^*c . It follows from (29.7) that this total cost is: $(N-1)/N$ times the prize value, V . Thus the fraction of the rent that is dissipated is a fraction, $(N-1)/N$, which is an increasing function of the number of contenders. With two contenders, half of the value is dissipated in a Nash equilibrium, and this fraction increases towards 1 as the number of contenders gets large. In general, *rent dissipation* is measured as the ratio of total expenditures on rent-seeking activity by all contenders to the value of the prize.

Notice that the fraction of rent dissipation is predicted to be independent of the cost of lobbying effort, c . For example, recall the previous example with four contenders, where the reduction in c from $\$3,000$ to $\$1,000$ raised the Nash equilibrium lobbying effort from 1 to 3, thereby maintaining a constant total level of expenditures.

These predictions were tested with a *Veconlab* experiment run on a single class with a “2x2” design with high and low lobbying costs, and with high and low numbers of competitors. There were 20 rounds (5 for each treatment). The prize value was $\$16,000$ in all rounds, and each person began with an initial cash balance of $\$100,000$. One person was selected at random *ex post* to receive a small percentage of earnings. The treatments involved a per-unit cost of either $\$500$ or $\$1,000$, and either 2 or 4 contenders. The four treatment combinations are shown in the two columns on the left side of Table 29.1.

Table 29.1. A Classroom Rent-seeking Experiment with a Prize Value of $\$16,000$

Number of Contenders: N	Per-unit Cost: C	Nash Predictions:		Observed Averages:	
		x^*	total cost (Ncx^*)	x	total cost (Ncx)
4	$\$1,000$	3	$\$12,000$	6	$\$24,000$
4	$\$500$	6	$\$12,000$	8	$\$16,000$
2	$\$1,000$	4	$\$8,000$	6	$\$12,000$

2	\$500	8	\$8,000	10	\$10,000
---	-------	---	---------	----	----------

The Nash prediction for a contender's lobbying effort, as calculated from (29.7), is shown in the third column. Multiplying this by the number of contenders and then by the cost per unit of effort yields the predicted total cost of rent seeking activities, which is shown in the fourth column. These total costs are \$12,000 with $N = 4$ (top two rows) and \$8,000 for $N = 2$ (bottom two rows). Thus the Nash prediction is that an increase in the number of contenders makes the situation more competitive and increases rent-seeking activity and total costs. For each fixed level of N , however, a reduction in the per-unit lobbying cost will double the effort, leaving the predicted total cost unchanged.

The Nash predictions can be compared with data from the classroom experiment, shown in the final two columns. The average effort decisions have been rounded off to the nearest integer, to make the table easier to read. Several conclusions are apparent:

- 1) The total costs of rent seeking activity are significant, and are greater than 50% of the prize value in all treatments.
- 2) The total costs of rent-seeking activity are greater than the Nash predictions.
- 3) An increase in the number of contenders tends to increase the total costs of rent seeking.
- 4) A decrease in per-unit effort costs does raise efforts, but not enough to offset the cost reduction.

V. Extensions and Further Reading

Rent seeking models have been used in the studying of political lobbying (Hillman and Samet, 1987), which is a natural application because lobbying expenditures often involve real resources. In fact, any type of contest for a single prize may have similar strategic elements, e.g. political campaigns or research and development contests. There is a large and growing literature that uses the rent-seeking paradigm to study non-market allocation activities.

There have been a number of research experiments using the payoff structure in equation (21.1). Millner and Pratt (1989) report that rent dissipation was significantly higher than the Nash prediction. In a companion paper, Millner and Pratt (1991) used a lottery choice decision to separate people into groups according to their risk aversion. Groups with more risk-averse individuals tended to expend more on rent-seeking activity.

Anderson, Goeree, and Holt (1998) provide a theoretical analysis of a model of rent seeking where the prize goes to the person with the *highest* effort. This is like a race or auction, where it only takes a small advantage to obtain a

sure win (unlike the model contained in this chapter where a higher effort increases the probability of obtaining the prize but does not ensure a win). The Anderson, Goeree and Holt model is based on logit stochastic response functions that were introduced in chapter 11. This incorporation of randomness in the game is used to derive a theoretical prediction that rent-seeking efforts will be above the levels predicted in a Nash equilibrium.

Questions

1. Suppose that the cost per unit of effort is reduced from \$3,000 to \$1,000, with four competitors and a prize of \$16,000. Show that a common effort of 3 is a Nash equilibrium for the rent seeking game with payoffs in equation (21.1). Hint: Check to be sure that a unilateral deviation to an effort of 2 or 4 would not be profitable, given that the other three people keep choosing 3. Of course, to be complete, all possible deviations would have to be considered, but you can get the idea by considering deviations to 2 and 4. An alternative is to use calculus.
2. Now suppose that the effort cost is reduced to \$500. Show that a common effort of 6 is a Nash equilibrium (consider unilateral deviations to 4 and 7).
3. Calculate the total amount of money spent in rent-seeking activities for the setup in question 1. Did the reduction in the per-unit (marginal) cost of rent seeking from \$1,000 to \$500 reduce the predicted *total* cost of rent seeking activity.
4. Now suppose that the number of competitors for the setup in question 1 is reduced from 4 to 2. The per-unit effort cost is \$1000 and the prize value is \$16,000. Show that a common effort of 4 is a Nash equilibrium (for simplicity, only consider deviations to 3 and 5, or use calculus).
5. Did the reduction in competitors from 4 to 2 in question 4 reduce the predicted extent of rent dissipation? (Rent dissipation is the total cost of all rent-seeking activity as a proportion of the prize value.)

Part VIII. Information and Learning

Information specific to individuals is often unobserved by others. Such information may be conveyed at a cost, but misrepresentation and strategic non-revelation is sometimes a problem. Informational asymmetries yield rich economic models that may have multiple equilibria and unusual patterns of behavior. One example considered in an earlier part is the case where individuals may decide to rely on information inferred from others' decisions and "follow the crowd," which creates information cascades. A diversity of investor information is also important in finance, where an issue may be the extent to which asset prices aggregate diverse bits of information. Many other interesting issues may arise, since market processes may fail when information cannot be priced.

Chapter 30 provides a closer look at exactly how new information is used to revise probabilistic beliefs (Bayes' rule). Chapter 31 also pertains to a sequence of decisions, but in this case, they are made by different people. Those later in the sequence are able to learn from others' earlier decisions. This raises the possibility of a type of bandwagon effect, which is known as an "information cascade."

Statistical discrimination can arise if the employer uses past correlations between the fixed signal and worker productivity to make job assignment decisions. "Biased" expectations may be confirmed in equilibrium as workers of one type become discouraged and stop investing. The statistical discrimination model presented in Chapter 31 is implemented in the webgame SD by assigning a color (purple or green) to each person in the worker role and letting them make an unobserved investment decision. Those in the employer role must then make job assignments on the basis of an imperfect productivity "test." Experience-based discrimination is most likely to show up when the test is inconclusive. This exercise provides an excellent opportunity for informed and non-emotional class discussions of issues like race and gender discrimination. The equilibrium analysis in this chapter makes use of Bayes' rule.

Chapter 33 introduces the standard signaling game, where the proposer has a privately known attribute or "type." This type is not observed by the responder, who can only see a "signal" made by the proposer. For example, an employer may attempt to use some educational credential or detailed portfolio to infer something about a worker's productivity and energy. There are several different types of behavior that may emerge, depending on whether the signals effectively distinguish the different types of people. The webgame SG can be configured to generate some of these behavioral patterns, and the results are compared with Nash predictions. Signaling models have been widely applied to other situations in economics, politics, law, and finance.

Chapter 30. Bayes' Rule

Any careful study of human behavior has to deal with the issue of how people learn from new information. For example, an employer may have beliefs about a prospective employee's value to the firm based on prior experience with those of similar backgrounds. Then new information is provided by a job interview or a test, and the issue is how to combine initial ("prior") beliefs with new information to form final ("posterior") beliefs. This chapter uses a simple ball counting rule-of-thumb or "heuristic" to explain Bayes' rule, which is a mathematical formula for forming posterior beliefs. Actual behavior will never be perfectly described by a mathematical formula, and some typical patterns of deviation are discussed. The appendix contains instructions for a game in which the participant sees draws of colored balls from one of two cups and has to form an opinion about which of the two cups is being used. This experiment can also be run with the Veconlab Bayes' Rule game (BAY), which is much quicker than hand-run versions, and it provides automatic calculations and a graph of average elicited probabilities as a function of the Bayes' prediction, under alternative scenarios (all data, data for specific participants, data from one draw or signal, etc.).

I. Introduction

Suppose that you have just received a test result indicating that you have a rare disease. Unfortunately, the disease is life-threatening, but you have some hope because the test is capable of producing "false positives," and the disease is rare. Your doctor tells you that the test is fairly accurate, with a false positive rate of only 1 percent. The "base rate" or incidence of the disease for those in your socio-economic group is only one per 10,000. What are your chances of having the disease? (Please write down a guess, so that you do not forget.)

Confronted with this problem, most people conclude will that it is more likely than not that the person actually has the disease, but such a guess would be seriously incorrect. The 1 percent false positive rate means that testing 10,000 randomly selected people will generate about 100 positive results (1%), but on average only one person out of 10,000 actually has the disease. Thus the chances of having the disease are only about one in a hundred, *even after you have tested positive with a test that is correct 99 times out of 100*. This example illustrates the dramatic effect of prior information about the "base rate" of some attribute in the population. This example also indicates how one might set up a simple frequency-based counting rule that will provide approximately correct probability calculations:

- Select a hypothetical sample of people (say 10,000).
- Use the base rate to determine how many of those, on average, would have the disease by multiplying the base rate and the sample size (e.g. $0.0001 \times 10,000 = 1$).
- Next, next calculate the expected number of positive test results that would come from someone who has the disease. (In the example, one person has the disease and the test would pick this up, so we have 1 positive from an infected person.)
- Subtract the number of infected people from the sample size to determine the number that are not infected ($10,000 - 1 = 9,999$).
- Estimate the number of positive test results that would come from people who are not infected. (The false positive rate is 0.01 in the example, so the number of positives from people who are not infected is $9,999/100$, which is approximately 100.)
- Calculate the chances of having the disease given the positive test result by taking the ratio of the number of positives from infected people (1) to the numbers of positives from those who are infected (1) and those who are not (100), which is about $1/101$, or about 1 percent.

Calculations such as those given above are an example of the use of *Bayes' Rule*. This chapter introduces this rule, which is an optimal procedure for using prior information, like a population base rate, together with new information, like a test result. As the incorrect answers to the disease question may suggest, peoples' decisions and inferences in such situations may be seriously biased. A related issue is the extent to which people are able to correct for potential biases in market situations where the incentives are high and one is able to learn from past experience. And in some (but not all) situations, those who do not correct for biases lose business to those who do.

When acquiring new information, it is useful to think about the difference between the initial beliefs, the information obtained, and the new beliefs after seeing the information. If the initial (prior) beliefs are very strongly held, then the new (posterior) beliefs are not likely to change very much, unless the information is very good. For example, passing a lie detector test will not eliminate suspicion if the investigator is almost positive that the suspect is guilty. On the other hand, very good information may overwhelm prior beliefs, as would happen with the discovery of DNA evidence that clears a person who has already been convicted. Obviously, the final (posterior) beliefs will depend on the relative quality of the prior information and on the reliability of the test result. Bayes' rule provides a mathematical way of handling diverse sources of information. The Bayesian

perspective is useful because it dictates how to evaluate different sources of information, based on their reliability. This perspective may even be valuable when it is undesirable or illegal to use some prior information, e.g. in the use of demographic information about crime rates in a jury trial. In such cases, knowing how such prior information would be used (by a Bayesian) may make it easier to guard against a bias derived from such information. Finally, Bayesian calculations provide a benchmark or baseline from which biases can be measured.

II. A Simple Example and a Counting Heuristic (Anderson and Holt, 1997a)

The simplest informational problem is deciding which of two possible situations or “states of nature” is relevant, e.g. guilty or innocent, infected or not, defective or not, etc. To be specific, suppose that there are two cups or “urns” that contain equally sized Amber (*a*) and Blue (*b*) marbles, as shown in Table 30.1. Cup A has two Ambers and one Blue, and Cup B has one Amber and two Blues. We will use the flip of a fair coin to choose one of these cups. The cup selection is hidden, so your prior information is that each cup is equally likely. Then you are allowed to see several draws of marbles from the selected cup. Each time a draw is made and shown, it is then returned to the cup, so the draws are “with replacement” from a cup with contents that do not change.

Table 30.1. A Two Cup Example

Cup A	Cup B
<i>a, a, b</i>	<i>a, b, b</i>

Suppose the first draw is Amber and you are asked to report the probability that cup A is being used. An answer of $1/2$ is disappointingly common, sometimes even among graduate students, since it can be justified by the argument that each urn was equally likely to be selected. But if each cup was equally likely beforehand, what was learned from the draw? Another commonly reported probability for cup A following an Amber draw is $1/3$, since this cup has a one half chance of being used, and if used, there is a $2/3$ chance of drawing an Amber. Then we multiply $1/2$ and $2/3$ to get $1/3$. This $1/3$ probability is clearly wrong, since the cups were equally likely *ex ante*, and the Amber draw is more likely when cup A is used, so the probability of cup A should be greater than $1/2$ after seeing an Amber draw. Another problem with this answer is that analogous reasoning requires the probability of cup B to be $1/2$ times $1/3$, or $1/6$. This yields an inconsistency, since if the probability of cup A is $1/3$ and the probability of cup B is $1/6$, where does the rest of the probability go? These probabilities ($1/3$ and $1/6$) sum to $1/2$, so we should double them (to $2/3$ and $1/3$), which are the correct

probabilities for cups A and B after the draw of one Amber marble. This kind of scaling up of probabilities will be seen later as a part of the mathematical formula for Bayes' rule.

A close look at where the Amber marbles are in Table 30.1 makes it clear why the probability of cup A being used is $2/3$ after the draw of an Amber marble. There are two *a* marbles on the left and one on the right. Why does this suggest that the right answer is $2/3$? All six marbles are equally likely to be drawn before the die is thrown to select one of the cups. So, no Amber marble is more likely to be chosen than any other, and 2 of the 3 Amber marbles are in cup A. It follows that the posterior probability of cup A given an Amber draw is $2/3$.

The calculations in the previous paragraphs are a special case of Bayes' rule with equal prior probabilities for each cup. Now consider what can be done when the probabilities are not equal. In particular, suppose that the first marble drawn (Amber) is returned to the cup, and the decision maker is told that a second draw is to be made from the *same* cup originally selected by the throw of the die. Having already seen an Amber, the person's beliefs before the second draw are that the probability of cup A is $2/3$ and the probability of cup B is $1/3$, so the person thinks that it is twice as likely that cup A is being used after observing one *a* draw. The next step is to figure out how the previous paragraph's method of counting balls (when the two cups were initially equally likely) can be modified for the new situation where one cup is twice as likely as the other one. In order to create a situation that corresponds to the new beliefs, we want to somehow get twice as many possible draws as coming from the cup that is twice as likely. Even though the physical number of marbles in each cup has not changed, we can represent these beliefs by thinking of cup A as having twice as many marbles as cup B, *with each marble in either cup having the same chance of being drawn*. These posterior beliefs are represented in Table 30.2, where the proportions of Amber and Blue marbles are the same as they were in cups A and B respectively. Even though the physical number of marbles is unchanged at six, the prior corresponds to a case in which the imagined marbles in Table 30.2 are numbered from one to nine, with one of the nine marbles chosen randomly.

Table 30.2. A Mental Model after One Amber Draw
(real marbles in bold, imagined marbles not bold)

Cup A	Cup B
<i>a</i> , <i>a</i> , <i>b</i> <i>a</i> , <i>a</i> , <i>b</i>	<i>a</i> , <i>b</i> , <i>b</i>

When the posterior beliefs after an Amber draw are represented in Table 30.2, it is clear that a Blue on the second draw is equally likely to have come from either urn, since each cup contains two *b* marbles. Thus the posterior probability for cup A after a Blue on the second draw is 1/2. This is consistent with intuition based on symmetry, since the prior probabilities for each urn were initially 1/2, and the draws of an Amber (first) and a Blue (second) are balanced. A mixed sample in the opposite order (Blue first, then Amber) would, of course, have the same effect.

Table 30.3. A Mental Model of the Situation
After Two Amber Draws

Cup A		Cup B
<i>a, a, b</i>		<i>a, b, b</i>
<i>a, a, b</i>		
<i>a, a, b</i>		
<i>a, a, b</i>		

Suppose instead that the two draws were Amber, with the two cups being equally likely to be used *ex ante*. As before, the posterior belief after the first Amber draw can be represented by the cups in Table 30.2. Since four of the five (real or imagined) Amber marbles are in cup A, the posterior probability of cup A after seeing a second Amber draw is 4/5. After two Amber draws, cup A is, therefore, four times as likely as cup B, since 4/5 is four times as large as 1/5. To represent these posterior beliefs in terms of colored marbles that are equally likely to be drawn, we need to have four times as many rows on the Cup A side of Table 30.2 as there are on the Cup B side. Thus we would need to add two imagined more rows of three marbles under cup A in Table 30.2, holding the proportions of Amber and Blue marbles fixed. To test your understanding, you might consider the probability of cup A after seeing two Ambers and a Blue (exercise 1 at the end of the chapter). Up to this point, the analysis has been intuitive, but it is now time to be a little more analytical.

III. Relating the Counting Heuristic to Bayes' Rule

To make the connection with Bayes' rule, we will need a little notation. Suppose there are *N* marbles in each cup. The marbles will be Amber or Blue, and we will use the letter *C* to represent a specific color, so *C* can be either Amber or Blue. What we want to know is the probability of cup A given the draw of a marble of color *C*. When *C* is Amber and the contents are shown in Table 30.1, we already know the answer (2/3), but our goal here is to find a general

formula for the probability of cup A given a draw of a color C marble. This probability is denoted by $\Pr(A|C)$, which reads “the probability of A given C.” This formula should be general enough to allow for different proportions of colored marbles, and for differences in the prior probability of each cup.

Consider $\Pr(C|A)$, which reverses the order of the A and the C from the order used in the previous paragraph. Note that $\Pr(C|A)$ is read as “the probability of color C given cup A.” Thus $\Pr(C|A)$ denotes the fraction of marbles in cup A that are of color C, where C is either Amber or Blue. Similarly, $\Pr(C|B)$ is the fraction of marbles in cup B that are of color C. For example, if there are 10 marbles in cup A and if $\Pr(C|A) = 0.6$, then there must be six marbles of color C in the cup (calculated as 0.6 times 10). In general, there are a total of $\Pr(C|A)N$ marbles of color C in cup A, and there are $\Pr(C|B)N$ marbles of color C in cup B. If each cup is equally likely to be selected, then each of the $2N$ marbles in the two cups is equally likely to be drawn *ex ante* (before the cup is selected). Suppose the marble drawn is of color C. The posterior probability that a marble of color C was drawn from cup A, denoted $\Pr(A|C)$, is just the ratio of the number of color C marbles in cup A to the total number of marbles of this color in both cups:

$$(30.1) \quad \Pr(A|C) = \frac{\text{Number of Color C Marbles in Cup A}}{\text{Number of Color C Marbles in Both Cups}},$$

which can be expressed:

$$(30.2) \quad \Pr(A|C) = \frac{\Pr(C|A)N}{\Pr(C|A)N + \Pr(C|B)N}.$$

It is worth emphasizing that this formula is only valid for the case of equal prior probabilities and equal numbers of marbles in each urn. Nothing is changed if we divide both numerator and denominator of the right side on (30.2) by $2N$, which is the total number of balls in both cups, which yields a formula for calculating the posterior when the priors are $1/2$:

$$(30.3) \quad \Pr(A|C) = \frac{\Pr(C|A)(1/2)}{\Pr(C|A)(1/2) + \Pr(C|B)(1/2)} \quad (\text{for priors of } 1/2).$$

A person who has seen one or more draws may not have prior probabilities of $1/2$, so this formula must be generalized. This involves replacing the $(1/2)$ terms on the right side of the with the new prior probabilities, denoted $\Pr(A)$ and $\Pr(B)$. This is Bayes' rule:

$$(30.4) \quad \Pr(A|C) = \frac{\Pr(C|A) \Pr(A)}{\Pr(C|A) \Pr(A) + \Pr(C|B) \Pr(B)} \quad (\text{Bayes' rule}).$$

For the previous example with equal priors, $\Pr(A) = 1/2$, $\Pr(a|A) = 2/3$, and $\Pr(a|B) = 1/3$, so equation (30.4) implies that the posterior probability following an Amber draw is:

$$\Pr(A|a) = \frac{\frac{2}{3} \frac{1}{2}}{\frac{2}{3} \frac{1}{2} + \frac{1}{3} \frac{1}{2}} = \frac{\frac{1}{3}}{\frac{1}{2}} = \frac{2}{3}.$$

Similarly, the probability of cup B is calculated: $\Pr(B|a) = (1/6)/(2/6 + 1/6) = 1/3$. Notice that the denominators in both of the previous calculations are $1/2$, so dividing by $1/2$ scales up the numerator by a factor of 2, which makes the probabilities add up to one.

Having arrived at the Bayes' rule formula (for two events) as it appears in the textbooks, it is important to point out that the argument was intuitive but not rigorous. It helps, therefore, to relate the general formula for Bayes' rule back to the counting heuristic. Recall that the two cups were initially equally likely, and the first time we saw an a draw, the probability of cup A was $2/3$, so the ratio $\Pr(A)/\Pr(B)$ was 2. Then we took the N balls in cup A and imagined that there were twice as many, i.e. $2N$ balls in cup A (N real balls and N imagined balls). After seeing two a draws, the probability of Cup A was $4/5$, or four times as great as the $1/5$ probability of Cup B, so we imagined that there were $4N$ balls in Cup A instead of the original N balls. How can this approach be generalized for the case where $\Pr(A)/\Pr(B)$ is something other than 2 or 4? Obviously, we calculate the ratio $\Pr(A)/\Pr(B)$ and imagine that there are N times $\Pr(A)/\Pr(B)$ balls in Cup A, holding the original proportions constant. To relate the Bayes' rule formula in (30.4) to this counting heuristic, we need to figure out how to get terms involving N times $\Pr(A)/\Pr(B)$ terms into (30.4). This can be done by dividing both numerator and denominator of (30.4) by $\Pr(B)/N$, to obtain:

$$(30.5) \quad \Pr(A|C) = \frac{\Pr(C|A) \frac{\Pr(A)}{\Pr(B)} N}{\Pr(C|A) \frac{\Pr(A)}{\Pr(B)} N + \Pr(C|B) N}.$$

If there are N marbles in each cup, this equation shows that the basic Bayes' rule formula is equivalent to imagining that the N marbles in cup A are increased or decreased to a number M , which is N times the prior “odds ratio,” $\Pr(A)/\Pr(B)$.

Take the representation in Table 30.2 for example, where the prior for cup A is $2/3$ after seeing an a draw. Then $\Pr(A)$ is now $2/3$, $\Pr(B)$ is $1/3$, and the odds ratio is $(2/3)/(1/3) = 2$. With $N = 3$ marbles in each cup, the odds ratio times N is 2 times 3, or 6, which replaces the odds ratio times N in the numerator and the left side of the denominator of (30.5). This is as if we imagine that there are 6 marbles in cup A and only 3 in cup B, the 3 corresponding to the N term on the right side of the denominator in (30.5).

To summarize, if there is a prior probability of $1/2$ that each cup is used and if the cups contain equal numbers of colored marbles, then the posterior probabilities can be calculated as ratios of numbers of marbles of the color drawn, as in equation (30.1). If the marble drawn is of color C, then the posterior that the draw was from cup A is the number of color C marbles in cup A divided by the total number of color C marbles in both cups. When the prior probabilities or numbers of marbles in the cups are unequal, then the $1/2$ terms in (30.3) are replaced by the prior probabilities, as in Bayes' rule (30.4). Finally, the Bayes' rule formula is equivalent to imagining that the number N of marbles in cup A is changed by a factor $\Pr(A)/\Pr(B)$, holding the proportion of color C marbles fixed at $\Pr(C|A)$, with the counting heuristic that each of the actual and imagined marbles is equally likely to be chosen.

IV. Experimental Results

Nobody would expect that something so noisy as the formation of beliefs would adhere strictly to a mathematical formula, and experiments have been directed towards finding the nature of systematic biases. The disease example mentioned earlier suggests that, in some contexts, people may underweight prior information based on population base rates.

This *base rate bias* was the motivation behind some experiments reported by Kahneman and Tversky (1973), who gave subjects lists of brief descriptions of people who were either lawyers or engineers. The subjects were told that the descriptions had been selected at random from a sample that contained 70 percent lawyers and 30 percent engineers. Subjects were asked to report the chances out of 100 that the description pertained to a lawyer. A second group was given some of the same descriptions, but with the information that the descriptions had been selected from a sample that contained 30 percent lawyers and 70 percent engineers. Respondents had no trouble with descriptions that obviously described one occupation or another, but some were intentionally neutral with phrases like "he is highly motivated" or "will be successful in his career." The modal response for such neutral descriptions involved probabilities of near a half, regardless of the respondent's treatment group. This behavior is insensitive to the prior information about the proportions of each occupation, a type of base rate bias.

Grether (1980) pointed out several potential procedural problems with the Kahneman and Tversky experiment. There was deception to the extent that the descriptions had been made up, and even if people “bought into” the context, they would have no incentive to think about the problem carefully. Moreover, the information conveyed in the descriptions is hard to evaluate in terms of factors that comprise Bayes’ rule formula. In other words, it is hard to determine an appropriate guess about the probability of a particular description conditional on the occupation. Grether based his experiments on cups with two types of balls as described. One of the biases that he considered is known as *representativeness bias*. In the two cup example discussed earlier, a sample of three draws that yields two Ambers and one Blue has the same proportions as cup A, and in this sense the sample looks representative of cup A. We saw that the probability of cup A after such a sample would be $2/3$, and a person who reports a higher probability, say 80 percent, may be doing so due to representativeness bias.

Notice that a sample of two Ambers and one Blue makes cup A more likely. If you ask someone which cup is more likely, an answer of A cannot distinguish between Bayesian behavior and a strong representativeness bias, which also favors cup A. Grether cleverly got around this problem by introducing some asymmetries which make it possible for representativeness to indicate a cup that is less likely given Bayes’ rule. For example, if you lower the prior probability of cup A to $1/8$, then the prior odds ratio on the right side of (30.5) will be: $\Pr(A)/\Pr(B) = (1/8)/(7/8) = 1/7$, which will make $\Pr(A|C)$ lower for each pattern of draws represented by C. In this manner, a sample of two Ambers and one Blue would look like the contents of cup A, but if the prior probability of A is small enough, the Bayesian probability of cup A would be less than one half. In this manner, representativeness and Bayes’ rule would have differing predictions when a person is asked which cup is more likely.

A binary choice question about which cup is more likely makes it easy to provide incentives: simply offer a cash prize if the cup actually used turns out to be the one the person said is more likely. This is the procedure that Grether used, with a \$15 prize for a correct prediction and a \$5 prize otherwise. When representativeness and Bayes’ rule gave different predictions, subjects tended to follow the Bayesian prediction more often than not, but a lot less often than when the two criteria matched, as with the symmetric example discussed in this chapter. When representativeness and Bayesian calculations indicated the same answer, subjects tended to give the correct answer about 80 percent of the time (with some variation depending on the specific sample). This percentage fell to about 60 percent when representativeness and Bayesian calculations suggested different answers.

V. Bayes' Rule with Elicited Probabilities

Sometimes it is useful to ask subjects to report a probability instead of just saying which event is more likely. This can be phrased as a question about the “chances out of 100 that the cup used is A.” The issue here is how to provide incentives for people to think carefully. The instructions in the appendix provide one approach, which is complicated, but which is based on a simple idea. Suppose that you send your friends to a fruit stand, and they ask you whether you prefer red or yellow tomatoes in case both are available. You would have no incentive to lie about your preference, since telling the truth allows your friends to make the best decision on your behalf. This section describes a method of eliciting probabilities that is based on this intuition. Probability elicitation is useful when we need specific probability numbers to evaluate theoretical predictions.

In particular, suppose that you have seen some draws (say Amber and Blue) and have concluded that the probability of cup A is 0.5. If you are promised a \$1,000 payment if cup A is really used, then you essentially have a “cup A lottery” that pays \$1,000 with probability one half. The idea is to ask someone for the probability (in chances out of 100) that cup A is used, and to give them the incentive to tell the truth. This incentive will be provided by having the person running the experiment make a choice for the subject. In this context, the experimenter is like the friend going to the fruit stand; the subject needs to tell the truth about their preferences so that the experimenter will make the best choice on the subject's behalf. In order to set up the right incentives to tell the truth, we will have the experimenter choose between that cup A lottery and another one that is constructed randomly, by using throws of 10-sided dice to get a number, N , between 0 and 100, in a manner which ensures that this “dice lottery” pays \$1,000 with chances N out of 100. If the probability of Cup A is $1/2$, this dice lottery is preferred if $N > 50$, and the cup A lottery is preferred if $N < 50$. The subject should answer that the chances of Cup A are 50 out of 100 so that the experimenter can make the right choice. This is really like the red and yellow tomato example discussed above, the experimenter is choosing between two lotteries on the subject's behalf, and therefore needs to know the subject's true value of the cup A lottery to make the right decision.

To convince yourself that the subject is motivated to tell the truth in this situation, consider what might happen otherwise. Suppose that the subject incorrectly reports that the chances of cup A are 75 out of 100, when the person really believes that each cup is equally likely. Thus, the cup A lottery provides a one-half chance of \$1,000. If the experimenter then throws a 7 and a 0, then $N = 70$, and the dice lottery would yield a 70% chance of \$1,000, which is a much better prospect than the 50-50 chance based on the subject's actual beliefs that each cup is equally likely. But since the subject incorrectly reported the chances

for cup A to be 75 out of 100, the experimenter would reject the dice lottery and base the subject's earnings on the cup A lottery, which gives a 20 percent lower chance of winning. A symmetric argument can be made for why it is bad to report that the chances of cup A are less than would be indicated by the subjects' true beliefs (question 3). For a mathematical derivation of the result that it is optimal to reveal one's true probability, see question 5.

The advantage of direct probability elicitation is that you can often make stronger conclusions when specific numbers are available, as opposed to the qualitative data obtained by asking which event is more likely. The disadvantage is that the elicitation process itself is not perfect in the sense that the measurements may contain more "noise" or measurement error than we get with binary comparisons.

Table 30.4. Elicited Probabilities for Two Subjects in a Bayes' Rule Experiment
Cup A: *a, a, b* Cup B: *a, b, b*

Subject 1			Subject 2		
Round	Draw (Bayes)	Elicited Probability	Round	Draw	Elicited Probability
21	None (0.50)	0.49	21	None (0.50)	0.50
22	<i>a</i> (0.67)	0.65	22	<i>b</i> (0.33)	0.30
23	<i>bb</i> (0.20)	0.18	23	<i>ba</i> (0.50)	0.60
24	<i>bab</i> (0.33)	0.25	24	<i>aba</i> (0.67)	0.70
25	<i>a</i> (0.67)	0.65	25	<i>a</i> (0.67)	0.65
26	<i>ab</i> (0.50)	0.49	26	<i>ab</i> (0.50)	0.30
27	<i>bba</i> (0.33)	0.33	27	<i>aab</i> (0.67)	0.80

Table 30.4 shows some elicited probabilities for a research experiment conducted by the author using University of Virginia subjects, and prize amounts of \$1.00 instead of \$1,000. All earnings were paid in cash. The two cups, A and B, each contained three marbles with the contents as shown in Table 30.1. The experiment consisted of three parts. The first part was done largely to acquaint people with the procedures, which are admittedly complicated. Then there were 10 rounds with asymmetric probabilities (a two-thirds chance of using cup A) and 10 rounds with symmetric probabilities (a one-half chance of using cup A). The order of the symmetric and asymmetric treatments was reversed with different groups of people.

The information and decisions for subjects 1 and 2 are shown in the table, for rounds 21-27 where the prior probability was 1/2 for each cup. First consider subject 1 on the left. In round 21, there were no draws, and the None (0.50) in the Draw column indicates that the correct Bayesian probability of cup A is 0.50. The subject's response in the Elicited Probability column was 0.49, as is also the case for this person in round 26 where the draws, AB, cancelled each other out. The other predictions can be derived using the ball counting heuristic or with Bayes' rule, as described above. In round 22, for example, there was only one draw, *a*, and the Bayesian probability of cup A is 0.67, since two of the three *a* balls from both cups are located in cup A. Subject 1 reports a probability 0.65, which is quite accurate. This person was unusually accurate, with the largest deviation from the theoretical prediction being in round 24, where the reported probability is a little too low, 0.25 versus 0.33.

Subject 2 is considerably less predictable. The *ab* and *ba* samples, which leave each cup being equally likely, resulted in answers of 0.60 and 0.30. The average of these answers is not far off, but the dispersion is atypically large for such an easy inference task, as compared with others in the sample. This person was fairly accurate with single draws of *b* (round 22) and of *a* (round 25), but the three-draw samples show behavior that is consistent with "representativeness bias." In round 27, for example, the sample of *aab* looks like cup A, and the elicited probability of 0.80 for cup A is much higher than the actual probability of 0.67. Subject 1 also seems to fall prey to this bias in round 24, where the sample, *bab*, looks like cup B.

Figure 30.1 shows some aggregate results for the symmetric and asymmetric treatments, for all 22 subjects in the study. The horizontal axis plots the Bayesian probability of cup A, which varies from 0.05 for a sample of *bbbb* in the symmetric treatment (with equal priors) to 0.97 for a sample of *aaaa* in the asymmetric treatment (where the prior probability of cup A was 2/3). The elicited probabilities are shown in the vertical dimension. The 45-degree dotted line shows the Bayes' prediction, the solid line connects the average of the elicited probabilities, and the dashed line connects the medians.




Figure 30.1. Elicited Probabilities Versus Bayes' Predictions (Source: Holt, 2002)

In the aggregate, Bayes' rule does quite well. A slight upward bias on the left side of the figure and a slight downward bias on the right side might possibly be due to the fact that there is more “room” for random error in the upward direction on the left and in the downward direction on the right. This conjecture is motivated by the observation that the medians (solid dashed line) are generally closer to Bayesian predictions. For example, one person got confused and reported a probability of 0.01 for cup A after observing draws of *aa* in the symmetric treatment, which should lead to a posterior of 0.8. (The draws were of light and dark marbles, and were coded on the decision sheet as LL.) On the right side of the figure, there is more “room” for extreme errors in the downward direction. To see this, consider a vertical line above the point 0.8 on the horizontal axis of Figure 30.1. If you imagine some people with no clue who put marks more or less equally spaced on this vertical line, more of the marks will be below the 45-degree dashed line, since there is more room below. Random errors of this type will tend to pull average reported probabilities down below the 45-degree line on the right side of the figure, and the reverse effect (upward bias)

would occur on the left side. Averages are much more sensitive to extreme errors than are medians, which may explain why the averages show more of a deviation from the 45 degree dashed line.

The averages graphed in the figure mask some interesting patterns of deviation. In the symmetric treatment, for example, there are several posterior probability levels that can result from different numbers of draws. A posterior of 0.67 can be achieved by draws of either *a* or *aab*. The average posterior for the *a* draw alone is 0.61, which is a little too low relative to the Bayes' prediction of 0.67, and the "representativeness" pattern of *aab* is actually closer to the mark at 0.66. Similarly, the posterior of 0.33 can be reached by draw patterns of *b* or *bba*. The *b* pattern yielded an average of 0.42, whereas the *bba* pattern resulted in a lower average of 0.31. Representativeness would imply that the elicited probability should be lower than what is observed for the *bba* sample and higher than what is observed for the *aab* sample. It may be after several draws, some people are ignoring the prior information and reporting a probability that matches the sample average, which would be a natural type of heuristic.

VI. Extensions

It is well known that elicited beliefs may deviate dramatically from Bayes's rule predictions, especially in extreme cases like the disease example discussed in section I. Instruction in the mathematics of conditional probability calculations may not help much, and such skills are probably quickly forgotten. On the other hand, the counting heuristics, once learned, may help one make good probability assessments in other situations. From a modeling perspective, the main issue is where to begin, i.e. whether to throw out Bayes' rule and begin fresh with something that has a more behavioral and less mathematical foundation. The alternative to begin by making Bayesian calculations and then trying to assess possible sources of bias.

Although there seem to be some systematic deviations from the predictions of Bayes' rule, there is no widely accepted alternative model of how information is actually processed by people in these situations. Some biases seem to be due to the use of heuristics, and some of the bias may be due to asymmetries in the way the measurements are made. At present, economists tend to use Bayes' rule to derive predictions, although there is some renewed interest in non-Bayesian models like that of reinforcement learning discussed in the previous chapter. Finally, it is useful to know that Bayes' rule can be used to derive the previous chapter's model of belief learning by making a specific assumption about the nature of the prior beliefs. This derivation is beyond the scope of this book (see DeGroot, 1970), but it is not very important since most people who use these models start making additional changes to the belief learning model that cannot be justified from Bayes' rule.

Questions

1. What would the mental model with imagined balls in Table 30.3 look like after seeing two Amber draws and one Blue draw? What does this model imply that the posterior probability for cup A would be?
2. Answer question 1 for a sample of three Amber draws, and check your answer using Bayes' rule. (Hint: the probability of getting three Amber draws from cup A is $(2/3)(2/3)(2/3) = 8/27$. You will also need to calculate the probability of getting three Amber draws from cup B.)
3. Suppose we are using the elicitation scheme described in section V, and that the subject's beliefs are that cups A and B are equally likely. Show why it would be bad for the person to report that the chances for cup A are 25 out of 100.
4. In the asymmetric treatment discussed in Section V, the prior probability of cup A is $2/3$. Use the ball counting heuristic to calculate the probability of cup A after seeing a sample of: **b**. What is the probability of cup A after a sample of: **aaaa**?
5. Twenty-one University of Virginia students in an undergraduate Experimental Economics (Spring 2002) participated in the Bayes' rule elicitation experiment, using the instructions in the Appendix to Chapter 7 (but with only 7 rounds). The experiment lasted for about 20 minutes, and cash payments amounted to about \$3-\$4 per person, with all payments made in cash, despite the fact that this was a classroom experiment. Each cup was equally likely to be used. Participants claimed that they did have time to do mathematical calculations. The draw sequences and the average elicited probabilities are shown in the table below. A) Calculate the Bayes' posterior for each case, and write it in the relevant box. B) How would you summarize the deviations from Bayesian predictions? Is there evidence for the representativeness bias in this context? Discuss briefly.

Mean and Median Elicited Probabilities for Cup A

Cup A = {L, L, D} Cup B = {L, D, D}

Draw:	no draws	L	D	DD	LD	LDD	DDD
Mean	0.5	0.68	0.33	0.24	0.47	0.28	0.13
Median	0.5	0.67	0.33	0.2	0.50	0.30	0.11
Bayes'							

6. (Advanced and Tedious) The incentives for the method of eliciting probabilities discussed in Section V can be evaluated with calculus. Let P denote the person's true probability of Cup A, i.e. the probability that

represents their beliefs after seeing the draws of colored marbles. Let R denote the reported probability. (To simplify notation, both R and P are fractions between 0 and 1, not numbers between 0 and 100 as required for the “chances out of 100” discussion in the text.) **a)** Explain why the following statement is true for the elicitation mechanism discussed in Section V: “If the person reports R , then there is an R probability of ending up with the Cup A lottery, which pays \$1 with probability P .” **b)** Explain why the following statement is true: “Similarly, there is a $1-R$ probability of ending up with the dice lottery, which pays \$1 with a probability $N/100$, where N is the outcome of the throw of the ten-sided die twice.” **c)** We know that $N/100$ is greater than R , since the dice lottery is only relevant if its probability of paying \$1 is greater than the reported probability of Cup A. Use these observations to express the expected payoff as: $1/2 + PR - R^2/2$. **d)** Then show that this quadratic expression is maximized when $R = P$, i.e. when the reported probability equals the probability that represents the person’s beliefs. *Hint:* Since the dice lottery is only relevant if $N/100 > R$ and $N/100 < 1$, this dice lottery has an expected value that is halfway between R and 1, i.e. $(1+R)/2$. *Comment:* Probability elicitation methods of this type are sometimes called “scoring rules.” The particular method being discussed is a “quadratic scoring rule” since the function being maximized, $1/2 + PR - R^2/2$, is quadratic.

7. (Fun for a Change - This question was provided by Professor Brent Kreider.) Suppose that a cook produces three pancakes, one with one burnt side, one with two burnt sides, and one with no burnt sides. The cook throws a 6 sided die and chooses the pancake at random, with each one having equal probability of the one being served. Then the cook flips the pancake high so that each side is equally likely to be the one that shows. All you know, besides the way the pancake was selected, is that the one on your plate is showing a burnt side on top. What is the probability that you have the pancake with only one burnt side? Explain with Bayes’ rule or with the counting heuristic.
8. A test for breast cancer is given and comes back positive. The rate of previously undiagnosed breast cancer in the female population for her age group is 0.01. The test is accurate in the sense that it produces virtually no false negative results, so if a person has breast cancer, the test will always be positive. The rate of “false positives” on the test is 1/9, or about 0.11. What is the probability that the patient actually has the cancer?

Chapter 31. Information Cascades

Suppose that individuals may receive different information about some unknown event, like whether or not a newly patented drug will be effective. This information is then used in making a decision, like whether or not to invest in the company that developed the new drug. If decisions are made in sequence, then the second and subsequent decision makers can observe and learn from earlier decisions. The dilemma occurs when one's private information suggests a decision that is different from what others have done before. An "information cascade" forms when people follow the consensus decision regardless of their own private information. This game could be implemented with draws from cups and throws of dice, but the web version (CAS) is easier to administer in a manner that controls for unwanted informational signals.

I. "To do exactly as your neighbors do is the only sensible rule."

Conformity is a common occurrence in social situations, as the section title implies (from Emily Post, 1927, Chapter 33). People may follow others' decisions because they value conformity and fear social sanctions. For example, an economic forecaster may prefer the risk of being wrong along with others to the risk of having a deviant forecast that turns out to be inaccurate. Similarly, Keynes (1936) remarked: "Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally." Some experimenters have suggested that there is an irrational preference for the current state of affairs, i.e. a *status quo bias*. In some hypothetical choice experiments, Samuelson and Zeckhauser (1988) showed subjects two portfolios that could be used to invest money inherited from a rich uncle. There was a strong preference for the portfolio that was indicated to be the current investment profile, and a switch in the *status quo* designation caused a tendency to switch in preference (in a between-subjects design).

In some situations, people may be tempted to follow others' decisions because of the belief that there is some wisdom or experience implicit in an established pattern. Maybe subjects in the *status quo* bias experiments concluded that the deceased uncle was wealthy because he had selected a good portfolio. The possible effects of collective wisdom may be amplified in a large group. For example, someone may prefer to buy a Honda or Toyota sedan thinking that the large market shares for those models signal a lot of customer satisfaction. This raises the possibility that a choice pattern started by a few individuals may set a precedent that results in a string of incorrect decisions. At any given moment, for example, there are certain classes of stocks that are to be considered good investments, and herd effects can be amplified by efforts of some to anticipate

where the next fads will lead. Anyone who purchased a wide range of “tech” stocks several years ago can attest to the dangers of following the bulls.

A particularly interesting type of bandwagon effect can develop when individuals make decisions in a sequence and can observe others’ prior decisions. For example, suppose that a person is applying for a job in an industry with a few employers who know each other well. If the applicant makes a bad impression in the first couple of interviews and is not hired, then the third employer who is approached may hesitate even if the applicant makes a good impression on the third try. This third employer may reason: “anyone may have an off day, but I wonder what the other two people saw that I missed.” If the joint information implied by the two previous decisions is deemed to be more informative than one’s own information, it may be rational to follow the pattern set by others’ decisions, in spite of contradictory evidence. A chain reaction started in this manner may take on a life of its own as subsequent employers hesitate even more. This is why first impressions can be important in the workplace. The effect of information inferred from a sequential pattern of conforming decisions is referred to as an “information cascade” (Bikhchandani, Hirschleifer, and Welch, 1992).

Since first impressions can be wrong, the interviews or tests that determine initial decisions may start an “incorrect” cascade that implies false information to those who follow. This possibility was raised in an article in the *Economist* about the popular drug Prozac: “Can 10 Million People be Wrong?” John Dryden is somewhat more poetic: “Nor is the people’s judgment always true, the most may err as grossly as the few.” This chapter considers how cascades, incorrect or otherwise, may form even if individuals are good Bayesians in the way that they process information.

II. A Model of Rational Learning from Others’ Decisions

Following Anderson and Holt (1997, 2002), the discussion of cascades will be based on a very stylized model in which the two events are referred to as Cup A with contents *a, a, b*, and Cup B with contents *a, b, b*. Think of these cups as containing amber or blue marbles, with the proportions being correlated with the cup label. Each cup is equally likely to be selected, with its contents then being emptied into an opaque container from which draws are made privately. Each person sees one randomly drawn marble from the selected cup, and then must guess which cup is being used. Decisions are made in a pre-specified sequence, so that the first person has nothing to go on but the color of the marble drawn. Draws are made with replacement, so the second person sees a draw from the unknown cup, and must make a decision based on two things: the first decision and the second (own) draw. Like the employers who cannot sit in on others’ interviews, each person can see prior decisions but not the private information that may have affected those decisions. There is no external

incentive to conform in the sense that one's payoff depends only on guessing the cup being used; there is no benefit in conforming to others' (prior or subsequent) decisions.

The first person in the sequence has only the observation, *a* or *b*, and the prior information that each cup is *ex ante* equally likely. Since two of the three *a* marbles are in cup A, the person should assess the probability of cup A to be 2/3 if an *a* is observed, and similarly for cup B (cf. Chapter 30). Thus, the first decision should reveal that person's information. In the experiments discussed below, about 95% of the people made the decision that corresponded to their information when they were first in the sequence.

The second person then sees a private draw, *a* or *b*, and faces an easy choice if the draw matches the previous choice. A conflict is more difficult to analyze. For example, if the first person predicts cup A, the second person who sees a *b* draw may reason: "the cups are now equally likely, since that was the initial situation and the *a* that I think the first person observed cancels out the *b* that I just saw." In such cases, the second person should be indifferent, and might choose each cup with equal probability. On the other hand, the second person may be a little cautious, being more sure about what they just saw than about what they are inferring about what the other person saw. Even a slight chance that the other person made a mistake (deliberate or not) would cause the second person to "go with their own information" in the event of a conflict. In the experiments, about 95% of the second decision makers behaved in this manner in the experiment described below, and we will base subsequent discussion on the assumption that this is the case.

Table 31.1. Possible Inferences Made by the Third Decision Maker

Prior Decisions	Own Draw	Inferred Pr(cup A)	Decision
A, A	<i>a</i>	0.89	A (no dilemma)
A, A	<i>b</i>	0.67	A (start cascade)
A, B	<i>a</i>	0.67	A
A, B	<i>b</i>	0.33	B

Now the third person will have observed two decisions and one draw. There is no loss of generality in letting the first decision be labeled A, so the various possibilities are listed in Table 31.1. (An analogous table could be constructed when the first decision is B, and the same conclusions would apply.) In the top row of the table, all three pieces of information line up, and the standard Bayesian calculations would indicate that the probability of cup A is 0.89, making A the best choice (see question 1). The case in the second row is more interesting, since the person's *own* draw is at odds with the information inferred from other's decisions. When there are two others, who each receive independent

draw “signals” that are just as informative as the person’s own draw, it is rational to go with the decision implied by the others’ decisions. (Recall from chapter 30 that an imbalance of one draw in favor of cup A will raise the probability of cup A to 0.67.) The decision in the final two rows is less difficult because the preponderance of information favors the decision that corresponds to one’s own information.

The pattern of following others’ behavior (seen in the top two rows of Table 31.1) may have a domino effect: the next person would see three A decisions and would be tempted to follow the crowd regardless of their own draw. This logic applies to all subsequent people, so an initial pair of matching decisions (AA or BB) can start an information cascade that will be followed by all others, regardless of their information. This logic also applies when the first two decisions cancel each other out (AB or BA) and the next two form an imbalance that causes the fifth person to decide to follow the majority. An example of such a situation would be: ABAA, in which case the fifth person should choose A even if the draw observed is *b*. Again, the intuition is that the first two decisions cancel each other and that the next two matching decisions are more informative than the person’s own draw.

Notice that there is nothing in this discussion that implies that the first two people will guess correctly. There is a $1/3$ chance that the first person will see the odd marble drawn from the selected cup, and will guess incorrectly. There is a $1/3$ chance that the same thing will happen to the second person, so in theory, there is a $(1/3)(1/3) = 1/9$ chance that both of the first two people will guess incorrectly and spark an incorrect cascade. Of course, cascades (incorrect or not) may initially fail to form and then may later form when the imbalance of draws is 2 or more in one direction or the other. In either case, the aggregate information inferred from others’ decisions is greater than the information inherent in any single person’s draw.

III. Experimental Evidence

Anderson and Holt (1997) used this setup in a laboratory experiment in which people earned \$2 for a correct guess, nothing otherwise. Subjects were in isolated booths, so that they could not see others’ draws. Decisions (but not subject ID numbers) were announced by a third person to avoid having confidence or doubts communicated by the decision-maker’s tone of voice. The marbles were light or dark, with two lights and a dark in cup A, and with two darks and a light in cup B. The cup to be selected was determined by the throw of a six-sided die, with a 1, 2, or 3 determining cup A. The die was thrown by a “monitor” selected at random from among the participants at the start of the session. There were six other subjects, so each prediction sequence consisted of six private draws and six public predictions. For each group of participants, there

were 15 prediction sequences. The monitor used a random device to determine the order in which individuals saw their draws and made their predictions. The monitor announced the cup that had actually been used at the end of the sequence, and all participants who had guessed correctly added \$2 to their cumulative earnings. The monitor received a fixed payment, and all others were paid their earnings in cash. This “ball and urn” setup was designed to reduce or eliminate preferences for conformity that were not based on informational considerations.

It is possible that an imbalance of signals does not develop, e.g. alternating a and b draws, making a cascade unlikely. An imbalance did occur in about half of the prediction sequences, and cascades formed about 70% of the time in such cases. A typical cascade sequence is:

Table 31.2. A Cascade

Subject:	58	57	59	55	56	60
Draw:	b	b	a	b	a	a
Prediction:	B	B	B	B	B	B

Sometimes people do deviate from the behavior that would be implied by a Bayesian analysis. For example, consider the sequence:

Table 31.3. A Typical “Error”

Subject:	8	9	12	10	11	7
Draw:	a	a	b	a	b	a
Prediction:	A	A	B	A	A	A

The decision of subject 12 in the third position is the most commonly observed type of error, i.e. basing a decision on one’s own information even when it conflicts with the information implied by others’ prior decisions. This type of error occurred in about a fourth of the cases in which it was possible. Notice that the cascade starts later with subject 11 in the fifth position.

Once a cascade begins, the subsequent decisions convey no information about their signals, since these people are just following the crowd. In this sense, patterns of conformity are based on a few draws, and consequently, cascades may be very fragile. In particular, one person who breaks the pattern by revealing their own information may alter the decisions of those who follow.

Finally, a number of incorrect cascades were observed. Cup B was actually being used for the sequence shown in Table 31.4, but the first two individuals were unfortunate and received misleading a signals. Their matching predictions caused the others to follow with a string of incorrect A predictions,

which would have been frustrating given that these individuals had seen private draws that indicated the correct cup.

Table 31.4. An Incorrect Cascade (Cup B Was Used)

Subject:	11	12	8	9	7	10
Draw:	<i>a</i>	<i>a</i>	<i>B</i>	<i>b</i>	<i>b</i>	<i>b</i>
Prediction:	A	A	A	A	A	A

To summarize, the general tendency was for individuals to let an emerging pattern of others' decisions override their own private information. The resulting information cascades were common, but were sometimes broken by later deviant decisions.

V. Extensions and Further Reading

The model discussed in this chapter is based on a model presented in Bikhchandani, Hirschleifer, and Welch (1992), which contains a rich array of examples and applications. Hung and Plott (2001) discuss cascades in voting situations, e.g., when the payoff depends on whether the majority decision is correct. Some of the more interesting applications are in the area of finance. Keynes (1936) compared investment decisions with people in a guessing game who must predict which beauty contestant will receive the most votes. Each player in this game, therefore, must try to guess who is viewed as being attractive to the others, and on a deeper level, who the others will think that others will find more attractive. Similarly, investment decisions in the stock market may involve both an analysis of fundamentals and an attempt to guess what stocks will attract attention from other investors, and the result may be "herd effects" that may cause surges in prices and later corrections in the other direction. Some of these price movements may be due to psychological considerations, which Keynes compared with "animal spirits," but herd effects may also result from attempts to infer information from others' decisions. In such situations, it may not be irrational to follow others' decisions during upswings and downswings in prices (Christie and Huang, 1995). Bannerjee's (1992) model of herd behavior is motivated by an investment example. This model has been evaluated in the context of a laboratory experiment reported by Alsopp and Hey (2000). Other applications to finance are discussed in Devenow and Welch (1996) and Welch (1992).

It is very difficult to sort out all of the possible factors that affect prices in a stock market, so again laboratory experiments may be useful. Plott, Wit, and Yang (1997) ran some experiments with a "parimutuel" betting structure like that of a horse race where the purse is divided among those who bet on the winner, with payouts in proportion to the amounts of the winning bets. There were six

alternative assets, and only one of them would have a positive payout. The payout was divided equally among the investors in that asset. Each subject received a private signal about asset profitability, but the order of decision making was not exogenous as in the cascade experiments. Individuals could observe others' decisions as they were made. In this manner, the information dispersed among the investors could become aggregated and incorporated into the prices of the six assets. In most cases, the asset prices ended up signaling which asset would actually produce a positive payout. In some cases, however, an initial surge of purchases for a particular asset would stimulate others to follow, and the result was that the market price rose for an asset that had no payout in the end, as would be the case with an incorrect cascade.

Questions

1. Use the ball counting heuristic from chapter 30 to verify the Bayesian probabilities in the right-hand column of Table 31.1 under the assumption that the first two decisions correctly reveal the draws seen.
2. (Advanced) The discussion in this chapter was based on the assumption that the second person in a sequence would make a decision that reveals their own information, even when this information differs from the draw inferred from the first decision. The result is that the third person will always follow the pattern set by two matching decisions, because the information implied by the first two decisions is greater than the informational content of the third person's draw. In this question, consider what happens if we alter the assumption that the second person always makes a decision that reveals the second draw. Suppose instead that the second person chooses randomly (with probabilities of $1/2$) when the second draw does not match the draw inferred from the first decision. Use Bayes' rule to show that two matching decisions (AA or BB) should start a cascade even if the third draw does not match. (The intuition for this result is that the information implied by the first decision is just as good as the information implied by the third draw, so the second draw can break a "tie" if it contains at least some information.)
3. In the Anderson and Holt experiments, the marbles were not labeled a and b , but rather, were referred to as "light" and "dark," since they were either translucent or dark. What are the advantages of each labeling method? Which would you prefer to use in a research experiment, and why?

Chapter 32. Statistical Discrimination

Economists and other social scientists have long speculated that discrimination based on observable traits (race, gender, etc.) could become self-perpetuating in a cycle of low expectations and low achievement by workers of one type. In such a case, the employers may be reacting rationally, without bias, to bad employment experiences with one type of worker. Workers of that “disadvantaged” type, in turn, may correctly perceive that job opportunities are diminished, and may reduce their investments in human capital. The result may be a type of “statistical discrimination” that is based on experience, not on any underlying bias. The game implemented by the Veconlab program SD sets up this kind of situation, with each worker being assigned a color, Purple or Green. The game can be used to stimulate a class discussion of topics that might be too sensitive for many students.

I. “Brown-eyed People Are More Civilized”

As Jane Elliott approached her Riceville, Iowa classroom one Friday in April 1968, she was going over the provocative experiment that she had stayed up late planning. Martin Luther King had been murdered in Memphis the day before, and she anticipated a lot of questions and confusion. She was desperate to do something that might make a difference. Her ideas had begun to take shape when she recalled an argument about racial prejudice with her father that had caused his hazel eyes to flare. She remembered commenting to her roommate that her father would be in trouble if hazel eyes ever went out of favor.

As soon as her third grade class was seated, Ms. Elliott divided them into two groups based on eye color, as described in *A Class Divided* (Peters, 1971). At first, brown-eyed people were designated as being superior, with the understanding that the roles would be reversed the next day. She began: “What I mean is that brown-eyed people are better than blue-eyed people. They are cleaner ... more civilized And they are smarter than blue-eyed people.” As the brown-eyed people were seated in the front of the room and allowed to drink from the water fountain without using paper cups, the blue-eyed students slumped in their chairs and exhibited other submissive types of behavior. Objections and complaints were transformed into rhetorical questions about whether blue-eyed children were impolite, etc., which resulted in a chorus of enthusiastic replies from the brown-eyed people. When one student questioned Ms. Elliott about her own eye color (blue) and she tried to defend her intelligence, the reply was that she was not as smart as the brown-eyed teachers. What began as a role-playing exercise began to take on an eerie reality of its own. These patterns were reversed when blue-eyed people were designated as superior on the following day.

Psychologists soon began conducting experiments under more controlled conditions (e.g., Tajfel, 1970; Vaughan, Tajfel, and Williams, 1981). A typical setup involved a task of skill like a trivia quiz or estimation of the lengths of some lines. The subjects would be told that they were being divided into groups on the basis of their skill, but in fact the assignments were random. Then the subjects would be asked to perform a task like division of money or candy that might indicate differential status (see the review in Anderson, Fryer, and Holt, 2002). Such effects were often manifested as lower allocations to those with lower status. Moreover, people sometimes offered preferential treatment to people of their own group, which is known as an “in-group bias.”

Economists are typically interested in whether group and status effects will carry over into market settings. For example, Ball et al. (2001) examined the effects of status assignments in double auctions where the demand and supply functions were “box” shapes with a large vertical overlap, which produced a range of market-clearing prices. In the earned-status treatment, traders on one side of the market (e.g. buyers) were told that they obtained a gold star as a result of their performance on a trivia quiz. The random-assignment treatment began with some people being selected to receive stars that were awarded in a special ceremony. The subjects were not told that the star recipients were selected at random. In each treatment, the people with stars were all on one side of the market (all buyers or all sellers). This seemingly trivial status permitted them to earn more of the total surplus, whether or not they were buyers or sellers.

Group differences were largely exogenous in the status experiments, but economists have long worried about the possibility that differences in endogenously acquired skills may develop and persist in a vicious circle of self-confirming differential expectations. Formal models of “experienced-based” economic discrimination were developed independently by Arrow (1973) and Phelps (1972), and have been refined by others. The intuition is that if some historical differences in opportunity cause employment opportunities for one group (race, gender, etc.) to be less attractive, then members of that group may rationally choose not to invest as much in human capital. In response, employers may form lower expectations for members of that group. Expectations are based on statistics from past experience, so these have been called models of “statistical discrimination.”

The obvious question is why a member of the disadvantaged group could not invest and break out of the cycle. This strategy may not work if the employer attributes a good impression made in an interview to random factors, and then bases a decision not to hire on past experience and the fear that the interview did not reveal potential problems. (Think of the human capital being discussed as any aspect of productivity that is learned and that cannot be observed without error in the job application process. For example, the employer may be able to ascertain

that a person attended a particular software class, but not the extent to which the person mastered the details of working with spreadsheets.) Even if investment by members of the disadvantaged group results in good job placement some of the time, a lower success rate for people in this group will nevertheless result in a lower investment rate. After all, investment is costly in terms of lost income and lost opportunities to have children, etc. These observations raise the disturbing possibility that two groups with *ex ante* identical abilities will be treated differently because of differential rates of investment in acquired skills. In such a situation, the employers need not be prejudiced in any way other than reacting rationally to their own experiences. It is, of course, easy to imagine how real prejudice might arise (or continue) and accentuate the economic forces that perpetuate discrimination.

II. Being Purple or Green

The focus in this chapter is on experiments in which some group aspects are exogenous (color) and some are endogenous (investments). The participants are assigned a role, employer or worker. Workers are also assigned a color, Purple or Green. Workers are given a chance to make a costly investment in skills that would be valuable to an employer, but only if they were hired. The employers can observe workers' colors and the results of an imperfect pre-employment test before making the hiring decision. Investment is costly for the worker, but it increases the chances of avoiding a bad result on the pre-employment test, and hence, of getting hired. In this setup, the employers prefer to hire workers who invested, and not to hire those who did not.

Table 32.1. Experiment Parameters

Worker's payoffs:	\$1.50 if hired \$0.00 if not hired
Employer's payoffs:	\$1.50 if the hired worker invested -\$1.50 if the hired worker did not invest \$0.50 if the worker is not hired

This experimental setup, used by Fryer, Goeree, and Holt (2002), matches a theoretical model (Coate and Loury, 1993) in which statistical discrimination is a possible outcome. The experiment consisted of a number of rounds with random pairings of employers and workers. Each round began with workers finding out their own (randomly determined) costs of making an investment in some skill. Investment costs are uniform on the range from \$0.00 to \$1.00, except as noted below. The employer could observe the worker's color, but not the investment decision. A test given to each worker provided noisy information

about the worker's investment, and on the basis of this information the employer decided whether or not to hire the worker. The payoff numbers are summarized in Table 32.1.

Workers prefer to be hired, which yields a payoff of \$1.50 per period, as opposed to the zero payoff for not being hired. The investment cost (if any) is deducted from this payoff, whether or not the worker is hired. Investment is good in that it increases the probability of getting hired at \$1.50. Thus, a risk-neutral worker would want to invest if the investment cost is less than the wage times the increase in the hire probability that results from investing. If the investment cost is C , then the decision rule is:

Worker decision: invest if $C < \$1.50$ times the increase in hire probability.

Employers prefer to hire a worker who invested, which yields a payoff of \$1.50. Hiring someone who did not invest yields a payoff of $-\$1.50$, which is worse than the \$0.50 payoff the employer receives if no worker is hired. (Think of the fifty cents as what the manager can earn without any competent help.) The manager wishes to hire a worker as long as the probability of investment, p , is such that the expected payoff from hiring the worker, $p(1.50) + (1-p)(-1.50)$, is greater than the \$0.50 earnings from not hiring. It is straightforward to show that $p(1.50) + (1-p)(-1.50) > 0.50$ when $p > 2/3$:

Employer decision: hire when the probability of investment $> 2/3$.

The decision rules just derived, are of course, incomplete, since the probabilities of investment and being hired are determined by the interaction of worker and employer decisions.

The employer's beliefs about the probability that a worker invested are affected by a pre-employment test, which provides an informative but imperfect indication of the worker's decision. If the worker invests, then the employer sees two independent draws, with replacement, from the "invest cup" with 3 Blue marbles and 3 Red marbles. If the worker does not invest, the employer sees two draws with replacement from a cup with only 1 Blue and 5 Reds. Thus the cups are:



Obviously, the Invest Cup provides three times as great a chance that each draw will be Blue (B), so Red (R) is considered a bad signal. If the decision is to invest (Inv), the probabilities of the draw combinations are just products: $\Pr(BB | \text{Inv}) = (1/2)(1/2) = 1/4$, $\Pr(RR | \text{Inv}) = (1/2)(1/2) = 1/4$, and therefore, the residual probability of a mixed signal (RB or BR) is $1/2$. Similarly, the probabilities of the signal combinations for no investment (No) are: $\Pr(BB | \text{No}) = (1/6)(1/6) = 1/36$, $\Pr(RR | \text{No}) = (5/6)(5/6) = 25/36$, and $\Pr(RB \text{ or } BR) = 10/36$, which is the residual.

IV. An Unfair Equilibrium

The signal combination probabilities just calculated can be used to determine whether or not it is worthwhile for a worker to invest, but we must know how employers react to the signals. The discussion will pertain to an asymmetric equilibrium outcome in which one color gets preferential treatment in the hiring process. (Symmetric equilibria without discrimination will be discussed afterwards.) The theoretical question is whether asymmetric equilibria (with discriminatory hiring decisions) can exist, even if workers of each color have the same investment cost opportunities and payoffs. Consider the discriminatory hiring strategy:

(32.1)	Signal BB :	Hire both colors.
	Signal BR or RB :	Hire Green, not hire Purple.
	Signal RR :	Do not hire either color.

There are two steps to complete the analysis of this equilibrium: 1) to figure out the ranges of investment costs for which workers of each color will invest, and 2) to verify that the conjectured hiring strategy given above is optimal for employers.

Step 1: Worker Investment Decisions

Recall that the worker will invest if the cost is less than the wage \$1.50 times the increase in the probability of being hired due to a decision to invest. The Purple worker is only hired if both draws are blue, which happens with probability $(1/2)(1/2) = 1/4$ following investment and $(1/6)(1/6) = 1/36$ following no investment, so the increase in the probability of being hired is $1/4 - 1/36 = 9/36 - 1/36 = 8/36 = 2/9$. Hence $(2/9)\$1.50 = \0.33 is the expected net gain from investment, which is the optimal decision if the investment cost is less than \$0.33. In contrast, the favored Green workers are hired whenever the draw combination is not RR. So investment results in a hire with probability $1 - (1/2)(1/2) = 3/4$. Similarly, the probability of a Green worker being hired without investing is the probability of not getting two R draws from the No Invest Cup (BRRRRR), which is $1 - (5/6)(5/6) = 11/36$. The increase in the chances of getting a job due to

investment is $3/4 - 11/36 = 27/36 - 11/36 = 16/36 = 4/9$. Hence the expected net gain from investment by a Green worker is $(4/9)\$1.50 = \0.67 . To summarize, given the hiring strategy that discriminates in favor of Green workers, the Greens will invest whenever their investment costs are less than \$0.67, and the Purple workers will invest half as often, i.e. when their costs are less than \$0.33. Thus the discriminatory hiring rates induce more investment by the favored color.

Step 2: Employer Hiring Decisions

Recall that the employer's payoffs in Table 32.1 are such that it is better to hire whenever the probability that the worker invested is greater than $2/3$. Since Greens invest twice as often as Purples (a probability of $2/3$ versus $1/3$), the employer's posterior probability that a worker invested will be higher for Green, regardless of the combination of draws. These posterior probabilities are calculated with Bayes' rule and are shown in Table 32.2.

Table 32.2. Probabilities that the Worker Invested Conditional on Test and Color:
Decision is to Hire if the Probability $> 2/3$

	Green Workers	Purple Workers
Pr(worker invested BB)	18/19 (hire Green)	9/11 (hire Purple)
Pr(worker invested BR or RB)	18/23 (hire Green)	9/19 (not hire Purple)
Pr(worker invested RR)	18/43 (not hire Green)	9/59 (not hire Purple)

For example, consider the top row of the Green column, which shows the probability that a Green worker with a BB signal invested. This probability, $18/19$, is calculated by taking the ratio, with the Greens who invested and received the BB signal in the numerator and all Greens (who invested or not) who received the BB signal in the denominator. Since two-thirds of Greens invest and investment produces the BB signal with probability $(1/2)(1/2) = 1/4$, the numerator is $(2/3)(1/4)$. Since one-third of the Greens do not invest and not investment produces BB signal with probability $(1/6)(1/6) = 1/36$, the denominator also includes a term that is the product: $(1/3)(1/36)$. It follows that the probability of investment for a Green with the BB signal is: $(2/3)(1/4)$ divided by $(2/3)(1/4) + (1/3)(1/36)$, which reduces to $18/19$, as shown in the Table. The employer's optimal decision is to hire when the investment probability is greater than $2/3$, so the Green worker with a BB test result should be hired. The other probabilities and hire decision in the table are calculated in a similar manner (question 1). A comparison of the middle and right columns of the table show that these decisions discriminate against Purple in a manner that matches the original specification in (32.1).

It can be shown that there is also a symmetric equilibrium where both colors invest when the cost is less than \$0.67 and both colors are treated like the employer treated Greens in Table 32.1, i.e. always hire a worker of either color unless the test outcome is RR (question 2).

IV. Data

The advantage of running experiments is obvious in this case, since the model has both asymmetric and symmetric equilibria, and since exogenous sources of bias can be controlled. Moreover, the setup is sufficiently complex that behavior may not be drawn to any of the equilibria. Fryer, Goeree, and Holt (2002) have run a number of sessions using the statistical discrimination game. In some of the sessions there is little distinction between the way workers with different color designations are treated. This is not too surprising, given the symmetry of the model and the presence of a symmetric equilibrium.

Many cases where unequal treatment is likely to persist have evolved from social situations where discrimination was either directly sanctioned or indirectly subsidized by legal rules and social norms. Therefore, we began some of the sessions with 10 rounds of unequal investment cost distributions. In particular, one color (e.g. Green) might be favored initially by drawing investment costs from a distribution that is uniform on [\$0.00, \$0.50], and Purples may draw from a distribution that is uniform on [\$0.50, \$1.00]. The vertical dashed line at round 10 in Figure 32.1 shows the round where these unequal cost opportunities ended. Subjects were not told of this cost difference; they were just told that all cost draws would be between \$0.00 and \$1.00, which was the case. All draws after round 10 were from a common uniform distribution on [\$0.00, \$1.00].

Figure 32.1 shows five-period average investment and hiring rates. Notice that the initial cost asymmetry works in the right direction; Greens start out investing at about twice the rate as that for Purples, as can be seen from the left panel. There is a slight crossover in rounds 10-15, which seems to have been caused by a “relative cost effect,” i.e. the tendency for Purples to invest a lot when encountering costs that are lower than initially high levels that they were used to seeing. Similarly, investment rates for Greens fell briefly after period 10 when they began drawing from a higher cost distribution. As can be seen from the right panel in the figure, the employers seemed to notice the lower initial investment rates for Purples, and they continued hiring Greens at a higher rate. The inertia of this preference for Green eventually caused the investment rates for Green to rise above those for Purple again, and separation continued.

The color-based hiring preference was clear for a majority of the employers, although the separation was not as uniform as that predicted in the previous section’s equilibrium analysis. After an employer saw mixed signals (BR or RB), Greens were hired more frequently (96% versus 58% in the final 20

rounds). In fact, Greens were even hired more frequently after a negative RR signal (53% versus 0% in the final 20 rounds).




Figure 32.1. A Separation Effect: Five -Period Average Investment and Hire Rates by Color (Green is Solid, Purple is Dashed) with Statistical Discrimination Induced by an Initial Investment Cost Asymmetry (Source: Fryer, Goeree, and Holt, 2002)

The individual employer decisions for the final 15 rounds of this session are shown in Table 32.3. Most of the employers discriminate to some extent, either in the case of a mixed (RB or BR) signal, or in the case of a negative (RR) signal. Employers 1, 2, and 3 hired all Green workers encountered, even when the test result was RR. Employers 4 and 5 tended to discriminate after mixed signals, and employers 2 and 3 tended to discriminate after the negative signal. Employer 6 is essentially color-blind in the way workers are treated.

The experience of employer 2 indicated the problems facing Purple workers. All Purples encountered by this employer after round 40 had invested, but the employer was unable to spot this trend since no Purples were hired after round 40. (The investment decisions of the hired workers are shown in dark bold letters (**Inv** or **No**) when the worker was hired, and the unobserved investment decisions are shown in light gray when the worker was not hired.) Such differential treatment, when it occurs, is particularly interesting when the experiment is conducted in class for teaching purposes. In classroom experiments, one person once remarked “Purple workers just can’t be trusted...they won’t invest.” Another student remarked: “I invested every time, even when costs were high, because I felt confident that I would get the ...job – because *I am Green*.”

Table 32.3. Employer Information and Decisions for the Final 15 Rounds

Key: **Inv** = investment observed *ex post* by employer.

Inv = investment not observed *ex post* by employer.

Round	Employer 1	Employer 2	Employer 3	Employer 4	Employer 5	Employer 6
36	Purple BB Hire Inv	Green BR Hire Inv	Green BR Hire Inv	Purple BR Hire Inv	Purple BR Hire Inv	Green BB Hire Inv
37	Purple BR Hire Inv	Green RR Hire Inv	Green BB Hire Inv	Green BR No Hire Inv	Purple BR No Hire Inv	Purple BB Hire Inv
38	Green BR Hire No	Purple BB Hire Inv	Purple RR No Hire No	Green RR No Hire Inv	Green RR No Hire Inv	Purple BB Hire Inv
39	Purple RR No Hire No	Green RR Hire No	Purple RR No Hire Inv	Green RR No Hire Inv	Green BR Hire Inv	Purple BR Hire Inv
40	Purple BR No Hire Inv	Purple BB Hire Inv	Purple BB Hire Inv	Green BR Hire Inv	Green BR Hire Inv	Green BR Hire Inv
41	Purple RR No Hire Inv	Green RR Hire Inv	Green BR Hire No	Green BR Hire Inv	Purple RR No Hire Inv	Purple BR Hire No
42	Purple BB Hire Inv	Purple BR No Hire Inv	Green RR Hire Inv	Purple BR No Hire Inv	Green BB Hire Inv	Green BB Hire Inv
43	Purple BR Hire Inv	Purple RR No Hire Inv	Green RR Hire Inv	Purple BR No Hire Inv	Green BR Hire Inv	Green BB Hire Inv
44	Purple BB Hire Inv	Green BB Hire Inv	Purple BR Hire Inv	Purple BR No Hire Inv	Green BR Hire Inv	Green RR No Hire Inv
45	Green BR Hire Inv	Green BB Hire Inv	Purple BR Hire Inv	Purple BR No Hire No	Green BR Hire Inv	Purple RR No Hire No
46	Purple RR No Hire Inv	Purple BR No Hire Inv	Green RR Hire Inv	Green BR Hire Inv	Green RR No Hire Inv	Purple BR Hire Inv
47	Purple BB Hire No	Green RR Hire Inv	Green BR Hire Inv	Purple BB Hire Inv	Purple RR No Hire No	Green BR Hire Inv
48	Purple RR No Hire No	Green BB Hire Inv	Green BB Hire Inv	Green RR No Hire Inv	Purple RR No Hire Inv	Purple BR Hire Inv
49	Green BR Hire Inv	Green RR Hire Inv	Purple BB Hire Inv	Purple RR No Hire No	Purple BR Hire Inv	Green BB Hire Inv
50	Purple RR No Hire No	Green BB Hire Inv	Green RR Hire No	Purple BR Hire Inv	Green BR Hire Inv	Purple RR No Hire No

The initial cost differences produced the same initial patterns in a second experiment, shown in Figure 32.2. The vertical dashed line again indicates the point after which the cost distributions became symmetric. Notice that the surge caused by the relative cost effect after 10 rounds carries over, causing a fascinating reverse separation where the initially disadvantaged group (Purple) ends up investing more for the remainder of the session. This differential investment rate seems to cause the hiring rates to diverge in the final ten periods. A similar crossover effect was observed in a classroom experiment with 5 rounds of asymmetric costs followed by 15 rounds of symmetric costs.




Figure 32.2. A Cross-Over Effect: Five-Period Average Investment and Hire Rates by Color (Green is Solid, Purple is Dashed) with Reverse Statistical Discrimination Induced by an Initial Investment Cost Asymmetry (Source: Fryer, Goeree, and Holt, 2002)

Extensions and Further Reading

There are a number of related theoretical models of statistical discrimination. These models are presented with a uniform notation in Fryer (2001). The classroom experiments mentioned in the previous section are discussed in more detail in Fryer, Goeree, and Holt (2001). Anderson and Hupert (1999) describe a classroom experiment with exogenously determined worker skill levels. Davis (1987) provides an experimental test of the idea that perceptions about a group may be influenced by the highest skill level encountered in the past. In this case, a majority group produces more workers, so statistically speaking, the highest skill level from that group should be higher than from the other group.

Questions

1. Verify the Bayes' rule probability calculations for Purple workers in the right column of Table 32.2.
2. Consider a symmetric equilibrium where employers hire a worker regardless of color if the test result is BB, BR, or RB, and not otherwise. To do this, first show that workers of either color will invest as long as the cost is less than \$0.67. Then calculate the probabilities of investment conditional on the test results, as in Table 3.2. Finally, show that the employer's best response is to hire unless the test result is RR.

Chapter 33. Signaling (unfinished)

When individuals have traits that cannot be directly observed by others, they may engage in costly activities in order to “signal” the presence of these traits. Many educational credentials, for example, are thought to provide signals of a prospective employee’s productivity and enthusiasm. The simplest signaling game is one where the signals are not directly productive, but differential costs of signaling may allow one type of person to distinguish themselves from others. There may also be “pooling” equilibria in which both types of people send identical signals, so that the signal loses meaning. The Veconlab signaling game SG facilitates the discussion of how these behavioral patterns may arise.

Appendices: Instructions for Class Experiments

Pit Market Instructions (Chapter 2)

We are going to set up a market in which the people on my right will be buyers, and the people on my left will be sellers. There will be equal numbers of buyers and sellers, which is the reason that some of you had to switch to the other side of the room. Several assistants have been selected to help record prices. I will now give each buyer and seller a numbered playing card. Some cards have been removed from the deck(s), and all remaining cards have a number. Please hold your card so that others do not see the number. The buyers' cards are red (Hearts or Diamonds), and the sellers' cards are black (Clubs or Spades). Each card represents a "unit" of an unspecified commodity that can be bought by buyers or sold by sellers.

Trading: Buyers and sellers will meet in the center of the room (or other designated area) and negotiate during a 5 minute trading period. When a buyer and a seller agree on a price, they will come together to the front of the room to report the price, which will be announced to all. Then the buyer and seller will turn in their cards, return to their original seats, and wait for the trading period to end. There will be several market periods.

Sellers: You can each sell a single unit of the commodity during a trading period. The number on your card is the dollar cost that you incur if you make a sale. You will be required to sell at a price that is no lower than the cost number on the card. Your earnings on the sale are calculated as the difference between the price that you negotiate and the cost number on the card. If you do not make a sale, you do not earn anything or incur any cost in that period. Think of it this way: its as if you knew someone who would sell you the commodity for a price that equals your cost number, so you can keep the difference if you are able to resell the commodity for a price that is above the acquisition cost. Suppose that your card is a 2 of Clubs and you negotiate a sale price of \$3. Then you would earn: $3 - 2 = \$1$. You would not be allowed to sell at a price below \$2 with this card (2 of Clubs). If you mistakenly agree to a price that is below your cost, then the trade will be invalidated when you come to the front desk; your card will be returned and you can resume negotiations.

Buyers: You can each buy a single unit of the commodity during a trading period. The number on your card is the dollar value that you receive if you make a purchase. You will be required to buy at a price that is no higher than the value number on the card. Your earnings on the purchase are calculated as the difference between the value number on the card and the price that you negotiate. If you do not make a purchase, you do not earn anything in the period. Think of it this way: its as if you knew someone who would later buy the unit from you at a price that equals your value number, so you can keep the difference if you are able to buy the unit at a price that is below the resale value. Suppose that your card is a 9 of Diamonds and you negotiate a purchase price of \$4. Then you would earn: $9 - 4 = \$5$. You would not be allowed to buy at a price above \$9 with this card (9 of Diamonds). If you mistakenly agree to a price that is above your value, then the trade will be invalidated when you come to the front desk; your card will be returned and you can resume negotiations.

Recording Earnings: Some buyers and sellers may not be able to negotiate a trade, but do not be discouraged since new cards will be passed out at the beginning of the next period. Remember that earnings are zero for any unit not bought or sold (sellers incur no cost and buyers receive no value). When the period ends, I will collect cards for the units not traded, and you can calculate your earnings while I shuffle and redistribute the cards. Your total earnings equal the sum of earnings for units traded in all periods, and you can use the Earnings Record Form on the back of this sheet to keep track of your earnings. Sellers use the left side of the Earnings Record Form, and buyers use the right side. At this time, please draw a diagonal line through the side that you will *not* use. All earnings are hypothetical. Please do not talk with each other until the trading period begins. Are there any questions?

Final Observations: When a buyer and a seller agree on a price, both should *immediately* come to the front table to turn in their cards together, so that we can verify that the price is neither lower than the seller's cost nor higher than the buyer's value. If there is a line, please wait together. After the price is verified, the assistant at the board will write the price and announce it loudly. Then those two traders can return to their seats to calculate their earnings. The assistants should come to their positions in the front of the room. Buyers and sellers, please come to the central trading area NOW, and begin calling out prices at which you are willing to buy or sell. The market is open, and there are 5 minutes remaining.

Your Name: _____

Seller Earnings				Buyer Earnings		
(sellers use this side)				(buyers use this side)		
		=	First			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
	—	=	Second			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
	—	=	Third			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
	—	=	Fourth			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
	—	=	Fifth			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
	—	=	Sixth			=
<u> </u>	<u> </u>	<u> </u>	Period	<u> </u>	<u> </u>	<u> </u>
(price)	(cost)	(earnings)		(value)	(price)	(earnings)
total earnings, for all periods:		\$		total earnings, for all periods:		\$
		_____				_____

Game Instructions (Chapter 3)

We are going to play a card game in which everybody will be matched with someone on the opposite side of the room. I will now give each of you a pair of playing cards, one red card (Hearts or Diamonds) and one black card (Clubs or Spades). The numbers or faces on the cards will not matter, just the color. You will be asked to play one of these cards by holding it to your chest (so we can see that you have made your decision, but not what that decision is). Your earnings are determined by the card that you play and by the card played by the person who is matched with you. If you play your red card, then your earnings in dollars will increase by \$2, and the earnings of the person matched with you will not change. If you play your black card, your earnings do not change and the dollar earnings of the person matched with you go up by \$3. In other words, think of there being some dollar bills on the table between you and the other person. You can either "pull" \$2 to yourself by playing the red card, or you can "push" \$3 to the other person by playing the black card. If you each pull your red card, you will each earn \$2. If you each push the black card, you will each earn \$3. If you push your black card and the other person pulls his or her red card, then you earn zero and the other person earns the \$5. If you pull red and the other person pushes black, then you earn the \$5, and the other person earns zero. Neither of you will be able to see what the other does (push or pull) until both decisions have been made. All earnings are hypothetical, except as noted below.

After you choose which card to play, hold it to your chest. We then tell you who you are matched with, and you can each reveal the card that you played. Record your earnings in the space below. (Option: After we finish all periods, I will pick one person with a random throw of dice and pay that person 10% of his or her total earnings, in cash. All earnings for everyone else are hypothetical. To make this easier, please write your name: _____ and the identification number that I will now give each of you: _____. Afterwards, I will throw a 10-sided die twice, with the first throw determining the "tens" digit, until I obtain the ID number of one of you, who will then be paid 10% of his or her total earnings in cash.) Any questions?

To begin: Would the people in the row that I designate please choose which card to play and write the color (R for Red, or B for Black) in the first column. Show that you have made your decision by picking up the card you want to play and holding it to your chest. Has everyone finished? Now, I will pair you with another person, ask you to reveal your choice, and calculate your earnings. Remember to keep track of earnings in the space provided below. Finally, please note that you will be matched with a different person in round 2, and payoffs will change. In round 3 you will be matched with a different person and payoffs

change again, but you will be paired with this person in all three subsequent periods.

Round	Payoffs	Your card (R or B)	Other's card (R or B)	Your Earnings
1	Red: pull \$2, Black: push \$3			
2	Red: pull \$2 Black: push \$8			
3	Red: pull \$2 Black: push \$3			
4	Red: pull \$2 Black: push \$3			
5	Red: pull \$2 Black: push \$3			

Lottery Choice Instructions (Chapter 4)

Your decision sheet shows ten decisions listed on the left. Each decision is a paired choice between "Option A" and "Option B." You will make ten choices and record these in the final column, but only one of them will be used in the end to determine your earnings. Before you start making your ten choices, please let me explain how these choices will affect your earnings, which will be hypothetical unless otherwise indicated.

Here is a ten-sided die that will be used to determine payoffs; the faces are numbered from 1 to 10 (the "0" face of the die will serve as 10.) After you have made all of your choices, we will throw this die twice, once to select one of the ten decisions to be used, and a second time to determine what your payoff is for the option you chose, A or B, for the particular decision selected. Even though you will make ten decisions, only one of these will end up affecting your earnings, but you will not know in advance which decision will be used. Obviously, each decision has an equal chance of being used in the end.

Now, please look at Decision 1 at the top. Option A pays \$2.00 if the throw of the ten sided die is 1, and it pays \$1.60 if the throw is 2-10. Option B yields \$3.85 if the throw of the die is 1, and it pays \$0.10 if the throw is 2-10. The other Decisions are similar, except that as you move down the table, the chances of the higher payoff for each option increase. In fact, for Decision 10 in the bottom row, the die will not be needed since each option pays the highest payoff for sure, so your choice here is between \$2.00 and \$3.85.

To summarize, you will make ten choices: for each decision row you will have to choose between Option A and Option B. You may choose A for some decision rows and B for other rows, and you may change your decisions and make them in any order. When you are finished, we will come to your desk and throw the ten-sided die to select which of the ten Decisions will be used, i.e. which row in the table will be relevant. Then we will throw the die again to determine your money earnings for the Option you chose for that Decision. Earnings for this choice will be added to your previous earnings (if any).

So now please look at the empty boxes on the right side of the record sheet. You will have to write a decision, A or B in each of these boxes, and then the die throw will determine which one is going to count. We will look at the decision that you made for the choice that counts, and circle it, before throwing the die again to determine your earnings for this part. Then you will write your earnings in the blank at the bottom of the page. Are there any questions?

	Option A	Option B	Your Choice A or B
Decision 1	\$2.00 if throw of die is 1 \$1.60 if throw of die is 2-10	\$3.85 if throw of die is 1 \$0.10 if throw of die is 2-10	
Decision 2	\$2.00 if throw of die is 1-2 \$1.60 if throw of die is 3-10	\$3.85 if throw of die is 1-2 \$0.10 if throw of die is 3-10	
Decision 3	\$2.00 if throw of die is 1-3 \$1.60 if throw of die is 4-10	\$3.85 if throw of die is 1-3 \$0.10 if throw of die is 4-10	
Decision 4	\$2.00 if throw of die is 1-4 \$1.60 if throw of die is 5-10	\$3.85 if throw of die is 1-4 \$0.10 if throw of die is 5-10	
Decision 5	\$2.00 if throw of die is 1-5 \$1.60 if throw of die is 6-10	\$3.85 if throw of die is 1-5 \$0.10 if throw of die is 6-10	
Decision 6	\$2.00 if throw of die is 1-6 \$1.60 if throw of die is 7-10	\$3.85 if throw of die is 1-6 \$0.10 if throw of die is 7-10	
Decision 7	\$2.00 if throw of die is 1-7 \$1.60 if throw of die is 8-10	\$3.85 if throw of die is 1-7 \$0.10 if throw of die is 8-10	
Decision 8	\$2.00 if throw of die is 1-8 \$1.60 if throw of die is 9-10	\$3.85 if throw of die is 1-8 \$0.10 if throw of die is 9-10	
Decision 9	\$2.00 if throw of die is 1-9 \$1.60 if throw of die is 10	\$3.85 if throw of die is 1-9 \$0.10 if throw of die is 10	
Decision 10	\$2.00 if throw of die is 1-10	\$3.85 if throw of die is 1-10	

BS Game Instructions (Chapter 5)

We are going to play a card game in which everybody will be matched with someone on the opposite side of the room. I will now give each of you a pair of playing cards, one Red card (Hearts or Diamonds) and one Black card (Clubs or Spades). (Instead, these may be index cards with numbers written in red or black). The people on the left side of the room have a red 8 and a black 2, whereas the people on the right side of the room have a red 2 and a black 8.

Left Side	Right Side
Red 8, Black 2	Red 2, Black 8

You will be asked to play one of these cards by holding it to your chest (so we can see that you have made your decision, but not what that decision is).

Your earnings are determined by the card that you play and by the card played by the person who is matched with you:

*If the colors of the cards do not match (Red and Black), you each earn nothing.
If the colors match, earnings are equal to the number on your card:*

After you choose which card to play, hold it to your chest. We then tell you who you are matched with, and you can each reveal the card that you played. Record your earnings in the space below. All earnings are hypothetical, except as noted below. (Option: After we finish all rounds, I will pick one person with a random throw of dice and pay that person ___% of his or her total earnings, in cash. All earnings for everyone else are hypothetical. To make this easier, please write your name: _____ and the identification number that I will now give each of you: _____. Afterwards, I will throw a 10-sided die twice, with the first throw determining the "tens" digit, until I obtain the ID number of one of you, who will then be paid 10% of his or her total earnings in cash.) Any questions?

To begin: Would the people in the row that I designate please choose which card to play and write the color (R for red or B for black) in the "Your Card" column. Show that you have made your decision by picking up the card you want to play and holding it to your chest. Everyone finished? Now, I will pair you with another person, ask you to reveal your choice, and calculate your earnings. Remember to keep track of the other's card and of your earnings in the space provided below. Finally, please note that in round 2 you will be matched with a different person, and payoffs change. In round 3 you will be matched with

a different person and payoffs change again, but you will be paired *with the same person in all three subsequent rounds*.

BS Game Payoffs:

Color key: R for Red, B for Black.

If the cards match in color (RR or BB), you earn a dollar amount (to be announced) for the card you played.

If the cards do not match in color (RB or BR), you do not earn anything.

Initial payoffs for round 1:

You play R, other plays R, you earn \$ ____.

You play B, other plays B, you earn \$ ____.

You play R, other plays B, you earn \$0.00.

You play B, other plays R, you earn \$0.00.

Round	Payoffs	Your card (R or B)	Other's card (R or B)	Your Earnings
1	R card: \$ ____ B card: \$ ____			
2	R card: \$ ____ B card: \$ ____			
3	R card: \$ ____ B card: \$ ____			
4	R card: \$ ____ B card: \$ ____			
5	R card: \$ ____ B card: \$ ____			
6	R card: \$ ____ B card: \$ ____			
7	R card: \$ ____ B card: \$ ____			
8	R card: \$ ____ B card: \$ ____			

Binary Prediction Game (Chapter 6)

This is an exercise in which you will be asked which of two random events will occur. The events will be called L (left) and R (right). We will use the throw of dice to determine which event will occur in each period. Please look at the decision sheet below. The number of the round is on the left, and you will use the second column to record your prediction, L or R. You begin by recording your prediction *for round 1 only*, leaving all remaining rows blank. After everybody has written their prediction for round 1 into the top row, we will throw the dice in the front of the room to determine the event, L or R. This throwing of the dice will be done behind a screen so that you cannot see what method is being used. The main thing to remember here is that the person throwing the dice will not know your predictions, so these decisions cannot affect the likelihood of future events. When we announce the event (L or R) for the current period, please write in your earnings: +50 cents if you were correct, and -50 cents if you were incorrect. You can keep track of your cumulative earnings in the final column. Your instructor will make an announcement about whether or not the earnings will actually be paid in cash. To summarize, each period consists of 1) the prediction stage, 2) the throwing of the dice to determine the event, 3) the announcement of the event and the calculation of results. Are there any questions?

Round	Your Prediction (L or R)	Observed Event (L or R)	Your Earnings	Total Earnings
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				

Lottery Choice Instructions (Chapter 7)

Your decision sheet shows ten decisions listed on the left, which are labeled 0, 1, 2, .. 9. Each decision is a paired choice between a randomly determined payoff described on the left and another one described on the right. You will make ten choices and record these in the final column, but only one of them will be used in the end to determine your earnings. Before you start making your ten choices, please let me explain how these choices will affect your earnings. Unless otherwise indicated, all earnings are hypothetical.

Here is a ten-sided die that will be used to determine payoffs; the faces are numbered from 0 to 9. After you have made all of your choices, we will throw this die three times, once to select one of the ten decisions to be used, and then two more times to determine what your payoff is for the option you chose (on the left side or on the right side) for the particular decision selected. Even though you will make ten decisions, only one of these will end up affecting your earnings, but you will not know in advance which decision will be used. Obviously, each decision has an equal chance of being used in the end.

Now, please look at Decision 0 at the top. If this were the decision that we ended up using, we would throw the ten-sided die two more times. The two throws will determine a number from 0 to 99, with the first throw determining the “tens” digit and the second one determining the “ones” digit. The option on the left side pays \$6.00 if the throw of the ten-sided die is 0-99, and it pays \$0.00 otherwise. Since all throws are between 0 and 99, this option provides a sure \$6.00. The option on the right pays \$8.00 if the throw of the die is 0-79, and it pays \$0.00 otherwise. Thus the option on the right side of Decision 0 provides 80 chances out of 100 (a four-fifths probability) of getting \$8.00. The left and right options for the other decision rows are similar, but with differing payoffs and chances of getting each payoff. In addition, you will receive \$10.00 for participating, so any earnings will be added to this amount. Some of the decisions involve losses, indicated by minus signs. If the decision selected ends up with a loss, this loss will be subtracted from the initial \$10.00 payment to determine your final earnings.

To summarize, you will make ten choices: for each decision row you will have to choose between the option on the left and the option on the right. Make your choice by putting a check by the option you prefer. If you change your mind, cross out the check and put it on the other side. Thus there should be one check mark in each row. You may change your decisions and make them in any order. When you are finished, mark your final choices, L for left or R for right, in the far-right column, and we will come to your desk and throw the ten-sided die to select which of the ten Decisions will be used. We will circle that decision before

throwing the die again to determine your money earnings for the Option you chose for that Decision.

Before you begin, let me mention that different people may make different choices for the same decision, in the same manner that one person may purchase a sweater that differs from that purchased by someone else. We are interested *your* preferences, i.e. in which option you prefer to have, so please think carefully about each decision, and please do not talk with others in the room. Are there any questions?

	Left Side	Right Side	Your Choice L or R
Decision 0	\$6.00 if throw of die is 0-99	\$8.00 if throw of die is 0-79 \$0.00 if throw of die is 80-99	
Decision 1	\$2.00 if throw of die is 0-99	\$4.00 if throw of die is 0-79 −\$4.00 if throw of die is 80-99	
Decision 2	\$6.00 if throw of die is 0-24 \$0.00 if throw of die is 25-99	\$8.00 if throw of die is 0-19 \$0.00 if throw of die is 20-99	
Decision 3	\$30.00 if throw of die is 0-99	\$40.00 if throw of die is 0-79 \$0.00 if throw of die is 80-99	
Decision 4	\$4.00 if throw of die is 0-49 \$3.20 if throw of die is 50-99	\$7.70 if throw of die is 1-49 \$0.20 if throw of die is 50-99	
Decision 5	−\$6.00 if throw of die is 0-99	−\$8.00 if throw of die is 0-79 −\$0.00 if throw of die is 80-99	
Decision 6	\$2.00 if throw of die is 0-49 \$1.20 if throw of die is 50-99	\$5.70 if throw of die is 1-49 −\$1.80 if throw of die is 50-99	
Decision 7	−\$6.00 if throw of die is 0-24 −\$0.00 if throw of die is 25-99	−\$8.00 if throw of die is 0-19 −\$0.00 if throw of die is 20-99	
Decision 8	−\$4.00 if throw of die is 0-49 −\$3.20 if throw of die is 50-99	−\$7.70 if throw of die is 1-49 −\$0.20 if throw of die is 50-99	
Decision 9	\$30.00 if throw of die is 0-24 \$0.00 if throw of die is 25-99	\$40.00 if throw of die is 0-19 \$0.00 if throw of die is 20-99	

Search Instructions (Chapter 8)

In this game, we will use throws of dice to determine a series of money amounts, and you have to choose which amount to accept, with the understanding that each additional throw entails a cost that will be deducted from the money amount that you finally accept.

Please look at this 10-sided die. The sides are marked from 0 to 9, so if I throw it twice and use the first throw for the “tens” digit, I will determine a number from 00 to 99 cents. By ignoring the first draw if it is a 9 (and throwing again for a new first number, I will obtain a number from between (and including) 0 and 89. A number of pennies determined in this way will be yours to keep, but there is a cost of obtaining such a number.

In particular, if you pay a cost of 5 cents, I will throw the dice for you, and you can either keep the amount determined, or you can decide to pay another 5 cents and I'll throw again to get a second number in the range from 0 to 89 cents. (If the throw for the “tens” digit is a 9, I will throw the die again.) You can do this as many times as you wish (with no limit), but you have to pay 5 cents for each new number. When you decide to stop, you can use the highest 2-digit number that you have received up to that point, and your earnings will be that number minus the total cost of the search process, which is 5 cents times the number of times that you asked me to throw the dice. There is no limit on the number of times you can pay the 5 cent cost in search of a higher number.

Use the table below to keep records, and you can use another sheet if you need additional space. Each new 2-digit number will be called an “offer”. For the 1st offer that you pay to obtain, look at the “1st” column. After I throw the dice, write the number in the top row for “value of current offer”. The highest offer received up to now is written in the second row. (At first, the current offer is also the highest offer.) Moving down to the next row, the total search cost is the number of draws times 5 cents. The “earnings if you stop” are calculated by subtracting the total search cost from the highest draw thusfar. After calculating these earnings, you can decide whether to stop or to pay and search again. Write “S” for stop or “C” for continue. When you stop, please circle your earnings at that point and stop recording the new numbers. I will keep throwing the dice and announcing numbers until everyone had decided to stop. Then the whole process will be repeated, for another “search sequence”, etc....

All earnings are hypothetical unless otherwise indicated. Are there any questions before we begin?

Markets, Games, and Strategic Behavior – Charles A. Holt

Your Name: _____

Record of the First Search Sequence

order of offer	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10
value of current offer (in cents)										
highest offer thusfar										
total search cost (5 for each search)	5	10	15	20	25	30	35	40	45	50
earnings if you stop (use ! for losses)										
will you stop now? (Stop/Continue)										

Record of the Second Search Sequence

order of offer	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10
value of current offer (in cents)										
highest offer thusfar										
total search cost (5 for each search)	5	10	15	20	25	30	35	40	45	50
earnings if you stop (use ! for losses)										
will you stop now? (Stop/Continue)										

Please note: the cost of search for the next sequence has increased to 20 cents.

Record of the Third Search Sequence

order of offer	1st	2nd	3rd	4th	5 th	6th	7th	8th	9th	10
value of current offer (in cents)										
highest offer thusfar										
total search cost (20 for each search)	20	40	60	80	100	120	140	160	180	200
earnings if you stop (use ! for losses)										
will you stop now? (Stop/Continue)										

Finally, I will do this again for a search cost of **20 cents**, but there is a maximum of three offers, which are equally likely to be any amount between 0 and 89 cents. You can stop before the first draw and earn 0, or you can pay 20 cents for the first offer, and stop or pay 20 more for a second, and stop with the highest of the two or pay 20 more for a third, and take the highest of the three. Before doing this, please decide how high the first offer must be for you to stop after only one, and how high the best of the two initial offers must be for you to stop after only two. You may want to discuss this with your neighbor. Please record your decisions below.

On the first offer, I will stop if it is at least _____ .

On the 2nd draw, I will stop if the highest of the first two offers is at least _____ .

Now I will throw the dice for the 1st offer, etc., and you can record your earnings below:

Instructions: Traveler's Dilemma (Chapter 11)

I will now distribute one record sheet to each pair (or small group) of people in the class. In each period, your group will be randomly matched with another group of students. The decisions made by your group and by the other group will determine the amounts earned by each group.

At the beginning of each period, you will choose a number or "claim." Claims will be made by writing the claim on the record sheet. Your claim amount may be any amount between and including 80 and 200 cents. The other group will also make a claim between and including 80 and 200 cents. If the claims are equal, then your group and the other group each receive the amount claimed. If the claims are not equal, then each of you receives the lower of the two claims. In addition, the group which makes the lower claim earns a reward of 10 cents, and the group with the higher claim pays a penalty of 10 cents. Thus your group will earn an amount that equals the lower of the two claims, plus a 10 cent reward if you made the lower claim, or minus a 10 cent penalty if you made the higher claim. There is no penalty or reward if the two claims are exactly equal, in which case each group receives what they claimed.

Example: Suppose that your claim is X and the other claim is Y .
If $X = Y$, you get X , and the other gets Y .
If $X > Y$, you get Y minus 50 cents, and the other gets Y plus 10.
If $X < Y$, you get X plus 50 cents, and the other gets X minus 10.

Please look at your Record Sheet and write your name or assigned ID number in the upper right corner. Going from left to right, you will see columns for the "period," "your claim," "other's claim," "minimum claim," penalty or reward (if any), and "your earnings." Your group begins by writing down your own claim in the appropriate column, *for the current period only*. As mentioned above, this claim must be greater than or equal to 80 cents and less than or equal to 200 cents, and the claim may be any amount in this range. Please do not use decimals, e.g. write "xy" instead of "0.xy" to indicate xy cents.

After you record your claim, hold your paper up to your chest, and I will designate two groups and ask them to announce their decisions (please be honest). When you find out the other group's claim, please write it in the third column, followed by the minimum, the penalty or reward (if the claims are not equal), and your earnings. Then leave your paper on your desk so that I will know which people have not been matched yet. This same process is repeated.

Record Sheet for Minimum Claim Game

Your Name: _____

period	your claim	other claim	minimum	penalty/reward	earnings
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					

Instructions: Coordination Game (Chapter 12)

I will now distribute one record sheet to each of you; please write your name or ID number in the upper right-hand corner. In each round of this experiment, you will begin will choose a number or "effort" that can be 1, 2, 3, 4, 5, 6, or 7. Efforts will be made by writing the number you choose in the second column of your record sheet, *for the current round only*. When everyone is finished, the sheets will be collected. Then I will call out the numbers and ask one of you to write them on the board.

Your payoff will depend on your effort choice and on the minimum of all efforts, including your own. After we write all efforts on the blackboard, the lowest one will be circled. The second lowest effort will be marked with an asterisk. Your payoff will be determined in one of two ways:

- If your effort is not the lowest, then the circled number will determine the relevant column of the payoff table below, and your effort determines the row. The number in the cell that corresponds to the intersection of this row and column will be your earnings in cents.
- If your effort is the lowest, then the lowest of the other efforts is the number marked with an asterisk, and that number determines the column, while your number determines the row as before.

		Minimum of Other Efforts						
		1	2	3	4	5	6	7
Your Effort	7	10	30	50	70	90	110	130
	6	20	40	60	80	100	120	120
	5	30	50	70	90	110	110	110
	4	40	60	80	100	100	100	100
	3	50	70	90	90	90	90	90
	2	60	80	80	80	80	80	80
	1	70	70	70	70	70	70	70

When you have determined the minimum of the other efforts and your earnings, please write them in the relevant columns, and we will begin the next round.

Record Sheet for Effort Game

Your Name: _____

Round	Your Effort	Minimum of Other Efforts	Your Earnings
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			

Instructions: Market Game (Chapter 13)

I will now distribute one record sheet to each pair (or small group) of people in the class. In each market period (after the first one), your group will be matched with the same other group of students. The decisions made by your group and by the other group will determine the amounts earned by each group. (In the first period, your group will be the only seller in your market.)

At the beginning of each period, your group will choose a quantity to produce and sell. Decisions will be made by writing the quantity on your record sheet. Your quantity may be any amount between and including 1 and 13 units. The other group will also select a quantity on this interval. Each unit that you produce will cost your group \$1, and similarly for the other group. The price at which you can sell all units is determined by the total quantity, that is, the sum of your quantity and that of the other group, as shown in the table at the top of your record sheet. For example, if you each select a quantity of 1, the total is 2 and the price will be 11. If you each select 13, then the total will be 26, and the price will be 0, as will be the case whenever the total is greater than 12.

Example: Suppose that your quantity is 2 and the other group also selects 1.
The price will be _____.
Your total sales revenue (price times your quantity) will be _____.
Your total cost (\$1 times your quantity) will be _____.
Your profit (total revenue – total cost) will be _____.
You should have an answer of \$18 in profit.

Now please write your name or assigned ID number in the upper right corner of your record sheet. Going from left to right, you will see columns for the Period, Your Quantity, Other Quantity, Market Price, Total Revenue, Your Cost, and Your Earnings. Your group begins by writing down your own quantity in the appropriate column, *for the current period only*. As mentioned above, this quantity 0 and 12. The units are indivisible, so please use only integer amounts (no fractions). After you find out the other group's quantity, you fill out the information for the current period just as you did in the example. There will be a number of rounds, as determined by your instructor.

At this time, make your decision for the first period. Note that the other group's quantity has been set to 0, i.e. your group will be the only seller in this first period. After that, you will be matched with the same other group in all remaining periods.

Record Sheet for Market Game

Your Name: _____

Determination of Price as a Function of Total Quantity

Q	1	2	3	4	5	6	7	8	9	10	11	12	13+
Price	12	11	10	9	8	7	6	5	4	3	2	1	0

Period	Your Quantity	Other's Quantity	Market Price	Total Revenue	Your Cost	Your Earnings
1		0				
2						
3						
4						
5						
6						
7						
8						
9						
10						

Instructions: Price/Quality Market (Chapter 17)

(These instructions are based on Holt and Sherman, 1999.)

This is a market with 4 buyers and 3 sellers, and each of you will be assigned to a group that will have one of these roles. The sellers will begin by choosing a price and a quality "grade." We will collect these decisions and write them on the blackboard. Then we will give buyers the chance to purchase from one of the sellers at the grade and price listed. The grade can be any number from 1 to 3; a higher grade costs more to produce and is worth more to buyers. The table on your record sheet shows your costs of different grades if you are a seller, and it shows your money values of different grades if you are a buyer.

Each buyer can buy only 1 "unit" of the commodity during a period. Each seller can sell up to 2 units in a period, but the 2nd unit costs \$1 more to produce. If you are a seller, the top row of the table above shows the cost of the 1st unit that you actually sell in a period (for the grade you choose), the 2nd unit costs \$1 more than the 1st unit. Unsold units are not produced and hence incur no cost.

Buyers earn money by making a purchase at a price that is below the value, which depends on the quality grade. The value to the buyer depends only on the grade, not on whether it is the seller's 1st or 2nd unit in the period. A buyer's earnings are calculated as the difference between the value and the purchase price. If a buyer does not make a purchase, the buyer earns \$0.

Sellers earn money by making one or more sales at a price that is above the cost of the unit (determined from the table above). A seller's earnings are calculated as the sum of the earnings on the units actually sold, and such earnings are calculated as the difference between the sale price and the cost of the grade produced. A seller who does not make a sale in a period will earn \$0.

When all sellers have finished choosing their prices and grades for the period, we will collect these sheets and write the prices and grades on the blackboard under the seller numbers. Then I will draw lots to determine a buyer number, and that buyer can purchase a unit from one of the sellers or from no seller. Buyers are then chosen in order; if buyer 2 goes first, then buyer 3 is second, ... and buyer 1 is last. Once a seller has sold a unit, the 2nd unit costs \$1 more, so the seller will be asked whether or not the seller wishes to sell a 2nd unit at the advertised price and grade. If a 2nd unit is sold, it must be at the same price and grade as the 1st unit. If a seller refuses to sell or sells both units in a period, I will draw a line through that seller's price.

You can use the table below to calculate (hypothetical) earnings. Any questions? We will begin by having each seller choose a price and grade for period 1, which you should write in the top two rows of your record table.

Seller Record Sheet

Seller Number: _____

	grade 1	grade 2	grade 3
seller cost of 1st unit	\$1.40	\$4.60	\$11.00
seller cost of 2nd unit	\$2.40	\$5.60	\$12.00

	pd.1	pd.2	pd.3	pd.4	pd.5
1) grade for current period	_____	_____	_____	_____	_____
2) price for current period	_____	_____	_____	_____	_____
3) sales price on 1st unit	_____	_____	_____	_____	_____
4) cost of 1st unit	_____	_____	_____	_____	_____
5) profit on 1st unit: (3) – (4)	_____	_____	_____	_____	_____
6) sales price on 2nd unit	_____	_____	_____	_____	_____
7) cost of 2nd unit	_____	_____	_____	_____	_____
8) profit on 2nd unit: (3) – (4)	_____	_____	_____	_____	_____
9) total profit: (5) + (8)	_____	_____	_____	_____	_____
10) cumulative profit	_____	_____	_____	_____	_____

Buyer Record Sheet

Buyer Number _____

	grade 1	grade 2	grade 3		
buyer value	\$4.00	\$8.80	\$13.60		
	pd.1	pd.2	pd.3	pd.4	pd.5
1) ID of seller of product	_____	_____	_____	_____	_____
2) grade of product	_____	_____	_____	_____	_____
3) value to you (from table)	_____	_____	_____	_____	_____
4) purchase price	_____	_____	_____	_____	_____
5) earnings: (3) – (4)	_____	_____	_____	_____	_____
6) cumulative earnings	_____	_____	_____	_____	_____

Private Value Auction (Chapter 19)

We will conduct a series of auctions in which the highest bidder obtains the "prize". In each auction, you will be paired with another randomly selected member of the class. You and the person you are paired with will each see half of the prize value. The value of the prize will be determined by throws of 10 sided dice. Prior to the auction, we will come to each of you and throw the 10 sided die two times: the first throw determines the dimes digit, and the final throw determines the pennies digit. These two throws determine the prize value to you if you win. The other person will also see two throws of the dice that will determine his/her prize value, so the prize will generally be worth different amounts to each of you. You will make your bid decision after seeing your value, but without knowledge of the other person's value or the other person's bid. The higher bidder wins the prize and earns the difference between their own prize value and their own bid amount. The low bidder earns nothing. You will be paired with another person in the classroom in each period. This other person will be selected at random, unless your instructor indicates otherwise.

The Record Sheet should be used to keep track of your decisions. The ID number will be used to match you with another bidder each period. The period number is shown on the left side of each row. At the beginning of the period, I will come to your desk to throw the dice that determine your prize value, which can be recorded in column (1). After recording this number, you should decide on a bid for the period, which will be entered in column (2). Do this when you make your bid, before you hear what the other person's bid was. Then I will match you with someone else in the class, and you will each call out your bids. You should use column (3) to record the other person's bid. If you had the higher bid, your earnings will be the value (1) minus your bid in (2). If you had the lower bid, your earnings will be zero for the period. Earnings are entered in column (4). Any questions?

Record Sheet for Private Value Auction

Your ID: _____

Period	(1) Your Value	(2) Your Bid	(3) Other Bid	(4) Your Earnings	(5) Your Earnings
1					
2					
3					
4					
5					

Instructions for Part B:

This part will be the same as before, with values determined by the throws of 10 sided dice and with the high bidder being the winner. The only difference is that the winning bidder only has to pay the *second-highest bid price*. The high bidder earns the difference between their own value and the other person's bid, and the low bidder earns nothing. So, for example, if you value is 2 and you bid X and the other person bids Y , then you win with a bid of X but you only have to pay Y , so you earn $2 - Y$.

Record Sheet for Second Price Private Value Auction

Period	(1) Your Value	(2) Your Bid	(3) Other Bid	(4) Your Earnings	(5) Your Earnings
6					
7					
8					
9					
10					

Instructions for Takeover Game (Chapter 20)

In a minute I will divide the class into 2-person teams. Half of the teams will be buyers and half will be sellers in a market. I will choose teams and indicate your role, buyer or seller. Each team should have a single copy of these instructions, with the role assignment, buyer or seller, written next to the place for your names below.

Each seller team is the owner of a business. The money value of the business to the seller is only known by the seller. Each of the buyers will be matched with one of the sellers. The buyer will make a single bid to buy the business. The buyer does not know the value of the business to the seller, but the buyer is a better manager and knows that he or she can increase the profits to 1.5 times the current level. After receiving the buyer's bid, the seller must decide whether or not to accept it. The seller will earn the value of the business if it is not sold, and the seller will earn the amount of the accepted bid if the business is sold. The buyer will earn nothing if the bid is rejected. If the buyer's bid is accepted, the buyer will earn 1.5 times the seller's value, minus the bid amount.

A ten-sided die will be thrown twice to determine the value to the seller, in thousands of dollars. This will be done for each seller individually. The die is numbered from 0 to 9; the first throw determines the tens digit and the second determines the ones digit, so the seller value is equally likely to be any integer number of thousands of dollars, from 0 to 99 thousand dollars. If the business is purchased by the buyer, it will be worth 1.5 times the seller value, so the value to the buyer will be between 0 and 150 thousand dollars.

The form below can be used to record the outcomes. For simplicity, records will be in thousands of dollars; i.e. a 50 means 50 thousand. I will begin by coming to each seller's desk to throw the 10-sided die twice, and you can record the resulting seller value in column (1). Then each buyer, not knowing the seller's value, will decide on a bid and record it in column (2). I will then match the buyers and sellers randomly, and each buyer will communicate the bid to the corresponding seller, who will say yes (accept) or no (reject). The seller decision is recorded in column (3). If the bid is accepted, then the seller will communicate the seller's value to the buyer, who will multiply it by 1.5 and enter the sum in column (4). Finally, buyers and sellers calculate their earnings in the appropriate column (5) or (6). The buyer earns either 0 (if no purchase) or the difference between the buyer value (4) and the accepted bid. The seller either earns the seller value (if no sale) or the amount of the accepted bid. Each of you will begin with an initial cash balance of 500 thousand dollars; gains are added to this and

losses are subtracted. You can keep track of your cumulative cash balance in column (7). Are there any questions?

Now we will come around to the desks of sellers and throw the dice to determine the seller values. While we are doing this, those of you who are buyers should now decide on a bid and enter it in column (2). For those of you who are buyers, the seller value column (1) is blank. You only find out the seller value after you announce your bid. Those of you who are sellers have no decision to make at this time. When you hear the buyer bid, all you have to do is choose between keeping the value of your business (determined by the dice throws) or giving it up in exchange for the buyer bid amount. Please only write in the top row of blanks at this time. (After everyone has been matched and has calculated their earnings for the first set of decisions, you will switch roles, with buyers becoming sellers, and vice versa. The second set of decisions will be recorded in the bottom row.)

Record Sheet for Takeover Game

Your Role: _____ Your Name: _____

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Round	Seller's Value	Buyer's Bid	Seller's Decision (accept or reject)	Value to Buyer $1.5*(1)$	Seller's Earnings (1) or (2)	Buyer's Earnings $1.5*(1)-(2)$ or 0	Cash Balance 500 (thousand)
1							
2							

Common Value Auction (Chapter 21)

We will conduct a series of auctions in which the highest bidder obtains the "prize". In each auction, you will be paired with another randomly selected member of the class. You and the person you are paired with will each see half of the prize value. The value of the prize will be determined by throws of 10 sided dice. Prior to the auction, we will come to each of you and throw the 10 sided die two times: the first throw determines the dimes digit, and the final throw determines the pennies digit. These two throws determine your value component, which is equally likely to be any penny amount from \$0.00 to \$0.99. The other person will see two throws of the dice that will determine his/her value component in exactly the same manner. The value of the prize is just the sum of the two value components, so the prize value will be no less than \$0.00 and no greater than \$1.98. You will make your bid decision after seeing your value component, but without knowledge of the other person's value component or the other person's bid. The higher bidder wins the prize and earns the difference between the prize value (sum of the two value components) and the bid amount. This will be a gain if the bid is lower than the prize value, and this will be a loss if the bid is higher than the prize value. The low bidder earns nothing. You will be randomly paired with another person in the classroom in each period, so you only bid against one other person, but that person will generally be different in each successive auction "period".

The Record Sheet should be used to keep track of your decisions. The ID number will be used to match you with another bidder each period. The period number is shown on the left side of each row. At the beginning of the period, I will come to your desk to throw the dice that determine your value component, which can be recorded in column (1). After recording this number, you should decide on a bid for the period, which will be entered in column (2). Do this when you make your bid, before you hear what the other person's bid was. Then I will randomly match you with someone else in the class, and you will each call out your bids. You should use column (3) to record the other person's bid. At this time, you will know whether you won the auction or not, but you will not know the prize value. Finally, each person will announce their value components, so that you can record the other person's value component in column (4) and calculate the sum of the components in column (5). If you had the higher bid, your earnings will be the value sum in (5) minus your bid in (2). If you had the lower bid, your earnings will be zero for the period. Earnings are entered in column (6). You begin period 1 with an initial cash balance of \$5.00; gains are added to this and losses are subtracted. You can keep track of your cumulative cash balance in column (7). Any questions?

Record Sheet for Common Value Auction

Your ID: _____

Period	(1) Your Value Component	(2) Your Bid	(3) Other Bid	(4) Other's Value Component	(5) Sum of Value Components	(6) Your Earnings	(7) Cumulative Earnings \$5.00
1							
2							
3							
4							

Multi-Unit Auction (Chapter 22)

(These instructions are loosely adapted from a practice procedure in the instructions used for some of the experiments reported in Cummings, Holt, and Laury (2001). The number of units to be purchased in the auction is announced in advance. An alternative to the fixed purchase quantity is for the instructor to announce a budget amount that will be spent, which corresponds more closely to the irrigation reduction auction. We use a fixed quantity here to clarify the incentives in the uniform price auction. The instructor may wish to give the discriminatory auction instructions to half of the class, and to give the uniform price instructions that follow to the other half.)

A Discriminatory Auction

Each person has been given a single colored pen. This is yours to keep when you leave today, unless you decide to sell it back to us. Here I have an amount of money that is sufficient to purchase _____ of these pens. What I will do is let each of you write down an offer to sell your pen, using the form:

Your Name:

Sale offer price:
(between \$0.00 and \$1.00)

Please write your name and offer (between \$0.00 and \$1.00) in the two boxes. After you have turned in your offer, all offers will be ranked from low to high, regardless of pen color, and the _____ lowest offers will be accepted. Each person with an accepted offer will receive an amount that equals *that person's own offer* amount, and they must, in turn, give up their pens. People with rejected offers will keep their pens.

For example, if the quantity goal were two pens and the bids were 10 cents, 13 cents, 15 cents, and 20 cents, then two bids (10 and 13) would be accepted. This is because we start with the lowest bid and stop when the quantity goal is reached. The people with accepted offers get reimbursed an amount that equals their own offer, which means that one person receives 10 cents and the other receives 13 cents. Each of these two people must then give up the pen. The

others keep their pens and receive no payment. In the actual auction, the quantity goal will be _____ pens instead of the two pens used in this example.

A Uniform Price Auction

Each person has been given a single colored pen. This is yours to keep when you leave today, unless you decide to sell it back to us. Here I have an amount of money that is sufficient to purchase _____ of these pens. What I will do is let each of you write down an offer to sell your pen, using the form:

Your Name:

Sale offer price:
(between \$0.00 and \$1.00)

Please write your name and offer (between \$0.00 and \$1.00) in the two boxes. After you have turned in your offer, all offers will be ranked from low to high, regardless of pen color, and the _____ lowest offers will be accepted. Each person with an accepted offer will receive an amount that equals the *lowest rejected offer*, and they must, in turn, give up their pens. People with rejected offers will keep their pens.

For example, if the quantity goal were two pens and the bids were 10 cents, 13 cents, 15 cents, and 20 cents, then two bids (10 and 13) would be accepted. This is because we start with the lowest bid and stop when the quantity goal is reached. The people with accepted offers get reimbursed an amount that equals the lowest rejected bid, which is 15 cents in this case. These people must then turn each give up the pen. The others keep their pens and receive no payment. In the actual auction, the quantity goal will be _____ pens instead of the two pens used in this example.

Ultimatum Bargaining (Chapter 23)

Proposer Number: _____

I have \$10 to split. Those in the front of the room are “responders,” and those in the back are “proposers.” There are 11 ways that the \$10 can be divided between two people. The proposer must suggest one of these by circling one (and only one) of the listed options. The proposer sheets will then be collected and shuffled. One of the sheets will be given to each responder, who must either accept or reject. If the responder accepts, then the proposal is enacted. If the responder rejects, then the money is not divided and each person, proposer and responder, will earn nothing. Then I will collect the sheets and use the throw of a ten-sided die to select one of the proposer numbers, and the earnings (if any) determined by the decisions on that sheet will actually be paid to the proposer and responder who made those decisions. This payment will be in cash.

\$0 for the proposer, \$10 for the responder

\$1 for the proposer, \$9 for the responder

\$2 for the proposer, \$8 for the responder

\$3 for the proposer, \$7 for the responder

\$4 for the proposer, \$6 for the responder

\$5 for the proposer, \$5 for the responder

\$6 for the proposer, \$4 for the responder

\$7 for the proposer, \$3 for the responder

\$8 for the proposer, \$2 for the responder

\$9 for the proposer, \$1 for the responder

\$10 for the proposer, \$0 for the responder

Responder Number: _____

_____ I accept, and earnings will be determined by the proposal.

_____ I reject, and both of us will earn nothing.

Play-or-Keep Game (Chapter 26)

(These instructions adapted from those in Holt and Laury, 1997.)

Each of you will now be given four playing cards, two of which are red (hearts or diamonds), and two of which are black (clubs or spades). All of your cards will be the same number.

The exercise will consist of a number of rounds. At the start of a round, I will come to each of you in order, and you will play *two* of your four cards by placing these two cards face down on top of the stack in my hand.

Your earnings in dollars are determined by what you do with your red cards. For each red card that you keep in a round you will earn four dollars for the round, and for each black card that you keep you will earn nothing. Red cards that are placed on the stack affect everyone's earnings in the following manner. I will count up the total number of red cards in the stack, and everyone will earn this number of dollars. Black cards placed on the stack have no effect on the count. When the cards are counted, I will not reveal who made which decisions. To summarize, your earnings for the round will be calculated:

Your Earnings = \$4 times the # of red cards you kept
+ \$1 times the total # of red cards I collect.

At the end of the round, I will return your own cards to you by coming to each of you in reverse order and giving you the top two cards, face down, off the stack in my hand. Thus you begin the next round with two cards of each color, regardless of which cards you just played.

After round 5, I will announce a change in the earnings for each red card you keep. Even though the value of red cards kept will change, red cards placed on the stack will always earn one dollar for each person.

Use the space below to record your decisions, your earnings, and your cumulative earnings. (Optional: At the end of the game, one person will be selected at random and will be paid ____ % of his or her actual earnings, in cash.) All earnings are hypothetical for everyone else. Are there any questions?

Record Sheet: Earnings from Each Red Card Collected = \$1.00

Round	Number of Red Cards Kept	Value of Each Red Card Kept	Earnings from Red Cards Kept	Earnings from Red Cards Collected	Total Earnings	Cumulative Earnings
1		\$4				
2		\$4				
3		\$4				
4		\$4				
5		\$4				
6		\$2				
7		\$2				
8		\$2				
9		\$2				
10		\$2				

Volunteer's Dilemma Game (Chapter 27)

(The terminology in these instructions is designed for classroom use.)

Each of you will now be given two playing cards, one red (hearts or diamonds), and one black (clubs or spades). At the start of a round, I will ask each of you to play one of your cards by picking it up and holding it against your chest with the color hidden until I ask everyone to reveal which card they played.

After the card play decisions have been made, but before they have been revealed, I will select one or more others to be in your group and will ask all of you to reveal your cards at the same time. Your earnings in dollars are determined by the card that you play and by the cards played by the others in your group. Think of playing a red card as a decision to volunteer to perform some task. Volunteering has a cost for you (\$0.25), but everyone in the group will benefit (by receiving \$2.00) if at least one person volunteers. There is no additional benefit if there is more than one volunteer in your group. The payoffs are:

Payoff if there is no volunteer: \$0.00

Payoff if there is at least one volunteer and you DO NOT volunteer: \$2.00

Payoff if there is at least one volunteer and you DO volunteer: \$1.75

We will begin with groups of size 2 for the first 5 rounds, so when I point to you and to another randomly selected person in the room, show your cards, return your card to your desk, and calculate your earnings. The people who have not yet been matched should keep holding their cards to their chests so that I can continue to match people until nobody remains to be matched. Then the next round will begin. Each person will decide which card to play for that round by holding that card against their chest until they are matched. After round 5, I will divide people into randomly selected groups of size 4 in each round. Before beginning, I will select one person to record all decisions.

All earnings are hypothetical. (Optional: At the end of the game, one person will be selected at random and will be paid ____ % of his or her actual earnings, in cash.) Are there any questions?

Record Sheet for Volunteer's Dilemma Game					
Round	Group Size	Card Played (R or B)	Number of Red Cards Played in Your Group	Your Earnings	Cumulative Earnings
1	2				
2	2				
3	2				
4	2				
5	2				
6	4				
7	4				
8	4				
9	4				
10	4				

Payoff if there is no volunteer: \$0.00

Payoff if there is at least one volunteer and you DO NOT volunteer: \$2.00

Payoff if there is at least one volunteer and you DO volunteer: \$1.75

Lobbying Game Instructions (Chapter 29)

Source: These instructions are adapted from Goeree and Holt (1999).

This is a simple card game. Each of you has been assigned to a team of investors bidding for a local government communications license that is worth \$16,000. The government will allocate the license by choosing randomly from the applications received. The paperwork and legal fees associated with each application will cost your team \$3,000, regardless of whether you obtain the license or not. (Think of this \$3,000 as the opportunity cost of the time and materials used in completing the required paperwork.) Each team is permitted to submit any number of applications, up to a limit of thirteen per team. Each team begins with a working capital of \$100,000.

There will be four teams competing for each license, each of which is provided with thirteen playing cards of the same suit. Your team will play *any* number of these cards by placing them in an envelope provided. Each card you play is like a lottery ticket in a drawing for a prize of \$16,000. All cards that are played by your team and the other three teams will be placed on a stack and shuffled. Then one card will be drawn from the deck. If that card is one of your suit, your team will win \$16,000. Otherwise you receive nothing from the lottery. Whether or not you win, your earnings will decrease by \$3,000 for each card that you play. To summarize, your earnings are calculated:

$$\begin{aligned}\text{Earnings} = & \$16,000 \text{ if you win the lottery} \\ & - \$3,000 \text{ times the number of cards you play.}\end{aligned}$$

Earnings are negative for the teams that do not win the lottery, and negative earnings are indicated with a minus sign in the record table below. The cumulative earnings column on the right begins with \$100,000, reflecting your initial financial capital. Earnings should be added to or subtracted from this amount. Are there any questions?

Round	Number Cards Played	Cost per Card Played	Total Cost	License Value	Your Earnings	Cumulative Earnings \$100,000
1		\$3,000		\$16,000		

The next lottery is for a second license. Your team begins again with thirteen cards, but the cost of each card played is reduced to \$1,000, due to a government

efficiency move that requires less paperwork for each application. This license is worth \$16,000 as before, whether or not your team already acquired a license.

Round	Number Cards Played	Cost per Card Played	Total Cost	License Value	Your Earnings	Cumulative Earnings \$100,000
2		\$1,000		\$16,000		

In the next lottery, the value of the license may differ from team to team. Your team begins again with thirteen cards, and the cost of each card played remains at \$1,000. Your instructor will inform you of the license value, which you should write in the appropriate place in the table below. Each license is worth \$16,000 as before, whether or not your team already acquired a license.

Round	Number Cards Played	Cost per Card Played	Total Cost	License Value	Your Earnings	Cumulative Earnings \$100,000
3		\$1,000				

In the final round, the license will be worth the same to you as it was in round 3, but there is no lottery and no application fee. Instead, I will conduct an auction by starting with a low price of \$8,000 and calling out successively higher prices until there is only one team actively bidding. The winning team will have to pay the amount of its final bid. The losing teams do not have to pay anything for the license that they did not purchase; the winning team earns an amount that equals its license value minus the price paid. The revenue from the auction will be divided equally among the teams:

Round	Your Earnings (your license value minus your bid if you win, \$0 otherwise)	Cumulative Earnings \$100,000
-------	---	----------------------------------

4

Bayes' Rule Instructions (Chapter 30)

Note: The setup requires a 6-sided die, a 10-sided die, three light marbles, 3 dark marbles, a plastic cup for making draws, two paper envelopes marked A and B,

and a way to hide the cup selection process from view. The reading of the instructions and the experiment itself can be done in about 25 minutes.

In this experiment, you will observe "balls" (colored marbles) drawn from one of two possible "cups." You will see the balls drawn, but you will not know for sure which cup is being used. Then you will be asked to indicate your beliefs about which cup is being used. I will begin by choosing one of you to serve as a "monitor" who assists in setting things up and drawing the marbles.

The two cups will be called cup A and cup B. Cup A contains 2 light balls and 1 dark ball, and Cup B contain 1 light ball and 2 dark balls. The cup will be chosen by the throw of a 6 sided die: cup A is used if the roll of the die yields a 1, 2, or 3, and cup B is used if the roll of the die yields a 4, 5, or 6. Thus it is equally likely that either cup will be selected.

Cup A (used if the die is 1, 2, or 3)	Cup B (used if the die is 4, 5, or 6)
2 Light Balls 1 Dark Ball	1 Light Ball 2 Dark Balls

Once a cup is determined by the roll of the die, we will empty the contents of that cup into a container. The container is always the same, regardless of which cup is being used, so you cannot guess the cup by looking at the container. Then we will draw one or more balls from the container. (If more than one draw is to be made, then the first ball drawn will be put back into the container, which is then shaken before a second draw is made, etc.)

Recording Your Beliefs

After the draw has been made, we will ask you to tell us your beliefs about the chances that cup A is being used. You will indicate a number between 0 and 100, which we will call P , such that the chances that cup A is being used are " P out of 100." If you could be sure that cup A is being used, you should choose $P = 100$ to indicate that the chances are 100 out of 100 that cup A is being used. If you could be sure that cup A is not being used, you should choose $P = 0$ to indicate that the chances are 0 out of 100 that cup A is being used. Thus the magnitude of P corresponds to the chances that cup A is being used. For example, if the manner in which the cup is selected and the balls that you see drawn cause you think that cup A is just as likely as cup B, then you should choose $P = 50$, indicating that the chances of A are 50 out of 100.

You will use the attached Decision Sheets to record your information and decisions. At the start of each "period," the monitor will throw a 6-sided die to select the cup (A if the throw is 1, 2, or 3, and B if the throw is 4, 5, or 6). The monitor will then place the 3 balls for that cup in the plastic container from which we make the draws. The period number is shown in column (1) of the decision sheet. The results of the draws are to be recorded in column (2). Please look at the decision sheet for periods 1-7. In the first period, there will be no draws, as indicated in column (2), so the only information that you have is the information about how the cup is selected.

In subsequent periods, you will see one or more draws, and you can use column (2) to record the draw(s). Write L (for Light) or D (for Dark) in this column at the time the draw is made. In periods 2 and 3, there will be only a single draw. In periods 4 and 5, there will be two draws from the same cup (with replacement after the first draw). In periods 6 and 7, there will be 3 draws (with each ball drawn being put back into the cup before the next draw is made).

After seeing the draw(s), if any, for the period, you will be asked to indicate the chances that cup A is being used. Do this by writing a number, P , between 1 and 100 in column (3).

Earnings

Next, let me describe a procedure that will help you make your decision. The procedure is complicated, but the underlying idea is simple. It's as if you send your friend to a fruit stand for tomatoes, and they ask you what to do if there are both red and yellow tomatoes. You should tell them the truth about your preference, since only then can they make the best choice for you. We'll set up the incentives so that your report about the chances of cup A will enable us to choose a lottery that gives you the highest chance of winning a \$1 prize.

Suppose that, after seeing the draw or draws from the unknown cup, you think the chances are P out of 100 that cup A is being used. If you were to receive \$1 in the event that the cup actually used was A, then this determines a " P lottery" which pays \$1 with chances of P out of 100, nothing otherwise.

After you write down the number P that represents your beliefs about likelihood of cup A, we will use a ten-sided die to determine a second lottery, the "dice lottery." To do this, we throw the die twice, with the first throw giving the "tens" digit (0, 1, ...9) and the second throw giving the "ones" digit (0, 1, ...9). Thus the random number will be between 0 and 99. If the number is N , then the "dice lottery" will give you an N out of 100 chance of earning \$1. (To play the dice lottery for a given value of N , we would throw the ten-sided die twice more to get a number that is equally likely to be 0, 1, ... 99, and you would earn \$1 if this second throw is less than or equal to N .)

Up to this point, you will have written down a number P for what you think are the chances out of 100 that cup A is being used. The resulting P lottery will pay \$1 if cup A is used. Then we will have used a throw of dice to determine a dice lottery that will pay \$1 with chance N out of 100. Which would you rather have? It is better to have the dice lottery if N is greater than P , and it is better to have the P lottery otherwise. We'll give you the one that is best for you.

Case 1: $P < N$

So if you write down a number P and the number N determined subsequently by the dice throws is greater, you will get to play the dice lottery (throwing the dice again to see if you get the \$1, i.e. when the number determined by the second throw is less than or equal to N .)

Case 2: $P \geq N$

So if you write down a number P that is greater than or equal to the number N determined subsequently by the dice throws, then your earnings are determined by the P lottery, i.e. you will get \$1 if the cup used turns out to be cup A.

Cup A
(used if the die is 1, 2, or 3)

Cup B
(used if the die is 4, 5, or 6)

2 Light Balls
1 Dark Ball

1 Light Ball
2 Dark Balls

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period Number	Your draws L or D	Chances of urn A P	Dice throw N	if $P < N$, use Dice Lottery		if $P \geq N$, use the P Lottery	
				Dice throw B	Earnings = \$1 if $B < N$	Cup (A or B)	Earnings = \$1 if cup A
1	---						
2							
3							
4							
5	,						
6	,						
7	,						
8	, ,						
9	, ,						
10	, ,						

References

- Akerlof, George A., "The Market for 'Lemons:' Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, August 1970, 84, 488-500.
- Allais, M. (1953): "Le Comportement De L'homme Rationnel Devant Le Risque, Critique Des Postulates Et Axiomes De L'ecole Americaine," *Econometrica*, 21, 503-546.
- Alsopp, L., and J. D. Hey (2000): "Two Experiments to Test a Model of Herd Behavior," *Experimental Economics*, 3, 121-136.
- Anderson, D. and M. Hauptert (1999): Employment and Statistical Discrimination: A Hands-On Experiment, *Journal of Econometrics*, 25(1) 85-102.
- Anderson, L. R. (1999): "Payoff Effects in Information Cascade Experiments," College of William and Mary.
- Anderson, L. R., Fryer, R. M. and C. A. Holt (2002): "Experimental Studies of Discrimination," Discussion Paper, forthcoming in the W. Rogers, ed., *Handbook of Discrimination*.
- Anderson, L.R. and C. A. Holt (1996a): "Classroom Games: Understanding Bayes' Rule," *Journal of Economic Perspectives*, 10, 179-187.
- (1996b): "Classroom Games: Information Cascades," *Journal of Economic Perspectives*, 10, 187-193.
- (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87, 847-862.
- (1998): "Information Cascade Experiments," in *Handbook of Experimental Economics Results*, ed. by C. R. Plott, and V. L. Smith. New York: Elsevier Press, forthcoming.
- (2003): "Experimental Economics and Public Choice," in the *Encyclopedia of Public Choice*, C. Rowley and F. Schneider, eds., New York: Kluwer, ***-***.
- (2003): "Information Cascades and Rational Conformity," in L. Nadel, ed., *Encyclopedia of Cognitive Science*, Vol. 2, pp. 847-862. London: Nature Publishing Group, McMillan.
- Anderson, L. R., R. Fryer, and C. A. Holt (2004): "Experimental Studies of Discrimination," forthcoming in the *Handbook of Discrimination*, W. Rogers, ed.
- Anderson, L.R., C. A. Holt, and D. Reiley (2004) "Congestion and Social Welfare," Discussion Paper, College of William and Mary.
- Anderson, S. P., J. K. Goeree, and C. A. Holt (1998): "Rent Seeking with Bounded Rationality: An Analysis of the All Pay Auction," *Journal of Political Economy*, 106, 828-853.

- (2001a): "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games," *Journal of Public Economics*, 70, 297-323.
- (2001b): "Minimum Effort Coordination Games: Stochastic Potential and Logit Equilibrium," *Games and Economic Behavior*, 34, 177-199.
- (2002): "The Logit Equilibrium: A Unified Perspective on Intuitive Behavioral Anomalies in Games with Rank-based Payoffs," *Southern Economic Journal*, 68: 21-47.
- (2003): ***(change title) Stochastic Game Theory: Adjustment to Equilibrium Under Noisy Directional Learning," forthcoming, *Scandinavian Journal of Economics*.
- Andreoni, J. (1993): "An Experimental Test of the Public Goods Crowding-out Hypothesis," *American Economic Review*, 83, 1317-1327.
- Andreoni, J., M. Castillo, and R. Petrie (2003): "New Experiments on Bargaining: The Squishy Game," *American Economic Review*, *****.
- Andreoni, J., W. Harbaugh, and L. Vesterlund (2001): "The Carrot or the Stick: The Demands for Rewards and Punishments and Their Effects on Cooperation," University of Wisconsin.
- Andreoni, J., and J. H. Miller (1993): "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," *Economic Journal*, 103, 570-585.
- (1997): "Giving According to Garp: An Experimental Study of Rationality and Altruism," University of Wisconsin.
- Andreoni, J., and L. Vesterlund (1999): "Which Is the Fair Sex? Gender Differences in Altruism," *Quarterly Journal of Economics*, forthcoming.
- Arkes, H. R., L. T. Herren, and A. M. Isen (1988): "The Role of Potential Loss in the Influence of Affect on Risk-Taking Behavior," *Organizational Behavior and Human Decision Processes*, 42, 191-193.
- Arkes, H., C. Joyner, M. Pezzo, J. G. Nash, K. Siegel-Jacobs, and E. Stone (1994): "The Psychology of Windfall Gains," *Organizational Behavior and Human Decision Processes*, 59, 331-347.
- Arrow, K. J. (1985) *Applied Economics* (Collected Papers of Kenneth J. Arrow), Vol. 6, Cambridge: Harvard University Press.
- (1973): "The Theory of Discrimination" in *Discrimination in Labor Markets*. Orley Ashenfelter and Albert Rees, eds., New Jersey: Princeton University Press.
- Babcock, L., G. Loewenstein, C. F. Camerer, and S. Issacharoff (1995): "Biased Judgements of Fairness in Bargaining," *American Economic Review*, 85, 1337-1343.
- Badasyan, N., J. K. Goeree, M. Hartmann, C. A. Holt, J. Morgan, T. Rosenblat, M. Servatka, D. Yandell (2004): "Vertical Integration of Successive

- Monopolists: A Classroom Experiment," Discussion Paper, University of Virginia.
- Bagnoli, M. and M. McKee (1991): "Voluntary Contribution Games: Efficient Private Provision of Public Goods," *Economic Inquiry*, 29(2), 351-366.
- Bannerjee, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797-817.
- Ball, S. B. (1998): "Research, Teaching, and Practice in Experimental Economics: A Progress Report and Review: Review Article," *Southern Economic Journal*, 64, 772-779.
- Ball, S. B., M. H. Bazerman, and J. S. Carroll (1990): "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes*, 48, 1-22.
- Ball, S. B., and C. A. Holt (1998): "Classroom Games: Bubbles in an Asset Market," *Journal of Economic Perspectives*, 12, 207-218.
- Ball, S. B. and C. Eckel (1996) "Buying Status: Experimental Evidence on Status in Negotiation," *Psychology and Marketing*, 13(4), 381-405.
- Ball, S. B., Eckel, C. C., Grossman, P. J., and W. Zane (2001): "Status in Markets," *Quarterly Journal of Economics*, February, 101-138*.
- Banks, J. S., C. Camerer, and D. P. Porter (1994): "An Experimental Analysis of Nash Refinements in Signaling Games," *Games and Economic Behavior*, 6, 1-31.
- Basu, K. (1994): "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory," *American Economic Review*, 84(2), 391-95.
- Battalio, R. C., L. Green, and J. H. Kagel (1981): "Income-Leisure Tradeoffs of Animal Workers," *American Economic Review*, 71, 621-32.
- Battalio, R. C., J. H. Kagel, and L. Green (1979): "Labor Supply Behavior of Animal Workers: Towards an Experimental Analysis," in *Research in Experimental Economics*, Vol. 1, ed. by V. L. Smith. Greenwich, Conn.: JAI Press, 231-253.
- Battalio, Raymond C., John H. Kagel, and Komain Jiranyakul (1990) "Testing Between Alternative Models of Choice Under Uncertainty: Some Initial Results," *Journal of Risk and Uncertainty*, 3(1), 25-50.
- Battalio, R. C., J. H. Kagel, and D. N. MacDonald (1985): "Animals' Choices over Uncertain Outcomes: Some Initial Experimental Results," *American Economic Review*, 75, 597-613.
- Bazerman, M. H., and W. F. Samuelson (1983): "I Won the Auction but Don't Want the Prize," *Journal of Conflict Resolution*, 27, 618-634.
- Becker, G. M., M. H. DeGroot, and J. Marschak (1964): "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9, 226-232.

- Berg, J. E., L. A. Daley, J. W. Dickhaut, and J. R. O'Brien (1986): "Controlling Preferences for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics*, 101, 281-306.
- Berg, J. E., J. W. Dickhaut, and K. A. McCabe (1995): "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, 10, 122-142.
- Bergstrom, T., and E. Kwok (2000): "Extracting Data from Classroom Trading Pits," University of California at Santa Barbara.
- Bergstrom, T. C., and J. H. Miller (1997): *Experiments with Economic Principles*. New York: McGraw-Hill.
- Bernoulli, D. (1738): "Specimen Theoriae Novae De Mensura Sortis (Exposition on a New Theory on the Measurement of Risk," *Comentarii Academiae Scientiarum Imperialis Petropolitanae*, 5, 175-192, translated by L. Sommer in *Econometrica*, 1954, 22, 23-36.
- Bikhchandani, S., D. Hirschleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992-1026.
- Binmore, K., A. Shaked, and J. Sutton (1985): "Testing Noncooperative Bargaining Theory: A Preliminary Study," *American Economic Review*, 75, 1178-1180.
- Binswanger, H. P. (1980): "Attitudes toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics*, 62, 395-407.
- Bloomfield, R. J. (1994): "Learning a Mixed Strategy Equilibrium in the Laboratory," *Journal of Economic Behavior and Organization*, 25, 411-436.
- Blount-Lyon, S., and R. Larrick (1999): "Framing the Game," University of Chicago.
- Blume, A., D. DeJong, G. R. Neumann, and N. E. Savin (1998): "Learning in Sender-Receiver Games," University of Iowa.
- Bohm, P. (1972): "Estimating Demand for Public Goods: An Experiment," *European Economic Review*, 3, 111-130.
- Bohnet, I., and B. S. Frey (1999): "The Sound of Silence in Prisoner's Dilemma and Dictator Games," *Journal of Economic Behavior and Organization*, 38, 43-57.
- Bolton, G. E. (1991): "A Comparative Model of Bargaining: Theory and Evidence," *American Economic Review*, 81, 1096-1136.
- Bolton, G. E., and A. Ockenfels (1998): "A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, forthcoming.
- Bornstein, Gary, Tamar Kugler, and Anthony Ziegelmeyer (2002): "Individual and Group Decisions in the Centipede Game," Working Paper, Hebrew University.

- Bossaerts, P., D. Kleiman, and C. R. Plott (1998): "Experimental Tests of the Capm as a Model of Equilibrium in Financial Markets," California Institute of Technology.
- Brams, S. J., and P. C. Fishburn (1978): "Approval Voting," *American Political Science Review*, 72, 831-847.
- Brandts, J., and G. Charness (2000): "Hot Vs. Cold: Sequential Responses and Preference Stability in Experimental Games," *Experimental Economics*, 2, 227-238.
- Brandts, J., and C. A. Holt (1992): "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review*, 82, 1350-1365.
- (1993): "Adjustment Patterns and Equilibrium Selection in Experimental Signaling Games," *International Journal of Game Theory*, 22, 279-302.
- Brown-Kruse, J., and D. Hummels (1993): "Gender Effects in Laboratory Public Goods Contribution: Do Individuals Put Their Money Where Their Mouth Is?," *Journal of Economic Behavior and Organization*, 22, 255-267.
- Bryant, John (1983): "A Simple Rational Expectations Keynes-Type Model," *Quarterly Journal of Economics*, 98, 525-528.
- Camerer, C., and D. Lovo (1999): "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89, 306-318.
- Camerer, C. F. (1989): "An Experimental Test of Several Generalized Utility Theories," *Journal of Risk and Uncertainty*, 2, 61-104.
- (1995): "Individual Decision Making," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton, N.J.: Princeton University Press, 587-703.
- (2003***): *Behavioral Game Theory*, Princeton: Princeton University Press.
- Camerer, C. F., and T.-H. Ho (1999): "Experience Weighted Attraction Learning in Normal-Form Games," *Econometrica*, 67, 827-874.
- Capen, E. C., R. V. Clapp, and W. M. Campbell (1971): "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23, 641-653.
- Capra, C. M., J. K. Goeree, R. Gomez, and C. A. Holt (1999): "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89, 678-690.
- (2001): "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition," *International Economic Review*, forthcoming.
- Capra, C. M., R. Gomez, J. K. Goeree, and C. A. Holt (2000): "Predation, Asymmetric Information, and Strategic Behavior in the Classroom: An Experimental Approach to the Teaching of Industrial Organization," *International Journal of Industrial Organization*, 18, 205-225.

- Capra, C. M., and C. A. Holt (1999): "Coordination," *Southern Economic Journal*, 65, 630-636.
- (2000): "Classroom Experiments: A Prisoner's Dilemma," *Journal of Economic Education*, 21(3), 229-236.
- Cardenas, J.C., J. Stranlund, and C. Willis (2000) "Local Environmental Control and Institutional Crowding Out," *World Development*, 28(10) 1719-1733.
- Carlsson, Hans and Mattias Ganslandt (1998) "Noisy Equilibrium Selection in Coordination Games," *Economics Letters*, 60, 23-34.
- Carter, J. R., and M. D. Irons (1991): "Are Economists Different, and If So, Why?," *Journal of Economic Perspectives*, 5, 171-177.
- Cason, T. N. (1992): "Call Market Efficiency with Simple Adaptive Learning," *Economics Letters*, 40, 27-32.
- Cason, Timothy N. (1995) "Cheap Talk and Price Signaling in Laboratory Markets," *Information Economics and Policy*, 7, 183-204.
- Cason, Timothy N. (1997**) "The Opportunity for Conspiracy in Asset Markets Organized with Dealer Intermediaries," working paper, University of Southern California.
- Cason, Timothy N. and Douglas D. Davis (1995) "Price Communications in Laboratory Markets: An Experimental Investigation," *Review of Industrial Organization*, 10, 769-787.
- Cason, T. N., and D. Friedman (1996): "Price Formation in Double Auction Markets," *Journal of Economic Dynamics and Control*, 20, 1307-1337.
- (1997): "Price Formation in Single Call Markets," *Econometrica*, 65, 311-345.
- (1999): "Consumer Search and Market Power: Some Laboratory Evidence," in *Advances in Applied Microeconomics*, Vol. 8, ed. by M. Baye. Greenwich, Conn.: JAI Press, forthcoming.
- Chamberlin, E. H. (1948): "An Experimental Imperfect Market," *Journal of Political Economy*, 56, 95-108.
- Charness, G., and E. Haruvy (1999): "Altruism, Fairness, and Reciprocity: An Encompassing Approach," Pompeu Fabra University.
- Chen, Y. (1998): "Asynchronicity and Learning in Cost Sharing Mechanisms," University of Michigan.
- Chen, Y., and C. R. Plott (1996): "The Groves-Ledyard Mechanism: An Experimental Study of Institutional Design," *Journal of Public Economics*, 59, 335-364.
- Chen, Y., and T. So**nmez (1999): "An Experimental Study of House Allocation Mechanisms," University of Michigan.
- (1999): "Improving the Efficiency of on-Campus Housing: An Experimental Study," University of Michigan.

- Christie, W. G. and R. D. Huang (1995): "Following the Pied Piper: Do Individual Returns Herd Around the Market?" *Financial Analysts Journal*, 51, 31-37.
- Clauser, Laura, and Charles R. Plott (1993) "On the Anatomy of the 'Nonfacilitating' Features of the Double Auction Institution in Conspiratorial Markets," in D. Friedman, S. Genakopolos, D. Lave, and J. Rust, eds., *Double Auction Market: Institutions, Theories, and Laboratory Evidence*, Reading: Addison-Wesley.
- Coate S. and G. Loury (1993): "Will Affirmative Action Eliminate Negative Stereotypes?" *American Economic Review*, 83(5) 1220-1240.
- Cochard, F., N. Van Phu, and M. Willinger (2000): "Trust and Reciprocity in a Repeated Investment Game," Working Paper, BETA, Louis Pasteur University.
- Coller, M., and M. Williams (1999): "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 107-127.
- Conlisk, J. (1989): "Three Variants on the Allais Example," *American Economic Review*, 79, 392-407.
- Cooper, D. J., S. Garvin, and J. H. Kagel (1997): "Adaptive Learning Vs. Equilibrium Refinements in an Entry Limit Pricing Game," *Economic Journal*, 107, 553-575.
- Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1989): "Communication in the Battle of the Sexes Game: Some Experimental Results," *Rand Journal of Economics*, 20, 568-587.
- (1992): "Communication in Coordination Games," *Quarterly Journal of Economics*, 107, 739-771
- (1993): "Forward Induction in the Battle-of-the-Sexes Games," *American Economic Review*, 83, 1303-1316.
- (1996): "Cooperation without Reputation: Experimental Evidence from Prisoners' Dilemma Games," *Games and Economic Behavior* 12(2) February 187-218.
- Cooper, Russell, and Andrew John (1988): "Coordinating Coordination Failures in Keynesian Models," *The Quarterly Journal of Economics*, 103, 441-464.
- Coppinger, V. M., V. L. Smith, and J. A. Titus (1980): "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions," *Economic Inquiry*, 18, 1-22.
- Coursey, D. L., and E. A. Dyl (1986): "Trading Suspensions, Daily Price Limits, and Information Efficiency: A Laboratory Examination," in *Laboratory Market Research*, ed. by S. Moriaty. Norman, Oklahoma: University of Oklahoma, Center for Econometrics and Management Research, 153-168.

- Coursey, D. L., J. L. Hovis, and W. D. Schulze (1987): "The Disparity between Willingness to Accept and Willingness to Pay Measures of Value," *Quarterly Journal of Economics*, 102, 679-690.
- Cox, James (1999): "Trust and Reciprocity: Implications of Game Triads and Social Context," Working Paper, University of Arizona*****.
- Cox, J. C., and R. L. Oaxaca (1989): "Laboratory Experiments with a Finite Horizon Job Search Model," *Journal of Risk and Uncertainty*, 2, 301-329.
- (1996): "Testing Job Search Models: The Laboratory Approach," in *Research in Labor Economics*, Vol. 15, ed. by S. Polachek. Greenwich, Conn.: JAI Press, ***.
- (2000): "Good News and Bad News: Search from Unknown Wage Offer Distributions," *Experimental Economics*, 2, 197-225.
- Cox, J. C., B. Roberson, and V. L. Smith (1982): "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, Vol. 2, ed. by V. L. Smith. Greenwich, Conn: JAI Press, 1-43.
- Cox, J. C., V. L. Smith, and J. M. Walker (1985): "Expected Revenue in Discriminative and Uniform Price Sealed-Bid Auctions," in *Research in Experimental Economics*, Vol. 3, ed. by V. L. Smith. Greenwich, Conn.: JAI Press, 183-232.
- (1988): "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1, 61-99.
- Cox, J. C., and S. Vjollca (2001): "Risk Aversion and Expected Utility Theory: Coherence for Small and Large Scale Gambles," University of Arizona.
- Cox, J. C., and M. Walker (1998): "Learning to Play Cournot Duopoly Strategies," *Journal of Economic Behavior and Organization*, 36, 141-161.
- Crawford, V. P. (1991): "An `Evolutionary` Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games and Economic Behavior*, 3, 25-59.
- Crawford, Vincent P. (1995) "Adaptive Dynamics in Coordination Games," *Econometrica*, 63, 103-144.
- Croson, R. T. A. (1996): "Partners and Strangers Revisited," *Economics Letters*, 53, 25-32.
- Croson, R. T. A. and M. Marks (2000): "Step Returns in Threshold Public Goods" A Meta- and Experimental Analysis," *Experimental Economics*, 2(3), 239-259.
- Croson, R. T. A., J. Sundali, and E. Gold (2000): "The Gambler's Fallacy Versus the Hot Hand: Empirical Data from Casinos," University of Pennsylvania.
- Cummings, R., C. A. Holt, and S. K. Laury (2003): "The Georgia Irrigation Reduction Auction: Experiments and Implementation," forthcoming****, *Journal of Policy Analysis and Management*.

- Cummings, R. G., S. Elliott, G. Harrison, and J. Murphy (1997): "Are Hypothetical Referenda Incentive Compatible?," *Journal of Political Economy*, 105, 609-621.
- Davis, D. D. (1987): "Maximal Quality Selection and Discrimination in Employment," *Journal of Economic Behavior and Organization*, 8, 97-112.
- Davis, D. D., G. W. Harrison, and A. W. Williams (1993): "Convergence to Nonstationary Competitive Equilibria: An Experimental Analysis," *Journal of Economic Behavior and Organization*, 22, 305-326.
- Davis, D. D., and C. A. Holt (1992): "Capacity Constraints, Market Power, and Mergers in Markets with Posted Prices," *Investigaciones Economicas*, 2, 73-79.
- (1993): *Experimental Economics*. Princeton, N.J.: Princeton University Press.
- (1994): "Market Power and Mergers in Markets with Posted Prices," *RAND Journal of Economics*, 25, 467-487.
- (1995): "An Examination of the Diamond Paradox: Initial Laboratory Results," *Actas de las X Jornadas de Economia Industrial*, 67-70.
- (1996): "Consumer Search Costs and Market Performance," *Economic Inquiry*, 34, 133-151*.
- (1998): "Conspiracies and Secret Price Discounts," *Economic Journal*, 108, 736-756*.
- (1993***): "Equilibrium Cooperation in Three-Person, Choice-of-Partner Games," *Games and Economic Behavior*, forthcoming*****.
- (1996): "Price Rigidities and Institutional Variations in Markets with Posted Prices," *Economic Theory*, 9(1): 63-80.
- (1994****): "The Effects of Discounting Opportunities in Laboratory Posted-Offer Markets," forthcoming in *Economics Letters*.
- (2003a): "The Exercise of Market Power in Laboratory Experiments," forthcoming in C. Plott and V. Smith, eds., *Handbook of Experimental Economics Results*, New York: Elsevier Press.
- (2003b): "The Effects of Collusion in Laboratory Experiments," forthcoming in C. Plott and V. Smith, eds., *Handbook of Experimental Economics Results*, New York: Elsevier Press.
- Davis, D. D., and A. W. Williams (1991): "The Hayek Hypothesis in Experimental Auctions: Institutional Effects and Market Power," *Economic Inquiry*, 29, 261-274.
- Davis, Douglas D. and Bart Wilson (1998****) "The Effects of Synergies on the Exercise of Market Power," working paper, Middlebury College.
- Davis, Douglas D. and Bart Wilson (1998****) "Detecting Collusion from Bidding Patterns," draft, Middlebury College.

- Dawes, R. M. (1980): "Social Dilemmas," *Annual Review of Psychology*, 31, 169-193.
- DeGroot, M. H. (1970): *Optimal Statistical Decisions*, New York: McGraw Hill.
- DeJong, D.V., Robert Forsythe, and Russell Lundholm, "Ripoffs, Lemons, and Reputation Formation in Agency Relationships: A Laboratory Market Study," *Journal of Finance*, 1985, 40, 809-820.
- Devenow, A. and I. Welch (1996): Rational Herding in Financial Economics," *European Economic Review*, 40, 603-615.
- Diamond, Peter A. (1971), "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 158-168.
- Diekmann, A. (1985) "Volunteer's Dilemma," *Journal of Conflict Resolution*, 29, 605-610.
- Diekmann, A. (1986): "Volunteer's Dilemma: A Social Trap without a Dominant Strategy and Some Empirical Results," in *Paradoxical Effects of Social Behavior: Essays in Honor of Anatol Rapoport*, ed. by A. Diekmann, and P. Mitter. Heidelberg: Physica-Verlag, 187-197.
- (1993): "Cooperation in Asymmetric Volunteer's Dilemma Game: Theory and Experimental Evidence," *International Journal of Game Theory*, 22, 75-85.
- Diekmann, A., and P. Mitter (1986): "Paradoxical Effects of Social Behavior: Essays in Honor of Anatol Rapoport," Heidelberg and Vienna: Physica-Verlag, 341.
- Dolbear Jr., F. T., L. Lave, G. Bowman, A. Lieberman, E. Prescott, F. Reuter, and R. Sherman (1968): "Collusion in Oligopoly: An Experiment on the Effect of Numbers and Information," *Quarterly Journal of Economics*, 82, 240-259.
- Dorsey, R. and L. Razzolini (2002): "Auctions vs. Lotteries: An Experimental Comparison of Behavior in Two Strategically Isomorphic Institutions," Working Paper, University of Mississippi.
- Duffy, J., and R. Nagel (1997): "On the Robustness of Behavior in Experimental 'Beauty Contest' Games," *Economic Journal*, 107, 1687-1700*.
- Dufwenberg, M., and G. Kirchsteiger (1998): "A Theory of Sequential Reciprocity," Tilburg University.
- Durham, Yvonne (2000): "An Experimental Examination of Double Marginalization and Vertical Relationships." *Journal of Economic Behavior and Organization*, 42(2): 207-229.
- Dwyer Jr., G. P., A. W. Williams, R. C. Battalio, and T. I. Masson (1993): "Tests of Rational Expectations in a Stark Setting," *Economic Journal*, 103, 586-601.
- Eavey, C. L., and G. J. Miller (1984): "Experimental Evidence on the Fallability of the Core," *American Journal of Political Science*, 28, 570-586.

- Eckel, C. C., and P. Grossman (1998): "Are Women Less Selfish Than Men? Evidence from Dictator Games," *Economic Journal*, 108, 726-735.
- Eckel, C. C., and P. J. Grossman (1999): "Differences in the Economic Decisions of Men and Women: Experimental Evidence," in *Handbook of Experimental Economics Results*, ed. by C. R. Plott, and V. L. Smith. New York: Elsevier, forthcoming.
- Eckel, C. C., and C. A. Holt (1989): "Strategic Voting Behavior in Agenda-Controlled Committee Experiments," *American Economic Review*, 79, 763-773.
- Ellsberg, D. (1961): "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643-669.
- Ensminger, Jean (2001): Market Integration and Fairness: Evidence from Ultimatum, Dictator, and Public Goods Experiments in East Africa," Discussion Paper, California Institute of Technology.
- Erev, I., and A. Rapoport (1998): "Coordination, "Magic," and Reinforcement Learning in a Market Entry Game," *Games and Economic Behavior*, 23, 146-175.
- Erev, I., and A. E. Roth (1998): "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88, 848-881.
- Fantino, E (1998) "Behavior Analysis and Decision Making," *Journal of the Experimental Analysis of Behavior*, 69, 355-364.
- Fehr, E., and S. Gächter (1999): "Do Economic Incentives Destroy Voluntary Cooperation?," University of Zurich.
- Fehr, E., G. Kirchsteiger, and A. Riedl (1993): "Does Fairness Prevent Market Clearing? An Experimental Investigation," *Quarterly Journal of Economics*, 108, 437-459.
- (1998): "Gift Exchange and Reciprocity in Competitive Experimental Markets," *European Economic Review*, 42, 1-34.
- Fehr, E., **Klein, and K. Schmidt (2001) "Fairness, Incentives, and Contractual Incompleteness," Working Paper 72, Institute of Empirical Research in Economics, University of Zurich.
- Fehr, E., and K. Schmidt (1999): "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114, 769-816.
- Fey, M. Richard D. McKelvey, and Thomas R. Palfrey (1996): "An Experimental Study of Constant-Sum Centipede Games," *International Journal of Game Theory*, 25: 269-287.
- Fischbacher, U. (1998): "Z-Tree: A Tool for Readymade Economic Experiment," University of Zurich.
- Fischbascher, U., and C. Fong (1999): "How Much Competition Is Needed to Wipe out Fairness?," University of Zurich.

- Forrester, J. (1961): *Industrial Dynamics*, New York: MIT Press.
- Forsythe, R., J. L. Horowitz, N. E. Savin, and M. Sefton (1988): "Fairness in Simple Bargaining Games," *Games and Economic Behavior*, 6, 347-369.
- Forsythe, R., F. Nelson, G. R. Neumann, and J. Wright (1991): "Forecasting the 1988 Presidential Election: A Field Experiment," in *Research in Experimental Economics*, Vol. 4. Greenwich, Conn.: JAI Press, 1-44.
- Forsythe, R., T. R. Palfrey, and C. R. Plott (1982): "Asset Valuation in an Experimental Market," *Econometrica*, 50, 537-568.
- Fouraker, L. E., and S. Siegel (1963): *Bargaining Behavior*. New York: McGraw Hill.
- Fouraker, L. E., S. Siegel, and D. I. Harnett (1962): "An Experimental Disposition of Alternative Bilateral Monopoly Models under Conditions of Price Leadership," *Operations Research*, 10, 41-50.
- Frazen, A. (1995): "Group Size and One Shot Collective Action," *Rationality and Society*, 7, 183-200.
- Friedman, D. (1993): "How Trading Institutions Affect Financial Market Performance: Some Laboratory Evidence," *Economic Inquiry*, 31: 410-435.
- (1998): "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly," *American Economic Review*, 88, 933-946.
- Friedman, D., and S. Sunder (1994): *Experimental Methods*. Cambridge, U.K.: Cambridge University Press.
- Friedman, J. W. (1963): "Individual Behavior in Oligopolistic Markets: An Experimental Study," *Yale Economic Essays*, 3, 359-417.
- Friedman, M. and R. D. Friedman (1989) *Free to Choose*, New York: Harcourt Brace and Company.
- Fryer, R. M. (2001) *Economists' Models of Discrimination: An Analytical Survey*, Unpublished Monograph, Chicago: University of Chicago Press.
- Fryer, R. M. Goeree, J. K., and C. A. Holt (2002): "Experimenting with Discrimination," Discussion Paper, University of Virginia.
- (2004): "Classroom Games: Experience-Based Discrimination," forthcoming, *Journal of Economic Education*.
- Fudenberg, D., and D. K. Levine (1998): *Learning in Games*. Cambridge, Mass.: MIT Press.
- Gächter, S. (2000): "The Role of Rewards and Sanctions in Voluntary Cooperation Games."
- Gardner, R., E. Ostrom, and J. M. Walker (1990): "The Nature of Common Pool Resource Problems," *Rationality and Society*, 2, 335-358.
- Gillette, A., and T. H. Noe (1998): "Repeated Tender Offers," Georgia State University.

- Godby, Robert (1997): "Market Power in Emission Permit Double Auctions," *Research in Experimental Economics*, vol. 7 (edited by C. A. Holt and R. M. Isaac), Greenwich, Connecticut: JAI Press (forthcoming).
- Goeree, J. K., S. P. Anderson, and C. A. Holt (1998): "The War of Attrition with Noisy Players," in *Advances in Applied Microeconomics, Volume 7*, ed. by M. R. Baye. Greenwich, Conn.: JAI Press, 15-29.
- Goeree, J. K. and C. A. Holt (1999a): "Employment and Prices in a Simple Macro-Economy," *Southern Economic Journal*, 65, 637-647.
- (1999b): "An Experimental Study of Costly Coordination," University of Virginia.
- (1999c): "An Explanation of Anomalous Behavior in Binary-Choice Games: Entry, Voting, Public Goods, and the Volunteer's Dilemma," forthcoming, *American Political Science Review*.
- (1999e): "Rent Seeking and the Inefficiency of Non-Market Allocations," *Journal of Economic Perspectives*, 13, 217-226.
- (1999f): "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, 10564-10567.
- (2000): "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games," *European Economic Review*, 1079-1089.
- (2001): "Ten Little Treasures of Game Theory, and Ten Intuitive Contradictions," *American Economic Review*, 90(5), 1402-1422.
- (2003): "Learning in Economics Experiments," in L. Nadel, ed., *Encyclopedia of Cognitive Science*, Vol. 2, pp. 1060-1069. London: Nature Publishing Group, McMillan.
- (2003): "Coordination Games," in L. Nadel, ed., *Encyclopedia of Cognitive Science*, Vol. 2, pp. 204-208. London: Nature Publishing Group, McMillan.
- (2004): "A Model of Noisy Introspection," *Games and Economic Behavior*, 46, 365-382 (***from galley proofs).
- Goeree, J. K., C. A. Holt, and S. K. Laury (2001): "Private Costs and Public Benefits: Unraveling the Effects of Altruism and Noisy Behavior," *Journal of Public Economics*, 82, 257-278.
- (2002): "Altruism and Error in Public Goods Experiments: Implications for the Environment," forthcoming in *Frontiers in Environmental Economics*, J. List and A. de Zeeuw, eds.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey (2002): "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104(1), 247-272.

- (2003): "Risk Averse Behavior in Asymmetric Matching Pennies Games," (with Jacob Goeree and Thomas Palfrey), *Games and Economic Behavior*, 45, 97-113.
- Goodfellow, Jessica and Charles R. Plott (1990) "An Experimental Examination of the Simultaneous Determination of Input Prices and Output Prices," *Southern Economic Journal*, 56, 969-83.
- Grether, D. M. (1980): "Bayes' Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 95, 537-557.
- (1992): "Testing Bayes' Rule and the Representativeness Heuristic: Some Experimental Evidence," *Journal of Economic Behavior and Organization*, 17, 31-57.
- Grether, D. M., and C. R. Plott (1979): "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review*, 69, 623-638.
- (1984) "The Effects of Market Practices in Oligopolistic Markets: An Experimental Examination of the *Ethyl* Case," *Economic Inquiry*, 24, 479-507.
- Grether, David M., Alan Schwartz, and Louis L. Wilde (1988) "Uncertainty and Shopping Behavior: An Experimental Analysis," *Review of Economic Studies*, 55, 323-342.
- Güth, W., R. Schmittberger, and B. Schwarze (1982): "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, 3, 367-388.
- Guyer, M., and A. Rapoport (1972): "2x2 Games Played Once," *Journal of Conflict Resolution*, 16, 409-431.
- Guzik, Victor S. (2004) "Contextual Framing Effects in a Common Pool Resource Experiment," Economics Honors Thesis, Middlebury College.
- Harbaugh, W. T., and K. Krause (1998): "Children's Contributions in Public Good Experiments: The Development of Altruistic and Free-Riding Behaviors," University of Oregon.
- Harbaugh, W. T., K. Krause, and T. Berry (2000): "Garp for Kids: On the Development of Rational Choice Behavior," University of Oregon.
- Harbaugh, W. T., K. Krause, and L. Vesterlund (1998): "The Endowment Effect in Children," University of Oregon.
- (1999): "Are Adults Better Behaved Than Children? Age, Experience, and the Endowment Effect," University of Oregon.
- (2002***): "Probability Weighting in Children and Adults," ***.
- Hardin, Garrett (1968): "The Tragedy of the Commons," *Science*, 162, 1243-1248.

- Harless, D. W. (1989): "More Laboratory Evidence on the Disparity between Willingness to Pay and Compensation Demanded," *Journal of Economic Behavior and Organization*, 11, 359-379.
- Harrison, G. W. (1989): "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79, 749-762.
- Harrison, G. W., and J. Hirshleifer (1989): "An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods," *Journal of Political Economy*, 97, 201-225.
- Harrison, G. W., M. Lau, and M. Williams (1998): "Estimating Individual Discount Rates in Denmark," University of South Carolina.
- Harrison, G. W., and M. McKee (1985): "Experimental Evaluation of the Coase Theorem," *Journal of Law and Economics*, 28, 653-670.
- Haruvy, E., I. Erev, and D. Sonsino (2000): "The Medium Prize Paradox: Experimental Evidence," Technion-Israel.
- Hawkes, Kristen (1993) "Why Hunter-Gatherers Work," *Current Anthropology*, 34(4), 341-361 (with commentaries and author's reply).
- Hazlett, D. (1998): "Inflation Uncertainty and Investment in an Experimental Credit Market," Whitman College.
- Hazlett, T. W. and R. J. Michaels (1993): "The Cost of Rent-Seeking: Evidence from Cellular Telephone License Lotteries," *Southern Economic Journal*, 59(3), 425-435.
- Henrich, J. et al. (2001): "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies," *American Economic Review*, 91(2), 73-84.
- Herrnstein, R. J., and D. Prelec (1991): "Melioration: A Theory of Distributed Choice," *Journal of Economic Perspectives*, 5, 137-156.
- Hertwig, R. and A. Ortmann (2001): "Experimental Practices in Economics: A Methodological Challenge to Psychologists?" *Behavioral and Brain Sciences*, 24(3), 383-403.
- Hewett, R., C. A. Holt, G. Kosmopoulou, C. Kymn, C. X. Long, S. Mousavi, S. Sarangi (2002): "A Classroom Exercise: Voting by Ballots and Feet," Discussion Paper, University of Virginia.
- Hay, George A. and Daniel Kelley (1974) "An Empirical Survey of Price Fixing Conspiracies," *Journal of Law and Economics*, 17, 13-38.
- Hey, J. D. (1981): "Search for Rules of Search," *Journal of Economic Behavior and Organization*, 3, 65-81.
- (1987): "Still Searching," *Journal of Economic Behavior and Organization*, 8, 137-144.
- (1994): *Experimental Economics*, Heidelberg: Springer-Verlag.
- (1995): "Experimental Investigations of Errors in Decision Making under Risk," *European Economic Review*, 39, 633-640.

- Hillman, A. L. and D. Samet (1987): "Dissipation of Contestable Rents by Small Numbers of Contenders," *Public Choice*, 54(1), 63-82.
- Hoffman, E., K. McCabe, K. Shachat, and V. L. Smith (1994): "Preferences, Property Rights, and Anonymity in Bargaining Games," *Games and Economic Behavior*, 7, 346-380.
- Hoffman, E., K. McCabe, and V. L. Smith (1996): "On Expectations and Monetary Stakes in Ultimatum Games," *International Journal of Game Theory*, 25(3), 289-301.
- Hoffman, E., and M. Spitzer (1990): "The Divergence between Willingness-to-Pay and Willingness-to-Accept Measures of Values," California Institute of Technology.
- Hoffman, E., and M. L. Spitzer (1982): "The Coase Theorem: Some Experimental Tests," *Journal of Law and Economics*, 25, 73-98.
- (1985): "Entitlements, Rights and Fairness: An Experimental Examination of Subjects' Concepts of Distributive Justice," *Journal of Legal Studies*, 14, 259-297.
- (1985): "Experimental Law and Economics: An Introduction," *Columbia Law Review*, 85, 991-1036.
- Hoggatt, A. C. (1959): "An Experimental Business Game," *Behavioral Science*, 4, 192-203.
- Holt, C. A. (1985): "An Experimental Test of the Consistent-Conjectures Hypothesis," *American Economic Review*, 75, 314-325.
- (1986): "Preference Reversals and the Independence Axiom," *American Economic Review*, 76, 508-515.
- (1989): "The Exercise of Market Power in Laboratory Experiments," *Journal of Law and Economics*, 32, S107-S131.
- (1992): "ISO Probability Matching," University of Virginia.
- (1995): "Industrial Organization: A Survey of Laboratory Results," in *Handbook of Experimental Economics*, ed. by J. Kagel, and A. Roth. Princeton, N.J.: Princeton University Press, 349-443.
- (1996): "Classroom Games: Trading in a Pit Market," *Journal of Economic Perspectives*, 10, 193-203.
- (1999): "Teaching Economics with Classroom Experiments: A Symposium," *Southern Economics Journal*, 65, 603-610.
- (2003): "Economic Science: An Experimental Approach for Teaching and Research (2002 Presidential Address)," *Southern Economics Journal*, 69(4), 755-771.
- Holt, C. A., and L. R. Anderson (1999): "Agendas and Strategic Voting," *Southern Economic Journal*, 65, 622-629.

- Holt, Charles A. and Davis, Douglas D. (1990) "The Effects of Non-Binding Price Announcements on Posted Offer Markets," *Economics Letters*, 34, pp. 307-310.
- Holt, C. A., L. Langan, and A. Villamil (1986): "Market Power in Oral Double Auctions," *Economic Inquiry*, 24, 107-123.
- Holt, C. A., and S. K. Laury (1997): "Classroom Games: Voluntary Provision of a Public Good," *Journal of Economic Perspectives*, 11, 209-215.
- (1998): "Theoretical Explanations of Treatment Effects in Voluntary Contributions Experiments," in C. R. Plott and V. L. Smith, eds. *Handbook of Experimental Economics Results*. New York: Elsevier, forthcoming.
- (2001): "Varying the Scale of Financial Incentives under Real and Hypothetical Conditions," *Behavioral and Brain Sciences*, 24(3), p. 217-218.
- (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5) December, 1644-1655.
- (2002) "Further Reflections on Prospect Theory," Discussion Paper, University of Virginia.
- (2004) "Risk Aversion and Incentive Effects: New Data without Order Effects," Discussion Paper, University of Virginia.
- Holt, Charles A. and Angela Moore (2004a) "The Bullwhip Effect and Vertical Price Competition," unpublished working paper, University of Virginia.
- (2004b) "A Survey of Experimental Studies in Political Science," unpublished working paper, University of Virginia, to be presented at the Fall 2004 ESA Meetings in Tucson.
- Holt, C. A., and T. McDaniel (1998): "Experimental Economics in the Classroom," in *Teaching Undergraduate Economics: A Handbook for Instructors*, ed. by W. B. Walstad, and P. Saunders. New York: McGraw Hill, 257-268.
- Holt, C. A., and A. E. Roth (2004) "The Nash Equilibrium: A Perspective," *Proceedings of the National Academy of Sciences, U.S.A*, 101(12), 3999-4002.
- Holt, C. A., and D. Scheffman (1987): "Facilitating Practices: The Effects of Advance Notice and Best-Price Policies," *RAND Journal of Economics*, 18, 187-197.
- Holt, Charles A. and Roger Sherman (1990) "Advertising and Product Quality on Posted-Offer Experiments," *Economic Inquiry*, 28(3), 39-56.
- Holt, C. A., and R. Sherman (1994): "The Loser's Curse," *American Economic Review*, 84, 642-652.
- (1999): "Classroom Games: A. Market for Lemons," *Journal of Economic Perspectives*, 13, 205-214.

- (2000): "Risk Aversion and the Winner's Curse," Discussion Paper, University of Virginia.
- Holt, C. A., and F. Solis-Soberon (1992): "The Calculation of Equilibrium Mixed Strategies in Posted-Offer Auctions," in *Research in Experimental Economics*, Vol. 5, ed. by R. M. Isaac. Greenwich, Conn.: JAI Press, 189-229.
- Hong, James T. and Charles R. Plott (1982) "Rate Filing Policies for Inland Water Transportation: An Experimental Approach," *Bell Journal of Economics*, 13, 1-19.
- Hung, A. A. and C. R. Plott (2001) "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," *American Economic Review*, **, ****.
- Isaac, R. M., and C. A. Holt (1999): "Emissions Permit Experiments," Stamford, Conn.: JAI Press.
- Isaac, R. M., K. McCue, F., and C. R. Plott (1985): "Public Goods Provision in an Experimental Environment," *Journal of Public Economics*, 26, 51-74.
- Isaac, R. M., and C. R. Plott (1981): "The Opportunity for Conspiracy in Restraint of Trade," *Journal of Economic Behavior and Organization*, 2, 1-30.
- Isaac, R. Mark, Valerie Ramey, and Arlington Williams (1984) "The Effects of Market Organization on Conspiracies in Restraint of Trade," *Journal of Economic Behavior and Organization*, 5, 191-222.
- Isaac, R. M., and J. M. Walker (1985): "Information and Conspiracy in Sealed Bid Auctions," *Journal of Economic Behavior and Organization*, 6, 139-159.
- (1988a): "Communication and Free-Riding Behavior: The Voluntary Contributions Mechanism," *Economic Inquiry*, 26, 585-608.
- (1988b): "Group Size Hypotheses of Public Goods Provision: The Voluntary Contributions Mechanism," *Quarterly Journal of Economics*, 103, 179-200*.
- Isaac, R. M., J. M. Walker, and S. H. Thomas (1984): "Divergent Evidence on Free Riding: An Experimental Examination of Possible Explanations," *Public Choice*, 43, 113-149.
- Isaac, R. M., J. M. Walker, and A. W. Williams (1994): "Group Size and the Voluntary Provision of Public Goods: Experimental Evidence Utilizing Large Groups," *Journal of Public Economics*, 54, 1-36.
- Joseph, M. (1965): "Role Playing in Teaching Economics," *American Economic Review*, 55, 556-565.
- Kachelmeier, S. J., and M. Shehata (1992): "Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, 82, 1120-1141.

- Kagel, J. H. (1987): "Economics According to the Rats (and Pigeons Too); What Have We Learned and What Can We Hope to Learn?," in *Laboratory Experimentation in Economics: Six Points of View*, ed. by A. E. Roth. Cambridge: Cambridge University Press, 155-192.
- (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton: Princeton University Press, 501-585.
- Kagel, J. H., R. C. Battalio, and L. Green (1983): "Matching Versus Maximizing: Comments on Prelec's Paper," *Psychology Review*, 90, 380-381.
- Kagel, J. H., R. C. Battalio, H. Rachlin, and L. Green (1981): "Demand Curves for Animal Consumers," *Quarterly Journal of Economics*, 96, 1-16.
- Kagel, J. H., R. C. Battalio, and J. M. Walker (1979): "Volunteer Artifacts in Experiments in Economics: Specification of the Problem and Some Initial Data from a Small-Scale Field Experiment," in *Research in Experimental Economics, I*, ed. by V. L. Smith. Greenwich, Conn.: JAI Press, 169-197.
- Kagel, J. H., and D. Levin (1986): "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894-920.
- (1998): "Independent Private Value Multi-Unit Demand Auctions: An Experiment Comparing Uniform Price and Vickrey Auctions," Ohio State University.
- Kagel, J. H., D. Levin, R. C. Battalio, and D. J. Meyer (1989): "First-Price Common Value Auctions: Bidder Behavior and the "Winner's Curse"," *Economic Inquiry*, 27, 241-258.
- Kagel, J. H., and A. E. Roth (1995): *The Handbook of Experimental Economics*, Princeton: Princeton University Press, 721.
- Kagel, J. H., and W. Vogt (1993): "The Buyers' Bid Double Auction: Preliminary Experimental Results," in D. Friedman and J. Rust, eds., Redwood City, CA: Addison-Wesley, 285-305.
- Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991): "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies," *Journal of Economic Perspectives*, 5, 193-206.
- Kahneman, D., P. Slovic, and A. Tversky (1982): "Judgement under Uncertainty: Heuristics and Biases," Cambridge: Cambridge University Press.
- Kahneman, D., and A. Tversky (1973): "On the Psychology of Prediction," *Psychological Review*, 80, 237-251.
- Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263-291.
- Kachelmeier, S. J. and M. Shehata (1992): "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, 82, 1120-1141.

- Keser, C. (1996): "Voluntary Contributions to a Public Good When Partial Contribution Is a Dominant Strategy," *Economics Letters*, 50, 359-366.
- Ketcham, J., V. L. Smith, and A. W. Williams (1984): "A Comparison of Posted-Offer and Double-Auction Pricing Institutions," *Review of Economic Studies*, 51, 595-614.
- Keynes, J. M. (1936) *The General Theory of Employment, Interest, and Money*, London: Macmillan.
- Klemperer, Paul (2002): "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16(1), Winter, 169-189.
- Knetsch, J. L. (1989): "The Endowment Effect and Evidence of Non-Reversible Indifference Curves," *American Economic Review*, 79, 1277-1284**.
- Knez, Marc, and Colin Camerer (1994): "Creating Expectational Assets in the Laboratory: Coordination in 'Weakest-Link' Games," *Strategic Management Journal*, 15, 101-119.
- Kogut, C. A. (1990): "Consumer Search Behavior and Sunk Costs," *Journal of Economic Behavior and Organization*, 14, 381-392.
- (1992): "Recall in Consumer Search," *Journal of Economic Behavior and Organization*, 17, 141-151.
- Krause, K., and W. T. Harbaugh (1999): "Do Children Behave According to Prospect Theory," University of New Mexico.
- Kreps, David (1991) *A Course in Microeconomic Theory*, Princeton: Princeton University Press.
- Kreps, David and Jose Scheinkman (1983) "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14, 326-37.
- Krueger, A. O. (1974): "The Political Economy of the Rent-seeking Society," *American Economic Review*, 64, 291-303.
- Kruse, J. B., S. Rassenti, S. Reynolds, and V. Smith (1994): "Bertrand-Edgeworth Competition in Experimental Markets," *Econometrica*, 62, 343-372.
- Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443-477.
- Laury, S. K., and C. A. Holt (1999): "Making Money," *Journal of Economic Perspectives*, 14, 205-213.
- (1999): "Multi-Market Equilibrium and the Law of One Price," *Southern Economic Journal*, 65, 611-621.
- (2000): "Voluntary Provision of Public Goods: Experimental Results with Interior Nash Equilibria," in C. R. Plott and V. L. Smith, eds. *Handbook of Experimental Economics Results*. New York: Elsevier, forthcoming.
- Lave, L. B. (1962): "An Empirical Approach to the Prisoner's Dilemma," *Quarterly Journal of Economics*, 76, 424-436.

- (1965): "Factors Affecting Cooperation in the Prisoner's Dilemma," *Behavioral Science*, 10, 26-38.
- Ledyard, J. O. (1995): "Public Goods: A Survey of Experimental Research," in *A Handbook of Experimental Economics*, ed. by A. Roth, and J. Kagel. Princeton: Princeton University Press, 111-194.
- Ledyard, J.O., D. Porter, and R. Wessen (2000): "A Market-Based Mechanism for Allocating Space Shuttle Secondary Payload Priority," *Experimental Economics*, 2, 173-195.
- Lee, H.L., V. Padmanabhan, and S. Whang (1997a): "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, 43(4), 546-458.
- (1997b): "The Paralyzing Effect of the Bullwhip Effect in a Supply Chain," *Sloan Management Review* ***, ***.
- Levine, M. E., and C. R. Plott (1977): "Agenda Influence and Its Implications," *Virginia Law Review*, 63, 561-604.
- Lian, P., and C. R. Plott (1998): "General Equilibrium, Markets, Macroeconomics and Money in a Laboratory Experimental Environment," *Economic Theory*, 12, 21-76.
- List, J.A., and T. L. Cherry (2000): "Learning to Accept in Ultimatum Games: Evidence from an Experimental Design That Generates Low Offers," *Experimental Economics*, 3, 11-29.
- List, J.A., and D. Lucking-Reilly (2000): "Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment," *American Economic Review*, 90(4), 961-972.
- (2001): "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign," *Journal of Political Economy*, forthcoming.
- Loewenstein, G., and R. H. Thaler (1989): "Anomalies: Intertemporal Choice," *Journal of Economic Perspectives*, 3, 181-193.
- Luce, R. D. (1959): *Individual Choice Behavior*, New York: John Wiley & Sons.
- Luce, R. D., and H. Raiffa (1957): *Games and Decisions*. New York: John Wiley & Sons.
- Lucking-Reiley, David "Experimental Evidence on the Endogenous Entry of Bidders in Internet Auctions," Discussion Paper, University of Arizona, 1999.
- Lucking-Reiley, David (2000): "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First Century E-Commerce," *Journal of Economic Perspectives*, 14(3), Summer, 183-192.
- Lucking-Reiley, D., and J. A. List (2000): "Do Seed Money and Refunds Influence Charitable Giving? Experimental Evidence from a University Capital Campaign," Vanderbilt University.

- Lynch, M., R. M. Miller, C. R. Plott, and R. Porter (1986): "Product Quality, Consumer Information and 'Lemons' in Experimental Markets," in *Empirical Approaches to Consumer Protection Economics*, ed. by P. M. Ippolito, and D. T. Scheffman. Washington, D.C.: Federal Trade Commission, Bureau of Economics, 251-306.
- Mackay, C. (1995): *Extraordinary Popular Delusions and the Madness of Crowds*, Hertfordshire, England: Wordsworth Editions Ltd. (originally 1841).
- Malouf, M. W. K., and A. E. Roth (1981): "Disagreement in Bargaining: An Experimental Study," *Journal of Conflict Resolution*, 25, 329-348.
- Marwell, G., and R. E. Ames (1981): "Economists Free Ride, Does Anyone Else? Experiments on the Provision of Public Goods, IV," *Journal of Public Economics*, 15, 295-310.
- McCabe, K. A., S. J. Rassenti, and V. L. Smith (1989): "Designing 'Smart' Computer-Assisted Markets: An Experimental Auction for Gas Networks," *Journal of Political Economy*, 5, 259-283.
- (1992): "Designing Call Auction Institutions: Is Double Dutch the Best?," *Economic Journal*, 102, 9-23.
- (1993): "Designing a Uniform-Price Double Auction, An Experimental Evaluation," in D. Friedman and J. Rust, eds, *The Double Auction Market: Institutions, Theory, and Evidence*, SFI Studies in the Sciences of Complexity, Proceedings, 15, Reading, Mass.: Addison-Wesley.
- McDaniel, T., and C. Starmer (1998): "Experimental Economics and Deception: A Comment," *Journal of Economic Psychology*, 19, 403-409.
- McDowell, R. "Going Once, Going Twice...," *GMDA News*, 2(2), 1.
- McKelvey, R. D., and P. C. Ordeshook (1979): "An Experimental Test of Several Theories of Committee Decision-Making under Majority Rule," in *Applied Game Theory*, ed. by S. J. Brams, A. Schotter, and G. Schwodiauer. Wurzburg: Physica Verlag.
- (1990): "A Decade of Experimental Research on Spatial Models of Elections and Committees," in *Readings in the Spatial Theory of Voting*, ed. by J. M. Enlow, and M. J. Hinich. Cambridge, England: Cambridge University Press, 99-144.
- McKelvey, R. D., and T. R. Palfrey (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60, 803-836.
- (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- (1998): "An Experimental Study of Jury Decisions," California Institute of Technology.
- McMillan, J. (1994): "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8(3) Summer, 145-162.

- Mestelman, S., R. Moir, and R. A. Muller (1999): "A Laboratory Test of a Canadian Proposal for an Emissions Trading Program," in *Research in Experimental Economics, Volume 7, Emissions Permit Experiments*, ed. by R. M. Isaac, and C. A. Holt. Stamford, Conn.: JAI Press, 45-91.
- Mestelman, Stuart, and J. Douglas Welland (1993***) "Price Flexibility and Market Performance in Experimental Markets," *Economic Theory*, forthcoming.
- Miller, R. M., and C. R. Plott (1985): "Product Quality Signaling in Experimental Markets," *Econometrica*, 53, 837-872.
- Millner, E. L. and M. D. Pratt (1989): "An Experimental Investigation of Efficient Rent Seeking," *Public Choice*, 62 (August), 139-151.
- (1991): "Risk Aversion and Rent Seeking: An Extension and Some Experimental Evidence," *Public Choice*, 69 (February), 81-92.
- Morgan, D., A. M. Bell, and W. A. Sethares (1999): "An Experimental Study of the El Farol Problem," University of Wisconsin, Madison.
- Morgan, J., and M. Sefton (1996): "Funding Public Goods with Lotteries: An Experiment," Penn State.
- (1998): "An Experimental Investigation of Unprofitable Games," Princeton University.
- Murnighan, J. K., J. W. Kim, and A. R. Metzger (1993): "The Volunteer Dilemma," *Administrative Science Quarterly*, 38, 515-538.
- Nagel, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.
- (1997): "A Survey of Experimental Beauty-Contest Games," Universitat Pompeu Fabra.
- Nagel, Rosmarie, and Fang Fang Tang (1998): "Experimental Results on the Centipede Game in Normal Form: An Investigation of Learning," *Journal of Mathematical Psychology*, 42: 356-384.
- Nash, J. (1950): "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences, U.S.A.*, 36, 48-49.
- Nöth, M., C. Camerer, C. R. Plott, and M. Weber (1999): "Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs," California Institute of Technology.
- Nöth, M., and M. Weber (1998): "Information Aggregation with Random Ordering: Cascades and Overconfidence," University of Mannheim.
- Noussair, C. N., C. R. Plott, and R. Riezman (1997): "The Principles of Exchange Rate Determination in an International Finance Experiment," *Journal of Political Economy*, 105, 822-861.
- Noussair, C. N., C. R. Plott, and R. Riezmann (1995): "An Experimental Investigation of the Patterns of International Trade," *American Economic Review*, 85, 462-491.

- Ochs, J. (1994): "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202-217.
- (1995): "Coordination Problems," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton, N.J.: Princeton University Press, 195-249.
- Ochs, J., and A. E. Roth (1989): "An Experimental Study of Sequential Bargaining," *American Economic Review*, 79, 355-384.
- Orbell, J. M., and R. M. Dawes (1981): "Social Dilemmas," in *Progress in Applied Social Psychology*, ed. by J. M. Stephenson, and J. H. Davis. New York: Wiley, 117-133.
- Ordeshook, P. (1986): *Game Theory and Political Theory*. London: Cambridge University Press.
- Ortmann, A., J. Fitzgerald, and C. Boeing (2000): "Trust, Reciprocity, and Social History: A Re-Examination," *Experimental Economics*, 3, 81-100.
- Ostrom, E., and R. Gardner (1993): "Coping with Asymmetries in the Commons: Self-Governing Irrigation Systems Can Work," *Journal of Economic Perspectives*, 7(4), 93-112.
- Ostrom, E., R. Gardner, and J. K. Walker (1994): *Rules, Games, and Common-Pool Resources*. Ann Arbor: University of Michigan Press.
- Ostrom, E., and J. K. Walker (1991): "Communication in a Commons: Cooperation without External Enforcement," in *Laboratory Research in Political Economy*, ed. by T. Palfrey. Ann Arbor: University of Michigan Press, 289-322.
- Palfrey, T., and J. Prisbrey (1997): "Anomalous Behavior in Linear Public Goods Experiments: How Much and Why?," *American Economic Review*, 87, 829-846.
- Palfrey, T. R., and H. Rosenthal (1988): "Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism," *Journal of Public Economics*, 35, 309-332.
- Peters, W. (1971) *A Class Divided*, **Doubleday and Company.
- Phelps, E. (1972): "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62, 659-661.
- Plot, C.R. (1979): "The Application of Laboratory Experimental Methods to the Public Choice," in C.S. Russell, *Collective Decision Making: Applications from Public Choice Theory*, Baltimore: Johns Hopkins Press, 137-160.
- (1982): "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature*, 20, 1485-1527.
- (1983): "Externalities and Corrective Policies in Experimental Markets," *Economic Journal*, 93, 106-127.
- (1986): "The Posted-Offer Trading Institution," *Science*, 232, 732-738.

- (1989): "An Updated Review of Industrial Organization: Applications of Experimental Methods," in *Handbook of Industrial Organization, Volume II*, ed. by R. Schmalensee, and R. D. Willig. Amsterdam: Elsevier Science, 1111-1176.
- (1991): "Will Economics Become an Experimental Science?," *Southern Economic Journal*, 57, 901-919.
- Plott, C. R., and M. E. Levine (1978): "A Model of Agenda Influence on Committee Decisions," *American Economic Review*, 68, 146-160.
- Plott, C. R., and V. L. Smith (1978): "An Experimental Examination of Two Exchange Institutions," *Review of Economic Studies*, 45, 133-153.
- Plott, C. R., J. Wit, and W. C. Yang (1997): "Paramutuel Betting Markets as Information Aggregation Devices: Experimental Results," Working Paper, California Institute of Technology.
- Ponti, G. (2002): "Cycles of Learning in the Centipede Game," *Games and Economic Behavior*, 30: 115-141.
- Porter, Robert H. and J. Douglas Zona (1993) "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 101, 518-538.
- Post, E. (1927): *Etiquette in Society, in Business, in Politics, and at Home*, New York: Funk and Wagnalls.
- Potters, J., C. G. de Vries, and F. van Winden (1998): "An Experimental Examination of Rational Rent-Seeking," *European Journal of Political Economy*, 14, 783-800.
- Prelec, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497-527.
- Quiggin, J. (1993): *Generalized Expected Utility Theory: The Rank-Dependent Model*. Dordrecht: Kluwer.
- Rabin, M. (1993): "Incorporating Fairness into Game Theory and Economics," *American Economic Review*, 83, 1281-1302.
- (2000): "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281-1292.
- Rapoport, A., and P. Dale (1966): "Models to Prisoner's Dilemma," *Journal of Mathematical Psychology*, 3, 269-286.
- Rapoport, A., D. A. Seale, I. Erev, and J. A. Sundali (1998): "Equilibrium Play in Large Market Entry Games," *Management Science*, 44, 129*-141.
- Rassenti, S., V. L. Smith, and B. Wilson (2000): "Market Power in Electricity Markets," University of Arizona.
- Reynolds, S. S. (1999): "Durable Goods Monopoly: Laboratory Market and Bargaining Experiments," *RAND Journal of Economics*, forthcoming.
- Reynolds, Stanley S. and Bart J. Wilson (1997): "Market Power, Price Markups and Capacity Investment under Uncertain Demand," manuscript, University of Arizona.

- Rigdon, M., K. McCabe, and V. Smith (2000): "Positive Reciprocity and Intentions in Trust Games," University of Arizona.
- Romer, David (1996): *Advanced Macroeconomics*, New York: McGraw-Hill.
- Rosenthal, R. W. (1982): "Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox," *Journal of Economic Theory*, 25, 92–100.
- (1989): "A Bounded Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273–292.
- Roth, A. E. (1995): "Bargaining Experiments," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton: Princeton University Press, 253–348.
- Roth, A. E. and M. W. K. Malouf (1979): "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 574–594.
- Roth, A. E., and J. K. Murnighan (1982): "The Role of Information in Bargaining: An Experimental Study," *Econometrica*, 50, 1123–1142.
- Roth, A. E., J. K. Murnighan, and F. Schoumaker (1988): "The Deadline Effect in Bargaining: Some Experimental Evidence," *American Economic Review*, 78, 806–823.
- Roth, A. E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir (1991): "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review*, 81, 1068–1095.
- Rubinstein, Ariel (1991) "Comments on the Interpretation of Game Theory," *Econometrica*, 59, 909–924.
- Saijo, T., and H. Nakamura (1995): "The "Spite" Dilemma in Voluntary Contribution Mechanism Experiments," *Journal of Conflict and Resolution*, 39, 535–360.
- Samuelson, W., and R. Zeckhauser (1988): "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1, 7–59.
- Sauermann, H., and R. Selten (1959): "Ein Oligopolexperiment," *Zeitschrift für die Gesamte Staatswissenschaft*, 115, 427–471.
- Schelling, T. C. (1960): *The Strategy of Conflict*. Cambridge: Harvard University Press.
- Schotter, A., and Y. M. Braunstein (1981): "Economic Search: An Experimental Study," *Economic Inquiry*, 19, 1–25.
- Schotter, A., and K. Weigelt (1992): "Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results," *Quarterly Journal of Economics*, 107, 511–539.
- Schotter, A., K. Weigelt, and C. Wilson (1994): "A Laboratory Investigation of Multiperson Rationality and Presentation Effects," *Games and Economic Behavior*, 6, 445–468.

- Schram, A., and J. Sonnemans (1996): "Voter Turnout as a Participation Game: An Experimental Investigation," *International Journal of Game Theory*, 25, 385-406.
- (1996): "Why People Vote: Experimental Evidence," *Journal of Economic Psychology*, 17, 417-442.
- Sefton, M. (1999): "A Model of Behavior in Coordination Game Experiments," *Experimental Economics*, 2, 151-164.
- Selten, Reinhard (1965) "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," parts I - II, *Zeitschrift für die Gesamte Staatswissenschaft*, 121, 301-324 and 667-689.
- (1990): "Bounded Rationality," *Journal of Institutional and Theoretical Economics*, 146, 649-658.
- (1991): "Anticipatory Learning in Games," in *Game Equilibrium Models, Vol. 1*, ed. by R. Selten. New York: Springer-Verlag, Chapter 3**.
- Selten, R., and J. Bucta (1999): "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," in *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, ed. by I. E. D. Burescu, I. Erev, and R. Zwick. Hillsdale N.J.: Erlbaum Association, 101-116.
- Selten, R., K. Sadreih, and K. Abbink (1999): "Money Does Not Induce Risk Neutral Behavior, but Binary Lotteries Do Even Worse," *Theory and Decisions*, 46, 211-249.
- Selten, R., and R. Stoecker (1986): "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach," *Journal of Economic Behavior and Organization*, 7, 47-70.
- Sherman, R. (1971): "Experimental Oligopoly," *Kyklos*, 24, 30-49.
- Sherstyuk, K. (1999): "Collusion without Conspiracy: An Experimental Study of One-Sided Auctions," *Experimental Economics*, 2, 59-75.
- Shogren, J. F., S. Y. Shin, D. J. Hayes, and J. B. Kliebenstein (1994): "Resolving Differences in Willingness to Pay and Willingness to Accept," *American Economic Review*, 84, 255-270.
- Shubik, M. (1971): "The Dollar Auction Game: A Paradox in Non-Cooperative Behavior and Escalation," *Journal of Conflict Resolution*, 15, 109-111.
- Siegel, S. (1956): *Nonparametric Statistics for the Behavioral Sciences*. New York: McGraw-Hill.
- (1961): "Decision Making and Learning under Varying Conditions of Reinforcement," *Annals of the New York Academy of Sciences*, 89, 766-783.
- Siegel, S., and J. Castellan Jr. (1988): *Nonparametric Statistics for the Behavioral Sciences*. New York: McGraw-Hill.
- Siegel, S., and L. E. Fournaker (1960): *Bargaining and Group Decision Making: Experiments in Bilateral Monopoly*. New York: McGraw-Hill.

- Siegel, S., and D. A. Goldstein (1959): "Decision-Making Behavior in a Two-Choice Uncertain Outcome Situation," *Journal of Experimental Psychology*, 57, 37-42.
- Siegel, S., A. Siegel, and J. Andrews (1964): *Choice, Strategy, and Utility*. New York: McGraw-Hill.
- Simon, H. (1955): "A Behavioral Model of Rational Choices," *Quarterly Journal of Economics*, 69, 99-118.
- Simon, H. A. (1993): "Altruism and Economics," *American Economic Review*, 83, 156-161.
- Slonim, R., and A. E. Roth (1998): "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic," *Econometrica*, 66, 569-596.
- Smith, Adam (1976, originally 1776) *The Wealth of Nations*, E. Cannan, ed., Chicago: University of Chicago Press.
- Smith, V. L. (1962): "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, 70, 111-137.
- (1964): "The Effect of Market Organization on Competitive Equilibrium," *Quarterly Journal of Economics*, 78, 181-201.
- (1976): "Experimental Economics: Induced Value Theory," *American Economic Review*, 66, 274-279.
- (1979): *Research in Experimental Economics, Vol. 1*. Greenwich, Conn.: JAI Press.
- (1981): "An Empirical Study of Decentralized Institutions of Monopoly Restraint," in J. Quirk and G. Horwich, eds., *Essays in Contemporary Fields of Economics in Honor of E.T. Weiler, 1914-1979* (Purdue University Press, West Lafayette), 83-106.
- (1982): "Markets as Economizers of Information: Experimental Examination of the 'Hayek Hypothesis'," *Economic Inquiry*, 20, 165-179.
- (1982b) "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923-55.
- Smith, V. L., G. L. Suchanek, and A. W. Williams (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119-1151.
- Smith, V. L., and J. M. Walker (1993): "Monetary Rewards and Decision Cost in Experimental Economics," *Economic Inquiry*, 31, 245-261.
- Smith, V. L., and A. W. Williams (1982): "The Effects of Rent Asymmetries in Experimental Auction Markets," *Journal of Economic Behavior and Organization*, 3, 99-116.
- (1988) "The Boundaries of Competitive Price Theory: Convergence, Expectations, and Transactions Costs," in L. Green and J. Kagel, eds., *Advances in Behavioral Economics*, vol. 2. Norwood, N.J.: Ablex Publishing.

- Smith, V. L., A. W. Williams, W.K. Bratton, and Vannoni (1982): "Competitive Market Institutions: Double Auctions vs. Sealed Bid-Offer Auctions," *American Economic Review*, 72: 58-77.
- Sonnemans, J., C. Hommes, and H. van de Velden (1999): "Expectation Formation: Group Experiments," University of Amsterdam.
- Sonnemans, J., F. van Dijk, and F. van Winden (2000): "Group Formation in a Public Good Experiment: On the Dynamics of Social Ties Structures," University of Amsterdam.
- Stahl, D. O., and P. W. Wilson (1994): "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25, 309-327.
- Starmer, C., and R. Sugden (1989): "Violations of the Independence Axiom in Common Ratio Problems: An Experimental Test of Some Competing Hypotheses," *Annals of Operations Research*, 19, 79-102.
- (1991): "Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation," *American Economic Review*, 81, 971-978.
- Sterman, John D. (1989): "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment," *Management Science*, 35(3), 321-339.
- Steiglitz, Ken and Daniel Shapiro (1998): "Simulating the Madness of Crowds: Price Bubbles in an Auction-Mediated Robot Market," *Computational Economics*, 12: 35-59.
- Straub, Paul G. (1995): "Risk Dominance and Coordination Failures in Static Games," *Quarterly Review of Economics and Finance*, 35(4), Winter 1995, 339-363.
- Sugden, R. (1984): "Reciprocity: The Supply of Public Goods through Voluntary Contributions," *Economic Journal*, 94, 777-787.
- Sunder, S. (1995): "Experimental Asset Markets: A Survey," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton, N.J.: Princeton University Press, 445-500.
- Tajfel, H. (1970): "Experiments in Inter-Group Discrimination," *Scientific American*, November, 96-102.
- Thaler, R. (1989): "Anomalies: The Ultimatum Game," *Journal of Economic Perspectives*, 2, 195-206.
- Thaler, R. H. (1988): "Anomalies: The Winner's Curse," *Journal of Economic Perspectives*, 2, 191-202.
- (1992): *The Winners Curse*. New York: Free Press.
- Tietz, R., W. Albers, and R. Selten (1988): *Experimental Games and Markets*. Berlin: Springer-Verlag.

- Tullock, G. (1967): "The Welfare Costs of Tariffs, Monopolies, and Thefts," *Western Economic Journal*, June 1967, 5(3), 224-232.
- (1975): "The Paradox of Not Voting for Oneself," *American Political Science Review*, 69, 919.
- (1999): "Non-Prisoner's Dilemma," *Journal of Economic Behavior and Organization*, 39, 455-458.
- Tversky, A., and D. Kahneman (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124-1131.
- (1991): "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *Quarterly Journal of Economics*, 106, 1039-1061.
- (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.
- Tversky, A., and R. H. Thaler (1990): "Anomalies: Preference Reversals," *Journal of Economic Perspectives*, 4, 201-211.
- Van Boening, M., A. W. Williams, and S. Lamaster (1993): "Price Bubbles and Crashes in Experimental Call Markets," *Economics Letters*, 41, 179-185.
- van Damme, E. (1998): "Auctioning Dutch Airwaves," Tilburg University.
- van Dijk, F., J. Sonnemans, and V. van Winden (1997): "Social Ties in a Public Good Experiment," Discussion Paper, University of Amsterdam.
- Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1990): "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80, 234-248.
- Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1991): "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics*, 91, 885-910.
- Van Huyck, John B., Raymond C. Battalio, and Frederick Rankin (2002) "On the Origin of Convention: Evidence from Coordination Games," working paper, Texas A&M University.
- Van Huyck, J. B., R. C. Battalio, S. Mathur, and P. Van Huyck (1995): "On the Origin of Convention: Evidence from Symmetric Bargaining Games," *International Journal of Game Theory*, 24, 187-212.
- Van Huyck, J. B., J. P. Cook, and R. C. Battalio (1997): "Adaptive Behavior and Coordination Failure," *Journal of Economic Behavior and Organization*, 32, 483-503.
- van Winden, F., and A. Riedl (2000): "An Experimental Investigation of Taxation and Unemployment in a Closed and Small Open Economy," University of Amsterdam.
- Vaughan, G. M., Tajfel, H., and J. Williams (1981): "Bias in Reward Allocation in an Intergroup and an Interpersonal Context," *Social Psychology Quarterly*, 44(1) 37-42.

- Vickrey, W. (1962): "Counterspeculation and Competitive Sealed Tenders," *Journal of Finance*, 16(1), 8-37.
- von Neumann, J., and O. Morgenstern (1944): *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Vulkan, N. (1998) "An Economist's Perspective on Probability Matching," Working Paper, University of Bristol.
- Walker, J. M., R. Gardner, and E. Ostrom (1990): "Rent Dissipation in a Limited-Access Common-Pool Resource: Experimental Evidence," *Journal of Environmental Economics and Management*, 19, 203-211.
- Walker, J. M., V. L. Smith, and J. C. Cox (1987): "Bidding Behavior in First-Price Sealed-Bid Auctions: Use of Computerized Nash Competitors," *Economics Letters*, 23, 239-244.
- (1990): "Inducing Risk-Neutral Preferences: An Examination in a Controlled Market Environment," *Journal of Risk and Uncertainty*, 3, 5-24.
- Walker, M., and J. Wooders (1999): "Minimax Play at Wimbledon," University of Arizona.
- Weber, M., and C. Camerer (1998): "The Disposition Effect in Securities Trading," *Journal of Economic Behavior and Organization*, 33, 167-184.
- Weber, R., and C. Camerer (2000): "Why Mergers Fail? Evidence from Experiments," California Institute of Technology.
- Welch, I. (1992): "Sequential Sales, Learning, and Cascades," *Journal of Finance*, 47, 695-732.
- Wilcox, N. T. (1993): "Lottery Choice: Incentives, Complexity and Decision Time," *Economic Journal*, 103, 1397-1417.
- Williams, A. W. (1979): "Intertemporal Competitive Equilibrium: On Further Experimental Results," in *Research in Experimental Economics*, Vol. 1, ed. by V. L. Smith. Greenwich, Conn.: JAI Press, 255-278.
- (1980): "Computerized Double-Auction Markets: Some Initial Experimental Results," *Journal of Business*, 53, 235-258.
- (1999): "Price Bubbles in Large Financial Asset Markets," in *Handbook of Experimental Economics Results*, ed. by C. R. Plott, and V. L. Smith. New York: Elsevier Press, forthcoming.
- Willinger, M., P. Nguyen, and F. Cocharde (2000): "Trust and Reciprocity in a Repeated Investment Game," University Louis Pasteur.
- Wilson, Bart J. (1997): "What Collusion? Unilateral Market Power as a Catalyst for Countercyclical Pricing," Manuscript, University of Arizona.
- Wilson, R. B. (1969): "Competitive Bidding with Disparate Options," *Management Science*, 15, 446-448.

- Wilson, R. K. (1986): "Forward and Backward Agenda Procedures: Committee Experiments on Structurally Induced Equilibrium," *Journal of Politics*, 48, 390-409.
- Wu, G., and R. Gonzalez (1996): "Curvature of the Probability Weighting Function," *Management Science*, 42, 1676-1690.
- Yoder, R. D. (1986): "The Performance of Farmer-Managed Irrigation Systems in the Hills of Nepal," Ph.D. Dissertation, Cornell University.
- Zauner, K. G. (1999): "A Payoff Uncertainty Explanation of Results in Experimental Centipede Games," *Games and Economic Behavior*, 26, 157-185.
- Ziegelmeyer, A. (2000): "Endogenous Information Cascades and Beliefs," Louis Pasteur University, Strasbourg.

Missing references:

add old FTC Holt and Sherman on quality

p. 103 Thaler and Rabin risk aversion not in refs. And add cite to paper at
Brown that finds error in Rabin paper

Composti not in refs

p. 169 Vickrey (1962) check date, is in refs and in text as 62, is date 61??

p. 190 Wilson (1967, 1968) not in refs