

GAMES AND INFORMATION, FOURTH EDITION

An Introduction to Game Theory

Eric Rasmusen

Basil Blackwell

Contents¹
(starred sections are less important)

List of Figures

List of Tables

Preface

Contents and Purpose

Changes in the Second Edition

Changes in the Third Edition

Using the Book

The Level of Mathematics

Other Books

Contact Information

Acknowledgements

Introduction

History

Game Theory's Method

Exemplifying Theory

This Book's Style

Notes

PART 1: GAME THEORY

1 The Rules of the Game

1.1 Definitions

1.2 Dominant Strategies: The Prisoner's Dilemma

1.3 Iterated Dominance: The Battle of the Bismarck Sea

1.4 Nash Equilibrium: Boxed Pigs, The Battle of the Sexes, and Ranked Coordination

1.5 Focal Points

Notes

Problems

¹xxx February 2, 2000. Eric Rasmusen, Erasmuse@indiana.edu. December 12, 2003. Footnotes starting with xxx are the author's notes to himself. Comments are welcomed.

2 Information

2.1 The Strategic and Extensive Forms of a Game

2.2 Information Sets

2.3 Perfect, Certain, Symmetric, and Complete Information

2.4 The Harsanyi Transformation and Bayesian Games

2.5 Example: The Png Settlement Game

Notes

Problems

3 Mixed and Continuous Strategies

3.1 Mixed Strategies: The Welfare Game

3.2 Chicken, The War of Attrition, and Correlated Strategies

3.3 Mixed Strategies with General Parameters and N Players: The Civic Duty Game

3.4 Different Uses of Mixing and Randomizing: Minimax and the Auditing Game

3.5 Continuous Strategies: The Cournot Game

3.6 Continuous Strategies: The Bertrand Game, Strategic Complements, and Strategic Substitutes

3.7 Existence of Equilibrium

Notes

Problems

4 Dynamic Games with Symmetric Information

4.1 Subgame Perfectness

4.2 An Example of Perfectness: Entry Deterrence I

4.3 Credible Threats, Sunk Costs, and the Open-Set Problem in the Game of Nuisance Suits

*4.4 Recoordination to Pareto-dominant Equilibria in Subgames: Pareto Perfection

Notes

Problems

5 Reputation and Repeated Games with Symmetric Information

5.1 Finitely Repeated Games and the Chainstore Paradox

5.2 Infinitely Repeated Games, Minimax Punishments, and the Folk Theorem

5.3 Reputation: the One-sided Prisoner's Dilemma

5.4 Product Quality in an Infinitely Repeated Game

- *5.5 Markov Equilibria and Overlapping Generations in the Game of Customer Switching Costs
- *5.6 Evolutionary Equilibrium: The Hawk-Dove Game
 - Notes
 - Problems

6 Dynamic Games with Incomplete Information

- 6.1 Perfect Bayesian Equilibrium: Entry Deterrence II and III
- 6.2 Refining Perfect Bayesian Equilibrium: the PhD Admissions Game
- 6.3 The Importance of Common Knowledge: Entry Deterrence IV and V
- 6.4 Incomplete Information in the Repeated Prisoner's Dilemma: The Gang of Four Model
- 6.5 The Axelrod Tournament
- *6.6 Credit and the Age of the Firm: The Diamond Model
 - Notes
 - Problems

PART 2: ASYMMETRIC INFORMATION

7 Moral Hazard: Hidden Actions

- 7.1 Categories of Asymmetric Information Models
- 7.2 A Principal-Agent Model: The Production Game
- 7.3 The Incentive Compatibility, Participation, and Competition Constraints
- 7.4 Optimal Contracts: The Broadway Game
 - Notes
 - Problems

8 Further Topics in Moral Hazard

- 8.1 Efficiency Wages
- 8.2 Tournaments
- 8.3 Institutions and Agency Problems
- *8.4 Renegotiation: the Repossession Game
- *8.5 State-space Diagrams: Insurance Games I and II
- *8.6 Joint Production by Many Agents: the Holmstrom Teams Model
 - Notes
 - Problems

9 Adverse Selection

- 9.1 Introduction: Production Game VI
- 9.2 Adverse Selection under Certainty: Lemons I and II
- 9.3 Heterogeneous Tastes: Lemons III and IV
- 9.4 Adverse Selection under Uncertainty: Insurance Game III
- *9.5 Market Microstructure
- *9.6 A Variety of Applications
 - Notes
 - Problems
- 10 Mechanism Design in Adverse Selection and in Moral Hazard with Hidden Information
 - 10.1 The Revelation Principle and Moral Hazard with Hidden Knowledge
 - 10.2 An Example of Moral Hazard with Hidden Knowledge: the Salesman Game
 - *10.3 Price Discrimination
 - *10.4 Rate-of-return Regulation and Government Procurement
 - *10.5 The Groves Mechanism
 - Notes
 - Problems
- 11 Signalling
 - 11.1 The Informed Player Moves First: Signalling
 - 11.2 Variants on the Signalling Model of Education
 - 11.3 General Comments on Signalling in Education
 - 11.4 The Informed Player Moves Second: Screening
 - *11.5 Two Signals: the Game of Underpricing New Stock Issues
 - *11.6 Signal Jamming and Limit Pricing
 - Notes
 - Problems
- PART 3: APPLICATIONS
- 12 Bargaining
 - 12.1 The Basic Bargaining Problem: Splitting a Pie
 - 12.2 The Nash Bargaining Solution
 - 12.3 Alternating Offers over Finite Time
 - 12.4 Alternating Offers over Infinite Time
 - 12.5 Incomplete Information

*12.6 Setting up a Way to Bargain: the Myerson-Satterthwaite Mechanism

Notes

Problems

13 Auctions

13.1 Auction Classification and Private-Value Strategies

13.2 Comparing Auction Rules

13.3 Risk and Uncertainty over Values

13.4 Common-value Auctions and the Winner's Curse

13.5 Information in Common-value Auctions

Notes

Problems

14 Pricing

14.1 Quantities as Strategies: Cournot Equilibrium Revisited

14.2 Prices as Strategies

14.3 Location Models

*14.4 Comparative Statics and Supermodular Games

*14.5 Durable Monopoly

Notes

Problems

*A Mathematical Appendix

*A.1 Notation

*A.2 The Greek Alphabet

*A.3 Glossary

*A.4 Formulas and Functions

*A.5 Probability Distributions

*A.6 Supermodularity

*A.7 Fixed Point Theorems

*A.8 Genericity

*A.9 Discounting

*A.10 Risk

References and Name Index

Subject Index

xxx September 6, 1999; February 2, 2000. February 9, 2000. May 24, 2002. Ariel Kemper August 6, 2003. Eric Rasmusen, Erasmuse@indiana.edu; Footnotes starting with xxx are the author's notes to himself. Comments are welcomed.

Preface

Contents and Purpose

This book is about noncooperative game theory and asymmetric information. In the Introduction, I will say why I think these subjects are important, but here in the Preface I will try to help you decide whether this is the appropriate book to read if they do interest you.

I write as an applied theoretical economist, not as a game theorist, and readers in anthropology, law, physics, accounting, and management science have helped me to be aware of the provincialisms of economics and game theory. My aim is to present the game theory and information economics that currently exist in journal articles and oral tradition in a way that shows how to build simple models using a standard format. Journal articles are more complicated and less clear than seems necessary in retrospect; precisely because it is original, even the discoverer rarely understands a truly novel idea. After a few dozen successor articles have appeared, we all understand it and marvel at its simplicity. But journal editors are unreceptive to new articles that admit to containing exactly the same idea as old articles, just presented more clearly. At best, the clarification is hidden in some new article's introduction or condensed to a paragraph in a survey. Students, who find every idea as complex as the originators of the ideas did when they were new, must learn either from the confused original articles or the oral tradition of a top economics department. This book tries to help.

Changes in the Second Edition, 1994

By now, just a few years later after the First Edition, those trying to learn game theory have more to help them than just this book, and I will list a number of excellent books below. I have also thoroughly revised *Games and Information*. George Stigler used to say that it was a great pity Alfred Marshall spent so much time on the eight editions of *Principles of Economics* that appeared between 1890 and 1920, given the opportunity cost of the other books he might have written. I am no Marshall, so I have been willing to sacrifice a Rasmusen article or two for this new edition, though I doubt I will keep it up till 2019.

What I have done for the Second Edition is to add a number of new topics, increase the number of exercises (and provide detailed answers), update the references, change the terminology here and there, and rework the entire book for clarity. A book, like a poem, is never finished, only abandoned (which is itself a good example of a fundamental economic principle). The one section I have dropped is the somewhat obtrusive discussion of existence theorems; I recommend Fudenberg & Tirole (1991a) on that subject. The new

topics include auditing games, nuisance suits, recoordination in equilibria, renegotiation in contracts, supermodularity, signal jamming, market microstructure, and government procurement. The discussion of moral hazard has been reorganized. The total number of chapters has increased by two, the topics of repeated games and entry having been given their own chapters.

Changes in the Third Edition, 2001

Besides numerous minor changes in wording, I have added new material and reorganized some sections of the book.

The new topics are 10.3 “Price Discrimination”; 12.6 “Setting up a Way to Bargain: The Myerson-Satterthwaite Mechanism”; 13.3 “Risk and Uncertainty over Values” (for private-value auctions); A.7 “Fixed-Point Theorems”; and A.8 “Genericity”.

To accommodate the additions, I have dropped 9.5 “Other Equilibrium Concepts: Wilson Equilibrium and Reactive Equilibrium” (which is still available on the book’s website), and Appendix A, “Answers to Odd-Numbered Problems”. These answers are very important, but I have moved them to the website because most readers who care to look at them will have web access and problem answers are peculiarly in need of updating. Ideally, I would like to discuss all likely wrong answers as well as the right answers, but I learn the wrong answers only slowly, with the help of new generations of students.

Chapter 10, “Mechanism Design in Adverse Selection and in Moral Hazard with Hidden Information”, is new. It includes two sections from chapter 8 (8.1 “Pooling versus Separating Equilibrium and the Revelation Principle” is now section 10.1; 8.2 “An Example of Moral Hazard with Hidden Knowledge: the Salesman Game” is now section 10.2) and one from chapter 9 (9.6 “The Groves Mechanism” is now section 10.5).

Chapter 15 “The New Industrial Organization”, has been eliminated and its sections reallocated. Section 15.1 “Why Established Firms Pay Less for Capital: The Diamond Model” is now section 6.6; Section 15.2 “Takeovers and Greenmail” remains section 15.2; section 15.3 “Market Microstructure and the Kyle Model” is now section 9.5; and section 15.4 “Rate-of-return Regulation and Government Procurement” is now section 10.4.

Topics that have been extensively reorganized or rewritten include 14.2 “Prices as Strategies”; 14.3 “Location Models”; the Mathematical Appendix, and the Bibliography. Section 4.5 “Discounting” is now in the Mathematical Appendix; 4.6 “Evolutionary Equilibrium: The Hawk-Dove Game” is now section 5.6; 7.5 “State-space Diagrams: Insurance Games I and II” is now section 8.5 and the sections in Chapter 8 are reordered; 14.2 “Signal Jamming: Limit Pricing” is now section 11.6. I have recast 1.1 “Definitions”, taking out the OPEC Game and using an entry deterrence game instead, to illustrate the difference between game theory and decision theory. Every other chapter has also been revised in minor ways.

Some readers preferred the First Edition to the Second because they thought the extra topics in the Second Edition made it more difficult to cover. To help with this problem, I have now starred the sections that I think are skippable. For reference, I continue to have

those sections close to where the subjects are introduced.

The two most novel features of the book are not contained within its covers. One is the website, at

[Http://php.indiana.edu/~erasmuse/GI/index.html](http://php.indiana.edu/~erasmuse/GI/index.html)

The website includes answers to the odd-numbered problems, new questions and answers, errata, files from my own teaching suitable for making overheads, and anything else I think might be useful to readers of this book.

The second new feature is a Reader— a prettified version of the course packet I use when I teach this material. This is available from Blackwell Publishers, and contains scholarly articles, news clippings, and cartoons arranged to correspond with the chapters of the book. I have tried especially to include material that is somewhat obscure or hard to locate, rather than just a collection of classic articles from leading journals.

If there is a fourth edition, three things I might add are (1) a long discussion of strategic complements and substitutes in chapter 14, or perhaps even as a separate chapter; (2) Holmstrom & Milgrom's 1987 article on linear contracts; and (3) Holmstrom & Milgrom's 1991 article on multi-task agency. Readers who agree, let me know and perhaps I'll post notes on these topics on the website.

Using the Book

The book is divided into three parts: Part I on game theory; Part II on information economics; and Part III on applications to particular subjects. Parts I and II, but not Part III, are ordered sets of chapters.

Part I by itself would be appropriate for a course on game theory, and sections from Part III could be added for illustration. If students are already familiar with basic game theory, Part II can be used for a course on information economics. The entire book would be useful as a secondary text for a course on industrial organization. I teach material from every chapter in a semester-long course for first- and second-year doctoral students at Indiana University's Kelley School of Business, including more or fewer chapter sections depending on the progress of the class.

Exercises and notes follow the chapters. It is useful to supplement a book like this with original articles, but I leave it to my readers or their instructors to follow up on the topics that interest them rather than recommending particular readings. I also recommend that readers try attending a seminar presentation of current research on some topic from the book; while most of the seminar may be incomprehensible, there is a real thrill in hearing someone attack the speaker with "Are you sure that equilibrium is perfect?" after just learning the previous week what "perfect" means.

Some of the exercises at the end of each chapter put slight twists on concepts in the text while others introduce new concepts. Answers to odd-numbered questions are given at the website. I particularly recommend working through the problems for those trying to learn this material without an instructor.

The endnotes to each chapter include substantive material as well as recommendations for further reading. Unlike the notes in many books, they are not meant to be skipped, since many of them are important but tangential, and some qualify statements in the main text. Less important notes supply additional examples or list technical results for reference. A mathematical appendix at the end of the book supplies technical references, defines certain mathematical terms, and lists some items for reference even though they are not used in the main text.

The Level of Mathematics

In surveying the prefaces of previous books on game theory, I see that advising readers how much mathematical background they need exposes an author to charges of being out of touch with reality. The mathematical level here is about the same as in Luce & Raiffa (1957), and I can do no better than to quote the advice on page 8 of their book:

Probably the most important prerequisite is that ill-defined quality: mathematical sophistication. We hope that this is an ingredient not required in large measure, but that it is needed to some degree there can be no doubt. The reader must be able to accept conditional statements, even though he feels the suppositions to be false; he must be willing to make concessions to mathematical simplicity; he must be patient enough to follow along with the peculiar kind of construction that mathematics is; and, above all, he must have sympathy with the method — a sympathy based upon his knowledge of its past successes in various of the empirical sciences and upon his realization of the necessity for rigorous deduction in science as we know it.

If you do not know the terms “risk averse,” “first order condition,” “utility function,” “probability density,” and “discount rate,” you will not fully understand this book. Flipping through it, however, you will see that the equation density is much lower than in first-year graduate microeconomics texts. In a sense, game theory is less abstract than price theory, because it deals with individual agents rather than aggregate markets and it is oriented towards explaining stylized facts rather than supplying econometric specifications. Mathematics is nonetheless essential. Professor Wei puts this well in his informal and unpublished class notes:

My experience in learning and teaching convinces me that going through a proof (which does not require much mathematics) is *the most effective way in learning, developing intuition, sharpening technical writing ability, and improving creativity*. However it is an extremely painful experience for people with simple mind and narrow interests.

Remember that a good proof should be *smooth* in the sense that any serious reader can read through it like the way we read *Miami Herald*; should be *precise* such that no one can add/delete/change a word—like the way we enjoy Robert Frost’s poetry!

I wouldn't change a word of that.

Other Books

At the time of the first edition of this book, most of the topics covered were absent from existing books on either game theory or information economics. Older books on game theory included Davis (1970), Harris (1987), Harsanyi (1977), Luce & Raiffa (1957), Moulin (1986a, 1986b), Ordeshook (1986), Rapoport (1960, 1970), Shubik (1982), Szep & Forgo (1985), Thomas (1984), and Williams (1966). Books on information in economics were mainly concerned with decision making under uncertainty rather than asymmetric information. Since the First Edition, a spate of books on game theory has appeared. The stream of new books has become a flood, and one of the pleasing features of this literature is its variety. Each one is different, and both student and teacher can profit by owning an assortment of them, something one cannot say of many other subject areas. We have not converged, perhaps because teachers are still converting into books their own independent materials from courses not taught with texts. I only wish I could say I had been able to use all my competitors' good ideas in the present edition.

Why, you might ask in the spirit of game theory, do I conveniently list all my competitor's books here, giving free publicity to books that could substitute for mine? For an answer, you must buy this book and read chapter 11 on signalling. Then you will understand that only an author quite confident that his book compares well with possible substitutes would do such a thing, and you will be even more certain that your decision to buy the book was a good one. (But see problem 11.6 too.)

Some Books on Game Theory and its Applications

- 1988** **Tirole**, Jean, *The Theory of Industrial Organization*, Cambridge, Mass: MIT Press. 479 pages. Still the standard text for advanced industrial organization.
- 1989** **Eatwell**, John, Murray Milgate & Peter Newman, eds., *The New Palgrave: Game Theory*. 264 pages. New York: Norton. A collection of brief articles on topics in game theory by prominent scholars.
- Schmalensee**, Richard & Robert Willig, eds., *The Handbook of Industrial Organization*, in two volumes, New York: North-Holland. A collection of not-so-brief articles on topics in industrial organization by prominent scholars.
- Spulber**, Daniel *Regulation and Markets*, Cambridge, Mass: MIT Press. 690 pages. Applications of game theory to rate of return regulation.
- 1990** **Banks**, Jeffrey, *Signalling Games in Political Science*. Chur, Switzerland: Harwood Publishers. 90 pages. Out of date by now, but worth reading anyway.
- Friedman**, James, *Game Theory with Applications to Economics*, 2nd edition, Oxford: Oxford University Press (First edition, 1986). 322 pages. By a leading expert on repeated games.
- Kreps**, David, *A Course in Microeconomic Theory*. Princeton: Princeton University Press. 850 pages. A competitor to Varian's Ph.D. micro text, in a more conversational style, albeit a conversation with a brilliant economist at a level of detail that scares some students.

Kreps, David, *Game Theory and Economic Modeling*, Oxford: Oxford University Press. 195 pages. A discussion of Nash equilibrium and its problems.

Krouse, Clement, *Theory of Industrial Economics*, Oxford: Blackwell Publishers. 602 pages. A good book on the same topics as Tirole's 1989 book, and largely overshadowed by it.

1991 **Dixit**, Avinash K. & Barry J. Nalebuff, *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*. New York: Norton. 393 pages. A book in the tradition of popular science, full of fun examples but with serious ideas too. I use this for my MBA students' half-semester course, though newer books are offering competition for that niche.

Fudenberg, Drew & Jean Tirole, *Game Theory*. Cambridge, Mass: MIT Press. 579 pages. This has become the standard text for second-year PhD courses in game theory. (Though I hope the students are referring back to *Games and Information* for help in getting through the hard parts.)

Milgrom, Paul and John Roberts, *Economics of Organization and Management*. Englewood Cliffs, New Jersey: Prentice-Hall. 621 pages. A model for how to think about organization and management. The authors taught an MBA course from this, but I wonder whether that is feasible anywhere but Stanford Business School.

Myerson, Roger, *Game Theory: Analysis of Conflict*, Cambridge, Mass: Harvard University Press. 568 pages. At an advanced level. In revising for the third edition, I noticed how well Myerson's articles are standing the test of time.

1992 **Aumann**, Robert & Sergiu Hart, eds., *Handbook of Game Theory with Economic Applications*, Volume 1, Amsterdam: North-Holland. 733 pages. A collection of articles by prominent scholars on topics in game theory.

Binmore, Ken, *Fun and Games: A Text on Game Theory*. Lexington, Mass: D.C. Heath. 642 pages. No pain, no gain; but pain and pleasure can be mixed even in the study of mathematics.

Gibbons, Robert, *Game Theory for Applied Economists*,. Princeton: Princeton University Press. 267 pages. Perhaps the main competitor to *Games and Information*. Shorter and less idiosyncratic.

Hirshleifer, Jack & John Riley, *The Economics of Uncertainty and Information*, Cambridge: Cambridge University Press. 465 pages. An underappreciated book that emphasizes information rather than game theory.

McMillan, John, *Games, Strategies, and Managers: How Managers Can Use Game Theory to Make Better Business Decisions*,. Oxford: Oxford University Press. 252 pages. Largely verbal, very well written, and an example of how clear thinking and clear writing go together.

Varian, Hal, *Microeconomic Analysis*, Third edition. New York: Norton. (1st edition, 1978; 2nd edition, 1984.) 547 pages. Varian was the standard PhD micro text when I took the course in 1980. The third edition is much bigger, with lots of game theory and information economics concisely presented.

1993 **Basu**, Kaushik, *Lectures in Industrial Organization Theory*, . Oxford: Blackwell Publishers. 236 pages. Lots of game theory as well as I.O.

Eichberger, Jurgen, *Game Theory for Economists*, San Diego: Academic Press. 315 pages. Focus on game theory, but with applications along the way for illustration.

Laffont, Jean-Jacques & Jean Tirole, *A Theory of Incentives in Procurement and Regulation*, Cambridge, Mass: MIT Press. 705 pages. If you like section 10.4 of *Games and Information*, here is an entire book on the model.

Martin, Stephen, *Advanced Industrial Economics*, Oxford: Blackwell Publishers. 660 pages. Detailed and original analysis of particular models, and much more attention to empirical articles than Krouse, Shy, and Tirole.

1994 **Baird**, Douglas, Robert Gertner & Randal Picker, *Strategic Behavior and the Law: The Role of Game Theory and Information Economics in Legal Analysis*, Cambridge, Mass: Harvard University Press. 330 pages. A mostly verbal but not easy exposition of game theory using topics such as contracts, procedure, and tort.

Gardner, Roy, *Games for Business and Economics*, New York: John Wiley and Sons. 480 pages. Indiana University has produced not one but two game theory texts.

Morris, Peter, *Introduction to Game Theory*, Berlin: Springer Verlag. 230 pages. Not in my library yet.

Morrow, James, *Game Theory for Political Scientists*, Princeton, N.J. : Princeton University Press. 376 pages. The usual topics, but with a political science slant, and especially good on things such as utility theory.

Osborne, Martin and Ariel Rubinstein, *A Course in Game Theory*, Cambridge, Mass: MIT Press. 352 pages. Similar in style to Eichberger's 1993 book. See their excellent "List of Results" on pages 313-19 which summarizes the mathematical propositions without using specialized notation.

1995 **Mas-Colell**, Andreu Michael D. Whinston and Jerry R. Green, *Microeconomic Theory*, Oxford: Oxford University Press. 981 pages. This combines the topics of Varian's PhD micro text, those of *Games and Information*, and general equilibrium. Massive, and a good reference.

Owen, Guillermo, *Game Theory*, New York: Academic Press, 3rd edition. (1st edition, 1968; 2nd edition, 1982.) This book clearly lays out the older approach to game theory, and holds the record for longevity in game theory books.

1996 **Besanko**, David, David Dranove and Mark Shanley, *Economics of Strategy*, New York: John Wiley and Sons. This actually can be used with Indiana M.B.A. students, and clearly explains some very tricky ideas such as strategic complements.

Shy, Oz, *Industrial Organization, Theory and Applications*, Cambridge, Mass: MIT Press. 466 pages. A new competitor to Tirole's 1988 book which is somewhat easier.

1997 **Gates**, Scott and Brian Humes, *Games, Information, and Politics: Applying Game Theoretic Models to Political Science*, Ann Arbor: University of Michigan Press. 182 pages.

Ghemawat, Pankaj, *Games Businesses Play: Cases and Models*, Cambridge, Mass: MIT Press. 255 pages. Analysis of six cases from business using game theory at the MBA level. Good for the difficult task of combining theory with evidence.

Macho-Stadler, Ines and J. David Perez-Castillo, *An Introduction to the Economics of Information: Incentives and Contracts*, Oxford: Oxford University Press. 277 pages. Entirely on moral hazard, adverse selection, and signalling.

Romp, Graham, *Game Theory: Introduction and Applications*, Oxford: Oxford University Press. 284 pages. With unusual applications (chapters on macroeconomics, trade policy, and environmental economics) and lots of exercises with answers.

Salanie, Bernard, *The Economics of Contracts: A Primer*, Cambridge, Mass: MIT Press. 232 pages. Specialized to a subject of growing importance.

1998 **Bierman**, H. Scott & Luis Fernandez, *Game Theory with Economic Applications*. Reading, Massachusetts: Addison Wesley, Second edition. (1st edition, 1993.) 452 pages. A text for undergraduate courses, full of good examples.

Dugatkin, Lee and Hudson Reeve, editors, *Game Theory & Animal Behavior*, Oxford: Oxford University Press. 320 pages. Just on biology applications.

1999 **Aliprantis**, Charalambos & Subir Chakrabarti *Games and Decisionmaking*, Oxford: Oxford University Press. 224 pages. An undergraduate text for game theory, decision theory, auctions, and bargaining, the third game theory text to come out of Indiana.

Basar, Tamar & Geert Olsder *Dynamic Noncooperative Game Theory*, 2nd edition, revised, Philadelphia: Society for Industrial and Applied Mathematics (1st edition 1982, 2nd edition 1995). This book is by and for mathematicians, with surprisingly little overlap between its bibliography and that of the present book. Suitable for people who like differential equations and linear algebra.

Dixit, Avinash & Susan Skeath, *Games of Strategy*, New York: Norton. 600 pages. Nicely laid out with color and boldfacing. Game theory plus chapters on bargaining, auctions, voting, etc. Detailed verbal explanations of many games.

Dutta, Prajit, *Strategies and Games: Theory And Practice*, Cambridge, Mass: MIT Press. 450 pages.

Stahl, Saul, *A Gentle Introduction to Game Theory*, Providence, RI: American Mathematical Society. 176 pages. In the mathematics department tradition, with many exercises and numerical answers.

Forthcoming

Gintis, Herbert, *Game Theory Evolving*, Princeton: Princeton University Press. (May 12, 1999 draft at www-unix.oit.umass.edu/~gintis.) A wonderful book of problems and solutions, with much explanation and special attention to evolutionary biology.

Muthoo, Abhinay, *Bargaining Theory With Applications*, Cambridge: Cambridge University Press.

Osborne, Martin, *An Introduction to Game Theory*, Oxford: Oxford University Press. Up on the web via this book's website if you'd like to check it out.

Rasmusen, Eric, editor, *Readings in Games and Information*, Oxford: Blackwell Publishers. Journal and newspaper articles on game theory and information economics.

Rasmusen, Eric *Games and Information*. Oxford: Blackwell Publishers, Third edition. (1st edition, 1989; 2nd edition, 1994.) Read on.

Contact Information

The website for the book is at

[Http://php.indiana.edu/~erasmuse/GI/index.html](http://php.indiana.edu/~erasmuse/GI/index.html)

This site has the answers to the odd-numbered problems at the end of the chapters. For answers to even-numbered questions, instructors or others needing them for good reasons should email me at Erasmuse@Indiana.edu; send me snailmail at Eric Rasmusen, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, 1309 East 10th Street, Bloomington, Indiana 47405-1701; or fax me at (812)855-3354.

If you wish to contact the publisher of this book, the addresses are 108 Cowley Road, Oxford, England, OX4 1JF; or Blackwell Publishers, 350 Main Street, Malden, Massachusetts 02148.

The text files on the website are two forms (a) *.te, LaTeX, which uses only ASCII characters, but does not have the diagrams, and (b) *.pdf, Adobe Acrobat, which is formatted and can be read using a free reader program. I encourage readers to submit additional homework problems as well as errors and frustrations. They can be sent to me by e-mail at Erasmuse@Indiana.edu.

Acknowledgements

I would like to thank the many people who commented on clarity, suggested topics and references, or found mistakes. I've put affiliations next to their names, but remember that these change over time (A.B. was not a finance professor when he was my research assistant!).

First Edition: Dean Amel (Board of Governors, Federal Reserve), Dan Asquith (S.E.C.), Sushil Bikhchandani (UCLA business economics), Patricia Hughes Brennan (UCLA accounting), Paul Cheng, Luis Fernandez (Oberlin economics), David Hirshleifer (Ohio State finance), Jack Hirshleifer (UCLA economics), Steven Lippman (UCLA management science), Ivan Png (Singapore), Benjamin Rasmusen (Roseland Farm), Marilyn Rasmusen (Roseland Farm), Ray Renken (Central Florida physics), Richard Silver, Yoon Suh (UCLA accounting), Brett Trueman (Berkeley accounting), Barry Weingast (Hoover) and students in Management 200a made useful comments. D. Koh, Jeanne Lamotte, In-Ho Lee, Loi Lu, Patricia Martin, Timothy Opler (Ohio State finance), Sang Tran, Jeff Vincent, Tao Yang, Roy Zerner, and especially Emmanuel Petrakis (Crete economics) helped me with research assistance at one stage or another. Robert Boyd (UCLA anthropology), Mark Ramseyer (Harvard law), Ken Taymor, and John Wiley (UCLA law) made extensive comments in a reading group as each chapter was written.

Second Edition: Jonathan Berk (U. British Columbia commerce), Mark Burkey (Appalachian State economics), Craig Holden (Indiana finance), Peter Huang (Penn Law), Michael Katz (Berkeley business), Thomas Lyon (Indiana business economics), Steve Postrel (Northwestern business), Herman Quirnbach (Iowa State economics), H. Shifrin, George Tsebelis (UCLA poli sci), Thomas Voss (Leipzig sociology), and Jong-Shin Wei made useful comments, and Alexander Butler (Louisiana State finance) and An-Sing Chen provided research assistance. My students in Management 200 at UCLA and G601 at Indiana University provided invaluable help, especially in suffering through the first drafts of the homework problems.

Third Edition: Kyung-Hwan Baik (Sung Kyun Kwan), Patrick Chen, Robert Dimand (Brock economics), Mathias Erlei (Muenster), Francisco Galera, Peter-John Gordon (University of the West Indies), Erik Johannessen, Michael Mesterton-Gibbons (Pennsylvania), David Rosenbaum (Nebraska economics), Richard Tucker, Hal Wasserman (Berkeley), and Chad Zutter (Indiana finance) made comments that were helpful for the Third Edition. Blackwell supplied anonymous reviewers of superlative quality. Scott Fluhr, Pankaj Jain and John Spence provided research assistance and new generations of students in G601 were invaluable in helping to clarify my writing.

Eric Rasmusen

Professor of Business Economics and Public Policy and Sanjay Subhedar Fellow,
Kelley School of Business, Indiana University.

Introduction

1

History

Not so long ago, the scoffer could say that econometrics and game theory were like Japan and Argentina. In the late 1940s both disciplines and both economies were full of promise, poised for rapid growth and ready to make a profound impact on the world. We all know what happened to the economies of Japan and Argentina. Of the disciplines, econometrics became an inseparable part of economics, while game theory languished as a subdiscipline, interesting to its specialists but ignored by the profession as a whole. The specialists in game theory were generally mathematicians, who cared about definitions and proofs rather than applying the methods to economic problems. Game theorists took pride in the diversity of disciplines to which their theory could be applied, but in none had it become indispensable.

In the 1970s, the analogy with Argentina broke down. At the same time that Argentina was inviting back Juan Peron, economists were beginning to discover what they could achieve by combining game theory with the structure of complex economic situations. Innovation in theory and application was especially useful for situations with asymmetric information and a temporal sequence of actions, the two major themes of this book. During the 1980s, game theory became dramatically more important to mainstream economics. Indeed, it seemed to be swallowing up microeconomics just as econometrics had swallowed up empirical economics.

Game theory is generally considered to have begun with the publication of von Neumann & Morgenstern's *The Theory of Games and Economic Behaviour* in 1944. Although very little of the game theory in that thick volume is relevant to the present book, it introduced the idea that conflict could be mathematically analyzed and provided the terminology with which to do it. The development of the "Prisoner's Dilemma" (Tucker [unpub]) and Nash's papers on the definition and existence of equilibrium (Nash [1950b, 1951]) laid the foundations for modern noncooperative game theory. At the same time, cooperative game theory reached important results in papers by Nash (1950a) and Shapley (1953b) on bargaining games and Gillies (1953) and Shapley (1953a) on the core.

By 1953 virtually all the game theory that was to be used by economists for the next 20 years had been developed. Until the mid 1970s, game theory remained an autonomous field with little relevance to mainstream economics, important exceptions being Schelling's 1960 book, *The Strategy of Conflict*, which introduced the focal point, and a series of papers (of which Debreu & Scarf [1963] is typical) that showed the relationship of the core of a game to the general equilibrium of an economy.

In the 1970s, information became the focus of many models as economists started

¹June 30, 1993, July 24, 1999. August 28, 1999. February 2, 2000. May 27, 2002. Ariel Kemper. August 6, 2003. Eric Rasmusen, Erasmuse@indiana.edu. Web: Php.indiana.edu/~erasmuse. Footnotes starting with xxx are the author's notes to himself. Comments are welcomed. This section is zzz pages long.

to put emphasis on individuals who act rationally but with limited information. When attention was given to individual agents, the time ordering in which they carried out actions began to be explicitly incorporated. With this addition, games had enough structure to reach interesting and non-obvious results. Important “toolbox” references include the earlier but long-unapplied articles of Selten (1965) (on perfectness) and Harsanyi (1967) (on incomplete information), the papers by Selten (1975) and Kreps & Wilson (1982b) extending perfectness, and the article by Kreps, Milgrom, Roberts & Wilson (1982) on incomplete information in repeated games. Most of the applications in the present book were developed after 1975, and the flow of research shows no sign of diminishing.

Game Theory’s Method

Game theory has been successful in recent years because it fits so well into the new methodology of economics. In the past, macroeconomists started with broad behavioral relationships like the consumption function, and microeconomists often started with precise but irrational behavioral assumptions such as sales maximization. Now all economists start with primitive assumptions about the utility functions, production functions, and endowments of the actors in the models (to which must often be added the available information). The reason is that it is usually easier to judge whether primitive assumptions are sensible than to evaluate high-level assumptions about behavior. Having accepted the primitive assumptions, the modeller figures out what happens when the actors maximize their utility subject to the constraints imposed by their information, endowments, and production functions. This is exactly the paradigm of game theory: the modeller assigns payoff functions and strategy sets to his players and sees what happens when they pick strategies to maximize their payoffs. The approach is a combination of the “Maximization Subject to Constraints” of MIT and the “No Free Lunch” of Chicago. We shall see, however, that game theory relies only on the spirit of these two approaches: it has moved away from maximization by calculus, and inefficient allocations are common. The players act rationally, but the consequences are often bizarre, which makes application to a world of intelligent men and ludicrous outcomes appropriate.

Exemplifying Theory

Along with the trend towards primitive assumptions and maximizing behavior has been a trend toward simplicity. I called this “no-fat modelling” in the First Edition, but the term “exemplifying theory” from Fisher (1989) is more apt. This has also been called “modelling by example” or “MIT-style theory.” A more smoothly flowing name, but immodest in its double meaning, is “exemplary theory.” The heart of the approach is to discover the simplest assumptions needed to generate an interesting conclusion—the starkest, barest model that has the desired result. This desired result is the answer to some relatively narrow question. Could education be just a signal of ability? Why might bid-ask spreads exist? Is predatory pricing ever rational?

The modeller starts with a vague idea such as “People go to college to show they’re smart.” He then models the idea formally in a simple way. The idea might survive intact; it might be found formally meaningless; it might survive with qualifications; or its opposite might turn out to be true. The modeller then uses the model to come up with precise

propositions, whose proofs may tell him still more about the idea. After the proofs, he goes back to thinking in words, trying to understand more than whether the proofs are mathematically correct.

Good theory of any kind uses Occam's razor, which cuts out superfluous explanations, and the *ceteris paribus* assumption, which restricts attention to one issue at a time. Exemplifying theory goes a step further by providing, in the theory, only a narrow answer to the question. As Fisher says, "Exemplifying theory does not tell us what *must* happen. Rather it tells us what *can* happen."

In the same vein, at Chicago I have heard the style called "Stories That Might be True." This is not destructive criticism if the modeller is modest, since there are also a great many "Stories That Can't Be True," which are often used as the basis for decisions in business and government. Just as the modeller should feel he has done a good day's work if he has eliminated most outcomes as equilibria in his model, even if multiple equilibria remain, so he should feel useful if he has ruled out certain explanations for how the world works, even if multiple plausible models remain. The aim should be to come up with one or more stories that might apply to a particular situation and then try to sort out which story gives the best explanation. In this, economics combines the deductive reasoning of mathematics with the analogical reasoning of law.

A critic of the mathematical approach in biology has compared it to an hourglass (Slatkin [1980]). First, a broad and important problem is introduced. Second, it is reduced to a very special but tractable model that hopes to capture its essence. Finally, in the most perilous part of the process, the results are expanded to apply to the original problem. Exemplifying theory does the same thing.

The process is one of setting up "If-Then" statements, whether in words or symbols. To apply such statements, their premises and conclusions need to be verified, either by casual or careful empiricism. If the required assumptions seem contrived or the assumptions and implications contradict reality, the idea should be discarded. If "reality" is not immediately obvious and data is available, econometric tests may help show whether the model is valid. Predictions can be made about future events, but that is not usually the primary motivation: most of us are more interested in explaining and understanding than predicting.

The method just described is close to how, according to Lakatos (1976), mathematical theorems are developed. It contrasts sharply with the common view that the researcher starts with a hypothesis and proves or disproves it. Instead, the process of proof helps show how the hypothesis should be formulated.

An important part of exemplifying theory is what Kreps & Spence (1984) have called "blackboxing": treating unimportant subcomponents of a model in a cursory way. The game "Entry for Buyout" of section 15.4, for example, asks whether a new entrant would be bought out by the industry's incumbent producer, something that depends on duopoly pricing and bargaining. Both pricing and bargaining are complicated games in themselves, but if the modeller does not wish to deflect attention to those topics he can use the simple Nash and Cournot solutions to those games and go on to analyze buyout. If the entire

focus of the model were duopoly pricing, then using the Cournot solution would be open to attack, but as a simplifying assumption, rather than one that “drives” the model, it is acceptable.

Despite the style’s drive towards simplicity, a certain amount of formalism and mathematics is required to pin down the modeller’s thoughts. Exemplifying theory treads a middle path between mathematical generality and nonmathematical vagueness. Both alternatives will complain that exemplifying theory is too narrow. But beware of calls for more “rich,” “complex,” or “textured” descriptions; these often lead to theory which is either too incoherent or too incomprehensible to be applied to real situations.

Some readers will think that exemplifying theory uses too little mathematical technique, but others, especially noneconomists, will think it uses too much. Intelligent laymen have objected to the amount of mathematics in economics since at least the 1880s, when George Bernard Shaw said that as a boy he (1) let someone assume that $a = b$, (2) permitted several steps of algebra, and (3) found he had accepted a proof that $1 = 2$. Forever after, Shaw distrusted assumptions and algebra. Despite the effort to achieve simplicity (or perhaps because of it), mathematics is essential to exemplifying theory. The conclusions can be retranslated into words, but rarely can they be found by verbal reasoning. The economist Wicksteed put this nicely in his reply to Shaw’s criticism:

Mr Shaw arrived at the sapient conclusion that there “was a screw loose somewhere”—not in his own reasoning powers, but—“in the algebraic art”; and thenceforth renounced mathematical reasoning in favour of the literary method which enables a clever man to follow equally fallacious arguments to equally absurd conclusions *without seeing that they are absurd*. This is the exact difference between the mathematical and literary treatment of the pure theory of political economy. (Wicksteed [1885] p. 732)

In exemplifying theory, one can still rig a model to achieve a wide range of results, but it must be rigged by making strange primitive assumptions. Everyone familiar with the style knows that the place to look for the source of suspicious results is the description at the start of the model. If that description is not clear, the reader deduces that the model’s counterintuitive results arise from bad assumptions concealed in poor writing. Clarity is therefore important, and the somewhat inelegant Players-Actions-Payoffs presentation used in this book is useful not only for helping the writer, but for persuading the reader.

This Book’s Style

Substance and style are closely related. The difference between a good model and a bad one is not just whether the essence of the situation is captured, but also how much froth covers the essence. In this book, I have tried to make the games as simple as possible. They often, for example, allow each player a choice of only two actions. Our intuition works best with such models, and continuous actions are technically more troublesome. Other assumptions, such as zero production costs, rely on trained intuition. To the layman, the assumption that output is costless seems very strong, but a little experience with these models teaches that it is the constancy of the marginal cost that usually matters, not its level.

What matters more than what a model says is what we understand it to say. Just as an article written in Sanskrit is useless to me, so is one that is excessively mathematical or poorly written, no matter how rigorous it seems to the author. Such an article leaves me with some new belief about its subject, but that belief is not sharp, or precisely correct. Overprecision in sending a message creates imprecision when it is received, because precision is not clarity. The result of an attempt to be mathematically precise is sometimes to overwhelm the reader, in the same way that someone who requests the answer to a simple question in the discovery process of a lawsuit is overwhelmed when the other side responds with 70 boxes of tangentially related documents. The quality of the author's input should be judged not by some abstract standard but by the output in terms of reader processing cost and understanding.

In this spirit, I have tried to simplify the structure and notation of models while giving credit to their original authors, but I must ask pardon of anyone whose model has been oversimplified or distorted, or whose model I have inadvertently replicated without crediting them. In trying to be understandable, I have taken risks with respect to accuracy. My hope is that the impression left in the readers' minds will be more accurate than if a style more cautious and obscure had left them to devise their own errors.

Readers may be surprised to find occasional references to newspaper and magazine articles in this book. I hope these references will be reminders that models ought eventually to be applied to specific facts, and that a great many interesting situations are waiting for our analysis. The principal-agent problem is not found only in back issues of *Econometrica*: it can be found on the front page of today's *Wall Street Journal* if one knows what to look for.

I make the occasional joke here and there, and game theory is a subject intrinsically full of paradox and surprise. I want to emphasize, though, that I take game theory seriously, in the same way that Chicago economists like to say that they take price theory seriously. It is not just an academic artform: people do choose actions deliberately and trade off one good against another, and game theory will help you understand how they do that. If it did not, I would not advise you to study such a difficult subject; there are much more elegant fields in mathematics, from an aesthetic point of view. As it is, I think it is important that every educated person have some contact with the ideas in this book, just as they should have some idea of the basic principles of price theory.

I have been forced to exercise more discretion over definitions than I had hoped. Many concepts have been defined on an article-by-article basis in the literature, with no consistency and little attention to euphony or usefulness. Other concepts, such as "asymmetric information" and "incomplete information," have been considered so basic as to not need definition, and hence have been used in contradictory ways. I use existing terms whenever possible, and synonyms are listed.

I have often named the players Smith and Jones so that the reader's memory will be less taxed in remembering which is a player and which is a time period. I hope also to reinforce the idea that a model is a story made precise; we begin with Smith and Jones, even if we quickly descend to s and j . Keeping this in mind, the modeller is less likely to build mathematically correct models with absurd action sets, and his descriptions are more

pleasant to read. In the same vein, labelling a curve “ $U = 83$ ” sacrifices no generality: the phrase “ $U = 83$ and $U = 66$ ” has virtually the same content as “ $U = \alpha$ and $U = \beta$, where $\alpha > \beta$,” but uses less short-term memory.

A danger of this approach is that readers may not appreciate the complexity of some of the material. While journal articles make the material seem harder than it is, this approach makes it seem easier (a statement that can be true even if readers find this book difficult). The better the author does his job, the worse this problem becomes. Keynes (1933) says of Alfred Marshall’s *Principles*,

The lack of emphasis and of strong light and shade, the sedulous rubbing away of rough edges and salients and projections, until what is most novel can appear as trite, allows the reader to pass too easily through. Like a duck leaving water, he can escape from this douche of ideas with scarce a wetting. The difficulties are concealed; the most ticklish problems are solved in footnotes; a pregnant and original judgement is dressed up as a platitude.

This book may well be subject to the same criticism, but I have tried to face up to difficult points, and the problems at the end of each chapter will help to avoid making the reader’s progress too easy. Only a certain amount of understanding can be expected from a book, however. The efficient way to learn how to do research is to start doing it, not to read about it, and after reading this book, if not before, many readers will want to build their own models. My purpose here is to show them the big picture, to help them understand the models intuitively, and give them a feel for the modelling process.

NOTES

- Perhaps the most important contribution of von Neumann & Morgenstern (1944) is the theory of expected utility (see section 2.3). Although they developed the theory because they needed it to find the equilibria of games, it is today heavily used in all branches of economics. In game theory proper, they contributed the framework to describe games, and the concept of mixed strategies (see section 3.1). A good historical discussion is Shubik (1992) in the Weintraub volume mentioned in the next note.
- A number of good books on the history of game theory have appeared in recent years. Norman Macrae’s *John von Neumann* and Sylvia Nasar’s *A Beautiful Mind* (on John Nash) are extraordinarily good biographies of founding fathers, while *Eminent Economists: Their Life Philosophies* and *Passion and Craft: Economists at Work*, edited by Michael Szenberg, and *Toward a History of Game Theory*, edited by Roy Weintraub, contain autobiographical essays by many scholars who use game theory, including Shubik, Riker, Dixit, Varian, and Myerson. Dimand and Dimand’s *A History of Game Theory*, the first volume of which appeared in 1996, is a more intensive look at the intellectual history of the field. See also Myerson (1999).
- For articles from the history of mathematical economics, see the collection by Baumol & Goldfeld (1968), Dimand and Dimand’s 1997 *The Foundations of Game Theory* in three volumes, and Kuhn (1997).

- Collections of more recent articles include Rasmusen (2000a), Binmore & Dasgupta (1986), Diamond & Rothschild (1978), and the immense Rubinstein (1990).
- On method, see the dialogue by Lakatos (1976), or Davis, Marchisotto & Hersh (1981), chapter 6 of which is a shorter dialogue in the same style. Friedman (1953) is the classic essay on a different methodology: evaluating a model by testing its predictions. Kreps & Spence (1984) is a discussion of exemplifying theory.
- Because style and substance are so closely linked, how one writes is important. For advice on writing, see McCloskey (1985, 1987) (on economics), Basil Blackwell (1985) (on books), Bowersock (1985) (on footnotes), Fowler (1965), Fowler & Fowler (1949), Halmos (1970) (on mathematical writing), Rasmusen (forthcoming), Strunk & White (1959), Weiner (1984), and Wydick (1978).
- **A fallacious proof that $1=2$.** Suppose that $a = b$. Then $ab = b^2$ and $ab - b^2 = a^2 - b^2$. Factoring the last equation gives us $b(a - b) = (a + b)(a - b)$, which can be simplified to $b = a + b$. But then, using our initial assumption, $b = 2b$ and $1 = 2$. (The fallacy is division by zero.)

xxx Footnotes starting with xxx are the author's notes to himself. Comments are welcomed. August 28, 1999. May 27, 2002. November 29, 2003. Eric Rasmusen, Erasmuse@indiana.edu. Web: Mypage.iu.edu/~erasmuse.

PART I GAME THEORY

1 The Rules of the Game

1.1: Definitions

Game theory is concerned with the actions of decision makers who are conscious that their actions affect each other. When the only two publishers in a city choose prices for their newspapers, aware that their sales are determined jointly, they are players in a game with each other. They are not in a game with the readers who buy the newspapers, because each reader ignores his effect on the publisher. Game theory is not useful when decisionmakers ignore the reactions of others or treat them as impersonal market forces.

The best way to understand which situations can be modelled as games and which cannot is to think about examples like the following:

1. OPEC members choosing their annual output;
2. General Motors purchasing steel from USX;
3. two manufacturers, one of nuts and one of bolts, deciding whether to use metric or American standards;
4. a board of directors setting up a stock option plan for the chief executive officer;
5. the US Air Force hiring jet fighter pilots;
6. an electric company deciding whether to order a new power plant given its estimate of demand for electricity in ten years.

The first four examples are games. In (1), OPEC members are playing a game because Saudi Arabia knows that Kuwait's oil output is based on Kuwait's forecast of Saudi output, and the output from both countries matters to the world price. In (2), a significant portion of American trade in steel is between General Motors and USX, companies which realize that the quantities traded by each of them affect the price. One wants the price low, the other high, so this is a game with conflict between the two players. In (3), the nut and bolt manufacturers are not in conflict, but the actions of one do affect the desired actions of the other, so the situation is a game none the less. In (4), the board of directors chooses a stock option plan anticipating the effect on the actions of the CEO.

Game theory is inappropriate for modelling the final two examples. In (5), each individual pilot affects the US Air Force insignificantly, and each pilot makes his employment decision without regard for the impact on the Air Force's policies. In (6), the electric company faces a complicated decision, but it does not face another rational agent. These situations are more appropriate for the use of **decision theory** than game theory, decision theory being the careful analysis of how one person makes a decision when he may be

faced with uncertainty, or an entire sequence of decisions that interact with each other, but when he is not faced with having to interact strategically with other single decision makers. Changes in the important economic variables could, however, turn examples (5) and (6) into games. The appropriate model changes if the Air Force faces a pilots' union or if the public utility commission pressures the utility to change its generating capacity.

Game theory as it will be presented in this book is a modelling tool, not an axiomatic system. The presentation in this chapter is unconventional. Rather than starting with mathematical definitions or simple little games of the kind used later in the chapter, we will start with a situation to be modelled, and build a game from it step by step.

Describing a Game

The essential elements of a game are **players**, **actions**, **payoffs**, and **information**—PAPI, for short. These are collectively known as the **rules of the game**, and the modeller's objective is to describe a situation in terms of the rules of a game so as to explain what will happen in that situation. Trying to maximize their payoffs, the players will devise plans known as **strategies** that pick actions depending on the information that has arrived at each moment. The combination of strategies chosen by each player is known as the **equilibrium**. Given an equilibrium, the modeller can see what actions come out of the conjunction of all the players' plans, and this tells him the **outcome** of the game.

This kind of standard description helps both the modeller and his readers. For the modeller, the names are useful because they help ensure that the important details of the game have been fully specified. For his readers, they make the game easier to understand, especially if, as with most technical papers, the paper is first skimmed quickly to see if it is worth reading. The less clear a writer's style, the more closely he should adhere to the standard names, which means that most of us ought to adhere very closely indeed.

Think of writing a paper as a game between author and reader, rather than as a single-player production process. The author, knowing that he has valuable information but imperfect means of communication, is trying to convey the information to the reader. The reader does not know whether the information is valuable, and he must choose whether to read the paper closely enough to find out.¹

To define the terms used above and to show the difference between game theory and decision theory, let us use the example of an entrepreneur trying to decide whether to start a dry cleaning store in a town already served by one dry cleaner. We will call the two firms "NewCleaner" and "OldCleaner." NewCleaner is uncertain about whether the economy will be in a recession or not, which will affect how much consumers pay for dry cleaning, and must also worry about whether OldCleaner will respond to entry with a price war or by keeping its initial high prices. OldCleaner is a well-established firm, and it would survive any price war, though its profits would fall. NewCleaner must itself decide whether to

¹Once you have read to the end of this chapter: What are the possible equilibria of this game?

initiate a price war or to charge high prices, and must also decide what kind of equipment to buy, how many workers to hire, and so forth.

Players are the individuals who make decisions. Each player's goal is to maximize his utility by choice of actions.

In the Dry Cleaners Game, let us specify the players to be NewCleaner and OldCleaner. Passive individuals like the customers, who react predictably to price changes without any thought of trying to change anyone's behavior, are not players, but environmental parameters. Simplicity is the goal in modelling, and the ideal is to keep the number of players down to the minimum that captures the essence of the situation.

Sometimes it is useful to explicitly include individuals in the model called **pseudo-players** whose actions are taken in a purely mechanical way.

Nature is a pseudo-player who takes random actions at specified points in the game with specified probabilities.

In the Dry Cleaners Game, we will model the possibility of recession as a move by Nature. With probability 0.3, Nature decides that there will be a recession, and with probability 0.7 there will not. Even if the players always took the same actions, this random move means that the model would yield more than just one prediction. We say that there are different **realizations** of a game depending on the results of random moves.

An **action** or **move** by player i , denoted a_i , is a choice he can make.

Player i 's **action set**, $A_i = \{a_i\}$, is the entire set of actions available to him.

An **action combination** is an ordered set $a = \{a_i\}$, ($i = 1, \dots, n$) of one action for each of the n players in the game.

Again, simplicity is our goal. We are trying to determine whether Newcleaner will enter or not, and for this it is not important for us to go into the technicalities of dry cleaning equipment and labor practices. Also, it will not be in Newcleaner's interest to start a price war, since it cannot possibly drive out Oldcleaners, so we can exclude that decision from our model. Newcleaner's action set can be modelled very simply as $\{Enter, Stay Out\}$. We will also specify Oldcleaner's action set to be simple: it is to choose price from $\{Low, High\}$.

By player i 's **payoff** $\pi_i(s_1, \dots, s_n)$, we mean either:

- (1) The utility player i receives after all players and Nature have picked their strategies and the game has been played out; or
- (2) The expected utility he receives as a function of the strategies chosen by himself and the other players.

For the moment, think of "strategy" as a synonym for "action". Definitions (1) and (2) are distinct and different, but in the literature and this book the term "payoff" is used

for both the actual payoff and the expected payoff. The context will make clear which is meant. If one is modelling a particular real-world situation, figuring out the payoffs is often the hardest part of constructing a model. For this pair of dry cleaners, we will pretend we have looked over all the data and figured out that the payoffs are as given by Table 1a if the economy is normal, and that if there is a recession the payoff of each player who operates in the market is 60 thousand dollars lower, as shown in Table 1b.

Table 1a: The Dry Cleaners Game: Normal Economy

		OldCleaner	
		<i>Low price</i>	<i>High price</i>
NewCleaner	<i>Enter</i>	-100, -50	100, 100
	<i>Stay Out</i>	0, 50	0, 300

Payoffs to: (NewCleaner, OldCleaner) in thousands of dollars

Table 1b: The Dry Cleaners Game: Recession

		OldCleaner	
		<i>Low price</i>	<i>High price</i>
NewCleaner	<i>Enter</i>	-160, -110	40, 40
	<i>Stay Out</i>	0, -10	0, 240

Payoffs to: (NewCleaner, OldCleaner) in thousands of dollars

Information is modelled using the concept of the **information set**, a concept which will be defined more precisely in Section 2.2. For now, think of a player's information set as his knowledge at a particular time of the values of different variables. The elements of the information set are the different values that the player thinks are possible. If the information set has many elements, there are many values the player cannot rule out; if it has one element, he knows the value precisely. A player's information set includes not only distinctions between the values of variables such as the strength of oil demand, but also knowledge of what actions have previously been taken, so his information set changes over the course of the game.

Here, at the time that it chooses its price, OldCleaner will know NewCleaner's decision about entry. But what do the firms know about the recession? If both firms know about the recession we model that as Nature moving before NewCleaner; if only OldCleaner knows, we put Nature's move after NewCleaner; if neither firm knows whether there is a recession at the time they must make their decisions, we put Nature's move at the end of the game. Let us do this last.

It is convenient to lay out information and actions together in an **order of play**. Here is the order of play we have specified for the Dry Cleaners Game:

- 1 Newcleaner chooses its entry decision from $\{Enter, Stay Out\}$.
- 2 Oldcleaner chooses its price from $\{Low, High\}$.
- 3 Nature picks demand, D , to be *Recession* with probability 0.3 or *Normal* with probability 0.7.

The purpose of modelling is to explain how a given set of circumstances leads to a particular result. The result of interest is known as the outcome.

*The **outcome** of the game is a set of interesting elements that the modeller picks from the values of actions, payoffs, and other variables after the game is played out.*

The definition of the outcome for any particular model depends on what variables the modeller finds interesting. One way to define the outcome of the Dry Cleaners Game would be as either *Enter* or *Stay Out*. Another way, appropriate if the model is being constructed to help plan NewCleaner's finances, is as the payoff that NewCleaner realizes, which is, from Tables 1a and 1b, one element of the set $\{0, 100, -100, 40, -160\}$.

Having laid out the assumptions of the model, let us return to what is special about the way game theory models a situation. Decision theory sets up the rules of the game in much the same way as game theory, but its outlook is fundamentally different in one important way: there is only one player. Return to NewCleaner's decision about entry. In decision theory, the standard method is to construct a **decision tree** from the rules of the game, which is just a graphical way to depict the order of play.

Figure 1 shows a decision tree for the Dry Cleaners Game. It shows all the moves available to NewCleaner, the probabilities of states of nature (actions that NewCleaner cannot control), and the payoffs to NewCleaner depending on its choices and what the environment is like. Note that although we already specified the probabilities of Nature's move to be 0.7 for *Normal*, we also need to specify a probability for OldCleaner's move, which is set at probability 0.5 of *Low price* and probability 0.5 of *High price*.




Figure 1: The Dry Cleaners Game as a Decision Tree

Once a decision tree is set up, we can solve for the optimal decision which maximizes the expected payoff. Suppose NewCleaner has entered. If OldCleaner chooses a high price, then NewCleaner's expected payoff is 82, which is $0.7(100) + 0.3(40)$. If OldCleaner chooses a low price, then NewCleaner's expected payoff is -118, which is $0.7(-100) + 0.3(-160)$. Since there is a 50-50 chance of each move by OldCleaner, NewCleaner's overall expected payoff from *Enter* is -18. That is worse than the 0 which NewCleaner could get by choosing *stay out*, so the prediction is that NewCleaner will stay out.

That, however, is wrong. This is a game, not just a decision problem. The flaw in the reasoning I just went through is the assumption that OldCleaner will choose *High price* with probability 0.5. If we use information about OldCleaner's payoffs and figure out what moves OldCleaner will take in solving its own profit maximization problem, we will come to a different conclusion.

First, let us depict the order of play as a **game tree** instead of a decision tree. Figure 2 shows our model as a game tree, with all of OldCleaner's moves and payoffs.




Figure 2: The Dry Cleaners Game as a Game Tree

Viewing the situation as a game, we must think about both players' decision making. Suppose NewCleaner has entered. If OldCleaner chooses *High price*, OldCleaner's expected profit is 82, which is $0.7(100) + 0.3(40)$. If OldCleaner chooses *Low price*, OldCleaner's expected profit is -68, which is $0.7(-50) + 0.3(-110)$. Thus, OldCleaner will choose *High price*, and with probability 1.0, not 0.5. The arrow on the game tree for *High price* shows this conclusion of our reasoning. This means, in turn, that NewCleaner can predict an expected payoff of 82, which is $0.7(100) + 0.3(40)$, from *Enter*.

Suppose NewCleaner has not entered. If OldCleaner chooses *High price*, OldCleaner's expected profit is 282, which is $0.7(300) + 0.3(240)$. If OldCleaner chooses *Low price*, OldCleaner's expected profit is 32, which is $0.7(50) + 0.3(-10)$. Thus, OldCleaner will choose *High price*, as shown by the arrow on *High price*. If NewCleaner chooses *Stay out*, NewCleaner will have a payoff of 0, and since that is worse than the 82 which NewCleaner can predict from *Enter*, NewCleaner will in fact enter the market.

This switching back from the point of view of one player to the point of view of another is characteristic of game theory. The game theorist must practice putting himself in *everybody* else's shoes. (Does that mean we become kinder, gentler people? – Or do we just get trickier?)

Since so much depends on the interaction between the plans and predictions of different players, it is useful to go a step beyond simply setting out actions in a game. Instead, the modeller goes on to think about **strategies**, which are action plans.

Player i 's strategy s_i is a rule that tells him which action to choose at each instant of the game, given his information set.

*Player i 's **strategy set** or **strategy space** $S_i = \{s_i\}$ is the set of strategies available to him.*

*A **strategy combination** $s = (s_1, \dots, s_n)$ is an ordered set consisting of one strategy for each of the n players in the game.²*

Since the information set includes whatever the player knows about the previous actions of other players, the strategy tells him how to react to their actions. In the Dry Cleaners Game, the strategy set for NewCleaner is just $\{Enter, Stay Out\}$, since NewCleaner moves first and is not reacting to any new information. The strategy set for OldCleaner, though, is

$$\left\{ \begin{array}{l} \text{High Price if NewCleaner Enters, Low Price if NewCleaner Stays Out} \\ \text{Low Price if NewCleaner Enters, High Price if NewCleaner Stays Out} \\ \text{High Price No Matter What} \\ \text{Low Price No Matter What} \end{array} \right\}$$

The concept of the strategy is useful because the action a player wishes to pick often depends on the past actions of Nature and the other players. Only rarely can we predict a player's actions unconditionally, but often we can predict how he will respond to the outside world.

Keep in mind that a player's strategy is a complete set of instructions for him, which tells him what actions to pick in every conceivable situation, even if he does not expect to reach that situation. Strictly speaking, even if a player's strategy instructs him to commit suicide in 1989, it ought also to specify what actions he takes if he is still alive in 1990. This kind of care will be crucial in Chapter 4's discussion of "subgame perfect" equilibrium. The completeness of the description also means that strategies, unlike actions, are unobservable. An action is physical, but a strategy is only mental.

Equilibrium

To predict the outcome of a game, the modeller focusses on the possible strategy combinations, since it is the interaction of the different players' strategies that determines what happens. The distinction between strategy combinations, which are sets of strategies, and outcomes, which are sets of values of whichever variables are considered interesting, is a common source of confusion. Often different strategy combinations lead to the same outcome. In the Dry Cleaners Game, the single outcome of *NewCleaner Enters* would result from either of the following two strategy combinations:

²The term "strategy combination" is not standard. The first edition of this book used it; I switched to the more common but jargonish "strategy profile" in the second edition, and now I have switched back.

$$\left\{ \begin{array}{l} \text{High Price if NewCleaner Enters, Low Price if NewCleaner Stays Out} \\ \text{Enter} \end{array} \right\}$$

$$\left\{ \begin{array}{l} \text{Low Price if NewCleaner Enters, High Price if NewCleaner Stays Out} \\ \text{Enter} \end{array} \right\}$$

Predicting what happens consists of selecting one or more strategy combinations as being the most rational behavior by the players acting to maximize their payoffs.

An **equilibrium** $s^* = (s_1^*, \dots, s_n^*)$ is a strategy combination consisting of a best strategy for each of the n players in the game.

The **equilibrium strategies** are the strategies players pick in trying to maximize their individual payoffs, as distinct from the many possible strategy combinations obtainable by arbitrarily choosing one strategy per player. Equilibrium is used differently in game theory than in other areas of economics. In a general equilibrium model, for example, an equilibrium is a set of prices resulting from optimal behavior by the individuals in the economy. In game theory, that set of prices would be the **equilibrium outcome**, but the equilibrium itself would be the strategy combination— the individuals' rules for buying and selling— that generated the outcome.

People often carelessly say “equilibrium” when they mean “equilibrium outcome,” and “strategy” when they mean “action.” The difference is not very important in most of the games that will appear in this chapter, but it is absolutely fundamental to thinking like a game theorist. Consider Germany's decision on whether to remilitarize the Rhineland in 1936. France adopted the strategy: *Do not fight*, and Germany responded by remilitarizing, leading to World War II a few years later. If France had adopted the strategy: *Do not fight unless Germany remilitarizes; otherwise, fight*, the outcome would still have been that France would not have fought and no war would have ensued, but Germany would not have remilitarized the Rhineland. Perhaps it was because he thought along these lines that John von Neumann was such a hawk in the Cold War, as MacRae describes in his biography (MacRae [1992]). This difference between actions and strategies, outcomes and equilibria, is one of the hardest ideas to teach in a game theory class, even though it is trivial to state.

To find the equilibrium, it is not enough to specify the players, strategies, and payoffs, because the modeller must also decide what “best strategy” means. He does this by defining an equilibrium concept.

An **equilibrium concept** or **solution concept** $F : \{S_1, \dots, S_n, \pi_1, \dots, \pi_n\} \rightarrow s^*$ is a rule that defines an equilibrium based on the possible strategy combinations and the payoff functions.

We have implicitly already used an equilibrium concept in the analysis above, which picked one strategy for each of the two players as our prediction for the game (what we implicitly used is the concept of **subgame perfectness** which will reappear in chapter 4). Only a few equilibrium concepts are generally accepted, and the remaining sections of this chapter are devoted to finding the equilibrium using the two best-known of them: dominant strategy and Nash equilibrium.

Uniqueness

Accepted solution concepts do not guarantee uniqueness, and lack of a unique equilibrium is a major problem in game theory. Often the solution concept employed leads us to believe that the players will pick one of the two strategy combinations A or B, not C or D, but we cannot say whether A or B is more likely. Sometimes we have the opposite problem and the game has no equilibrium at all. By this is meant either that the modeller sees no good reason why one strategy combination is more likely than another, or that some player wants to pick an infinite value for one of his actions.

A model with no equilibrium or multiple equilibria is underspecified. The modeller has failed to provide a full and precise prediction for what will happen. One option is to admit that his theory is incomplete. This is not a shameful thing to do; an admission of incompleteness like Section 5.2's Folk Theorem is a valuable negative result. Or perhaps the situation being modelled really is unpredictable. Another option is to renew the attack by changing the game's description or the solution concept. Preferably it is the description that is changed, since economists look to the rules of the game for the differences between models, and not to the solution concept. If an important part of the game is concealed under the definition of equilibrium, in fact, the reader is likely to feel tricked and to charge the modeller with intellectual dishonesty.

1.2 Dominated and Dominant Strategies: The Prisoner's Dilemma

In discussing equilibrium concepts, it is useful to have shorthand for "all the other players' strategies."

For any vector $y = (y_1, \dots, y_n)$, denote by y_{-i} the vector $(y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n)$, which is the portion of y not associated with player i .

Using this notation, s_{-Smith} , for instance, is the combination of strategies of every player except player *Smith*. That combination is of great interest to Smith, because he uses it to help choose his own strategy, and the new notation helps define his best response.

*Player i 's **best response** or **best reply** to the strategies s_{-i} chosen by the other players is the strategy s_i^* that yields him the greatest payoff; that is,*

$$\pi_i(s_i^*, s_{-i}) \geq \pi_i(s'_i, s_{-i}) \quad \forall s'_i \neq s_i^*. \quad (1)$$

The best response is strongly best if no other strategies are equally good, and weakly best otherwise.

The first important equilibrium concept is based on the idea of **dominance**.

*The strategy s_i^d is a **dominated strategy** if it is strictly inferior to some other strategy no matter what strategies the other players choose, in the sense that whatever strategies they pick, his payoff is lower with s_i^d . Mathematically, s_i^d is dominated if there exists a single s'_i such that*

$$\pi_i(s_i^d, s_{-i}) < \pi_i(s'_i, s_{-i}) \quad \forall s_{-i}. \quad (2)$$

Note that s_i^d is not a dominated strategy if there is no s_{-i} to which it is the best response, but sometimes the better strategy is s'_i and sometimes it is s''_i . In that case, s_i^d could have the redeeming feature of being a good compromise strategy for a player who cannot predict what the other players are going to do. A dominated strategy is unambiguously inferior to some single other strategy.

There is usually no special name for the superior strategy that beats a dominated strategy. In unusual games, however, there is some strategy that beats *every* other strategy. We call that a “dominant strategy”.

*The strategy s_i^* is a **dominant strategy** if it is a player’s strictly best response to any strategies the other players might pick, in the sense that whatever strategies they pick, his payoff is highest with s_i^* . Mathematically,*

$$\pi_i(s_i^*, s_{-i}) > \pi_i(s'_i, s_{-i}) \quad \forall s_{-i}, \quad \forall s'_i \neq s_i^*. \quad (3)$$

*A **dominant strategy equilibrium** is a strategy combination consisting of each player’s dominant strategy.*

A player’s dominant strategy is his strictly best response even to wildly irrational actions by the other players. Most games do not have dominant strategies, and the players must try to figure out each others’ actions to choose their own.

The Dry Cleaners Game incorporated considerable complexity in the rules of the game to illustrate such things as information sets and the time sequence of actions. To illustrate equilibrium concepts, we will use simpler games, such as the Prisoner’s Dilemma. In the Prisoner’s Dilemma, two prisoners, Messrs Row and Column, are being interrogated separately. If both confess, each is sentenced to eight years in prison; if both deny their involvement, each is sentenced to one year.³ If just one confesses, he is released but the other prisoner is sentenced to ten years. The Prisoner’s Dilemma is an example of a **2-by-2 game**, because each of the two players—Row and Column—has two possible actions in his action set: *Confess* and *Deny*. Table 2 gives the payoffs (The arrows represent a player’s preference between actions, as will be explained in Section 1.4).

³Another way to tell the story is to say that if both deny, then with probability 0.1 they are convicted anyway and serve ten years, for an expected payoff of $(-1, -1)$.

Table 2: The Prisoner's Dilemma

		Column	
		<i>Deny</i>	<i>Confess</i>
Row	<i>Deny</i>	-1, -1	-10, 0
	<i>Confess</i>	0, -10	-8, -8

Payoffs to: (Row, Column)

Each player has a dominant strategy. Consider Row. Row does not know which action Column is choosing, but if Column chooses *Deny*, Row faces a *Deny* payoff of -1 and a *Confess* payoff of 0 , whereas if Column chooses *Confess*, Row faces a *Deny* payoff of -10 and a *Confess* payoff of -8 . In either case Row does better with *Confess*. Since the game is symmetric, Column's incentives are the same. The dominant strategy equilibrium is (*Confess*, *Confess*), and the equilibrium payoffs are $(-8, -8)$, which is worse for both players than $(-1, -1)$. Sixteen, in fact, is the greatest possible combined total of years in prison.

The result is even stronger than it seems, because it is robust to substantial changes in the model. Because the equilibrium is a dominant strategy equilibrium, the information structure of the game does not matter. If Column is allowed to know Row's move before taking his own, the equilibrium is unchanged. Row still chooses *Confess*, knowing that Column will surely choose *Confess* afterwards.

The Prisoner's Dilemma crops up in many different situations, including oligopoly pricing, auction bidding, salesman effort, political bargaining, and arms races. Whenever you observe individuals in a conflict that hurts them all, your first thought should be of the Prisoner's Dilemma.

The game seems perverse and unrealistic to many people who have never encountered it before (although friends who are prosecutors assure me that it is a standard crime-fighting tool). If the outcome does not seem right to you, you should realize that very often the chief usefulness of a model is to induce discomfort. Discomfort is a sign that your model is not what you think it is—that you left out something essential to the result you expected and didn't get. Either your original thought or your model is mistaken; and finding such mistakes is a real if painful benefit of model building. To refuse to accept surprising conclusions is to reject logic.

Cooperative and Noncooperative Games

What difference would it make if the two prisoners could talk to each other before making their decisions? It depends on the strength of promises. If promises are not binding, then although the two prisoners might agree to *Deny*, they would *Confess* anyway when the time came to choose actions.

A **cooperative game** is a game in which the players can make binding commitments, as opposed to a **noncooperative game**, in which they cannot.

This definition draws the usual distinction between the two theories of games, but the real difference lies in the modelling approach. Both theories start off with the rules of the game, but they differ in the kinds of solution concepts employed. Cooperative game theory is axiomatic, frequently appealing to pareto-optimality,⁴ fairness, and equity. Noncooperative game theory is economic in flavor, with solution concepts based on players maximizing their own utility functions subject to stated constraints. Or, from a different angle: cooperative game theory is a reduced-form theory, which focusses on properties of the outcome rather than on the strategies that achieve the outcome, a method which is appropriate if modelling the process is too complicated. Except for Section 12.2 in the chapter on bargaining, this book is concerned exclusively with noncooperative games. For a good defense of the importance of cooperative game theory, see the essay by Aumann (1996).

In applied economics, the most commonly encountered use of cooperative games is to model bargaining. The Prisoner's Dilemma is a noncooperative game, but it could be modelled as cooperative by allowing the two players not only to communicate but to make binding commitments. Cooperative games often allow players to split the gains from cooperation by making **side-payments**—transfers between themselves that change the prescribed payoffs. Cooperative game theory generally incorporates commitments and side-payments via the solution concept, which can become very elaborate, while noncooperative game theory incorporates them by adding extra actions. The distinction between cooperative and noncooperative games does *not* lie in conflict or absence of conflict, as is shown by the following examples of situations commonly modelled one way or the other:

A cooperative game without conflict. Members of a workforce choose which of equally arduous tasks to undertake to best coordinate with each other.

A cooperative game with conflict. Bargaining over price between a monopolist and a monopsonist.

A noncooperative game with conflict. The Prisoner's Dilemma.

A noncooperative game without conflict. Two companies set a product standard without

⁴If outcome X **strongly pareto-dominates** outcome Y , then all players have higher utility under outcome X . If outcome X **weakly pareto-dominates** outcome Y , some player has higher utility under X , and no player has lower utility. A zero-sum game does not have outcomes that even weakly pareto-dominate other outcomes. All of its equilibria are pareto-efficient, because no player gains without another player losing.

It is often said that strategy combination x “pareto dominates” or “dominates” strategy combination y . Taken literally, this is meaningless, since strategies do not necessarily have any ordering at all— one could define *Deny* as being bigger than *Confess*, but that would be arbitrary. The statement is really shorthand for “The payoff combination resulting from strategy combination x pareto-dominates the payoff combination resulting from strategy y .”

communication.

1.3 Iterated Dominance: The Battle of the Bismarck Sea

Very few games have a dominant strategy equilibrium, but sometimes dominance can still be useful even when it does not resolve things quite so neatly as in the Prisoner's Dilemma. The Battle of the Bismarck Sea, a game I found in Haywood (1954), is set in the South Pacific in 1943. General Imamura has been ordered to transport Japanese troops across the Bismarck Sea to New Guinea, and General Kenney wants to bomb the troop transports. Imamura must choose between a shorter northern route or a longer southern route to New Guinea, and Kenney must decide where to send his planes to look for the Japanese. If Kenney sends his planes to the wrong route he can recall them, but the number of days of bombing is reduced.

The players are Kenney and Imamura, and they each have the same action set, $\{North, South\}$, but their payoffs, given by Table 3, are never the same. Imamura loses exactly what Kenney gains. Because of this special feature, the payoffs could be represented using just four numbers instead of eight, but listing all eight payoffs in Table 3 saves the reader a little thinking. The 2-by-2 form with just four entries is a **matrix game**, while the equivalent table with eight entries is a **bimatrix game**. Games can be represented as matrix or bimatrix games even if they have more than two moves, as long as the number of moves is finite.

Table 3: The Battle of the Bismarck Sea

		Imamura	
		North	South
Kenney	North	2,-2 $\leftrightarrow$	2, -2
	South	1, -1 $\leftarrow$	3, -3
<i>Payoffs to: (Kenney, Imamura)</i>			

Strictly speaking, neither player has a dominant strategy. Kenney would choose *North* if he thought Imamura would choose *North*, but *South* if he thought Imamura would choose *South*. Imamura would choose *North* if he thought Kenney would choose *South*, and he would be indifferent between actions if he thought Kenney would choose *North*. This is what the arrows are showing. But we can still find a plausible equilibrium, using the concept of “weak dominance”.

*Strategy s'_i is **weakly dominated** if there exists some other strategy s''_i for player i which is possibly better and never worse, yielding a higher payoff in some strategy combination and never yielding a lower payoff. Mathematically, s'_i is weakly dominated if there exists s''_i such that*

$$\begin{aligned} \pi_i(s''_i, s_{-i}) &\geq \pi_i(s'_i, s_{-i}) \quad \forall s_{-i}, \text{ and} \\ \pi_i(s''_i, s_{-i}) &> \pi_i(s'_i, s_{-i}) \quad \text{for some } s_{-i}. \end{aligned} \tag{4}$$

One might define a **weak dominance equilibrium** as the strategy combination found by deleting all the weakly dominated strategies of each player. Eliminating weakly dominated strategies does not help much in the Battle of the Bismarck Sea, however. Imamura's strategy of *South* is weakly dominated by the strategy *North* because his payoff from *North* is never smaller than his payoff from *South*, and it is greater if Kenney picks *South*. For Kenney, however, neither strategy is even weakly dominated. The modeller must therefore go a step further, to the idea of the iterated dominance equilibrium.

An iterated dominance equilibrium is a strategy combination found by deleting a weakly dominated strategy from the strategy set of one of the players, recalculating to find which remaining strategies are weakly dominated, deleting one of them, and continuing the process until only one strategy remains for each player.

Applied to the Battle of the Bismarck Sea, this equilibrium concept implies that Kenney decides that Imamura will pick *North* because it is weakly dominant, so Kenney eliminates "Imamura chooses *South*" from consideration. Having deleted one column of Table 3, Kenney has a strongly dominant strategy: he chooses *North*, which achieves payoffs strictly greater than *South*. The strategy combination (*North*, *North*) is an iterated dominance equilibrium, and indeed (*North*, *North*) was the outcome in 1943.

It is interesting to consider modifying the order of play or the information structure in the Battle of the Bismarck Sea. If Kenney moved first, rather than simultaneously with Imamura, (*North*, *North*) would remain an equilibrium, but (*North*, *South*) would also become one. The payoffs would be the same for both equilibria, but the outcomes would be different.

If Imamura moved first, (*North*, *North*) would be the only equilibrium. What is important about a player moving first is that it gives the other player more information before he acts, not the literal timing of the moves. If Kenney has cracked the Japanese code and knows Imamura's plan, then it does not matter that the two players move literally simultaneously; it is better modelled as a sequential game. Whether Imamura literally moves first or whether his code is cracked, Kenney's information set becomes either {Imamura moved *North*} or {Imamura moved *South*} after Imamura's decision, so Kenney's equilibrium strategy is specified as (*North* if Imamura moved *North*, *South* if Imamura moved *South*).

Game theorists often differ in their terminology, and the terminology applied to the idea of eliminating dominated strategies is particularly diverse. The equilibrium concept used in the Battle of the Bismarck Sea might be called **iterated dominance equilibrium** or **iterated dominant strategy equilibrium**, or one might say that the game is **dominance solvable**, that it can be **solved by iterated dominance**, or that the equilibrium strategy combination is **serially undominated**. Sometimes the terms are used to mean deletion of strictly dominated strategies and sometimes to mean deletion of weakly dominated strategies.

The significant difference is between strong and weak dominance. Everyone agrees that no rational player would use a strictly dominated strategy, but it is harder to argue against weakly dominated strategies. In economic models, firms and individuals are often indifferent about their behavior in equilibrium. In standard models of perfect competition, firms earn zero profits but it is crucial that some firms be active in the market and some stay out and produce nothing. If a monopolist knows that customer Smith is willing to pay up to ten dollars for a widget, the monopolist will charge exactly ten dollars to Smith in equilibrium, which makes Smith indifferent about buying and not buying, yet there is no equilibrium unless Smith buys. It is impractical, therefore, to rule out equilibria in which a player is indifferent about his actions. This should be kept in mind later when we discuss the “open-set problem” in Section 4.3.

Another difficulty is multiple equilibria. The dominant strategy equilibrium of any game is unique if it exists. Each player has at most one strategy whose payoff in any strategy combination is strictly higher than the payoff from any other strategy, so only one strategy combination can be formed out of dominant strategies. A strong iterated dominance equilibrium is unique if it exists. A weak iterated dominance equilibrium may not be, because the order in which strategies are deleted can matter to the final solution. If all the weakly dominated strategies are eliminated simultaneously at each round of elimination, the resulting equilibrium is unique, if it exists, but possibly no strategy combination will remain.

Consider Table 4’s Iteration Path Game. The strategy combination (r_1, c_1) and (r_1, c_3) are both iterated dominance equilibria, because each of those strategy combination can be found by iterated deletion. The deletion can proceed in the order (r_3, c_3, c_2, r_2) or in the order (r_2, c_2, c_1, r_3) .

Table 4: The Iteration Path Game

		Column		
		c_1	c_2	c_3
	r_1	2,12	1,10	1,12
Row	r_2	0,12	0,10	0,11
	r_3	0,12	1,10	0,13

Payoffs to: (Row, Column)

Despite these problems, deletion of weakly dominated strategies is a useful tool, and it is part of more complicated equilibrium concepts such as Section 4.1’s “subgame perfectness”.

If we may return to the Battle of the Bismarck Sea, that game is special because the payoffs of the players always sum to zero. This feature is important enough to deserve a name.

*A **zero-sum game** is a game in which the sum of the payoffs of all the players is zero whatever strategies they choose. A game which is not zero-sum is **nonzero-sum game** or **variable-sum**.*

In a zero-sum game, what one player gains, another player must lose. The Battle of the Bismarck Sea is a zero-sum game, but the Prisoner's Dilemma and the Dry Cleaners Game are not, and there is no way that the payoffs in those games can be rescaled to make them zero-sum without changing the essential character of the games.

If a game is zero-sum the utilities of the players can be represented so as to sum to zero under any outcome. Since utility functions are to some extent arbitrary, the sum can also be represented to be non-zero even if the game is zero-sum. Often modellers will refer to a game as zero-sum even when the payoffs do not add up to zero, so long as the payoffs add up to some constant amount. The difference is a trivial normalization.

Although zero-sum games have fascinated game theorists for many years, they are uncommon in economics. One of the few examples is the bargaining game between two players who divide a surplus, but even this is often modelled nowadays as a nonzero-sum game in which the surplus shrinks as the players spend more time deciding how to divide it. In reality, even simple division of property can result in loss— just think of how much the lawyers take out when a divorcing couple bargain over dividing their possessions.

Although the 2-by-2 games in this chapter may seem facetious, they are simple enough for use in modelling economic situations. *The Battle of the Bismarck Sea*, for example, can be turned into a game of corporate strategy. Two firms, Kenney Company and Imamura Incorporated, are trying to maximize their shares of a market of constant size by choosing between the two product designs *North* and *South*. Kenney has a marketing advantage, and would like to compete head-to-head, while Imamura would rather carve out its own niche. The equilibrium is (*North*, *North*).

1.4 Nash Equilibrium: *Boxed Pigs*, the *Battle of the Sexes*, and *Ranked Coordination*

For the vast majority of games, which lack even iterated dominance equilibria, modellers use Nash equilibrium, the most important and widespread equilibrium concept. To introduce Nash equilibrium we will use the game *Boxed Pigs* from Baldwin & Meese (1979). Two pigs are put in a box with a special control panel at one end and a food dispenser at the other end. When a pig presses the panel, at a utility cost of 2 units, 10 units of food are dispensed at the dispenser. One pig is “dominant” (let us assume he is bigger), and if he gets to the dispenser first, the other pig will only get his leavings, worth 1 unit. If, instead, the small pig is at the dispenser first, he eats 4 units, and even if they arrive at the same

time the small pig gets 3 units. Table 5 summarizes the payoffs for the strategies *Press* the panel and *Wait* by the dispenser at the other end.

Table 5: Boxed Pigs

		Small Pig	
		Press	Wait
Big Pig	Press	5, 1 → 4 , 4	
	Wait	9 , -1 → 0, 0	

Payoffs to: (Big Pig, Small Pig)

Boxed Pigs has no dominant strategy equilibrium, because what the big pig chooses depends on what he thinks the small pig will choose. If he believed that the small pig would press the panel, the big pig would wait by the dispenser, but if he believed that the small pig would wait, the big pig would press the panel. There does exist an iterated dominance equilibrium, $(Press, Wait)$, but we will use a different line of reasoning to justify that outcome: Nash equilibrium.

Nash equilibrium is the standard equilibrium concept in economics. It is less obviously correct than dominant strategy equilibrium but more often applicable. Nash equilibrium is so widely accepted that the reader can assume that if a model does not specify which equilibrium concept is being used it is Nash or some refinement of Nash.

*The strategy combination s^* is a **Nash equilibrium** if no player has incentive to deviate from his strategy given that the other players do not deviate. Formally,*

$$\forall i, \pi_i(s_i^*, s_{-i}^*) \geq \pi_i(s_i', s_{-i}^*), \quad \forall s_i'. \quad (5)$$

The strategy combination $(Press, Wait)$ is a Nash equilibrium. The way to approach Nash equilibrium is to propose a strategy combination and test whether each player's strategy is a best response to the others' strategies. If the big pig picks *Press*, the small pig, who faces a choice between a payoff of 1 from pressing and 4 from waiting, is willing to wait. If the small pig picks *Wait*, the big pig, who has a choice between a payoff of 4 from pressing and 0 from waiting, is willing to press. This confirms that $(Press, Wait)$ is a Nash equilibrium, and in fact it is the unique Nash equilibrium.⁵

It is useful to draw arrows in the tables when trying to solve for the equilibrium, since the number of calculations is great enough to soak up quite a bit of mental RAM.

⁵This game, too, has its economic analog. If Bigpig, Inc. introduces granola bars, at considerable marketing expense in educating the public, then Smallpig Ltd. can imitate profitably without ruining Bigpig's sales completely. If Smallpig introduces them at the same expense, however, an imitating Bigpig would hog the market.

Another solution tip, illustrated in Boxed Pigs, is to circle payoffs that dominate other payoffs (or box, them, as is especially suitable here). Double arrows or dotted circles indicate weakly dominant payoffs. Any payoff combination in which every payoff is circled, or which has arrows pointing towards it from every direction, is a Nash equilibrium. I like using arrows better in 2-by-2 games, but circles are better for bigger games, since arrows become confusing when payoffs are not lined up in order of magnitude in the table (see Chapter 2's Table 2).

The pigs in this game have to be smarter than the players in the Prisoner's Dilemma. They have to realize that the only set of strategies supported by self-consistent beliefs is $(Press, Wait)$. The definition of Nash equilibrium lacks the " $\forall s_{-i}$ " of dominant strategy equilibrium, so a Nash strategy need only be a best response to the other Nash strategies, not to all possible strategies. And although we talk of "best responses," the moves are actually simultaneous, so the players are predicting each others' moves. If the game were repeated or the players communicated, Nash equilibrium would be especially attractive, because it is even more compelling that beliefs should be consistent.

Like a dominant strategy equilibrium, a Nash equilibrium can be either weak or strong. The definition above is for a weak Nash equilibrium. To define strong Nash equilibrium, make the inequality strict; that is, require that no player be indifferent between his equilibrium strategy and some other strategy.

Every dominant strategy equilibrium is a Nash equilibrium, but not every Nash equilibrium is a dominant strategy equilibrium. If a strategy is dominant it is a best response to *any* strategies the other players pick, including their equilibrium strategies. If a strategy is part of a Nash equilibrium, it need only be a best response to the other players' *equilibrium* strategies.

The Modeller's Dilemma of Table 6 illustrates this feature of Nash equilibrium. The situation it models is the same as the Prisoner's Dilemma, with one major exception: although the police have enough evidence to arrest the prisoner's as the "probable cause" of the crime, they will not have enough evidence to convict them of even a minor offense if neither prisoner confesses. The northwest payoff combination becomes (0,0) instead of $(-1, -1)$.

		Column	
		<i>Deny</i>	<i>Confess</i>
Table 6: The Modeller's Dilemma	<i>Deny</i>	$\boxed{0}, \boxed{0}$	$\leftrightarrow -10, \boxed{0}$
	<i>Confess</i>	$\boxed{0}, -10$	$\rightarrow \boxed{-8}, \boxed{-8}$
		Payoffs to: (Row, Column)	

The Modeller's Dilemma does not have a dominant strategy equilibrium. It does have what might be called a weak dominant strategy equilibrium, because *Confess* is still

a weakly dominant strategy for each player. Moreover, using this fact, it can be seen that $(Confess, Confess)$ is an iterated dominance equilibrium, and it is a strong Nash equilibrium as well. So the case for $(Confess, Confess)$ still being the equilibrium outcome seems very strong.

There is, however, another Nash equilibrium in the Modeller's Dilemma: $(Deny, Deny)$, which is a weak Nash equilibrium. This equilibrium is weak and the other Nash equilibrium is strong, but $(Deny, Deny)$ has the advantage that its outcome is pareto-superior: $(0, 0)$ is uniformly greater than $(-8, -8)$. This makes it difficult to know which behavior to predict.

The Modeller's Dilemma illustrates a common difficulty for modellers: what to predict when two Nash equilibria exist. The modeller could add more details to the rules of the game, or he could use an **equilibrium refinement**, adding conditions to the basic equilibrium concept until only one strategy combination satisfies the refined equilibrium concept. There is no single way to refine Nash equilibrium. The modeller might insist on a strong equilibrium, or rule out weakly dominated strategies, or use iterated dominance. All of these lead to $(Confess, Confess)$ in the Modeller's Dilemma. Or he might rule out Nash equilibria that are pareto-dominated by other Nash equilibria, and end up with $(Deny, Deny)$. Neither approach is completely satisfactory. In particular, do not be misled into thinking that weak Nash equilibria are to be despised. Often, no Nash equilibrium at all will exist unless the players have the expectation that player B chooses X when he is indifferent between X and Y. It is not that we are picking the equilibrium in which it is assumed B does X when he is indifferent. Rather, we are finding the *only* set of consistent expectations about behavior. (You will read more about this in connection with the "open-set problem" of Section 4.2.)

The Battle of the Sexes

The third game we will use to illustrate Nash equilibrium is the Battle of the Sexes, a conflict between a man who wants to go to a prize fight and a woman who wants to go to a ballet. While selfish, they are deeply in love, and would, if necessary, sacrifice their preferences to be with each other. Less romantically, their payoffs are given by Table 7.

Table 7: The Battle of the Sexes ⁶

⁶Political correctness has led to bowdlerized versions of this game being presented in many game theory books. This is the original, unexpurgated game.

		Woman	
		<i>Prize Fight</i>	<i>Ballet</i>
Man	<i>Prize Fight</i>	2,1 ↑	← 0, 0 ↓
	<i>Ballet</i>	0, 0	→ 1,2

Payoffs to: (Man, Woman)

The Battle of the Sexes does not have an iterated dominance equilibrium. It has two Nash equilibria, one of which is the strategy combination (*Prize Fight*, *Prize Fight*). Given that the man chooses *Prize Fight*, so does the woman; given that the woman chooses *Prize Fight*, so does the man. The strategy combination (*Ballet*, *Ballet*) is another Nash equilibrium by the same line of reasoning.

How do the players know which Nash equilibrium to choose? Going to the fight and going to the ballet are both Nash strategies, but for different equilibria. Nash equilibrium assumes correct and consistent beliefs. If they do not talk beforehand, the man might go to the ballet and the woman to the fight, each mistaken about the other's beliefs. But even if the players do not communicate, Nash equilibrium is sometimes justified by repetition of the game. If the couple do not talk, but repeat the game night after night, one may suppose that eventually they settle on one of the Nash equilibria.

Each of the Nash equilibria in the Battle of the Sexes is pareto-efficient; no other strategy combination increases the payoff of one player without decreasing that of the other. In many games the Nash equilibrium is not pareto-efficient: (*Confess*, *Confess*), for example, is the unique Nash equilibrium of the Prisoner's Dilemma, although its payoffs of $(-8, -8)$ are pareto- inferior to the $(-1, -1)$ generated by (*Deny*, *Deny*).

Who moves first is important in the Battle of the Sexes, unlike any of the three previous games we have looked at. If the man could buy the fight ticket in advance, his commitment would induce the woman to go to the fight. In many games, but not all, the player who moves first (which is equivalent to commitment) has a **first-mover advantage**.

The Battle of the Sexes has many economic applications. One is the choice of an industrywide standard when two firms have different preferences but both want a common standard to encourage consumers to buy the product. A second is to the choice of language used in a contract when two firms want to formalize a sales agreement but they prefer different terms. Both sides might, for example, want to add a "liquidated damages" clause which specifies damages for breach, rather than trust to the courts to estimate a number later, but one firm wants the value to be \$10,000 and the other firm wants \$12,000.

Coordination Games

Sometimes one can use the size of the payoffs to choose between Nash equilibria. In the following game, players Smith and Jones are trying to decide whether to design the

computers they sell to use large or small floppy disks. Both players will sell more computers if their disk drives are compatible, as shown in Table 8.

Table 8: Ranked Coordination

		Jones	
		<i>Large</i>	<i>Small</i>
Smith	<i>Large</i>	2,2	← −1, −1
	<i>Small</i>	−1, −1	→ 1,1

Payoffs to: (Smith, Jones)

The strategy combinations (*Large, Large*) and (*Small, Small*) are both Nash equilibria, but (*Large, Large*) pareto-dominates (*Small, Small*). Both players prefer (*Large, Large*), and most modellers would use the pareto- efficient equilibrium to predict the actual outcome. We could imagine that it arises from pre-game communication between Smith and Jones taking place outside of the specification of the model, but the interesting question is what happens if communication is impossible. Is the pareto-efficient equilibrium still more plausible? The question is really one of psychology rather than economics.

Ranked Coordination is one of a large class of games called **coordination games**, which share the common feature that the players need to coordinate on one of multiple Nash equilibria. Ranked Coordination has the additional feature that the equilibria can be pareto ranked. Section 3.2 will return to problems of coordination to discuss the concepts of “correlated strategies” and “cheap talk.” These games are of obvious relevance to analyzing the setting of standards; see, e.g., Michael Katz & Carl Shapiro (1985) and Joseph Farrell & Garth Saloner (1985). They can be of great importance to the wealth of economies— just think of the advantages of standard weights and measures (or read Charles Kindleberger (1983) on their history). Note, however, that not all apparent situations of coordination on pareto-inferior equilibria turn out to be so. One oft-cited coordination problem is that of the QWERTY typewriter keyboard, developed in the 1870s when typing had to proceed slowly to avoid jamming. QWERTY became the standard, although it has been claimed that the faster speed possible with the Dvorak keyboard would amortize the cost of retraining full-time typists within ten days (David [1985]). Why large companies would not retrain their typists is difficult to explain under this story, and Liebowitz & Margolis (1990) show that economists have been too quick to accept claims that QWERTY is inefficient. English language spelling is a better example.

Table 9 shows another coordination game, Dangerous Coordination, which has the same equilibria as Ranked Coordination, but differs in the off-equilibrium payoffs. If an experiment were conducted in which students played Dangerous Coordination against each other, I would not be surprised if (*Small, Small*), the pareto-dominated equilibrium, were the one that was played out. This is true even though (*Large, Large*) is still a Nash

equilibrium; if Smith thinks that Jones will pick *Large*, Smith is quite willing to pick *Large* himself. The problem is that if the assumptions of the model are weakened, and Smith cannot trust Jones to be rational, well-informed about the payoffs of the game, and unconfused, then Smith will be reluctant to pick *Large* because his payoff if Jones picks *Small* is then -1,000. He would play it safe instead, picking *Small* and ensuring a payoff of at least -1 . In reality, people do make mistakes, and with such an extreme difference in payoffs, even a small probability of a mistake is important, so $(Large, Large)$ would be a bad prediction.

Table 9: Dangerous Coordination

		Jones	
		<i>Large</i>	<i>Small</i>
Smith	<i>Large</i>	2,2 ←	-1000, -1
	<i>Small</i>	-1, -1 →	1,1

Payoffs to: (Smith, Jones)

Games like Dangerous Coordination are a major concern in the 1988 book by Harsanyi and Selten, two of the giants in the field of game theory. I will not try to describe their approach here, except to say that it is different from my own. I do not consider the fact that one of the Nash equilibria of Dangerous Coordination is a bad prediction as a heavy blow against Nash equilibrium. The bad prediction is based on two things: using the Nash equilibrium concept, and using the game Dangerous Coordination. If Jones might be confused about the payoffs of the game, then the game actually being played out is not Dangerous Coordination, so it is not surprising that it gives poor predictions. The rules of the game ought to describe the probabilities that the players are confused, as well as the payoffs if they take particular actions. If confusion is an important feature of the situation, then the two-by-two game of Table 9 is the wrong model to use, and a more complicated game of incomplete information of the kind described in Chapter 2 is more appropriate. Again, as with the Prisoner's Dilemma, the modeller's first thought on finding that the model predicts an odd result should not be "Game theory is bunk," but the more modest "Maybe I'm not describing the situation correctly" (or even "Maybe I should not trust my 'common sense' about what will happen").

Nash equilibrium is more complicated but also more useful than it looks. Jumping ahead a bit, consider a game slightly more complex than the ones we have seen so far. Two firms are choosing outputs Q_1 and Q_2 simultaneously. The Nash equilibrium is a pair of numbers (Q_1^*, Q_2^*) such that neither firm would deviate unilaterally. This troubles the beginner, who says to himself,

"Sure, Firm 1 will pick Q_1^* if it thinks Firm 2 will pick Q_2^* . But Firm 1 will realize that if it makes Q_1 bigger, then Firm 2 will react by making Q_2 smaller. So the

situation is much more complicated, and (Q_1^*, Q_2^*) is not a Nash equilibrium. Or, if it is, Nash equilibrium is a bad equilibrium concept.”

If there is a problem in this model, it is not Nash equilibrium but the model itself. Nash equilibrium makes perfect sense as a stable outcome in this model. The beginner’s hypothetical is false because if Firm 1 chooses something other than Q_1^* , Firm 2 would not observe the deviation till it was too late to change Q_2 —remember, this is a simultaneous move game. The beginner’s worry is really about the rules of the game, not the equilibrium concept. He seems to prefer a game in which the firms move sequentially, or maybe a repeated version of the game. If Firm 1 moved first, and then Firm 2, then Firm 1’s strategy would still be a single number, Q_1 , but Firm 2’s strategy—its action rule—would have to be a function, $Q_2(Q_1)$. A Nash equilibrium would then consist of an equilibrium number, Q_1^{**} , and an equilibrium function, $Q_2^{**}(Q_1)$. The two outputs actually chosen, Q_1^{**} and $Q_2^{**}(Q_1^{**})$, will be different from the Q_1^* and Q_2^* in the original game. And they should be different—the new model represents a very different real-world situation. Look ahead, and you will see that these are the Cournot and Stackelberg models of Chapter 3.

One lesson to draw from this is that it is essential to figure out the mathematical form the strategies take before trying to figure out the equilibrium. In the simultaneous move game, the strategy combination is a pair of non-negative numbers. In the sequential game, the strategy combination is one nonnegative number and one function defined over the nonnegative numbers. Students invariably make the mistake of specifying Firm 2’s strategy as a number, not a function. This is a far more important point than any beginner realizes. Trust me—you’re going to make this mistake sooner or later, so it’s worth worrying about.

1.5 Focal Points

Schelling’s book, *The Strategy of Conflict* (1960) is a classic in game theory, even though it contains no equations or Greek letters. Although it was published more than 40 years ago, it is surprisingly modern in spirit. Schelling is not a mathematician but a strategist, and he examines such things as threats, commitments, hostages, and delegation that we will examine in a more formal way in the remainder of this book. He is perhaps best known for his coordination games. Take a moment to decide on a strategy in each of the following games, adapted from Schelling, which you win by matching your response to those of as many of the other players as possible.

1 Circle one of the following numbers: 100, 14, 15, 16, 17, 18.

2 Circle one of the following numbers 7, 100, 13, 261, 99, 666.

3 Name Heads or Tails.

4 Name Tails or Heads.

5 You are to split a pie, and get nothing if your proportions add to more than 100 percent.

6 You are to meet somebody in New York City. When? Where?

Each of the games above has many Nash equilibria. In example (1), if each player thinks every other player will pick 14, he will too, and this is self-confirming; but the same is true if each player thinks every other player will pick 15. But to a greater or lesser extent they also have Nash equilibria that seem more likely. Certain of the strategy combinations are **focal points**: Nash equilibria which for psychological reasons are particularly compelling.

Formalizing what makes a strategy combination a focal point is hard and depends on the context. In example (1), 100 is a focal point, because it is a number clearly different from all the others, it is biggest, and it is first in the listing. In example (2), Schelling found 7 to be the most common strategy, but in a group of Satanists, 666 might be the focal point. In repeated games, focal points are often provided by past history. Examples (3) and (4) are identical except for the ordering of the choices, but that ordering might make a difference. In (5), if we split a pie once, we are likely to agree on 50:50. But if last year we split a pie in the ratio 60:40, that provides a focal point for this year. Example (6) is the most interesting of all. Schelling found surprising agreement in independent choices, but the place chosen depended on whether the players knew New York well or were unfamiliar with the city.

The **boundary** is a particular kind of focal point. If player Russia chooses the action of putting his troops anywhere from one inch to 100 miles away from the Chinese border, player China does not react. If he chooses to put troops from one inch to 100 miles *beyond* the border, China declares war. There is an arbitrary discontinuity in behavior at the boundary. Another example, quite vivid in its arbitrariness, is the rallying cry, “Fifty-Four Forty or Fight!,” which refers to the geographic parallel claimed as the boundary by jingoist Americans in the Oregon dispute between Britain and the United States in the 1840s.⁷

Once the boundary is established it takes on additional significance because behavior with respect to the boundary conveys information. When Russia crosses an established boundary, that tells China that Russia intends to make a serious incursion further into China. Boundaries must be sharp and well known if they are not to be violated, and a large part of both law and diplomacy is devoted to clarifying them. Boundaries can also arise in business: two companies producing an unhealthful product might agree not to mention relative healthfulness in their advertising, but a boundary rule like “Mention unhealthfulness if you like, but don’t stress it,” would not work.

Mediation and **communication** are both important in the absence of a clear focal point. If players can communicate, they can tell each other what actions they will take, and sometimes, as in Ranked Coordination, this works, because they have no motive to lie. If the players cannot communicate, a mediator may be able to help by suggesting an equilibrium to all of them. They have no reason not to take the suggestion, and they would

⁷The threat was not credible: that parallel is now deep in British Columbia.

use the mediator even if his services were costly. Mediation in cases like this is as effective as arbitration, in which an outside party imposes a solution.

One disadvantage of focal points is that they lead to inflexibility. Suppose the pareto-superior equilibrium (*Large, Large*) were chosen as a focal point in Ranked Coordination, but the game was repeated over a long interval of time. The numbers in the payoff matrix might slowly change until (*Small, Small*) and (*Large, Large*) both had payoffs of, say, 1.6, and (*Small, Small*) started to dominate. When, if ever, would the equilibrium switch?

In Ranked Coordination, we would expect that after some time one firm would switch and the other would follow. If there were communication, the switch point would be at the payoff of 1.6. But what if the first firm to switch is penalized more? Such is the problem in oligopoly pricing. If costs rise, so should the monopoly price, but whichever firm raises its price first suffers a loss of market share.

NOTES

N1.2 Dominant Strategies: The Prisoner's Dilemma

- Many economists are reluctant to use the concept of cardinal utility (see Starmer [2000]), and even more reluctant to compare utility across individuals (see Cooter & Rappoport [1984]). Noncooperative game theory never requires interpersonal utility comparisons, and only ordinal utility is needed to find the equilibrium in the Prisoner's Dilemma. So long as each player's rank ordering of payoffs in different outcomes is preserved, the payoffs can be altered without changing the equilibrium. In general, the dominant strategy and pure strategy Nash equilibria of games depend only on the ordinal ranking of the payoffs, but the mixed strategy equilibria depend on the cardinal values. Compare Section 3.2's Chicken game with Section 5.6's Hawk-Dove.
- If we consider only the ordinal ranking of the payoffs in 2-by-2 games, there are 78 distinct games in which each player has strict preference ordering over the four outcomes and 726 distinct games if we allow ties in the payoffs. Rapoport, Guyer & Gordon's 1976 book, *The 2x2 Game*, contains an exhaustive description of the possible games.
- The Prisoner's Dilemma was so named by Albert Tucker in an unpublished paper, although the particular 2-by-2 matrix, discovered by Dresher and Flood, was already well known. Tucker was asked to give a talk on game theory to the psychology department at Stanford, and invented a story to go with the matrix, as recounted in Straffin (1980), pp. 101-18 of Poundstone (1992), and pp. 171-3 of Raiffa (1992).
- In the Prisoner's Dilemma the notation *cooperate* and *defect* is often used for the moves. This is bad notation, because it is easy to confuse with *cooperative* games and with *deviations*. It is also often called the Prisoners' Dilemma (rs', not r's) ; whether one looks at from the point of the individual or the group, the prisoners have a problem.
- The Prisoner's Dilemma is not always defined the same way. If we consider just ordinal payoffs, then the game in Table 10 is a Prisoner's Dilemma if $T(\text{temptation}) > R(\text{revolt}) > P(\text{punishment}) > S(\text{Sucker})$, where the terms in parentheses are mnemonics. This is standard notation; see, for example, Rapoport, Guyer & Gordon (1976), p. 400. If the game is repeated, the cardinal values of the payoffs can be important. The requirement $2R > T + S > 2P$ should be added if the game is to be a standard Prisoner's Dilemma, in which (*Deny, Deny*) and (*Confess, Confess*) are the best and worst possible outcomes in terms of the sum of payoffs. Section 5.3 will show that an asymmetric game called the One-Sided Prisoner's Dilemma has properties similar to the standard Prisoner's Dilemma, but does not fit this definition.

Sometimes the game in which $2R < T + S$ is also called a prisoner's dilemma, but in it the sum of the players' payoffs is maximized when one confesses and the other denies. If the game were repeated or the prisoners could use the correlated equilibria defined in Section 3.2, they would prefer taking turns being confessed against, which would make the game a coordination game similar to the Battle of the Sexes. David Shimko has suggested the name "Battle of the Prisoners" for this (or, perhaps, the "Sex Prisoners' Dilemma").

Table 10: A General Prisoner's Dilemma

		Column	
		<i>Deny</i>	<i>Confess</i>
Row	<i>Deny</i>	R, R →	S, T ↓
	<i>Confess</i>	T, S →	P, P

Payoffs to: (Row, Column)

- Herodotus (429 B.C., III-71) describes an early example of the reasoning in the Prisoner's Dilemma in a conspiracy against the Persian emperor. A group of nobles met and decided to overthrow the emperor, and it was proposed to adjourn till another meeting. One of them named Darius then spoke up and said that if they adjourned, he knew that one of them would go straight to the emperor and reveal the conspiracy, because if nobody else did, he would himself. Darius also suggested a solution— that they immediately go to the palace and kill the emperor.

The conspiracy also illustrates a way out of coordination games. After killing the emperor, the nobles wished to select one of themselves as the new emperor. Rather than fight, they agreed to go to a certain hill at dawn, and whoever's horse neighed first would become emperor. Herodotus tells how Darius's groom manipulated this randomization scheme to make him the new emperor.

- Philosophers are intrigued by the Prisoner's Dilemma: see Campbell & Sowden (1985), a collection of articles on the Prisoner's Dilemma and the related Newcombe's paradox. Game theory has even been applied to theology: if one player is omniscient or omnipotent, what kind of equilibrium behavior can we expect? See Brams (1983).

N1.4 Nash Equilibrium: Boxed Pigs, the Battle of the Sexes, and Ranked Coordination

- I invented the payoffs for Boxed Pigs from the description of one of the experiments in Baldwin & Meese (1979). They do *not* think of this as an experiment in game theory, and they describe the result in terms of “reinforcement.” The Battle of the Sexes is taken from p. 90 of Luce & Raiffa (1957). I have changed their payoffs of $(-1, -1)$ to $(-5, -5)$ to fit the story.
- Some people prefer the term “equilibrium point” to “Nash equilibrium,” but the latter is more euphonious, since the discoverer's name is “Nash” and not “Mazurkiewicz.”
- Bernheim (1984a) and Pearce (1984) use the idea of mutually consistent beliefs to arrive at a different equilibrium concept than Nash. They define a **rationalizable strategy** to be a strategy which is a best response for some set of rational beliefs in which a player believes that the other players choose their best responses. The difference from Nash is that not all players need have the same beliefs concerning which strategies will be chosen, nor need their beliefs be consistent.

This idea is attractive in the context of Bertrand games (see Section 3.6). The Nash equilibrium in the Bertrand game is weakly dominated— by picking any other price above marginal cost, which yields the same profit of zero as does the equilibrium. Rationalizability rules that out.

- Jack Hirshleifer (1982) uses the name “the Tender Trap” for a game essentially the same as Ranked Coordination, and the name “the Assurance Game” has also been used for it.

- O. Henry's story, "The Gift of the Magi" is about a coordination game noteworthy for the reason communication is ruled out. A husband sells his watch to buy his wife combs for Christmas, while she sells her hair to buy him a watch fob. Communication would spoil the surprise, a worse outcome than discoordination.
- Macroeconomics has more game theory in it than is readily apparent. The macroeconomic concept of *rational expectations* faces the same problems of multiple equilibria and consistency of expectations as Nash equilibrium. Game theory is now often explicitly used in macroeconomics; see the books by Canzoneri & Henderson (1991) and Cooper (1999).

N1.5 Focal Points

- Besides his 1960 book, Schelling has written books on diplomacy (1966) and the oddities of aggregation (1978). Political scientists are now looking at the same issues more technically; see Brams & Kilgour (1988) and Ordeshook (1986). Riker (1986) and Muzzio's 1982 book, *Watergate Games* are absorbing examples of how game theory can be used to analyze specific historical episodes.
- In Chapter 12 of *The General Theory*, Keynes (1936) suggests that the stock market is a game with multiple equilibria, like a contest in which a newspaper publishes the faces of 20 girls, and contestants submit the name of the one they think most people would submit as the prettiest. When the focal point changes, big swings in predictions about beauty and value result.
- Not all of what we call boundaries have an arbitrary basis. If the Chinese cannot defend themselves as easily once the Russians cross the boundary at the Amur River, they have a clear reason to fight there.
- Crawford & Haller (1990) take a careful look at focalness in repeated coordination games by asking which equilibria are objectively different from other equilibria, and how a player can learn through repetition which equilibrium the other players intend to play. If on the first repetition the players choose strategies that are Nash with respect to each other, it seems focal for them to continue playing those strategies, but what happens if they begin in disagreement?

Problems

1.1. Nash and Iterated Dominance

- (a) Show that every iterated dominance equilibrium s^* is Nash.
- (b) Show by counterexample that not every Nash equilibrium can be generated by iterated dominance.
- (c) Is every iterated dominance equilibrium made up of strategies that are not weakly dominated?

1.2. 2-by-2 Games

Find examples of 2-by-2 games with the following properties:

- (a) No Nash equilibrium (you can ignore mixed strategies).
- (b) No weakly pareto-dominant strategy combination.
- (c) At least two Nash equilibria, including one equilibrium that pareto-dominates all other strategy combinations.
- (d) At least three Nash equilibria.

1.3. Pareto Dominance (from notes by Jong-Shin Wei)

- (a) If a strategy combination s^* is a dominant strategy equilibrium, does that mean it weakly pareto-dominates all other strategy combinations?
- (b) If a strategy combination s strongly pareto-dominates all other strategy combinations, does that mean it is a dominant strategy equilibrium?
- (c) If s weakly pareto-dominates all other strategy combinations, then must it be a Nash equilibrium?

1.4. Discoordination

Suppose that a man and a woman each choose whether to go to a prize fight or a ballet. The man would rather go to the prize fight, and the woman to the ballet. What is more important to them, however, is that the man wants to show up to the same event as the woman, but the woman wants to avoid him.

- (a) Construct a game matrix to illustrate this game, choosing numbers to fit the preferences described verbally.
- (b) If the woman moves first, what will happen?
- (c) Does the game have a first-mover advantage?

- (d) Show that there is no Nash equilibrium if the players move simultaneously.

1.5. Drawing Outcome Matrices

It can be surprisingly difficult to look at a game using new notation. In this exercise, redraw the outcome matrix in a different form than in the main text. In each case, read the description of the game and draw the outcome matrix as instructed. You will learn more if you do this from the description, without looking at the conventional outcome matrix.

- (a) The Battle of the Sexes (Table 7). Put (*Prize Fight*, *Prize Fight*) in the northwest corner, but make the woman the row player.
- (b) The Prisoner's Dilemma (Table 2). Put (*Confess*, *Confess*) in the northwest corner.
- (c) The Battle of the Sexes (Table 7). Make the man the row player, but put (*Ballet*, *Prize Fight*) in the northwest corner.

1.6. Finding Nash Equilibria

Find the Nash equilibria of the game illustrated in Table 11. Can any of them be reached by iterated dominance?

Table 11: An Abstract Game

		Column		
		<i>Left</i>	<i>Middle</i>	<i>Right</i>
	<i>Up</i>	10,10	0,0	-1,15
Row:	<i>Sideways</i>	-12,1	8,8	-1,-1
	<i>Down</i>	15,1	8,-1	0,0

Payoffs to: (Row, Column).

1.7. Finding More Nash Equilibria

Find the Nash equilibria of the game illustrated in Table 12. Can any of them be reached by iterated dominance?

Table 12: Flavor and Texture

		Brydox	
		<i>Flavor</i>	<i>Texture</i>
	<i>Flavor</i>	-2,0	0,1
Apex:	<i>Texture</i>	-1,-1	0,-2

Payoffs to: (Apex, Brydox).

1.8. Which Game?

Table 13 is like the payoff matrix for what game that we have seen?

- (a) a version of the Battle of the Sexes.
- (b) a version of the Prisoner's Dilemma.
- (c) a version of Pure Coordination.
- (d) a version of the Legal Settlement Game.
- (e) none of the above.

Table 13: Which Game?

		COL	
		A	B
ROW	A	3,3	0,1
	B	5,0	-1,-1

1.9. Choosing Computers

The problem of deciding whether to adopt IBM or HP computers by two offices in a company is most like which game that we have seen?

1.10. Finding Equilibria

Find the pure-strategy Nash equilibria of the game in Table 14.

1.11. Campaign Contributions

The large Wall Street investment banks have recently agreed not to make campaign contributions to state treasurers, which up till now has been a common practice. What was the game in the past, and why can the banks expect this agreement to hold fast?

1.12. Three-by-Three Equilibria

Identify any dominated strategies and any Nash equilibria in pure strategies in the game of Table 15.

Table 15: A Three-By-Three Game

		Column		
		<i>Left</i>	Middle	<i>Right</i>
Row:	<i>Up</i>	1,4	5, -1	0, 1
	Sideways	-1, 0	-2,-2	-3, 4
	<i>Down</i>	0, 3	9,-1	5, 0

Payoffs to: (Row, Column).

1.13. A Sequential Prisoner's Dilemma

Suppose Row moves first, then Column, in the Prisoner's Dilemma. What are the possible actions?

What are the possible strategies? Construct a normal form, showing the relationship between strategy profiles and payoffs.

Hint: The normal form is *not* a two-by-two matrix here.

2 Information

6 September 1999. January 3, 2000. December 7, 2003. xxx Footnotes starting with xxx are the author's notes to himself. Comments are welcomed.

2.1 The Strategic and Extensive Forms of a Game

If half of strategic thinking is predicting what the other player will do, the other half is figuring out what he knows. Most of the games in Chapter 1 assumed that the moves were simultaneous, so the players did not have a chance to learn each other's private information by observing each other. Information becomes central as soon as players move in sequence. The important difference, in fact, between simultaneous-move games and sequential-move games is that in sequential-move games the second player acquires the information on how the first player moved before he must make his own decision.

Section 2.1 shows how to use the strategic form and the extensive form to describe games with sequential moves. Section 2.2 shows how the extensive form, or game tree, can be used to describe the information available to a player at each point in the game. Section 2.3 classifies games based on the information structure. Section 2.4 shows how to redraw games with incomplete information so that they can be analyzed using the Harsanyi transformation, and derives Bayes's Rule for combining a player's prior beliefs with information which he acquires in the course of the game. Section 2.5 concludes the chapter with the Png Settlement Game, an example of a moderately complex sequential-move game.

The Strategic Form and the Outcome Matrix

Games with moves in sequence require more care in presentation than single-move games. In Section 1.4 we used the 2-by-2 form, which for the game *Ranked Coordination* is shown in Table 1.

Table 1: Ranked Coordination

		Jones	
		<i>Large</i>	<i>Small</i>
Smith	<i>Large</i>	2,2	← −1, −1
	<i>Small</i>	↑ −1, −1	↓ 1,1

Payoffs to: (Smith, Jones)

Because strategies are the same as actions in *Ranked Coordination* and the outcomes are simple, the 2-by-2 form in Table 1 accomplishes two things: it relates strategy combinations to payoffs, and action combinations to outcomes. These two mappings are called the strategic form and the outcome matrix, and in more complicated games they are distinct from each other. The strategic form shows what payoffs result from each possible strategy

combination, while the outcome matrix shows what outcome results from each possible action combination. The definitions below use n to denote the number of players, k the number of variables in the outcome vector, p the number of strategy combinations, and q the number of action combinations.

The **strategic form** (or **normal form**) consists of

1 All possible strategy combinations $s^1, s^2, \dots, s^p$.

2 Payoff functions mapping s^i onto the payoff n -vector π^i , ($i = 1, 2, \dots, p$).

The **outcome matrix** consists of

1 All possible action combinations $a^1, a^2, \dots, a^q$.

2 Outcome functions mapping a^i onto the outcome k -vector z^i , ($i = 1, 2, \dots, q$).

Consider the following game based on Ranked Coordination, which we will call Follow-the-Leader I since we will create several variants of the game. The difference from Ranked Coordination is that Smith moves first, committing himself to a certain disk size no matter what size Jones chooses. The new game has an outcome matrix identical to Ranked Coordination, but its strategic form is different because Jones's strategies are no longer single actions. Jones's strategy set has four elements,

$$\left\{ \begin{array}{l} \text{(If Smith chose } Large, \text{ choose } Large; \text{ if Smith chose } Small, \text{ choose } Large), \\ \text{(If Smith chose } Large, \text{ choose } Large; \text{ if Smith chose } Small, \text{ choose } Small), \\ \text{(If Smith chose } Large, \text{ choose } Small; \text{ if Smith chose } Small, \text{ choose } Large), \\ \text{(If Smith chose } Large, \text{ choose } Small; \text{ if Smith chose } Small, \text{ choose } Small) \end{array} \right\}$$

which we will abbreviate as

$$\left\{ \begin{array}{l} (L|L, L|S), \\ (L|L, S|S), \\ (S|L, L|S), \\ (S|L, S|S) \end{array} \right\}$$

Follow-the-Leader I illustrates how adding a little complexity can make the strategic form too obscure to be very useful. The strategic form is shown in Table 2, with equilibria boldfaced and labelled E_1 , E_2 , and E_3 .

		Jones			
		J_1	J_2	J_3	J_4
		$L L, L S$	$L L, S S$	$S L, L S$	$S L, S S$
Smith	$S_1 : Large$	[2], [2] (E_1)	[2], [2] (E_2)	~~[1], [1]~~ -1	-1, -1
	$S_2 : Small$	-1, -1	1, ~~[1]~~	~~[1], [1]~~ -1	[1], [1] (E_3)

Payoffs to: (Smith, Jones)

Table 2: Follow-the-Leader I

Equilibrium	Strategies	Outcome
E_1	$\{Large, (L L, L S)\}$	Both pick <i>Large</i>
E_2	$\{Large, (L L, S S)\}$	Both pick <i>Large</i>
E_3	$\{Small, (S L, S S)\}$	Both pick <i>Small</i>

Consider why E_1 , E_2 , and E_3 are Nash equilibria. In Equilibrium E_1 , Jones will respond with *Large* regardless of what Smith does, so Smith quite happily chooses *Large*. Jones would be irrational to choose *Large* if Smith chose *Small* first, but that event never happens in equilibrium. In Equilibrium E_2 , Jones will choose whatever Smith chose, so Smith chooses *Large* to make the payoff 2 instead of 1. In Equilibrium E_3 , Smith chooses *Small* because he knows that Jones will respond with *Small* whatever he does, and Jones is willing to respond with *Small* because Smith chooses *Small* in equilibrium. Equilibria E_1 and E_3 are not completely sensible, but, except for a little discussion in the context of the game tree, we will defer to Chapter 4 the discussion of how to redefine the equilibrium concept to rule them out.

The Order of Play

The “normal form” is rarely used in modelling games of any complexity. Already, in Section 1.1, we have seen an easier way to model a sequential game: the *order of play*. For it Follow the Leader I, this would be:

- 1 Smith chooses his disk size to be either *Large* or *Small*.
- 2 Jones chooses his disk size to be either *Large* or *Small*.

The reason I have retained the concept of the normal form in this edition is that it reinforces the idea of laying out all the possible strategies and comparing their payoffs. The order of play, however, gives us a better way to describe games, as I will explain next.

The Extensive Form and the Game Tree

Two other ways to describe a game are the extensive form and the game tree. First we need to define their building blocks. As you read the definitions, you may wish to refer to Figure 1 as an example.

A **node** is a point in the game at which some player or Nature takes an action, or the game ends.

A **successor** to node X is a node that may occur later in the game if X has been reached.

A **predecessor** to node X is a node that must be reached before X can be reached.

A **starting node** is a node with no predecessors.

An **end node** or **end point** is a node with no successors.

A **branch** is one action in a player’s action set at a particular node.

A **path** is a sequence of nodes and branches leading from the starting node to an end node.

These concepts can be used to define the extensive form and the game tree.

The **extensive form** is a description of a game consisting of

- 1 A configuration of nodes and branches running without any closed loops from a single starting node to its end nodes.
- 2 An indication of which node belongs to which player.
- 3 The probabilities that Nature uses to choose different branches at its nodes.
- 4 The information sets into which each player's nodes are divided.
- 5 The payoffs for each player at each end node.

The **game tree** is the same as the extensive form except that (5) is replaced with
 5' The outcomes at each end node.

“Game tree” is a looser term than “extensive form.” If the outcome is defined as the payoff combination, one payoff for each player, then the extensive form is the same as the game tree.

The extensive form for Follow-the-Leader I is shown in Figure 1. We can see why Equilibria E_1 and E_3 of Table 2 are unsatisfactory even though they are Nash equilibria. If the game actually reached nodes J_1 or J_2 , Jones would have dominant actions, *Small* at J_1 and *Large* at J_2 , but E_1 and E_3 specify other actions at those nodes. In Chapter 4 we will return to this game and show how the Nash concept can be refined to make E_2 the only equilibrium.




Figure 1: Follow-the-Leader I in Extensive Form

The extensive form for *Ranked Coordination*, shown in Figure 2, adds dotted lines to the extensive form for *Follow-the-Leader I*. Each player makes a single decision between two actions. The moves are simultaneous, which we show by letting Smith move first, but not letting Jones know how he moved. The dotted line shows that Jones's knowledge stays the same after Smith moves. All Jones knows is that the game has reached some node within the information set defined by the dotted line; he does not know the exact node reached.




Figure 2: Ranked Coordination in Extensive Form

The Time Line

The **time line**, a line showing the order of events, is another way to describe games. Time lines are particularly useful for games with continuous strategies, exogenous arrival of information, and multiple periods, games that are frequently used in the accounting and finance literature. A typical time line is shown in Figure 3a, which represents a game that will be described in Section 11.5.




Figure 3: The Time Line for Stock Underpricing: (a) A Good Time Line; (b) A Bad Time Line

The time line illustrates the order of actions and events, not necessarily the passage of time. Certain events occur in an instant, others over an interval. In Figure 3a, events 2 and 3 occur immediately after event 1, but events 4 and 5 might occur ten years later. We sometimes refer to the sequence in which decisions are made as **decision time** and the interval over which physical actions are taken as **real time**. A major difference is that players put higher value on payments received earlier in real time because of time preference (on which see the appendix).

A common and bad modelling habit is to restrict the use of the dates on the time line to separating events in real time. Events 1 and 2 in Figure 2.3a are not separated by real time: as soon as the entrepreneur learns the project's value, he offers to sell stock. The modeller might foolishly decide to depict his model by a picture like Figure 3b in which both events happen at date 1. Figure 3b is badly drawn, because readers might wonder which event occurs first or whether they occur simultaneously. In more than one seminar, 20 minutes of heated and confusing debate could have been avoided by 10 seconds care to delineate the order of events.

2.2: Information Sets

A game's information structure, like the order of its moves, is often obscured in the strategic form. During the Watergate affair, Senator Baker became famous for the question "How much did the President know, and when did he know it?". In games, as in scandals, these are the big questions. To make this precise, however, requires technical definitions so that one can describe who knows what, and when. This is done using the "information set," the set of nodes a player thinks the game might have reached, as the basic unit of knowledge.

*Player i 's **information set** ω_i at any particular point of the game is the set of different nodes in the game tree that he knows might be the actual node, but between which he cannot distinguish by direct observation.*

As defined here, the information set for player i is a set of nodes belonging to one player but on different paths. This captures the idea that player i knows whose turn it is to move, but not the exact location the game has reached in the game tree. Historically, player i 's information set has been defined to include only nodes at which player i moves, which is appropriate for single-person decision theory, but leaves a player's knowledge undefined for most of any game with two or more players. The broader definition allows comparison of information across players, which under the older definition is a comparison of apples and oranges.

In the game in Figure 4, Smith moves at node S_1 in 1984 and Jones moves at nodes J_1, J_2, J_3 , and J_4 in 1985 or 1986. Smith knows his own move, but Jones can tell only whether Smith has chosen the moves which lead to J_1, J_2 , or "other"; he cannot distinguish between J_3 and J_4 . If Smith has chosen the move leading to J_3 , his own information set is simply $\{J_3\}$, but Jones's information set is $\{J_3, J_4\}$.

One way to show information sets on a diagram is to put dashed lines around or between nodes in the same information set. The resulting diagrams can be very cluttered, so it is often more convenient to draw dashed lines around the information sets of the

player making the move at a node. The dashed lines in Figure 4 show that J_3 and J_4 are in the same information set for Jones, even though they are in different information sets for Smith. An expressive synonym for information set which is based on the appearance of these diagrams is **cloud**; one would say that nodes J_3 and J_4 are in the same cloud, so that while Jones can tell that the game has reached that cloud, he cannot pierce the fog to tell exactly which node has been reached.




Figure 4: Information Sets and Information Partitions.

One node cannot belong to two different information sets of a single player. If node J_3 belonged to information sets $\{J_2, J_3\}$ and $\{J_3, J_4\}$ (unlike in Figure 4), then if the game reached J_3 , Jones would not know whether he was at a node in $\{J_2, J_3\}$ or a node in $\{J_3, J_4\}$ —which would imply that they were really the same information set.

If the nodes in one of Jones’s information sets are nodes at which he moves, his action set must be the same at each node, because he knows his own action set (though his actions might differ later on in the game depending on whether he advances from J_3 or J_4). Jones has the same action sets at nodes J_3 and J_4 , because if he had some different action available at J_3 he would know he was there and his information set would reduce to just $\{J_3\}$. For the same reason, nodes J_1 and J_2 could not be put in the same information set; Jones must know whether he has three or four moves in his action set. We also require end nodes to be in different information sets for a player if they yield him different payoffs.

With these exceptions, we do not include in the information structure of the game any information acquired by a player’s rational deductions. In Figure 4, for example, it seems clear that Smith would choose *Bottom*, because that is a dominant strategy — his payoff is 8 instead of the 4 from *Lower*, regardless of what Jones does. Jones should be able to deduce this, but even though this is an uncontroversial deduction, it is none the less a deduction, not an observation, so the game tree does not split J_3 and J_4 into separate information sets.

Information sets also show the effects of unobserved moves by Nature. In Figure 4, if the initial move had been made by Nature instead of by Smith, Jones's information sets would be depicted the same way.

*Player i 's **information partition** is a collection of his information sets such that*

- 1 Each path is represented by one node in a single information set in the partition, and*
- 2 The predecessors of all nodes in a single information set are in one information set.*

The information partition represents the different positions that the player knows he will be able to distinguish from each other at a given stage of the game, carving up the set of all possible nodes into the subsets called information sets. One of Smith's information partitions is $(\{J_1\}, \{J_2\}, \{J_3\}, \{J_4\})$. The definition rules out information set $\{S_1\}$ being in that partition, because the path going through S_1 and J_1 would be represented by two nodes. Instead, $\{S_1\}$ is a separate information partition, all by itself. The information partition refers to a stage of the game, not chronological time. The information partition $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$ includes nodes in both 1985 and 1986, but they are all immediate successors of node S_1 .

Jones has the information partition $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$. There are two ways to see that his information is worse than Smith's. First is the fact that one of his information sets, $\{J_3, J_4\}$, contains *more* elements than Smith's, and second, that one of his information partitions, $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$, contains *fewer* elements.

Table 3 shows a number of different information partitions for this game. Partition I is Smith's partition and partition II is Jones's partition. We say that partition II is **coarser**, and partition I is **finer**. A combination of two or more of the information sets in a partition, which reduces the number of information sets and increases the numbers of nodes in one or more of them is a **coarsening**. A splitting of one or more of the information sets in a partition, which increases the number of information sets and reduces the number of nodes in one or more of them, is a **refinement**. Partition II is thus a coarsening of partition I, and partition I is a refinement of partition II. The ultimate refinement is for each information set to be a **singleton**, containing one node, as in the case of partition I. As in bridge, having a singleton can either help or hurt a player. The ultimate coarsening is for a player not to be able to distinguish between any of the nodes, which is partition III in Table 3.¹

A finer information partition is the formal definition for "better information." Not all information partitions are refinements or coarsenings of each other, however, so not all information partitions can be ranked by the quality of their information. In particular, just because one information partition contains more information sets does not mean it is a refinement of another information partition. Consider partitions II and IV in Figure 3. Partition II separates the nodes into three information sets, while partition IV separates them into just two information sets. Partition IV is not a coarsening of partition II, however, because it cannot be reached by combining information sets from partition II, and one cannot say that a player with partition IV has worse information. If the node reached is J_1 , partition II gives more precise information, but if the node reached is J_4 ,

¹Note, however, that partitions III and IV are not really allowed in this game, because Jones could tell the node from the actions available to him, as explained earlier.

partition IV gives more precise information.

<i>Nodes</i>	I	II	III	IV
J_1	$\{J_1\}$	$\{J_1\}$	$\left\{ \begin{matrix} J_1 \\ J_2 \end{matrix} \right\}$	$\left\{ \begin{matrix} J_1 \\ J_2 \end{matrix} \right\}$
J_2	$\{J_2\}$	$\{J_2\}$	$\left\{ \begin{matrix} J_2 \\ J_3 \end{matrix} \right\}$	$\left\{ \begin{matrix} J_2 \\ J_3 \end{matrix} \right\}$
J_3	$\{J_3\}$	$\left\{ \begin{matrix} J_3 \\ J_4 \end{matrix} \right\}$	$\left\{ \begin{matrix} J_3 \\ J_4 \end{matrix} \right\}$	$\left\{ \begin{matrix} J_3 \\ J_4 \end{matrix} \right\}$
J_4	$\{J_4\}$	$\left\{ \begin{matrix} J_4 \end{matrix} \right\}$	$\left\{ \begin{matrix} J_4 \end{matrix} \right\}$	$\{J_4\}$

Table 3: Information Partitions

Information quality is defined independently of its utility to the player: it is possible for a player's information to improve and for his equilibrium payoff to fall as a result. Game theory has many paradoxical models in which a player prefers having worse information, not a result of wishful thinking, escapism, or blissful ignorance, but of cold rationality. Coarse information can have a number of advantages. It may permit a player to engage in trade because other players do not fear his superior information. It may give a player a stronger strategic position because he usually has a strong position and is better off not knowing that in a particular realization of the game his position is weak. Or, as in the more traditional economics of uncertainty, poor information may permit players to insure each other.

Discussion of the first two points will be deferred to later (models of entry deterrence in Section 6.3 and used car markets in Chapter 9). The point about insurance, however, is best discussed here, as an example showing that even when information is symmetric and behavior nonstrategic, it may still may be the case that better information, in the sense of a finer information partition, can actually reduce the utility of all the players in a game — a Pareto worsening.

To see this, suppose that Smith and Jones, both risk averse, work for the same employer, and that both know that one of them will be randomly fired at the end of the year, while the other will be promoted. The one who is fired will end with a wealth of 0 and the one who is not fired will end with a wealth of 100. Smith and Jones will agree to insure each other by pooling their wealth. That is, they will agree that whoever is promoted will pay 50 to whoever is fired. If someone offers to tell them who will be fired before they make this agreement, they should cover their ears and refuse to listen. This refinement of their information would make both worse off, in expectation, because it would ruin their insurance arrangement. Once they knew who would be promoted, they would no longer

agree to pool their wealth. Better information would reduce the expected utility of both players.

Common Knowledge

We have been implicitly assuming that the players know what the game tree looks like. In fact, we have assumed that the players also know that the other players know what the game tree looks like. The term “common knowledge” is used to avoid spelling out the infinite recursion to which this leads.

*Information is **common knowledge** if it is known to all the players, if each player knows that all the players know it, if each player knows that all the players know that all the players know it, and so forth ad infinitum.*

Because of this recursion (the importance of which will be seen in Section 6.3), the assumption of common knowledge is stronger than the assumption that players have the same beliefs about where they are in the game tree. Hirshleifer & Riley (1992, p. 169) use the term **concordant beliefs** to describe a situation where players share the same belief about the probabilities that Nature has chosen different states of the world, but where they do not necessarily know they share the same beliefs. (Brandenburger [1992] uses the term **mutual knowledge** for the same idea.)

For clarity, models are set up so that information partitions are common knowledge. Every player knows how precise the other players’ information is, however ignorant he himself may be as to which node the game has reached. Modelled this way, the information partitions are independent of the equilibrium concept. Making the information partitions common knowledge is important for clear modelling, and restricts the kinds of games that can be modelled less than one might think. This will be illustrated in Section 2.4 when the assumption will be imposed on a situation in which one player does not even know which of three games he is playing.

2.3 Perfect, Certain, Symmetric, and Complete Information

We categorize the information structure of a game in four different ways, so a particular game might have perfect, complete, certain, and symmetric information. The categories are summarized in Table 4.

Information category	Meaning
Perfect	Each information set is a singleton
Certain	Nature does not move after any player moves
Symmetric	No player has information different from other players when he moves, or at the end nodes
Complete	Nature does not move first, or her initial move is observed by every player

Table 4: Information Categories

The first category divides games into those with perfect and those with imperfect information.

*In a game of **perfect information** each information set is a singleton. Otherwise the game is one of **imperfect information**.*

The strongest informational requirements are met by a game of perfect information, because in such a game each player always knows exactly where he is in the game tree. No moves are simultaneous, and all players observe Nature’s moves. Ranked Coordination is a game of imperfect information because of its simultaneous moves, but Follow-the-Leader I is a game of perfect information. Any game of incomplete or asymmetric information is also a game of imperfect information.

*A game of **certainty** has no moves by Nature after any player moves. Otherwise the game is one of **uncertainty**.*

The moves by Nature in a game of uncertainty may or may not be revealed to the players immediately. A game of certainty can be a game of perfect information if it has no simultaneous moves. The notion “game of uncertainty” is new with this book, but I doubt it would surprise anyone. The only quirk in the definition is that it allows an initial move by Nature in a game of certainty, because in a game of incomplete information Nature moves first to select a player’s “type.” Most modellers do not think of this situation as uncertainty.

We have already talked about information in Ranked Coordination, a game of imperfect, complete, and symmetric information with certainty. The Prisoner’s Dilemma falls into the same categories. Follow-the-Leader I, which does not have simultaneous moves, is a game of perfect, complete, and symmetric information with certainty.

We can easily modify Follow-the-Leader I to add uncertainty, creating the game Follow-the-Leader II (Figure 5). Imagine that if both players pick *Large* for their disks, the market yields either zero profits or very high profits, depending on the state of demand, but demand would not affect the payoffs in any other strategy combination. We can quantify

this by saying that if $(Large, Large)$ is picked, the payoffs are $(10,10)$ with probability 0.2, and $(0,0)$ with probability 0.8, as shown in Figure 5.




Figure 5: Follow-the-Leader II

When players face uncertainty, we need to specify how they evaluate their uncertain future payoffs. The obvious way to model their behavior is to say that the players maximize the expected values of their utilities. Players who behave in this way are said to have **von Neumann-Morgenstern utility functions**, a name chosen to underscore von Neumann & Morgenstern's (1944) development of a rigorous justification of such behavior.

Maximizing their expected utilities, the players would behave exactly the same as in Follow-the-Leader I. Often, a game of uncertainty can be transformed into a game of certainty without changing the equilibrium, by eliminating Nature's moves and changing the payoffs to their expected values based on the probabilities of Nature's moves. Here we could eliminate Nature's move and replace the payoffs 10 and 0 with the single payoff 2 $(= 0.2[10] + 0.8[0])$. This cannot be done, however, if the actions available to a player depend on Nature's moves, or if information about Nature's move is asymmetric.

The players in Figure 5 might be either risk averse or risk neutral. Risk aversion is implicitly incorporated in the payoffs because they are in units of utility, not dollars. When players maximize their expected utility, they are not necessarily maximizing their expected dollars. Moreover, the players can differ in how they map money to utility. It could be that $(0,0)$ represents $(\$0, \$5,000)$, $(10,10)$ represents $(\$100,000, \$100,000)$, and $(2,2)$, the expected utility, could here represent a non-risky $(\$3,000, \$7,000)$.

*In a game of **symmetric information**, a player's information set at*
1 any node where he chooses an action, or
2 an end node

*contains at least the same elements as the information sets of every other player. Otherwise the game is one of **asymmetric information**.*

In a game of asymmetric information, the information sets of players differ in ways relevant to their behavior, or differ at the end of the game. Such games have imperfect information, since information sets which differ across players cannot be singletons. The definition of “asymmetric information” which is used in the present book for the first time is intended for capturing a vague meaning commonly used today. The essence of asymmetric information is that some player has useful **private information**: an information partition that is different and not worse than another player’s.

A game of symmetric information can have moves by Nature or simultaneous moves, but no player ever has an informational advantage. The one point at which information may differ is when the player *not* moving has superior information because he knows what his own move *was*; for example, if the two players move simultaneously. Such information does not help the informed player, since by definition it cannot affect his move.

A game has asymmetric information if information sets differ at the end of the game because we conventionally think of such games as ones in which information differs, even though no player takes an action after the end nodes. The principal-agent model of Chapter 7 is an example. The principal moves first, then the agent, and finally Nature. The agent observes the agent’s move, but the principal does not, although he may be able to deduce it. This would be a game of symmetric information except for the fact that information continues to differ at the end nodes.

*In a game of **incomplete information**, Nature moves first and is unobserved by at least one of the players. Otherwise the game is one of **complete information**.*

A game with incomplete information also has imperfect information, because some player’s information set includes more than one node. Two kinds of games have complete but imperfect information: games with simultaneous moves, and games where, late in the game, Nature makes moves not immediately revealed to all players.

Many games of incomplete information are games of asymmetric information, but the two concepts are not equivalent. If there is no initial move by Nature, but Smith takes a move unobserved by Jones, and Smith moves again later in the game, the game has asymmetric but complete information. The principal-agent games of Chapter 7 are again examples: the agent knows how hard he worked, but his principal never learns, not even at the end nodes. A game can also have incomplete but symmetric information: let Nature, unobserved by either player, move first and choose the payoffs for (*Confess*, *Confess*) in the Prisoner’s Dilemma to be either $(-6, -6)$ or $(-100, -100)$.

Harris & Holmstrom (1982) have a more interesting example of incomplete but symmetric information: Nature assigns different abilities to workers, but when workers are young their ability is known neither to employers nor to themselves. As time passes, the abilities become common knowledge, and if workers are risk averse and employers are risk neutral, the model shows that equilibrium wages are constant or rising over time.

Poker Examples of Information Classification

In the game of poker, the players make bets on who will have the best hand of cards at the

end, where a ranking of hands has been pre-established. How would the following rules for behavior before betting be classified? (Answers are in note N2.3)

1. All cards are dealt face up.
2. All cards are dealt face down, and a player cannot look even at his own cards before he bets.
3. All cards are dealt face down, and a player can look at his own cards.
4. All cards are dealt face up, but each player then scoops up his hand and secretly discards one card.
5. All cards are dealt face up, the players bet, and then each player receives one more card face up.
6. All cards are dealt face down, but then each player scoops up his cards without looking at them and holds them against his forehead so all the *other* players can see them (Indian poker).

2.4 The Harsanyi Transformation and Bayesian Games

The Harsanyi Transformation: Follow-the-Leader III

The term “incomplete information” is used in two quite different senses in the literature, usually without explicit definition. The definition in Section 2.3 is what economists commonly *use*, but if asked to *define* the term, they might come up with the following, older, definition.

Old definition

*In a game of **complete information**, all players know the rules of the game. Otherwise the game is one of **incomplete information**.*

The old definition is not meaningful, since the game itself is ill defined if it does not specify exactly what the players’ information sets are. Until 1967, game theorists spoke of games of incomplete information to say that they could not be analyzed. Then John Harsanyi pointed out that any game that had incomplete information under the old definition could be remodelled as a game of complete but imperfect information without changing its essentials, simply by adding an initial move in which Nature chooses between different sets of rules. In the transformed game, all players know the new meta-rules, including the fact that Nature has made an initial move unobserved by them. Harsanyi’s suggestion trivialized the definition of incomplete information, and people began using

the term to refer to the transformed game instead. Under the old definition, a game of incomplete information was transformed into a game of complete information. Under the new definition, the original game is ill defined, and the transformed version is a game of incomplete information.

Follow-the-Leader III serves to illustrate the Harsanyi transformation. Suppose that Jones does not know the game's payoffs precisely. He does have some idea of the payoffs, and we represent his beliefs by a subjective probability distribution. He places a 70 percent probability on the game being game (A) in Figure 6 (which is the same as Follow-the-Leader I), a 10 percent chance on game (B), and a 20 percent on game (C). In reality the game has a particular set of payoffs, and Smith knows what they are. This is a game of incomplete information (Jones does not know the payoffs), asymmetric information (when Smith moves, Smith knows something Jones does not), and certainty (Nature does not move after the players do.)

Figure 6: Follow-the-Leader III: Original

The game cannot be analyzed in the form shown in Figure 6. The natural way to approach such a game is to use the Harsanyi transformation. We can remodel the game to look like Figure 7, in which Nature makes the first move and chooses the payoffs of game (A), (B), or (C), in accordance with Jones's subjective probabilities. Smith observes Nature's move, but Jones does not. Figure 7 depicts the same game as Figure 6, but now we can analyze it. Both Smith and Jones know the rules of the game, and the difference between them is that Smith has observed Nature's move. Whether Nature actually makes the moves with the indicated probabilities or Jones just imagines them is irrelevant, so long as Jones's initial beliefs or fantasies are common knowledge.




Figure 7: Follow-the-Leader III: After the Harsanyi Transformation

Often what Nature chooses at the start of a game is the strategy set, information partition, and payoff function of one of the players. We say that the player can be any of several “types,” a term to which we will return in later chapters. When Nature moves, especially if she affects the strategy sets and payoffs of both players, it is often said that Nature has chosen a particular “state of the world.” In Figure 7 Nature chooses the state of the world to be (A), (B), or (C).

*A player’s **type** is the strategy set, information partition, and payoff function which Nature chooses for him at the start of a game of incomplete information.*

*A **state of the world** is a move by Nature.*

As I have already said, it is good modelling practice to assume that the structure of the game is common knowledge, so that though Nature’s choice of Smith’s type may really just represent Jones’s opinions about Smith’s possible type, Smith knows what Jones’s possible opinions are and Jones knows that they are just opinions. The players may have different beliefs, but that is modelled as the effect of their observing different moves by Nature. All players begin the game with the same beliefs about the probabilities of the moves Nature will make—the same priors, to use a term that will shortly be introduced. This modelling assumption is known as the **Harsanyi doctrine**. If the modeller is following it, his model can never reach a situation where two players possess exactly the same information but disagree as to the probability of some past or future move of Nature. A model cannot, for example, begin by saying that Germany believes its probability of winning a war against France is 0.8 and France believes it is 0.4, so they are both willing to go to war. Rather, he must assume that beliefs begin the same but diverge because of private information. Both players initially think that the probability of a German victory is 0.4 but that if General Schmidt is a genius the probability rises to 0.8, and then Germany discovers that Schmidt is indeed a genius. If it is France that has the initiative to declare war, France’s mistaken

beliefs may lead to a conflict that would be avoidable if Germany could credibly reveal its private information about Schmidt's genius.

An implication of the Harsanyi doctrine is that players are at least slightly open-minded about their opinions. If Germany indicates that it is willing to go to war, France must consider the possibility that Germany has discovered Schmidt's genius and update the probability that Germany will win (keeping in mind that Germany might be bluffing). Our next topic is how a player updates his beliefs upon receiving new information, whether it be by direct observation of Nature or by observing the moves of another player who might be better informed.

Updating Beliefs with Bayes's Rule

When we classify a game's information structure we do not try to decide what a player can deduce from the other players' moves. Player Jones might deduce, upon seeing Smith choose *Large*, that Nature has chosen state (A), but we do not draw Jones's information set in Figure 7 to take this into account. In drawing the game tree we want to illustrate only the exogenous elements of the game, uncontaminated by the equilibrium concept. But to find the equilibrium we do need to think about how beliefs change over the course of the game.

One part of the rules of the game is the collection of **prior beliefs** (or **priors**) held by the different players, beliefs that they update in the course of the game. A player holds prior beliefs concerning the types of the other players, and as he sees them take actions he updates his beliefs under the assumption that they are following equilibrium behavior.

The term **Bayesian equilibrium** is used to refer to a Nash equilibrium in which players update their beliefs according to Bayes's Rule. Since Bayes's Rule is the natural and standard way to handle imperfect information, the adjective, "Bayesian," is really optional. But the two-step procedure of checking a Nash equilibrium has now become a three-step procedure:

- 1 Propose a strategy combination.
- 2 See what beliefs the strategy combination generates when players update their beliefs in response to each others' moves.
- 3 Check that given those beliefs together with the strategies of the other players each player is choosing a best response for himself.

The rules of the game specify each player's initial beliefs, and Bayes's Rule is the rational way to update beliefs. Suppose, for example, that Jones starts with a particular prior belief, $Prob(\text{Nature chose } (A))$. In Follow-the-Leader III, this equals 0.7. He then observes Smith's move — *Large*, perhaps. Seeing *Large* should make Jones update to the **posterior** belief, $Prob(\text{Nature chose } (A)|\text{Smith chose Large})$, where the symbol " $|$ " denotes "conditional upon" or "given that."

Bayes's Rule shows how to revise the prior belief in the light of new information such as Smith's move. It uses two pieces of information, the likelihood of seeing Smith choose *Large* given that Nature chose state of the world (A), $Prob(\text{Large}|(A))$, and the

likelihood of seeing Smith choose *Large* given that Nature did not choose state (A), $Prob(Large|(B) \text{ or } (C))$. From these numbers, Jones can calculate $Prob(\text{Smith chooses Large})$, the **marginal likelihood** of seeing *Large* as the result of one or another of the possible states of the world that Nature might choose.

$$\begin{aligned} Prob(\text{Smith chooses Large}) &= Prob(Large|A)Prob(A) + Prob(Large|B)Prob(B) \\ &\quad + Prob(Large|C)Prob(C). \end{aligned} \tag{1}$$

To find his posterior, $Prob(\text{Nature chose (A)}|\text{Smith chose Large})$, Jones uses the likelihood and his priors. The joint probability of both seeing Smith choose *Large* and Nature having chosen (A) is

$$Prob(Large, A) = Prob(A|Large)Prob(Large) = Prob(Large|A)Prob(A). \tag{2}$$

Since what Jones is trying to calculate is $Prob(A|Large)$, rewrite the last part of (2) as follows:

$$Prob(A|Large) = \frac{Prob(Large|A)Prob(A)}{Prob(Large)}. \tag{3}$$

Jones needs to calculate his new belief — his posterior — using $Prob(Large)$, which he calculates from his original knowledge using (1). Substituting the expression for $Prob(Large)$ from (1) into equation (3) gives the final result, a version of Bayes's Rule.

$$Prob(A|Large) = \frac{Prob(Large|A)Prob(A)}{Prob(Large|A)Prob(A) + Prob(Large|B)Prob(B) + Prob(Large|C)Prob(C)}. \tag{4}$$

More generally, for Nature's move x and the observed data,

$$Prob(x|data) = \frac{Prob(data|x)Prob(x)}{Prob(data)} \tag{5}$$

Equation (6) is a verbal form of Bayes's Rule, which is useful for remembering the terminology, summarized in Table 5.

$$(Posterior \text{ for Nature's Move}) = \frac{(Likelihood \text{ of Player's Move}) \cdot (Prior \text{ for Nature's Move})}{(Marginal likelihood \text{ of Player's Move})}. \tag{6}$$

Bayes's Rule is not purely mechanical. It is the only way to rationally update beliefs. The derivation is worth understanding, because Bayes's Rule is hard to memorize but easy to rederive.

Table 5: Bayesian terminology

Name	Meaning
Likelihood	$Prob(data event)$
Marginal likelihood	$Prob(data)$
Conditional Likelihood	$Prob(data \text{ X} data \text{ Y}, event)$
Prior	$Prob(event)$
Posterior	$Prob(event data)$

Updating Beliefs in Follow-the-Leader III

Let us now return to the numbers in Follow-the-Leader III to use the belief- updating rule that was just derived. Jones has a prior belief that the probability of event “Nature picks state (A)” is 0.7 and he needs to update that belief on seeing the data “Smith picks *Large*”. His prior is $Prob(A) = 0.7$, and we wish to calculate $Prob(A|Large)$.

To use Bayes’s Rule from equation (4), we need the values of $Prob(Large|A)$, $Prob(Large|B)$, and $Prob(Large|C)$. These values depend on what Smith does in equilibrium, so Jones’s beliefs cannot be calculated independently of the equilibrium. This is the reason for the three-step procedure suggested above, for what the modeller must do is propose an equilibrium and then use it to calculate the beliefs. Afterwards, he must check that the equilibrium strategies are indeed the best responses given the beliefs they generate.

A candidate for equilibrium in Follow-the-Leader III is for Smith to choose *Large* if the state is (A) or (B) and *Small* if it is (C), and for Jones to respond to *Large* with *Large* and to *Small* with *Small*. This can be abbreviated as $(L|A, L|B, S|C; L|L, S|S)$. Let us test that this is an equilibrium, starting with the calculation of $Prob(A|Large)$.

If Jones observes *Large*, he can rule out state (C), but he does not know whether the state is (A) or (B). Bayes’s Rule tells him that the posterior probability of state (A) is

$$\begin{aligned} Prob(A|Large) &= \frac{(1)(0.7)}{(1)(0.7)+(1)(0.1)+(0)(0.2)} \\ &= 0.875. \end{aligned} \tag{7}$$

The posterior probability of state (B) must then be $1 - 0.875 = 0.125$, which could also be calculated from Bayes’s Rule, as follows:

$$\begin{aligned} Prob(B|Large) &= \frac{(1)(0.1)}{(1)(0.7)+(1)(0.1)+(0)(0.2)} \\ &= 0.125. \end{aligned} \tag{8}$$

Figure 8 shows a graphic intuition for Bayes’s Rule. The first line shows the total probability, 1, which is the sum of the prior probabilities of states (A), (B), and (C). The second line shows the probabilities, summing to 0.8, which remain after *Large* is observed and state (C) is ruled out. The third line shows that state (A) represents an amount 0.7 of that probability, a fraction of 0.875. The fourth line shows that state (B) represents an amount 0.1 of that probability, a fraction of 0.125.




Figure 8: Bayes's Rule

Jones must use Smith's strategy in the proposed equilibrium to find numbers for $Prob(Large|A)$, $Prob(Large|B)$, and $Prob(Large|C)$. As always in Nash equilibrium, the modeller assumes that the players know which equilibrium strategies are being played out, even though they do not know which particular actions are being chosen.

Given that Jones believes that the state is (A) with probability 0.875 and state (B) with probability 0.125, his best response is *Large*, even though he knows that if the state were actually (B) the better response would be *Small*. Given that he observes *Large*, Jones's expected payoff from *Small* is -0.625 ($= 0.875[-1] + 0.125[2]$), but from *Large* it is 1.875 ($= 0.875[2] + 0.125[1]$). The strategy combination $(L|A, L|B, S|C; L|L, S|S)$ is a Bayesian equilibrium.

A similar calculation can be done for $Prob(A|Small)$. Using Bayes's Rule, equation (4) becomes

$$Prob(A|Small) = \frac{(0)(0.7)}{(0)(0.7) + (0)(0.1) + (1)(0.2)} = 0. \quad (9)$$

Given that he believes the state is (C), Jones's best response to *Small* is *Small*, which agrees with our proposed equilibrium.

Smith's best responses are much simpler. Given that Jones will imitate his action, Smith does best by following his equilibrium strategy of $(L|A, L|B, S|C)$.

The calculations are relatively simple because Smith uses a nonrandom strategy in equilibrium, so, for instance, $Prob(Small|A) = 0$ in equation (9). Consider what happens if Smith uses a random strategy of picking *Large* with probability 0.2 in state (A), 0.6 in state (B), and 0.3 in state (C) (we will analyze such "mixed" strategies in Chapter 3). The equivalent of equation (7) is

$$Prob(A|Large) = \frac{(0.2)(0.7)}{(0.2)(0.7) + (0.6)(0.1) + (0.3)(0.2)} = 0.54 \text{ (rounded)}. \quad (10)$$

If he sees *Large*, Jones's best guess is still that Nature chose state (A), even though in state (A) Smith has the smallest probability of choosing *Large*, but Jones's subjective posterior probability, $Pr(A|Large)$, has fallen to 0.54 from his prior of $Pr(A) = 0.7$.

The last two lines of Figure 8 illustrate this case. The second-to-last line shows the total probability of *Large*, which is formed from the probabilities in all three states and sums to 0.26 ($=0.14 + 0.06 + 0.06$). The last line shows the component of that probability arising from state (A), which is the amount 0.14 and fraction 0.54 (rounded).

Regression to the Mean

Regression to the mean is an old statistical idea that has a Bayesian interpretation. Suppose that each student's performance on a test results partly from his ability and partly from random error because of his mood the day of the test. The teacher does not know the individual student's ability, but does know that the average student will score 70 out of 100. If a student scores 40, what should the teacher's estimate of his ability be?

It should not be 40. A score of 30 points below the average score could be the result of two things: (1) the student's ability is below average, or (2) the student was in a bad mood the day of the test. Only if mood is completely unimportant should the teacher use 40 as his estimate. More likely, both ability and luck matter to some extent, so the teacher's best guess is that the student has an ability below average but was also unlucky. The best estimate lies somewhere between 40 and 70, reflecting the influence of both ability and luck. Of the students who score 40 on the test, more than half can be expected to score above 40 on the next test. Since the scores of these poorly performing students tend to float up towards the mean of 70, this phenomenon is called "regression to the mean." Similarly, students who score 90 on the first test will tend to score less well on the second test.

This is "regression to the mean" ("towards" would be more precise) not "regression beyond the mean." A low score does indicate low ability, on average, so the predicted score on the second test is still below average. Regression to the mean merely recognizes that both luck and ability are at work.

In Bayesian terms, the teacher in this example has a prior mean of 70, and is trying to form a posterior estimate using the prior and one piece of data, the score on the first test. For typical distributions, the posterior mean will lie between the prior mean and the data point, so the posterior mean will be between 40 and 70.

In a business context, regression to the mean can be used to explain business conservatism. It is sometimes claimed that businesses pass up profitable investments because they have an excessive fear of risk. Let us suppose that the business is risk neutral, because the risk associated with the project and the uncertainty over its value are nonsystematic — that is, they are risks that a widely held corporation can distribute in such a way that each shareholder's risk is trivial. Suppose that the firm will not spend \$100,000 on an investment with a present value of \$105,000. This is easily explained if the \$105,000 is an estimate and the \$100,000 is cash. If the average value of a new project of this kind is less than \$100,000 — as is likely to be the case since profitable projects are not easy to find — the best estimate of the value will lie between the measured value of \$105,000 and that

average value, unless the staffer who came up with the \$105,000 figure has already adjusted his estimate. Regressing the \$105,000 to the mean may regress it past \$100,000. Put a bit differently, if the prior mean is, let us say, \$80,000, and the data point is \$105,000, the posterior may well be less than \$100,000.

It is important to keep regression to the mean in mind as an alternative to strategic behavior in explaining odd phenomena. In analyzing test scores, one might try to explain the rise in the scores of poor students by changes in their effort level in an attempt to achieve a target grade in the course with minimum work. In analyzing business decisions, one might try to explain why apparently profitable projects are rejected because of managers' dislike for innovations that would require them to work harder. These explanations might well be valid, but models based on Bayesian updating or regression to the mean might explain the situation just as well and with fewer hard-to-verify assumptions about the utility functions of the individuals involved.

2.5: An Example: The Png Settlement Game

The Png (1983) model of out-of-court settlement is an example of a game with a fairly complicated extensive form.² The plaintiff alleges that the defendant was negligent in providing safety equipment at a chemical plant, a charge which is true with probability q . The plaintiff files suit, but the case is not decided immediately. In the meantime, the defendant and the plaintiff can settle out of court.

What are the moves in this game? It is really made up of two games: the one in which the defendant is liable for damages, and the one in which he is blameless. We therefore start the game tree with a move by Nature, who makes the defendant either liable or blameless. At the next node, the plaintiff takes an action: *Sue* or *Grumble*. If he decides on *Grumble* the game ends with zero payoffs for both players. If he decides to *Sue*, we go to the next node. The defendant then decides whether to *Resist* or *Offer* to settle. If the defendant chooses *Offer*, then the plaintiff can *Settle* or *Refuse*; if the defendant chooses to *Resist*, the plaintiff can *Try* the case or *Drop* it. The following description adds payoffs to this model.

The Png Settlement Game

Players

The plaintiff and the defendant.

The Order of Play

0 Nature chooses the defendant to be Liable for injury to the plaintiff with probability $q = 0.13$ and Blameless otherwise. The defendant observes this but the plaintiff does not.

1 The plaintiff decides to Sue or just to Grumble.

2 The defendant Offers a settlement amount of $S = 0.15$ to the plaintiff, or Resist, setting $S = 0$.

3a If the defendant offered $S = 0.15$, the plaintiff agrees to Settle or he Refuses and goes to trial.

²"Png," by the way, is pronounced the same way it is spelt.

3b If the defendant offered $S = 0$, the plaintiff Drops the case, for legal costs of $P = 0$ and $D = 0$ for himself and the defendant, or chooses to Try it, creating legal costs of $P = 0.1$ and $D = 0.2$

4 If the case goes to trial, the plaintiff wins damages of $W = 1$ if the defendant is Liable and $W = 0$ if the defendant is Blameless. If the case is dropped, $W = 0$.

Payoffs

The plaintiff's payoff is $S + W - P$. The defendant's payoff is $-S - W - D$.

We can also depict this on a game tree, as in Figure 9.




Figure 9: The Game Tree for the Png Settlement Game

This model assumes that the settlement amount, $S = 0.15$, and the amounts spent on legal fees are exogenous. Except in the infinitely long games without end nodes that will appear in Chapter 5, an extensive form should incorporate all costs and benefits into the payoffs at the end nodes, even if costs are incurred along the way. If the court required a \$100 filing fee (which it does not in this game, although a fee will be required in the similar game of Nuisance Suits in Section 4.3), it would be subtracted from the plaintiff's payoffs at every end node except those resulting from his choice of *Grumble*. Such consolidation makes it easier to analyze the game and would not affect the equilibrium strategies unless payments along the way revealed information, in which case what matters is the information, not the fact that payoffs change.

We assume that if the case reaches the court, justice is done. In addition to his legal fees D , the defendant pays damages $W = 1$ only if he is liable. We also assume that the players are risk neutral, so they only care about the expected dollars they will receive, not the variance. Without this assumption we would have to translate the dollar payoffs into utility, but the game tree would be unaffected.

This is a game of certain, asymmetric, imperfect, and incomplete information. We

have assumed that the defendant knows whether he is liable, but we could modify the game by assuming that he has no better idea than the plaintiff of whether the evidence is sufficient to prove him so. The game would become one of symmetric information and we could reasonably simplify the extensive form by eliminating the initial move by Nature and setting the payoffs equal to the expected values. We cannot perform this simplification in the original game, because the fact that the defendant, and only the defendant, knows whether he is liable strongly affects the behavior of both players.

Let us now find the equilibrium. Using dominance we can rule out one of the plaintiff's strategies immediately — *Grumble* — which is dominated by *(Sue, Settle, Drop)*.

Whether a strategy combination is a Nash equilibrium depends on the parameters of the model— S, W, P, D and q , which are the settlement amount, the damages, the court costs for the plaintiff and defendant, and the probability the defendant is liable. Depending on the parameter values, three outcomes are possible: settlement (if the settlement amount is low), trial (if expected damages are high and the plaintiff's court costs are low), and the plaintiff dropping the action (if expected damages minus court costs are negative). Here, I have inserted the parameter values $S = 0.15, D = 0.2, W = 1, q = 0.13$, and $P = 0.1$. Two Nash equilibria exist for this set of parameter values, both weak.

One equilibrium is the strategy combination $\{(Sue, Settle, Try), (Offer, Offer)\}$. The plaintiff sues, the defendant offers to settle (whether liable or not), and the plaintiff agrees to settle. Both players know that if the defendant did not offer to settle, the plaintiff would go to court and try the case. Such **out-of-equilibrium** behavior is specified by the equilibrium, because the threat of trial is what induces the defendant to offer to settle, even though trials never occur in equilibrium. This is a Nash equilibrium because given that the plaintiff chooses *(Sue, Settle, Try)*, the defendant can do no better than *(Offer, Offer)*, settling for a payoff of -0.15 whether he is liable or not; and, given that the defendant chooses *(Offer, Offer)*, the plaintiff can do no better than the payoff of 0.15 from *(Sue, Settle, Try)*.

The other equilibrium is $\{(Sue, Refuse, Try), (Resist, Resist)\}$. The plaintiff sues, the defendant resists and makes no settlement offer, the plaintiff would refuse any offer that was made, and goes to trial. Given that he foresees that the plaintiff will refuse a settlement offer of $S = 0.15$ and will go to trial no matter what, the defendant is willing to resist because it makes no difference what he does.

One final observation on the Png Settlement Game: the game illustrates the Harsanyi doctrine in action, because while the plaintiff and defendant differ in their beliefs as to the probability the plaintiff will win, they do so because the defendant has different information, not because the modeller assigns them different beliefs at the start of the game. This seems awkward compared to the everyday way of approaching this problem in which we simply note that potential litigants have different beliefs, and will go to trial if they both think they can win. It is very hard to make the story consistent, however, because if the differing beliefs are common knowledge, both players know that one of them is wrong, and each has to believe that he is correct. This may be fine as a “reduced form,” in which the attempt is to simply describe what happens without explaining it in any depth. After all, even in The Png Settlement Game, if a trial occurs it is because the players differ in their beliefs, so

one could simply chop off the first part of the game tree. But that is also the problem with violating the Harsanyi doctrine: one cannot analyze how the players react to each other's moves if the modeller simply assigns them inflexible beliefs. In the Png Settlement Game, a settlement is rejected and a trial can occur under certain parameters because the plaintiff weighs the probability that the defendant knows he will win versus the probability that he is bluffing, and sometimes decides to risk a trial. Without the Harsanyi doctrine it is very hard to evaluate such an explanation for trials.

NOTES

N2.1 The strategic and extensive forms of a game

- The term “outcome matrix” is used in Shubik (1982, p. 70), but never formally defined there.
- The term “node” is sometimes defined to include only points at which a player or Nature makes a decision, which excludes the end points.

N2.2 Information Sets

- If you wish to depict a situation in which a player does not know whether the game has reached node A_1 or A_2 and he has different action sets at the two nodes, restructure the game. If you wish to say that he has action set (X,Y,Z) at A_1 and (X,Y) at A_2 , first add action Z to the information set at A_2 . Then specify that at A_2 , action Z simply leads to a new node, A_3 , at which the choice is between X and Y .
- The term “common knowledge” comes from Lewis (1969). Discussions include Brandenburger (1992) and Geanakoplos (1992). For rigorous but nonintuitive definitions of common knowledge, see Aumann (1976) and Milgrom (1981a). Following Milgrom, let (Ω, p) be a probability space, let P and Q be partitions of Ω representing the information of two agents, and let R be the finest common coarsening of P and Q . Let ω be an event (an item of information) and $R(\omega)$ be that element of R which contains ω .

*An event A is **common knowledge** at ω if $R(\omega) \subset A$.*

N2.3 Perfect, Certain, Symmetric, and Complete Information

- Tirole (1988, p. 431) (and more precisely Fudenberg & Tirole [1991a, p. 82]) have defined games of *almost perfect* information. They use this term to refer to repeated simultaneous-move games (of the kind studied here in Chapter 5) in which at each repetition all players know the results of all the moves, including those of Nature, in previous repetitions. It is a pity they use such a general-sounding term to describe so narrow a class of games; it could be usefully extended to cover all games which have perfect information except for simultaneous moves.
- **Poker Classifications.** (1) Perfect, certain. (2) Incomplete, symmetric, certain. (3) Incomplete, asymmetric, certain. (4) Complete, asymmetric, certain. (5) Perfect, uncertain. (6) Incomplete, asymmetric, certain.
- For an explanation of von Neumann-Morgenstern utility, see Varian (1992, chapter 11) or Kreps (1990a, Chapter 3). For other approaches to utility, see Starmer (2000). Expected utility and Bayesian updating are the two foundations of standard game theory, partly because they seem realistic and partly because they are so simple to use. Sometimes they do not explain people’s behavior well, however, and there exist extensive literatures (a) pointing out anomalies, and (b) suggesting alternatives. So far no alternatives have proven to be big enough improvements to justify replacing the standard techniques, given the tradeoff between descriptive realism and added complexity in modelling. The standard response is to admit and ignore the anomalies in theoretical work, and to not press any

theoretical models too hard in situations where the anomalies are likely to make a significant difference. On anomalies, see Kahneman, Slovic & Tversky (1982) (an edited collection); Thaler (1992) (essays from his *Journal of Economic Perspectives* feature); and Dawes (1988) (a good mix of psychology and business).

- Mixed strategies (to be described in Section 3.1) are allowed in a game of perfect information because they are an aspect of the game’s equilibrium, not of its exogenous structure.
- Although the word “perfect,” appears in both “perfect information” (Section 2.3) and “perfect equilibrium” (Section 4.1), the concepts are unrelated.
- An unobserved move by Nature in a game of symmetric information can be represented in any of three ways: (1) as the last move in the game; (2) as the first move in the game; or (3) by replacing the payoffs with the expected payoffs and not using any explicit moves by Nature.

N2.4 The Harsanyi Transformation and Bayesian Games

- Mertens & Zamir (1985) probes the mathematical foundations of the Harsanyi transformation. The transformation requires the extensive form to be common knowledge, which raises subtle questions of recursion.
- A player always has some idea of what the payoffs are, so we can always assign him a subjective probability for each possible payoff. What would happen if he had no idea? Such a question is meaningless, because people always have some notion, and when they say they do not, they generally mean that their prior probabilities are low but positive for a great many possibilities. You, for instance, probably have as little idea as I do of how many cups of coffee I have consumed in my lifetime, but you would admit it to be a nonnegative number less than 3,000,000, and you could make a much more precise guess than that. On the topic of subjective probability, the classic reference is Savage (1954).
- The term “marginal likelihood” is confusing for economists, since it refers to an unconditional likelihood. Statisticians came up with it because they start with $Prob(a, b)$ and then move to $Prob(a)$. That is like going to the margin of a graph—the a -axis—and asking how probable each value of a is.
- If two players have common priors and their information partitions are finite, but they each have private information, iterated communication between them will lead to the adoption of a common posterior. This posterior is not always the posterior they would reach if they directly pooled their information, but it is almost always that posterior (Geanakoplos & Polemarchakis [1982]).
- For formal analysis of regression to the mean and business conservatism, see Rasmusen (1992b). This can also explain why even after discounting revenues further in the future, businesses favor projects that offer quicker returns, if more distant revenue forecasts are less accurate.

Problems

2.1. The Monty Hall Problem

You are a contestant on the TV show, “Let’s Make a Deal.” You face three curtains, labelled A, B and C. Behind two of them are toasters, and behind the third is a Mazda Miata car. You choose A, and the TV showmaster says, pulling curtain B aside to reveal a toaster, “You’re lucky you didn’t choose B, but before I show you what is behind the other two curtains, would you like to change from curtain A to curtain C?” Should you switch? What is the exact probability that curtain C hides the Miata?

2.2. Elmer’s Apple Pie

Mrs Jones has made an apple pie for her son, Elmer, and she is trying to figure out whether the pie tasted divine, or merely good. Her pies turn out divinely a third of the time. Elmer might be ravenous, or merely hungry, and he will eat either 2, 3, or 4 pieces of pie. Mrs Jones knows he is ravenous half the time (but not which half). If the pie is divine, then, if Elmer is hungry, the probabilities of the three consumptions are (0, 0.6, 0.4), but if he is ravenous the probabilities are (0, 0, 1). If the pie is just good, then the probabilities are (0.2, 0.4, 0.4) if he is hungry and (0.1, 0.3, 0.6) if he is ravenous.

Elmer is a sensitive, but useless, boy. He will always say that the pie is divine and his appetite weak, regardless of his true inner feelings.

- (a) What is the probability that he will eat four pieces of pie?
- (b) If Mrs Jones sees Elmer eat four pieces of pie, what is the probability that he is ravenous and the pie is merely good?
- (b) If Mrs Jones sees Elmer eat four pieces of pie, what is the probability that the pie is divine?

2.3. Cancer Tests (adapted from McMillan [1992:211])

Imagine that you are being tested for cancer, using a test that is 98 percent accurate. If you indeed have cancer, the test shows positive (indicating cancer) 98 percent of the time. If you do not have cancer, it shows negative 98 percent of the time. You have heard that 1 in 20 people in the population actually have cancer. Now your doctor tells you that you tested positive, but you shouldn’t worry because his last 19 patients all died. How worried should you be? What is the probability you have cancer?

2.4. The Battleship Problem (adapted from Barry Nalebuff, “Puzzles,” *Journal of Economic Perspectives*, 2:181-82 [Fall 1988])

The Pentagon has the choice of building one battleship or two cruisers. One battleship costs the same as two cruisers, but a cruiser is sufficient to carry out the navy’s mission — if the cruiser survives to get close enough to the target. The battleship has a probability of p of carrying out its mission, whereas a cruiser only has probability $p/2$. Whatever the outcome, the war ends and any surviving ships are scrapped. Which option is superior?

2.5. Joint Ventures

Software Inc. and Hardware Inc. have formed a joint venture. Each can exert either high or

low effort, which is equivalent to costs of 20 and 0. Hardware moves first, but Software cannot observe his effort. Revenues are split equally at the end, and the two firms are risk neutral. If both firms exert low effort, total revenues are 100. If the parts are defective, the total revenue is 100; otherwise, if both exert high effort, revenue is 200, but if only one player does, revenue is 100 with probability 0.9 and 200 with probability 0.1. Before they start, both players believe that the probability of defective parts is 0.7. Hardware discovers the truth about the parts by observation before he chooses effort, but Software does not.

- (a) Draw the extensive form and put dotted lines around the information sets of Software at any nodes at which he moves.
- (b) What is the Nash equilibrium?
- (c) What is Software's belief, in equilibrium, as to the probability that Hardware chooses low effort?
- (d) If Software sees that revenue is 100, what probability does he assign to defective parts if he himself exerted high effort and he believes that Hardware chose low effort?

2.6. California Drought (Adapted from McMillan [1992], page xxx)

California is in a drought and the reservoirs are running low. The probability of rainfall in 1991 is $1/2$, but with probability 1 there will be heavy rainfall in 1992 and any saved water will be useless. The state uses rationing rather than the price system, and it must decide how much water to consume in 1990 and how much to save till 1991. Each Californian has a utility function of $U = \log(w_{90}) + \log(w_{91})$. Show that if the discount rate is zero the state should allocate twice as much water to 1990 as to 1991.

2.7. Smith's Energy Level

The boss is trying to decide whether Smith's energy level is high or low. He can only look in on Smith once during the day. He knows if Smith's energy is low, he will be yawning with a 50 percent probability, but if it is high, he will be yawning with a 10 percent probability. Before he looks in on him, the boss thinks that there is an 80 percent probability that Smith's energy is high, but then he sees him yawning. What probability of high energy should the boss now assess?

2.8. Two Games

Suppose that Column gets to choose which of the following two payoff structures applies to the simultaneous-move game he plays with Row. Row does not know which of these Column has chosen.

Table 5: Payoffs A, The Prisoner's Dilemma

		Column	
		<i>Deny</i>	<i>Confess</i>
Row:	<i>Deny</i>	-1,-1	-10, 0
	<i>Confess</i>	0,-10	-8,-8

Payoffs to: (Row, Column).

Table 6: Payoffs B, A Confession Game

		Column	
		<i>Deny</i>	<i>Confess</i>
Row:	<i>Deny</i>	-4,-4	-12,-200
	<i>Confess</i>	-200,-12	-10,-10

Payoffs to: (Row,Column).

- What is one example of a strategy for each player?
- Find a Nash equilibrium. Is it unique? Explain your reasoning.
- Is there a dominant strategy for Column? Explain why or why not.
- Is there a dominant strategy for Row? Explain why or why not.
- Does Row's choice of strategy depend on whether Column is rational or not? Explain why or why not.

2.9. Bank Runs, Coordination, and Asymmetric Information

A recent article has suggested that during the Chicago bank run of 1932, only banks that actually had negative net worth failed, even though depositors tried to take their money out of all the banks in town. (Charles Calomiris and Joseph Mason (1998), "Contagion and Bank Failures During the Great Depression: The June 1932 Chicago Banking Panic", *American Economic Review*, December 1997, 87: 863-883.) A bank run occurs when many depositors all try to withdraw their deposits simultaneously, which creates a cash crunch for the bank since banks ordinarily do not keep much cash on hand, and have lent out most of it in business and home loans.

- Explain why some people might model bank runs as coordination games.
- Why would the prisoner's dilemma be an inappropriate model for bank runs?
- Suppose that some banks are owned by embezzlers who each year secretly steal some of the funds deposited in the banks, and that these thefts will all be discovered in 1940. The present year is 1931. Some depositors learn in 1932 which banks are owned by embezzlers and which are not, and the other depositors know who these depositors are. Construct a game to capture this situation and predict what would happen.
- How would your model change if the government introduced deposit insurance in 1931, which would pay back all depositors if the bank were unable to do so?

3 Mixed and Continuous Strategies

3.1 Mixed Strategies: The Welfare Game

The games we have looked at have so far been simple in at least one respect: the number of moves in the action set has been finite. In this chapter we allow a continuum of moves, such as when a player chooses a price between 10 and 20 or a purchase probability between 0 and 1. Chapter 3 begins by showing how to find mixed- strategy equilibria for a game with no pure-strategy equilibrium. In Section 3.2 the mixed-strategy equilibria are found by the payoff- equating method, and mixed strategies are applied to a dynamic game, The War of Attrition. Section 3.3 takes a more general look at mixed strategy equilibria and distinguishes between mixed strategies and random actions in auditing games. Section 3.4 begins the analysis of continuous action spaces, and this is continued in Section 3.5 in the Cournot duopoly model, where the discussion focusses on the example of two firms choosing output from the continuum between zero and infinity. These sections introduce other ideas that will be built upon in later chapters— dynamic games in Chapter 4, auditing and agency in Chapters 7 and 8, and Cournot oligopoly in Chapter 14. Section 3.6 looks at the Bertrand model and strategic substitutes. Section 3.7 switches gears a bit and talks about four reasons why a Nash equilibrium might not exist.

We invoked the concept of Nash equilibrium to provide predictions of outcomes without dominant strategies, but some games lack even a Nash equilibrium. It is often useful and realistic to expand the strategy space to include random strategies, in which case a Nash equilibrium almost always exists. These random strategies are called “mixed strategies.”

A **pure strategy** maps each of a player's possible information sets to one action. $s_i : \omega_i \rightarrow a_i$.

A **mixed strategy** maps each of a player's possible information sets to a probability distribution over actions.

$$s_i : \omega_i \rightarrow m(a_i), \text{ where } m \geq 0 \text{ and } \int_{A_i} m(a_i) da_i = 1.$$

A **completely mixed strategy** puts positive probability on every action, so $m > 0$.

The version of a game expanded to allow mixed strategies is called the **mixed extension** of the game.

A pure strategy constitutes a rule that tells the player what action to choose, while a mixed strategy constitutes a rule that tells him what dice to throw in order to choose an action. If a player pursues a mixed strategy, he might choose any of several different actions

in a given situation, an unpredictability which can be helpful to him. Mixed strategies occur frequently in the real world. In American football games, for example, the offensive team has to decide whether to pass or to run. Passing generally gains more yards, but what is most important is to choose an action not expected by the other team. Teams decide to run part of the time and pass part of the time in a way that seems random to observers but rational to game theorists.

The Welfare Game

The Welfare Game models a government that wishes to aid a pauper if he searches for work but not otherwise, and a pauper who searches for work only if he cannot depend on government aid. This is a well-known problem in public policy, called the Samaritan's Dilemma by Tullock (1983, p. 59), who attributes it to James Buchanan. The same problem arises on a private level when parents decide how much to help a lazy child.

Table 1 shows payoffs which represent the situation. Neither player has a dominant strategy, and with a little thought we can see that no Nash equilibrium exists in pure strategies either.

Table 1: The Welfare Game

		Pauper	
		<i>Work</i> (γ_w)	<i>Loaf</i> ($1 - \gamma_w$)
Government	<i>Aid</i> (θ_a)	3, 2	→ -1, 3
	<i>No Aid</i> ($1 - \theta_a$)	↑ -1, 1	← ↓ 0, 0
<i>Payoffs to: (Government, Pauper)</i>			

Each strategy combination must be examined in turn to check for Nash equilibria.

1 The strategy combination (*Aid*, *Work*) is not a Nash equilibrium, because the Pauper would respond with *Loaf* if the Government picked *Aid*.

2 (*Aid*, *Loaf*) is not Nash, because the Government would switch to *No Aid*.

3 (*No Aid*, *Loaf*) is not Nash, because the Pauper would switch to *Work*.

4 (*No Aid*, *Work*) is not Nash, because the Government would switch to *Aid*, which brings us back to (1).

The Welfare Game does have a mixed-strategy Nash equilibrium, which we can calculate. The players' payoffs are the expected values of the payments from Table 1. If the Government plays *Aid* with probability θ_a and the Pauper plays *Work* with probability γ_w , the Government's expected payoff is

$$\begin{aligned}
\pi_{\text{Government}} &= \theta_a[3\gamma_w + (-1)(1 - \gamma_w)] + [1 - \theta_a][-1\gamma_w + 0(1 - \gamma_w)] \\
&= \theta_a[3\gamma_w - 1 + \gamma_w] - \gamma_w + \theta_a\gamma_w \\
&= \theta_a[5\gamma_w - 1] - \gamma_w.
\end{aligned} \tag{1}$$

If only pure strategies are allowed, θ_a equals zero or one, but in the mixed extension of the game, the Government's action of θ_a lies on the continuum from zero to one, the pure strategies being the extreme values. Following the usual procedure for solving a maximization problem, we would differentiate the payoff function with respect to the choice variable to obtain the first-order condition:

$$\begin{aligned} 0 &= \frac{d\pi_{Government}}{d\theta_a} = 5\gamma_w - 1 \\ \Rightarrow \gamma_w &= 0.2. \end{aligned} \tag{2}$$

In the mixed-strategy equilibrium, the Pauper selects *Work* 20 percent of the time. This is a bit strange, though: we obtained the Pauper's strategy by differentiating the Government's payoff! That is because we have not used maximization in the standard way. The problem has a corner solution, because depending on the Pauper's strategy, one of three strategies maximizes the Government's payoff: (i) Do not aid ($\theta_a = 0$) if the Pauper is unlikely enough to try to work; (ii) Definitely aid ($\theta_a = 1$) if the Pauper is likely enough to try to work; (iii) any probability of aid, if the Government is indifferent because the Pauper's probability of work is right on the border line of $\gamma_w = 0.2$.

It is possibility (iii) which allows a mixed strategy equilibrium to exist. To see this, go through the following 4 steps:

- 1 I assert that an optimal mixed strategy exists for the government.
- 2 If the Pauper selects *Work* more than 20 percent of the time, the Government always selects *Aid*. If the Pauper selects *Work* less than 20 percent of the time, the Government never selects *Aid*.
- 3 If a mixed strategy is to be optimal for the Government, the pauper must therefore select *Work* with probability exactly 20 percent.

To obtain the probability of the Government choosing *Aid*, we must turn to the Pauper's payoff function, which is

$$\begin{aligned} \pi_{Pauper} &= \theta_a(2\gamma_w + 3[1 - \gamma_w]) + (1 - \theta_a)(1\gamma_w + 0[1 - \gamma_w]), \\ &= 2\theta_a\gamma_w + 3\theta_a - 3\theta_a\gamma_w + \gamma_w - \theta_a\gamma_w, \\ &= -\gamma_w(2\theta_a - 1) + 3\theta_a. \end{aligned} \tag{3}$$

The first order condition is

$$\begin{aligned} \frac{d\pi_{Pauper}}{d\gamma_w} &= -(2\theta_a - 1) = 0, \\ \Rightarrow \theta_a &= 1/2. \end{aligned} \tag{4}$$

If the Pauper selects *Work* with probability 0.2, the Government is indifferent among selecting *Aid* with probability 100 percent, 0 percent, or anything in between. If the strategies are to form a Nash equilibrium, however, the Government must choose $\theta_a = 0.5$.

In the mixed-strategy Nash equilibrium, the Government selects *Aid* with probability 0.5 and the Pauper selects *Work* with probability 0.2. The equilibrium outcome could be any of the four entries in the outcome matrix. The entries having the highest probability of occurrence are *(No Aid, Loaf)* and *(Aid, Loaf)*, each with probability 0.4 ($= 0.5[1 - 0.2]$).

Interpreting Mixed Strategies

Mixed strategies are not as intuitive as pure strategies, and many modellers prefer to restrict themselves to pure-strategy equilibria in games which have them. One objection to mixed strategies is that people in the real world do not take random actions. That is not a compelling objection, because all that a model with mixed strategies requires to be a good description of the world is that the actions appear random to observers, even if the player himself has always been sure what action he would take. Even explicitly random actions are not uncommon, however—the Internal Revenue Service randomly selects which tax returns to audit, and telephone companies randomly monitor their operators' conversations to discover whether they are being polite.

A more troubling objection is that a player who selects a mixed strategy is always indifferent between two pure strategies. In the Welfare Game, the Pauper is indifferent between his two pure strategies and a whole continuum of mixed strategies, given the Government's mixed strategy. If the Pauper were to decide not to follow the particular mixed strategy $\gamma_w = 0.2$, the equilibrium would collapse because the Government would change its strategy in response. Even a small deviation in the probability selected by the Pauper, a deviation that does not change his payoff if the Government does not respond, destroys the equilibrium completely because the Government does respond. A mixed-strategy Nash equilibrium is weak in the same sense as the *(North, North)* equilibrium in the Battle of the Bismarck Sea: to maintain the equilibrium, a player who is indifferent between strategies must pick a particular strategy from out of the set of strategies.

One way to reinterpret The Welfare Game is to imagine that instead of a single pauper there are many, with identical tastes and payoff functions, all of whom must be treated alike by the Government. In the mixed-strategy equilibrium, each of the paupers chooses *Work* with probability 0.2, just as in the one-pauper game. But the many-pauper game has a pure-strategy equilibrium: 20 percent of the paupers choose the pure strategy *Work* and 80 percent choose the pure strategy *Loaf*. The problem persists of how an individual pauper, indifferent between the pure strategies, chooses one or the other, but it is easy to imagine that individual characteristics outside the model could determine which actions are chosen by which paupers.

The number of players needed so that mixed strategies can be interpreted as pure strategies in this way depends on the equilibrium probability γ_w , since we cannot speak of a fraction of a player. The number of paupers must be a multiple of five in The Welfare Game in order to use this interpretation, since the equilibrium mixing probability is a multiple of $1/5$. For the interpretation to apply no matter how we vary the parameters of a model we would need a *continuum* of players.

Another interpretation of mixed strategies, which works even in the single-pauper game, assumes that the pauper is drawn from a population of paupers, and the Government

does not know his characteristics. The Government only knows that there are two types of paupers, in the proportions (0.2, 0.8): those who pick *Work* if the Government picks $\theta_a = 0.5$, and those who pick *Loaf*. A pauper drawn randomly from the population might be of either type. Harsanyi (1973) gives a careful interpretation of this situation.

3.2 Chicken, The War of Attrition, and Correlated Strategies

Chicken and the Payoff-Equating Method

The next game illustrates why we might decide that a mixed-strategy equilibrium is best even if pure-strategy equilibria also exist. In the game of *Chicken*, the players are two Malibu teenagers, Smith and Jones. Smith drives a hot rod south down the middle of Route 1, and Jones drives north. As collision threatens, each decides whether to *Continue* in the middle or *Swerve* to the side. If a player is the only one to *Swerve*, he loses face, but if neither player picks *Swerve* they are both killed, which has an even lower payoff. If a player is the only one to *Continue*, he is covered with glory, and if both *Swerve* they are both embarrassed. (We will assume that to *Swerve* means by convention to *Swerve* right; if one swerved to the left and the other to the right, the result would be both death and humiliation.) Table 2 assigns numbers to these four outcomes.

Table 2: *Chicken*

		Jones	
		<i>Continue</i> (θ)	<i>Swerve</i> ($1 - \theta$)
Smith:	<i>Continue</i> (θ)	$-3, -3$	$\rightarrow$ $2, 0$
	<i>Swerve</i> ($1 - \theta$)	$\downarrow$ $0, 2$	$\uparrow$ $1, 1$
<i>Payoffs to: (Smith, Jones)</i>			

Chicken has two pure-strategy Nash equilibria, (*Swerve*, *Continue*) and (*Continue*, *Swerve*), but they have the defect of asymmetry. How do the players know which equilibrium is the one that will be played out? Even if they talk before the game started, it is not clear how they could arrive at an asymmetric result. We encountered the same dilemma in choosing an equilibrium for *Battle of the Sexes*. As in that game, the best prediction in *Chicken* is perhaps the mixed-strategy equilibrium, because its symmetry makes it a focal point of sorts, and does not require any differences between the players.

The **payoff-equating** method used here to calculate the mixing probabilities for *Chicken* will be based on the logic followed in Section 3.1, but it does not use the calculus of maximization. In the mixed strategy equilibrium, Smith must be indifferent between *Swerve* and *Continue*. Moreover, *Chicken*, unlike the *Welfare Game*, is a symmetric game, so we can guess that in equilibrium each player will choose the same mixing probability. If that is the case, then, since the payoffs from each of Jones' pure strategies must be equal in a mixed-strategy equilibrium, it is true that

$$\begin{aligned}
 \pi_{\text{Jones}}(\text{Swerve}) &= (\theta_{\text{Smith}}) \cdot (0) + (1 - \theta_{\text{Smith}}) \cdot (1) \\
 &= (\theta_{\text{Smith}}) \cdot (-3) + (1 - \theta_{\text{Smith}}) \cdot (2) = \pi_{\text{Jones}}(\text{Continue}).
 \end{aligned} \tag{5}$$

From equation (5) we can conclude that $1 - \theta_{Smith} = 2 - 5\theta_{Smith}$, so $\theta_{Smith} = 0.25$. In the symmetric equilibrium, both players choose the same probability, so we can replace θ_{Smith} with simply θ . As for the question which represents the greatest interest to their mothers, the two teenagers will survive with probability $1 - (\theta \cdot \theta) = 0.9375$.

The payoff-equating method is easier to use than the calculus method if the modeller is sure which strategies will be mixed, and it can also be used in asymmetric games. In the Welfare Game, it would start with $V_g(Aid) = V_g(No Aid)$ and $V_p(Loaf) = V_p(Work)$, yielding two equations for the two unknowns, θ_a and γ_w , which when solved give the same mixing probabilities as were found earlier for that game. The reason why the payoff-equating and calculus maximization methods reach the same result is that the expected payoff is linear in the possible payoffs, so differentiating the expected payoff equalizes the possible payoffs. The only difference from the symmetric-game case is that two equations are solved for two different mixing probabilities instead of a single equation for the one mixing probability that both players use.

It is interesting to see what happens if the payoff of -3 in the northwest corner of Table 2 is generalized to x . Solving the analog of equation (5) then yields

$$\theta = \frac{1}{1 - x}. \quad (6)$$

If $x = -3$, this yields $\theta = 0.25$, as was just calculated, and if $x = -9$, it yields $\theta = 0.10$. This makes sense; increasing the loss from crashes reduces the equilibrium probability of continuing down the middle of the road. But what if $x = 0.5$? Then the equilibrium probability of continuing appears to be $\theta = 2$, which is impossible; probabilities are bounded by zero and one.

When a mixing probability is calculated to be greater than one or less than zero, the implication is either that the modeller has made an arithmetic mistake or, as in this case, that he is wrong in thinking that the game has a mixed-strategy equilibrium. If $x = 0.5$, one can still try to solve for the mixing probabilities, but, in fact, the only equilibrium is in pure strategies— (*Continue*, *Continue*) (the game has become a Prisoner's Dilemma). The absurdity of probabilities greater than one or less than zero is a valuable aid to the fallible modeller because such results show that he is wrong about the qualitative nature of the equilibrium— it is pure, not mixed. Or, if the modeller is not sure whether the equilibrium is mixed or not, he can use this approach to prove that the equilibrium is not in mixed strategies.

The War of Attrition

The War of Attrition is a game something like Chicken stretched out over time, where both players start with *Continue*, and the game ends when the first one picks *Swerve*. Until the game ends, both earn a negative amount per period, and when one exits, he earns zero and the other player earns a reward for outlasting him.

We will look at a war of attrition in discrete time. We will continue with Smith and Jones, who have both survived to maturity and now play games with more expensive toys: they control two firms in an industry which is a natural monopoly, with demand strong

enough for one firm to operate profitably, but not two. The possible actions are to *Exit* or to *Continue*. In each period that both *Continue*, each earns -1 . If a firm exits, its losses cease and the remaining firm obtains the value of the market's monopoly profit, which we set equal to 3. We will set the discount rate equal to $r > 0$, although that is inessential to the model, even if the possible length of the game is infinite (discount rates will be discussed in detail in Section 4.3).

The War of Attrition has a continuum of Nash equilibria. One simple equilibrium is for Smith to choose (*Continue* regardless of what Jones does) and for Jones to choose (*Exit* immediately), which are best responses to each other. But we will solve for a symmetric equilibrium in which each player chooses the same mixed strategy: a constant probability θ that the player picks *Exit* given that the other player has not yet exited.

We can calculate θ as follows, adopting the perspective of Smith. Denote the expected discounted value of Smith's payoffs by V_{stay} if he stays and V_{exit} if he exits immediately. These two pure strategy payoffs must be equal in a mixed strategy equilibrium (which was the basis for the payoff-equating method). If Smith exits, he obtains $V_{exit} = 0$. If Smith stays in, his payoff depends on what Jones does. If Jones stays in too, which has probability $(1 - \theta)$, Smith gets -1 currently and his expected value for the following period, which is discounted using r , is unchanged. If Jones exits immediately, which has probability θ , then Smith receives a payment of 3. In symbols,

$$V_{stay} = \theta \cdot (3) + (1 - \theta) \left(-1 + \left[\frac{V_{stay}}{1 + r} \right] \right), \quad (7)$$

which, after a little manipulation, becomes

$$V_{stay} = \left(\frac{1 + r}{r + \theta} \right) (4\theta - 1). \quad (8)$$

Once we equate V_{stay} to V_{exit} , which equals zero, equation (8) tells us that $\theta = 0.25$ in equilibrium, and that this is independent of the discount rate r .

Returning from arithmetic to ideas, why does Smith *Exit* immediately with positive probability, given that Jones will exit first if Smith waits long enough? The reason is that Jones might choose to continue for a long time and both players would earn -1 each period until Jones exited. The equilibrium mixing probability is calculated so that both of them are likely to stay in long enough so that their losses soak up the gain from being the survivor. Papers on the War of Attrition include Fudenberg & Tirole (1986b), Ghemawat & Nalebuff (1985), Maynard Smith (1974), Nalebuff & Riley (1985), and Riley (1980). All are examples of “rent-seeking” welfare losses. As Posner (1975) and Tullock (1967) have pointed out, the real costs of acquiring rents can be much bigger than the second-order triangle losses from allocative distortions, and the war of attrition shows that the big loss from a natural monopoly might be not the reduced trade that results from higher prices, but the cost of the battle to gain the monopoly.

In The War of Attrition, the reward goes to the player who does not choose the move which ends the game, and a cost is paid each period that both players refuse to end it. Various other **timing games** also exist. The opposite of a war of attrition is a **pre-emption game**, in which the reward goes to the player who chooses the move which ends

the game, and a cost is paid if both players choose that move, but no cost is incurred in a period when neither player chooses it. The game of **Grab the Dollar** is an example. A dollar is placed on the table between Smith and Jones, who each must decide whether to grab for it or not. If both grab, both are fined one dollar. This could be set up as a one-period game, a T period game, or an infinite-period game, but the game definitely ends when someone grabs the dollar. Table 3 shows the payoffs.

Table 3: Grab the Dollar

		Jones	
		<i>Grab</i>	<i>Don't Grab</i>
Smith:	<i>Grab</i>	$-1, -1$ $\downarrow$	$\rightarrow$ 1,0 $\uparrow$
	<i>Don't Grab</i>	0,1 $\leftarrow$	$0, 0$

Payoffs to: (Smith, Jones)

Like The War of Attrition, Grab the Dollar has asymmetric equilibria in pure strategies, and a symmetric equilibrium in mixed strategies. In the infinite-period version, the equilibrium probability of grabbing is 0.5 per period in the symmetric equilibrium.

Still another class of timing games are duels, in which the actions are discrete occurrences which the players locate at particular points in continuous time. Two players with guns approach each other and must decide when to shoot. In a **noisy duel**, if a player shoots and misses, the other player observes the miss and can kill the first player at his leisure. An equilibrium exists in pure strategies for the noisy duel. In a **silent duel**, a player does not know when the other player has fired, and the equilibrium is in mixed strategies. Karlin (1959) has details on duelling games, and Chapter 4 of Fudenberg & Tirole (1991a) has an excellent discussion of games of timing in general. See also Shubik (1954) on the rather different problem of who to shoot first in a battle with three or more sides.

We will go through one more game of timing to see how to derive a continuous mixed strategies probability distribution, instead of just the single number derived earlier. In presenting this game, a new presentation scheme will be useful. If a game has a continuous strategy set, it is harder or impossible to depict the payoffs using tables or the extensive form using a tree. Tables of the sort we have been using so far would require a continuum of rows and columns, and trees a continuum of branches. A new format for game descriptions of the players, actions, and payoffs will be used for the rest of the book. The new format will be similar to the way the rules of the Dry Cleaners Game were presented in Section 1.1.

Patent Race for a New Market

Players

Three identical firms, Apex, Brydox, and Central.

The Order of Play

Each firm simultaneously chooses research spending $x_i \geq 0$, ($i = a, b, c$).

Payoffs

Firms are risk neutral and the discount rate is zero. Innovation occurs at time $T(x_i)$ where $T' < 0$. The value of the patent is V , and if several players innovate simultaneously they share its value.

$$\pi_i = \begin{cases} V - x_i & \text{if } T(x_i) < T(x_j), (\forall j \neq i) & \text{(Firm } i \text{ gets the patent)} \\ \frac{V}{1+m} - x_i & \text{if } T(x_i) = T(x_k), & \text{(Firm } i \text{ shares the patent with} \\ & m = 1 \text{ or } 2 \text{ other firms)} \\ -x_i & \text{if } T(x_i) > T(x_j) \text{ for some } j & \text{(Firm } i \text{ does not get the patent)} \end{cases}$$

The format first assigns the game a title, after which it lists the players, the order of play (together with who observes what), and the payoff functions. Listing the players is redundant, strictly speaking, since they can be deduced from the order of play, but it is useful for letting the reader know what kind of model to expect. The format includes very little explanation; that is postponed, lest it obscure the description. This exact format is not standard in the literature, but every good article begins its technical section by specifying the same information, if in a less structured way, and the novice is strongly advised to use all the structure he can.

The game Patent Race for a New Market does not have any pure strategy Nash equilibria, because the payoff functions are discontinuous. A slight difference in research by one player can make a big difference in the payoffs, as shown in Figure 1 for fixed values of x_b and x_c . The research levels shown in Figure 1 are not equilibrium values. If Apex chose any research level x_a less than V , Brydax would respond with $x_a + \varepsilon$ and win the patent. If Apex chose $x_a = V$, then Brydax and Central would respond with $x_b = 0$ and $x_c = 0$, which would make Apex want to switch to $x_a = \varepsilon$.

Figure 1: The Payoffs in Patent Race for a New Market

There does exist a symmetric mixed strategy equilibrium. Denote the probability that firm i chooses a research level less than or equal to x as $M_i(x)$. This function describes the firm's mixed strategy. In a mixed-strategy equilibrium a player is indifferent between any of the pure strategies among which he is mixing. Since we know that the pure strategies $x_a = 0$ and $x_a = V$ yield zero payoffs, if Apex mixes over the support $[0, V]$ then the expected payoff for every strategy mixed between must also equal zero. The expected payoff from the pure strategy x_a is the expected value of winning minus the cost of research. Letting x stand for nonrandom and X for random variables, this is

$$\pi_a(x_a) = V \cdot \Pr(x_a \geq X_b, x_a \geq X_c) - x_a = 0 = \pi_a(x_a = 0), \quad (9)$$

which can be rewritten as

$$V \cdot \Pr(X_b \leq x_a) \Pr(X_c \leq x_a) - x_a = 0, \quad (10)$$

or

$$V \cdot M_b(x_a) M_c(x_a) - x_a = 0. \quad (11)$$

We can rearrange equation (11) to obtain

$$M_b(x_a) M_c(x_a) = \frac{x_a}{V}. \quad (12)$$

If all three firms choose the same mixing distribution M , then

$$M(x) = \left(\frac{x}{V}\right)^{1/2} \text{ for } 0 \leq x \leq V. \quad (13)$$

What is noteworthy about a patent race is not the nonexistence of a pure strategy equilibrium but the overexpenditure on research. All three players have expected payoffs of zero, because the patent value V is completely dissipated in the race. As in Brecht's *Threepenny Opera*, "When all race after happiness/Happiness comes in last."¹ To be sure, the innovation is made earlier than it would have been by a monopolist, but hurrying the innovation is not worth the cost, from society's point of view, a result that would persist even if the discount rate were positive. Rogerson (1982) uses a game very similar to Patent Race for a New Market to analyze competition for a government monopoly franchise.

Correlated Strategies

One example of a war of attrition is setting up a market for a new security, which may be a natural monopoly for reasons to be explained in Section 8.5. Certain stock exchanges have avoided the destructive symmetric equilibrium by using lotteries to determine which of them would trade newly listed stock options under a system similar to the football draft.²

¹Act III, scene 7 of the *Threepenny Opera*, translated by John Willett (Berthold Brecht, *Collected Works*, London: Eyre Methuen (1987).

²"Big Board will Begin Trading of Options on 4 Stocks it Lists," *Wall Street Journal*, p. 15 (4 October 1985).

Rather than waste resources fighting, these exchanges use the lottery as a coordinating device, even though it might not be a binding agreement.

Aumann (1974) has pointed out that it is often important whether players can use the same randomizing device for their mixed strategies. If they can, we refer to the resulting strategies as **correlated strategies**. Consider the game of *Chicken*. The only mixed-strategy equilibrium is the symmetric one in which each player chooses *Continue* with probability 0.25 and the expected payoff is 0.75. A correlated equilibrium would be for the two players to flip a coin and for Smith to choose *Continue* if it comes up heads and for Jones to choose *Continue* otherwise. Each player's strategy is a best response to the other's, the probability of each choosing *Continue* is 0.5, and the expected payoff for each is 1.0, which is better than the 0.75 achieved without correlated strategies.

Usually the randomizing device is not modelled explicitly when a model refers to correlated equilibrium. If it is, uncertainty over variables that do not affect preferences, endowments, or production is called **extrinsic uncertainty**. Extrinsic uncertainty is the driving force behind **sunspot models**, so called because the random appearance of sunspots might cause macroeconomic changes via correlated equilibria (Maskin & Tirole [1987]) or bets made between players (Cass & Shell [1983]).

One way to model correlated strategies is to specify a move in which Nature gives each player the ability to commit first to an action such as *Continue* with equal probability. This is often realistic because it amounts to a zero probability of both players entering the industry at exactly the same time without anyone knowing in advance who will be the lucky starter. Neither firm has an a priori advantage, but the outcome is efficient.

The population interpretation of mixed strategies cannot be used for correlated strategies. In ordinary mixed strategies, the mixing probabilities are statistically independent, whereas in correlated strategies they are not. In *Chicken*, the usual mixed strategy can be interpreted as populations of Smiths and Joneses, each population consisting of a certain proportion of pure swervers and pure stayers. The correlated equilibrium has no such interpretation.

Another coordinating device, useful in games that, like *Battle of the Sexes*, have a coordination problem, is **cheap talk** (Crawford & Sobel [1982], Farrell [1987]). Cheap talk refers to costless communication before the game proper begins. In *Ranked Coordination*, cheap talk instantly allows the players to make the desirable outcome a focal point. In *Chicken*, cheap talk is useless, because it is dominant for each player to announce that he will choose *Continue*. But in *The Battle of the Sexes*, coordination and conflict are combined. Without communication, the only symmetric equilibrium is in mixed strategies. If both players know that making inconsistent announcements will lead to the wasteful mixed-strategy outcome, then they are willing to mix announcing whether they will go to the ballet or the prize fight. With many periods of announcements before the final decision, their chances of coming to an agreement are high. Thus communication can help reduce inefficiency even if the two players are in conflict.

3.3 Mixed Strategies with General Parameters and N Players: The Civic Duty Game

Having looked at a number of specific games with mixed-strategy equilibria, let us now apply the method to the general game of Table 4.

Table 4: The General 2-by-2 Game

		Column	
		<i>Left</i> (θ)	<i>Right</i> ($1 - \theta$)
Row:	<i>Up</i> (γ)	a, w	b, x
	<i>Down</i> ($1 - \gamma$)	c, y	d, z
<i>Payoffs to: (Row, Column)</i>			

To find the game's equilibrium, equate the payoffs from the pure strategies. For Row, this yields

$$\pi_{Row}(Up) = \theta a + (1 - \theta)b \quad (14)$$

and

$$\pi_{Row}(Down) = \theta c + (1 - \theta)d. \quad (15)$$

Equating (14) and (15) gives

$$\theta(a + d - b - c) + b - d = 0, \quad (16)$$

which yields

$$\theta^* = \frac{d - b}{(d - b) + (a - c)}. \quad (17)$$

Similarly, equating the payoffs for Column gives

$$\pi_{Column}(Left) = \gamma w + (1 - \gamma)y = \pi_{Column}(Right) = \gamma x + (1 - \gamma)z, \quad (18)$$

which yields

$$\gamma^* = \frac{z - y}{(z - y) + (w - x)}. \quad (19)$$

The equilibrium represented by (17) and (19) illustrates a number of features of mixed strategies.

First, it is possible, but wrong, to follow the payoff-equating method for finding a mixed strategy even if no mixed strategy equilibrium actually exists. Suppose, for example, that *Down* is a strongly dominant strategy for Row, so $c > a$ and $d > b$. Row is unwilling to mix, so the equilibrium is not in mixed strategies. Equation (17) would be misleading, though some idiocy would be required to stay misled for very long since the equation implies that $\theta^* > 1$, or $\theta^* \leq 0$ in this case.

Second, the exact features of the equilibrium in mixed strategies depend heavily on the cardinal values of the payoffs, not just on their ordinal values like the pure strategy equilibria in other 2-by-2 games. Ordinal rankings are all that is needed to know that an equilibrium exists in mixed strategies, but cardinal values are needed to know the exact mixing probabilities. If the payoff to Column from (*Confess*, *Confess*) is changed slightly

in the Prisoner's Dilemma it makes no difference at all to the equilibrium. If the payoff of z to Column from (*Down*, *Right*) is increased slightly in the General 2-by-2 Game, equation (19) says that the mixing probability γ^* will change also.

Third, the payoffs can be changed by affine transformations without changing the game substantively, even though cardinal payoffs do matter (which is to say that monotonic but non-affine transformations do make a difference). Let each payoff π in Table 4 become $\alpha + \beta\pi$. Equation (19) then becomes

$$\begin{aligned}\gamma^* &= \frac{\alpha + \beta z - \alpha - \beta y}{(\alpha + \beta z - \alpha - \beta y) + (\alpha + \beta w - \alpha - \beta x)} \\ &= \frac{z - y}{(z - y) + (w - x)}.\end{aligned}\tag{20}$$

The affine transformation has left the equilibrium strategy unchanged.

Fourth, as was mentioned earlier in connection with the Welfare Game, each player's mixing probability depends only on the payoff parameters of the other player. Row's strategy γ^* in equation (19) depends on the parameters w, x, y and z , which are the payoff parameters for Column, and have no direct relevance for Row.

Categories of Games with Mixed Strategies

Table 5 uses the players and actions of Table 4 to depict three major categories of 2-by-2 games in which mixed-strategy equilibria are important. Some games fall in none of these categories—those with tied payoffs, such as the Swiss Cheese Game in which all eight payoffs equal zero—but the three games in Table 5 encompass a wide variety of economic phenomena.

Table 5: 2-by-2 Games with Mixed Strategy Equilibria ³

$\begin{array}{cc} a, w & \rightarrow & b, x \\ \uparrow & & \downarrow \\ c, y & \leftarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \leftarrow & b, x \\ \downarrow & & \uparrow \\ c, y & \rightarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \leftarrow & b, x \\ \uparrow & & \downarrow \\ c, y & \rightarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \rightarrow & b, x \\ \downarrow & & \uparrow \\ c, y & \leftarrow & d, z \end{array}$
Discoordination Games		Coordination Games	Contribution Games

Discoordination games have a single equilibrium, in mixed strategies. The payoffs are such that either (a) $a > c$, $d > b$, $x > w$, and $y > z$, or (b) $c > a$, $b > d$, $w > x$, and $z > y$. The Welfare Game is a discoordination game, as is Auditing Game I in the next section and Matching Pennies in problem 3.3.

Coordination games have three equilibria: two symmetric equilibria in pure strategies and one symmetric equilibrium in mixed strategies. The payoffs are such that $a > c$, $d > b$, $w > x$, and $z > y$. Ranked Coordination and the Battle of the Sexes are two varieties of coordination games in which the players have the same and opposite rankings of the pure-strategy equilibria.

³xxx Put equilibria in bold fonts in Coord and Cont games

Contribution games have three equilibria: two asymmetric equilibria in pure strategies and one symmetric equilibrium in mixed strategies. The payoffs are such that $c > a$, $b > d$, $x > w$, and $y > z$. Also, it must be true that $c < b$ and $y > x$.

I have invented the name “contribution game” for the occasion, since the type of game described by this term is often used to model a situation in which two players each have a choice of taking some action that contributes to the public good, but would each like the other player to bear the cost. The difference from the Prisoner’s Dilemma is that each player in a contribution game is willing to bear the cost alone if necessary.

Contribution games appear to be quite different from the Battle of the Sexes, but they are essentially the same. Both of them have two pure-strategy equilibria, ranked oppositely by the two players. In mathematical terms, the fact that contribution games have the equilibria in the southwest and northeast corners of the outcome matrix whereas coordination games have them in the northwest and southeast, is unimportant; the location of the equilibria could be changed by just switching the order of Row’s strategies. We do view real situations differently, however, depending on whether players choose the same actions or different actions in equilibrium.

Let us take a look at a particular contribution game to show how to extend two-player games to games with several players. A notorious example in social psychology is the murder of Kitty Genovese, who was killed in New York City in 1964 despite the presence of numerous neighbors. “For more than half an hour 38 respectable, law-abiding citizens in Queens watched a killer stalk and stab a woman in three separate attacks in Kew Gardens.... Twice the sound of their voices and the sudden glow of their bedroom lights interrupted him and frightened him off. Each time he returned, sought her out, and stabbed her again. Not one person telephoned the police during the assault; one witness called after the woman was dead.” (Martin Gansberg, “38 Who Saw Murder Didn’t Call Police,” *The New York Times*, March 27, 1964, p. 1.) Even as hardened an economist as myself finds it somewhat distasteful to call this a “game,” but game theory does explain what happened.

I will use a slightly less appalling situation for our model. In the Civic Duty Game of Table 6, Smith and Jones observe a burglary taking place. Each would like someone to call the police and stop the burglary because having it stopped adds 10 to his payoff, but neither wishes to make the call himself because the effort subtracts 3. If Smith can be assured that Jones will call, Smith himself will ignore the burglary. Table 6 shows the payoffs.

Table 6: The Civic Duty Game

		Jones	
		<i>Ignore</i> (γ)	<i>Telephone</i> ($1 - \gamma$)
Smith:	<i>Ignore</i> (γ)	0, 0	→ 10, 7
	<i>Telephone</i> ($1 - \gamma$)	↓ 7, 10	← ↑ 7, 7
<i>Payoffs to: (Row, Column)</i>			

The Civic Duty Game has two asymmetric pure-strategy equilibria and a symmetric mixed-strategy equilibrium. In solving for the mixed-strategy equilibrium, let us move from two players to N players. In the N -player version of the game, the payoff to Smith is 0 if nobody calls, 7 if he himself calls, and 10 if one or more of the other $N - 1$ players calls. This game also has asymmetric pure-strategy and a symmetric mixed-strategy equilibrium. If all players use the same probability γ of *Ignore*, the probability that the other $N - 1$ players besides Smith all choose *Ignore* is γ^{N-1} , so the probability that one or more of them chooses *Telephone* is $1 - \gamma^{N-1}$. Thus, equating Smith's pure-strategy payoffs using the payoff-equating method of equilibrium calculation yields

$$\pi_{Smith}(Telephone) = 7 = \pi_{Smith}(Ignore) = \gamma^{N-1}(0) + (1 - \gamma^{N-1})(10). \quad (21)$$

Equation (21) tells us that

$$\gamma^{N-1} = 0.3 \quad (22)$$

and

$$\gamma^* = 0.3^{\frac{1}{N-1}}. \quad (23)$$

If $N = 2$, Smith chooses *Ignore* with a probability of 0.30. As N increases, Smith's expected payoff remains equal to 7 whether $N = 2$ or $N = 38$, since his expected payoff equals his payoff from the pure strategy of *Telephone*. The probability of *Ignore*, γ^* , however, increases with N . If $N = 38$, the value of γ^* is about 0.97. When there are more players, each player relies more on somebody else calling.

The probability that nobody calls is γ^{*N} . Equation (22) shows that $\gamma^{*N-1} = 0.3$, so $\gamma^{*N} = 0.3\gamma^*$, which is increasing in N because γ^* is increasing in N . If $N = 2$, the probability that neither player phones the police is $\gamma^{*2} = 0.09$. When there are 38 players, the probability rises to γ^{*38} , about 0.29. The more people that watch a crime, the less likely it is to be reported.

As in the Prisoner's Dilemma, the disappointing result in the Civic Duty Game suggests a role for active policy. The mixed-strategy outcome is clearly bad. The expected payoff per player remains equal to 7 whether there is 1 player or 38, whereas if the equilibrium played out was the equilibrium in which one and only one player called the police, the average payoff would rise from 7 with 1 player to about 9.9 with 38 ($= [1(7) + 37(10)]/38$). A situation like this requires something to make one of the pure-strategy equilibria a focal point. The problem is divided responsibility. One person must be made responsible for calling the police, whether by tradition (e.g., the oldest person on the block always calls the police) or direction (e.g., Smith shouts to Jones: "Call the police!").

3.4 Different Uses of Mixing and Randomizing: Minimax and the Auditing Game

A pure strategy may be strictly dominated by a mixed strategy, even if it is not dominated by any of the other pure strategies in the game.

Table 7: Pure Strategies Dominated by a Mixed Strategy

		Column	
		North	South
Row	North	0,0	4,-4
	South	4,-4	0,0
	Defense	1,-1	1,-1

Payoffs to: (Row, Column)

In the zero-sum game of Table 7, Row's army can attack in the North, attack in the South, or remain on the defense. An unexpected attack gives Row a payoff of 4, an expected attack a payoff of 0, and defense a payoff of 1. Column can respond by preparing to defend in the North or in the South.

Row could guarantee himself a payoff of 1 if he chose Defense. But suppose he plays North with probability .5 and South with probability .5. His expected payoff from this mixed strategy if Column plays North with probability N is $.5(N)(0) + .5(1 - N)(4) + .5(N)(4) + .5(1 - N)(0) = 2$, so whatever response Column picks, Row's expected payoff is higher than his payoff of 1 from Defense. Defense is thus strictly dominated for Row by (.5 North, .5 South).^{4 5}

The next three games will illustrate the difference between mixed strategies and random actions, a subtle but important distinction. In all three games, the Internal Revenue Service must decide whether to audit a certain class of suspect tax returns to discover whether they are accurate or not. The goal of the IRS is to either prevent or catch cheating at minimum cost. The suspects want to cheat only if they will not be caught. Let us assume that the benefit of preventing or catching cheating is 4, the cost of auditing is C , where $C < 4$, the cost to the suspects of obeying the law is 1, and the cost of being caught is the fine $F > 1$.

Even with all of this information, there are several ways to model the situation. Table 8 shows one way: a 2-by-2 simultaneous-move game.

Table 8: Auditing Game I

		Suspects	
		Cheat (θ)	Obey ($1 - \theta$)
IRS:	Audit (γ)	$4 - C, -F$	$4 - C, -1$
	Trust ($1 - \gamma$)	0,0	4, -1

Payoffs to: (IRS, Suspects)

⁴xxx In the unique Nash equilibrium, Column would choose North with probability $N=.5$ and south with probability .5; this is his unique equilibrium action because any other choice of N would cause Row to deviate to whatever direction Column is not guarding as heavily.

⁵xxx Use this as an example of the minimax theorem at work. It really is good for 2-person zero sum games. Fundamental insight of the Minimax theorem: With mixing, you can protect yourself even if you aren't smart. Mention this.

Auditing Game I is a discoordination game, with only a mixed strategy equilibrium. Equations (17) and (19) or the payoff-equating method tell us that

$$\begin{aligned} \text{Probability}(\text{Cheat}) = \theta^* &= \frac{4-(4-C)}{(4-(4-C))+((4-C)-0)} \\ &= \frac{C}{4} \end{aligned} \quad (24)$$

and

$$\begin{aligned} \text{Probability}(\text{Audit}) = \gamma^* &= \frac{-1-0}{(-1-0)+(-F-1)} \\ &= \frac{1}{F}. \end{aligned} \quad (25)$$

Using (24) and (25), the payoffs are

$$\begin{aligned} \pi_{IRS}(\text{Audit}) = \pi_{IRS}(\text{Trust}) &= \theta^*(0) + (1 - \theta^*)(4) \\ &= 4 - C. \end{aligned} \quad (26)$$

and

$$\begin{aligned} \pi_{Suspect}(\text{Obey}) = \pi_{Suspect}(\text{Cheat}) &= \gamma^*(-F) + (1 - \gamma^*)(0) \\ &= -1. \end{aligned} \quad (27)$$

A second way to model the situation is as a sequential game. Let us call this Auditing Game II. The simultaneous game implicitly assumes that both players choose their actions without knowing what the other player has decided. In the sequential game, the IRS chooses government policy first, and the suspects react to it. The equilibrium in Auditing Game II is in pure strategies, a general feature of sequential games of perfect information. In equilibrium, the IRS chooses *Audit*, anticipating that the suspect will then choose *Obey*. The payoffs are $4 - C$ for the IRS and -1 for the suspects, the same for both players as in Auditing Game I, although now there is more auditing and less cheating and fine-paying.

We can go a step further. Suppose the IRS does not have to adopt a policy of auditing or trusting every suspect, but instead can audit a random sample. This is not necessarily a mixed strategy. In Auditing Game I, the equilibrium strategy was to audit all suspects with probability $1/F$ and none of them otherwise. That is different from announcing in advance that the IRS will audit a random sample of $1/F$ of the suspects. For Auditing Game III, suppose the IRS move first, but let its move consist of the choice of the proportion α of tax returns to be audited.

We know that the IRS is willing to deter the suspects from cheating, since it would be willing to choose $\alpha = 1$ and replicate the result in Auditing Game II if it had to. It chooses α so that

$$\pi_{suspect}(\text{Obey}) \geq \pi_{suspect}(\text{Cheat}), \quad (28)$$

i.e.,

$$-1 \geq \alpha(-F) + (1 - \alpha)(0). \quad (29)$$

In equilibrium, therefore, the IRS chooses $\alpha = 1/F$ and the suspects respond with *Obey*. The IRS payoff is $4 - \alpha C$, which is better than the $4 - C$ in the other two games, and the suspect's payoff is -1 , exactly the same as before.

The equilibrium of Auditing Game III is in pure strategies, even though the IRS's action is random. It is different from Auditing Game I because the IRS must go ahead with the costly audit even if the suspect chooses *Obey*. Auditing Game III is different in another way also: its action set is continuous. In Auditing Games I and Auditing Game II the action set is $\{Audit, Trust\}$, although the strategy set becomes $\gamma \in [0, 1]$ once mixed strategies are allowed. In Auditing Game III, the action set is $\alpha \in [0, 1]$, and the strategy set would allow mixing of any of the elements in the action set, although mixed strategies are pointless for the IRS because the game is sequential.

Games with mixed strategies are like games with continuous strategies since a probability is drawn from the continuum between zero and one. Auditing Game III also has a strategy drawn from the interval between zero and one, but it is not a mixed strategy to pick an audit probability of, say, 70 percent. A mixed strategy in that game would be something like a choice of a probability 0.5 of an audit probability of 60 percent and 0.5 of 80 percent. The big difference between the pure strategy choice of an audit probability of .70 and the mixed strategy of (.7-audit, .3-don't audit) is that the pure strategy is an irreversible choice that might be used even when the player is not indifferent between pure strategies, but the mixed strategy is the result of a player who in equilibrium is indifferent as to what he does. The next section will show another difference between mixed strategies and continuous strategies: the payoffs are linear in the mixed-strategy probability, as is evident from payoff equations (14) and (15), but they can be nonlinear in continuous strategies generally.

I have used auditing here mainly to illustrate what mixed strategies are and are not, but auditing is interesting in itself and optimal auditing schemes have many twists to them. An example is the idea of **cross-checking**. Suppose an auditor is supposed to check the value of some variable $x \in [0, 1]$, but his employer is worried that he will not report the true value. This might be because the auditor will be lazy and guess rather than go to the effort of finding x , or because some third party will bribe him, or that certain values of x will trigger punishments or policies the auditor dislikes (this model applies even if x is the auditor's own performance on some other task). The idea of cross-checking is to hire a second auditor and ask him to simultaneously report x . Then, if both auditors report the same x , they are both rewarded, but if they report different values they are both punished. There will still be multiple equilibria, because anything in which they report the same value is an equilibrium. But at least truthful reporting becomes a possible equilibrium. See Kandori & Matsushima (1998) for details.

3.5 Continuous Strategies: The Cournot Game

Most of the games so far in the book have had discrete strategy spaces: *Aid* or *No Aid*, *Confess* or *Deny*. Quite often when strategies are discrete and moves are simultaneous, no pure-strategy equilibrium exists. The only sort of compromise possible in the Welfare Game, for instance, is to choose *Aid* sometimes and *No Aid* sometimes, a mixed strategy. If "*A Little Aid*" were a possible action, maybe there would be a pure-strategy equilibrium. The simultaneous-move game we discuss next, the Cournot Game, has a continuous strategy space even without mixing. It models a duopoly in which two firms choose output levels in competition with each other.

The Cournot Game

Players

Firms Apex and Brydax

The Order of Play

Apex and Brydax simultaneously choose quantities q_a and q_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $Q = q_a + q_b$, and we will assume it to be linear (for generalization see Chapter 14), and, in fact, will use the following specific function:

$$p(Q) = 120 - q_a - q_b. \quad (30)$$

Payoffs are profits, which are given by a firm's price times its quantity minus its costs, i.e.,

$$\begin{aligned} \pi_{Apex} &= (120 - q_a - q_b)q_a - cq_a = (120 - c)q_a - q_a^2 - q_aq_b; \\ \pi_{Brydax} &= (120 - q_a - q_b)q_b = (120 - c)q_b - q_aq_b - q_b^2. \end{aligned} \quad (31)$$




Figure 2: Reaction Curves in The Cournot Game

If this game were cooperative (see Section 1.2), firms would end up producing somewhere on the 45° line in Figure 2, where total output is the monopoly output and maximizes the sum of the payoffs. The monopoly output maximizes $pQ - cQ = (120 - Q - c)Q$ with respect to the total output of Q , resulting in the first-order condition

$$120 - c - 2Q = 0, \quad (32)$$

which implies a total output of $Q = 54$ and a price of 66. Deciding how much of that output of 54 should be produced by each firm—where the firm's output should be located on the

45 ° line—would be a zero-sum cooperative game, an example of bargaining. But since the Cournot Game is noncooperative, the strategy combinations such that $q_a + q_b = 54$ are not necessarily equilibria despite their Pareto optimality (where Pareto optimality is defined from the point of view of the two players, not of consumers, and under the implicit assumption that price discrimination cannot be used).

Cournot noted in Chapter 7 of his 1838 book that this game has a unique equilibrium when demand curves are linear. To find that “Cournot-Nash” equilibrium, we need to refer to the **best response functions** for the two players. If Brydoux produced 0, Apex would produce the monopoly output of 54. If Brydoux produced $q_b = 108$ or greater, the market price would fall to 12 and Apex would choose to produce zero. The best response function is found by maximizing Apex’s payoff, given in equation (31), with respect to his strategy, q_a . This generates the first order condition $120 - c - 2q_a - q_b = 0$, or

$$q_a = 60 - \frac{q_b + c}{2} = 54 - \frac{q_b}{2}. \quad (33)$$

Another name for the best response function, the name usually used in the context of the Cournot Game, is the **reaction function**. Both names are somewhat misleading since the players move simultaneously with no chance to reply or react, but they are useful in imagining what a player would do if the rules of the game did allow him to move second. The reaction functions of the two firms are labelled R_a and R_b in Figure 2. Where they cross, point E, is the **Cournot-Nash equilibrium**, which is simply the Nash equilibrium when the strategies consist of quantities. Algebraically, it is found by solving the two reaction functions for q_a and q_b , which generates the unique equilibrium, $q_a = q_b = 40 - c/3 = 36$. The equilibrium price is then 48 ($= 120 - 36 - 36$).

In the Cournot Game, the Nash equilibrium has the particularly nice property of **stability**: we can imagine how starting from some other strategy combination the players might reach the equilibrium. If the initial strategy combination is point X in Figure 2, for example, Apex’s best response is to decrease q_a and Brydoux’s is to increase q_b , which moves the combination closer to the equilibrium. But this is special to The Cournot Game, and Nash equilibria are not always stable in this way.

Stackelberg Equilibrium

There are many ways to model duopoly. The three most prominent are Cournot, Stackelberg, and Bertrand. Stackelberg equilibrium differs from Cournot in that one firm gets to choose its quantity first. If Apex moved first, what output would it choose? Apex knows how Brydoux will react to its choice, so it picks the point on Brydoux’s reaction curve that maximizes Apex’s profit (see Figure 3).

The Stackelberg Game

Players

Firms Apex and Brydoux

The Order of Play

1. Apex chooses quantity q_a from the set $[0, \infty)$.
2. Brydax chooses quantity q_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $Q = q_a + q_b$:

$$p(Q) = 120 - q_a - q_b. \quad (34)$$

Payoffs are profits, which are given by a firm's price times its quantity minus its costs, i.e.,

$$\begin{aligned} \pi_{Apex} &= (120 - q_a - q_b)q_a - cq_a = (120 - c)q_a - q_a^2 - q_aq_b; \\ \pi_{Brydax} &= (120 - q_a - q_b)q_b = (120 - c)q_b - q_aq_b - q_b^2. \end{aligned} \quad (35)$$




Figure 3: Stackelberg Equilibrium

Apex, moving first, is called the **Stackelberg leader** and Brydax is the **Stackelberg follower**. The distinguishing characteristic of a Stackelberg equilibrium is that one player gets to commit himself first. In Figure 3, Apex moves first intertemporally. If moves were simultaneous but Apex could commit himself to a certain strategy, the same equilibrium would be reached as long as Brydax was not able to commit himself. Algebraically, since Apex forecasts Brydax's output to be $q_b = 60 - \frac{q_a + c}{2}$ from the analog of equation (33), Apex can substitute this into his payoff function in (31), obtaining

$$\pi_a = (120 - c)q_a - q_a^2 - q_a\left(60 - \frac{q_a + c}{2}\right). \quad (36)$$

Maximizing with respect to q_a yields the first order condition

$$(120 - c) - 2q_a - 60 + q_a + \frac{c}{2} = 0, \quad (37)$$

which generates $q_a = 60 - \frac{c}{2} = 54$ (which only equals the monopoly output by coincidence, due to the particular numbers in this example). Once Apex chooses this output, Brydax chooses his output to be $q_b = 27$. (That Brydax chooses exactly half the monopoly output is also accidental.) The market price is $120 - 54 - 27 = 39$ for both firms, so Apex has benefited from his status as Stackelberg leader, but industry profits have fallen compared to the Cournot equilibrium.

3.6 Continuous Strategies: The Bertrand Game, Strategic Complements, and Strategic Substitutes (formerly 14.2, new)

A natural alternative to a duopoly model in which the two firms pick outputs simultaneously is a model in which they pick prices simultaneously. This is known as **Bertrand equilibrium**, because the difficulty of choosing between the two models was stressed in Bertrand (1883), a review discussion of Cournot's book. We will use the same two-player linear-demand world as before, but now the strategy spaces will be the prices, not the quantities. We will also use the same demand function, equation (30), which implies that if p is the lowest price, $q = 120 - p$. In the Cournot model, firms chose quantities but allowed the market price to vary freely. In the Bertrand model, they choose prices and sell as much as they can.

The Bertrand Game

Players

Firms Apex and Brydax

The Order of Play

Apex and Brydax simultaneously choose prices p_a and p_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $p(Q) = 120 - Q$. The payoff function for Apex (and analogously for Brydax) is

$$\pi_a = \begin{cases} p_a(120 - p_a - c) & \text{if } p_a < p_b \\ \frac{p_a(120 - p_a - c)}{2} & \text{if } p_a = p_b \\ 0 & \text{if } p_a > p_b \end{cases}$$

The Bertrand Game has a unique Nash equilibrium: $p_a = p_b = c = 12$. That this is a weak Nash equilibrium is clear: if either firm deviates to a higher price, it loses all its customers and so fails to increase its profits to above zero. (In fact, this is an example of a Nash equilibrium in weakly dominated strategies.) That it is unique is less clear. To see why it is, divide the strategy combinations into four groups:

$(p_a < c \text{ or } p_b < c)$ In either of these cases, the firm with the lowest price will earn negative profits, and could profitably deviate to a price high enough to reduce its demand to zero.

- $p_b > c$ or $p_b > p_a > c$) In either of these cases the firm with the higher price could deviate to a price below its rival and increase its profits from zero to some positive value.
- $(p_a = p_b > c)$ In this case, Apex could deviate to a price ϵ less than Brydrex and its profit would rise, because it would go from selling half the market quantity to selling all of it with an infinitesimal decline in profit per unit sale.
- $p_b = c$ or $p_b > p_a = c$) In this case, the firm with the price of c could move from zero profits to positive profits by increasing its price slightly while keeping it below the other firm's price.

This proof is a good example of one common method of proving uniqueness of equilibrium in game theory: partition the strategy combination space and show area by area that deviations would occur. It is such a good example that I recommend it to anyone teaching from this book as a good test question.⁶

Like the surprising outcome of Prisoner's Dilemma, the Bertrand equilibrium is less surprising once one thinks about the model's limitations. What it shows is that duopoly profits do not arise just because there are two firms. Profits arise from something else, such as multiple periods, incomplete information, or differentiated products.

Both the Bertrand and Cournot models are in common use. The Bertrand model can be awkward mathematically because of the discontinuous jump from a market share of 0 to 100 percent after a slight price cut. The Cournot model is useful as a simple model that avoids this problem and which predicts that the price will fall gradually as more firms enter the market. There are also ways to modify the Bertrand model to obtain intermediate prices and gradual effects of entry. Let us proceed to look at one such modification.⁷

The Differentiated Bertrand Game

The Bertrand model generates zero profits because only slight price discounts are needed to bid away customers. The assumption behind this is that the two firms sell identical goods, so if Apex's price is slightly higher than Brydrex's all the customers go to Brydrex. If customers have brand loyalty or poor price information, the equilibrium is different. Let us now move to a different duopoly market, where the demand curves facing Apex and Brydrex are

$$q_a = 24 - 2p_a + p_b \quad (38)$$

and

$$q_b = 24 - 2p_b + p_a, \quad (39)$$

and they have constant marginal costs of $c = 3$.

The greater the difference in the coefficients on prices in demand curves like these, the less substitutable are the products. As with standard demand curves like (30), we have

⁶Is it still a good question given that I have just provided a warning to the students? Yes. First, it will prove a filter for discovering which students have even skimmed the assigned reading. Second, questions like this are not always easy even if one knows they are on the test. Third, and most important, even if in equilibrium every student answers the question correctly, that very fact shows that the incentive to learn this particular item has worked – and that is our main goal, is it not?

⁷xxx Here or in the notes discuss the sequential undifferentiated Bertrand Game. No Difference.

made implicit assumptions about the extreme points of (38) and (39). These equations only apply if the quantities demanded turn out to be nonnegative, and we might also want to restrict them to prices below some ceiling, since otherwise the demand facing one firm becomes infinite as the other's price rises to infinity. With those restrictions, the payoffs are

$$\pi_a = (p_a - c)(24 - 2p_a + p_b) \quad (40)$$

and

$$\pi_b = (p_b - c)(24 - 2p_b + p_a). \quad (41)$$

The order of play is the same as in The Bertrand Game (or Undifferentiated Bertrand Game, as we will call it when that is necessary to avoid confusion): Apex and Brydax simultaneously choose prices p_a and p_b from the set $[0, \infty)$.

Maximizing Apex's payoff by choice of p_a , we obtain the first-order condition

$$\frac{d\pi_a}{dp_a} = 24 - 4p_a + p_b + 2c = 0, \quad (42)$$

and the reaction function

$$p_a = 6 + \frac{c}{2} + \frac{1}{4}p_b = 7.5 + \frac{1}{4}p_b. \quad (43)$$

Since Brydax has a parallel first-order condition, the equilibrium occurs where $p_a = p_b = 10$. The quantity each firm produces is 14, which is below the 21 each would produce at prices of $p_a = p_b = c = 3$. Figure 4 shows that the reaction functions intersect. Apex's demand curve has the elasticity

$$\left(\frac{\partial q_a}{\partial p_a}\right) \cdot \left(\frac{p_a}{q_a}\right) = -2 \left(\frac{p_a}{q_a}\right), \quad (44)$$

which is finite even when $p_a = p_b$, unlike in the undifferentiated-good Bertrand model.




Figure 4: Bertrand Reaction Functions with Differentiated Products

The differentiated-good Bertrand model is very important, because it is what most economists would consider the most descriptively realistic model of oligopoly.⁸ A basic taxonomy of marketing is that selling something depends on “The Four P’s”: Product, Place, Promotion, and Price. Economists have concentrated heavily on price differences between products, but we realize that differences in product quality and characteristics, where something is sold, and how the sellers get information about it to the buyers also matter. Sellers use their prices as control variables more often than their quantities, but the seller with the lowest price does not get all the customers.

Why, then, did I bother to even describe the Cournot and Undifferentiated Bertrand models? Aren’t they obsolete? No, because descriptive realism is not the *summum bonum* of modelling. Simplicity matters a lot too. The Cournot and Undifferentiated Bertrand models are simpler, especially when we go to three or more firms, so they are better models in many applications.

Strategic Substitutes and Strategic Complements

You may have noticed an interesting difference between the Cournot and Differentiated Bertrand reaction curves in Figures 2 and 4: the reactions curves have opposite slopes. Figure 5 puts the two together for easier comparison. Note that I have changed the Differentiated Bertrand model a little in Figure 5—I now assume that if a firm’s price exceeds some value $\bar{p}$ (the **reservation price**) then it attracts zero customers regardless of how high the other firm’s price may be.

Figure 5: Cournot vs. Differentiated Bertrand Reaction Functions (Strategic Substitutes vs. Strategic Complements)

⁸xxx Maybe I should do Dixit-Stiglitz in Chapter 14, taking off from this.

In both models, the reaction curves cross once, so there is a unique Nash equilibrium. Off the equilibrium path, though, there is an interesting difference. If a Cournot firm increases its output, its rival will do the opposite and reduce its output. If a Bertrand firm increases its price, its rival will do the same thing, and increase its price too.

We can ask of any game: “If the other players do more of their strategy, will I do more of my own strategy, or less?” The dichotomy between games in which the answer is “more” and games in which the answer is “less” turns out to be very useful. Jeremy Bulow, John Geanakoplos & Paul Klemperer (1985) call the strategies in the first (“more”) kind of game, “strategic complements,” because when Player 1 does more of his strategy it increases Player 2’s marginal payoff from his own strategy, just as when I buy more bread it increases my marginal utility from buying more butter. If strategies are strategic complements, then their reaction curves are upward sloping, as in The Differentiated Bertrand Game.

On the other hand, in the “less” kind of game when Player 1 does more of his strategy it *reduces* Player 2’s marginal payoff from his own strategy, just as when I buy more potato chips it reduces my marginal utility from buying more corn chips. We therefore call the strategies “strategic substitutes” and their reaction curves are downward sloping, as in The Cournot Game.

Which way the reaction curves slope has important implications for player behavior. Esther Gal-Or (1985) notes that if reaction curves slope down (as with strategic substitutes and Cournot) there is a first-mover advantage, whereas if they slope upwards (as with strategic complements and Differentiated Bertrand) there is a second-mover advantage. We can see that in Figure 5. The Cournot Game in which Player 1 moves first is simply The Stackelberg Game, which we have already analyzed using Figure 3. The equilibrium moves from E to E^* in Figure 5a, Player 1’s payoff increases, and Player 2’s payoff falls.

We have not analyzed The Differentiated Bertrand Game when Player 1 moves first, but since price is a strategic complement, the effect of sequentiality is very different from in The Cournot Game. We actually cannot tell what Player 1’s optimal strategy is from the diagram alone, but Figure 5 illustrates one possibility. Player 1 chooses the reservation price, $\bar{p}$, predicting that Player 2’s response will be a price somewhat lower, but still greater than the simultaneous Bertrand price at E . The result is that Player 2’s payoff is higher than Player 1’s—a second-mover advantage—but both players are better off at E^* than at E .

Both sequential games could be elaborated further by adding moves beforehand which would determine which player would choose his price or quantity first, but I will leave that to you. The important point for now is that whether a game has strategic complements or strategic substitutes is hugely important to the incentives of the players.

The point is simple enough and important enough that I devote an entire session of my MBA game theory course to strategic complements and strategic substitutes. In the practical game theory that someone with a Master of Business Administration degree ought to know, the most important thing is to learn how to describe a situation in terms of players, actions, information, and payoffs. Often there is not enough data to use a specific functional form, but it is possible to figure out with a mixture of qualitative and

quantitative information whether the relation between actions and payoffs is one of strategic substitutes or strategic complements. The businessman then knows whether, for example, he should try to be a first mover or a second mover, and whether he should keep his action secret or proclaim his action to the entire world.

We will explore this further in Chapter 14. For now, though, try to think about how you would model situations like these (noting that there is no universally right answer for any of them):

1. Two firms are choosing their research and development budgets. Are the budgets strategic complements or strategic substitutes?
2. Smith and Jones are both trying to be elected President of the United States. Each must decide how much he will spend on advertising in California. Are the advertising budgets strategic complements or strategic substitutes?
3. Seven firms are each deciding whether to make their products more special, or more suited to the average consumer. Is the degree of specialness a strategic complement or a strategic substitute?
4. Iran and Iraq are each deciding whether to make their armies larger or smaller. Is army size a strategic complement or a strategic substitute?

3.7 Existence of Equilibrium

One of the strong points of Nash equilibria is that they exist in practically every game one is likely to encounter. There are four common reasons why an equilibrium might not exist or might only exist in mixed strategies.

(1) An unbounded strategy space.

Suppose in a stock market game that Smith can borrow money and buy as many shares x of stock as he likes, so his strategy set, the amount of stock he can buy, is $[0, \infty)$, a set which is unbounded above. (Note, by the way, that we thus assume that he can buy fractional shares, e.g. $x = 13.4$, but cannot sell short, e.g. $x = -100$.)

If Smith knows that the price is lower today than it will be tomorrow, his payoff function will be $\pi(x) = x$ and he will want to buy an infinite number of shares, which is not an equilibrium purchase. If the amount he buys is restricted to be less than or equal to 1,000, however, then the strategy set is bounded (by 1,000), and an equilibrium exists— $x = 0$.

Sometimes, as in the Cournot Game discussed earlier in this chapter, the unboundedness of the strategy sets does not matter because the optimum is an interior solution. In other games, though, it is important, not just to get a determinate solution but because the real world is a rather bounded place. The solar system is finite in size, as is the amount of human time past and future.

(2) An open strategy space. Again consider Smith. Let his strategy be $x \in [0, 1,000)$, which is the same as saying that $0 \leq x < 1,000$, and his payoff function be $\pi(x) = x$. Smith's strategy set is bounded (by 0 and 1,000), but it is open rather than closed, because he can choose any number less than 1,000, but not 1,000 itself. This means no equilibrium will exist, because he wants to buy 999.999... shares. This is just a technical problem; we ought to have specified Smith strategy space to be $[0, 1,000]$, and then an equilibrium would exist, at $x = 1,000$.

(3) A discrete strategy space (or, more generally, a nonconvex strategy space). Suppose we start with an arbitrary pair of strategies s_1 and s_2 for two players. If the players' strategies are strategic complements, then if player 1 increases his strategy in response to s_2 , then player 2 will increase his strategy in response to that. An equilibrium will occur where the players run into diminishing returns or increasing costs, or where they hit the upper bounds of their strategy sets. If, on the other hand, the strategies are strategic substitutes, then if player 1 increases his strategy in response to s_2 , player 2 will in turn want to reduce his strategy. If the strategy spaces are continuous, this can lead to an equilibrium, but if they are discrete, player 2 cannot reduce his strategy just a little bit—he has to jump down a discrete level. That could then induce Player 1 to increase his strategy by a discrete amount. This jumping of responses can be never-ending—there is no equilibrium.

That is what is happening in The Welfare Game of Table 1 in this chapter. No compromise is possible between a little aid and no aid, or between working and not working—until we introduce mixed strategies. That allows for each player to choose a continuous amount of his strategy.

This problem is not limited to games such as 2-by-2 games that have discrete strategy spaces. Rather, it is a problem of “gaps” in the strategy space. Suppose we had a game in which the Government was not limited to amount 0 or 100 of aid, but could choose any amount in the space $\{[0, 10], [90, 100]\}$. That is a continuous, closed, and bounded strategy space, but it is non-convex—there is gap in it. (For a space $\{x\}$ to be convex, it must be true that if x_1 and x_2 are in the space, so is $\theta x_1 + (1 - \theta)x_2$ for any $\theta \in [0, 1]$.) Without mixed strategies, an equilibrium to the game might well not exist.

(4) A discontinuous reaction function arising from nonconcave or discontinuous payoff functions.

Even if the strategy spaces are closed, bounded, and convex, a problem remains. For a Nash equilibrium to exist, we need for the reaction functions of the players to intersect. If the reaction functions are discontinuous, they might not intersect.

Figure 6 shows this for a two-player game in which each player chooses a strategy from the interval between 0 and 1. Player 1's reaction function, $s_1(s_2)$, must pick one or more value of s_1 for each possible value of s_2 , so it must cross from the bottom to the top of the diagram. Player 2's reaction function, $s_2(s_1)$, must pick one or more value of s_2 for each possible value of s_1 , so it must cross from the left to the right of the diagram. If the strategy sets were unbounded or open, the reaction functions might not exist, but that is

not a problem here: they do exist. And in Panel (a) a Nash equilibrium exists, at the point, E , where the two reaction functions intersect.

In Panel (b), however, no Nash equilibrium exists. The problem is that Firm 2's reaction function $s_2(s_1)$ is discontinuous at the point $s_1 = 0.5$. It jumps down from $s_2(0.5) = 0.6$ to $s_2(0.50001) = 0.4$. As a result, the reaction curves never intersect, and no equilibrium exists.

If the two players can use mixed strategies, then an equilibrium will exist even for the game in Panel (b), though I will not prove that here. I would, however, like to say why it is that the reaction function might be discontinuous. A player's reaction functions, remember, is derived by maximizing his payoff as a function of his own strategy given the strategies of the other players.

Thus, a first reason why Player 1's reaction function might be discontinuous in the other players' strategies is that his payoff function is discontinuous in either his own or the other players' strategies. This is what happens in the Hotelling Pricing Game, where if Player 1's price drops enough (or Player 2's price rises high enough), all of Player 2's customers suddenly rush to Player 1.

A second reason why Player 1's reaction function might be discontinuous in the other players' strategies is that his payoff function is not concave. The intuition is that if an objective function is not concave, then there might be a number of maxima that are local but not global, and as the parameters change, which maximum is the global one can suddenly change. This means that the reaction function will suddenly jump from one maximizing choice to another one that is far-distant, rather than smoothly changing as it would in a more nicely behaved problem.




Figure 6: Continuous and Discontinuous Reaction Functions

Problems (1) and (2) are really decision theory problems, not game theory problems,

because unboundedness and openness lead to nonexistence of the solution to a one-player maximization problem. Problems (3) and (4) are special to game theory. They arise because although each player has a best response to the other players, no profile of best choices is such that everybody has chosen his best response to everybody else. They are similar to the decision theory problem of nonexistence of an interior solution, but if only one player were involved, we would at least have a corner solution.

xxx ADD SOME NICE PARAGRAPH CONCLUDING FOR THE CHAPTER

NOTES

N3.1 Mixed Strategies: The Welfare Game

- Waldegrave (1713) is a very early reference to mixed strategies.
- The January 1992 issue of *Rationality and Society* is devoted to attacks on and defenses of the use of game theory in the social sciences, with considerable discussion of mixed strategies and multiple equilibria. Contributors include Harsanyi, Myerson, Rapaport, Tullock, and Wildavsky. The Spring 1989 issue of the *RAND Journal of Economics* also has an exchange on the use of game theory, between Franklin Fisher and Carl Shapiro. I also recommend the Peltzman (1991) attack on the game theory approach to industrial organization in the spirit of the “Chicago School”.
- In this book it will always be assumed that players remember their previous moves. Without this assumption of **perfect recall**, the definition in the text is not that for a mixed strategy, but for a **behavior strategy**. As historically defined, a player pursues a mixed strategy when he randomly chooses between pure strategies at the starting node, but he plays a pure strategy thereafter. Under that definition, the modeller cannot talk of random choices at any but the starting node. Kuhn (1953) showed that the definition of mixed strategy given in the text is equivalent to the original definition if the game has perfect recall. Since all important games have perfect recall and the new definition of mixed strategy is better in keeping with the modern spirit of sequential rationality, I have abandoned the old definition.

The classic example of a game without perfect recall is **bridge**, where the four players of the actual game can be cutely modelled as two players who forget what half their cards look like at any one time in the bidding. A more useful example is a game that has been simplified by restricting players to Markov strategies (see Section 5.4), but usually the modeller sets up such a game with perfect recall and then rules out non-Markov equilibria after showing that the Markov strategies form an equilibrium for the general game.

- For more examples of calculating mixed strategy equilibria, see Sections 5.6, 11.6, 12.5, and 12.6.
- It is *not* true that when two pure-strategy equilibria exist a player would be just as willing to use a strategy mixing the two even when the other player is using a pure strategy. In Battle of the Sexes, for instance, if the man knows the woman is going to the ballet he is not indifferent between the ballet and the prize fight.
- A continuum of players is useful not only because the modeller need not worry about fractions of players, but because he can use more modelling tools from calculus— taking the integral of the quantities demanded by different consumers, for example, rather than the sum. But using a continuum is also mathematically more difficult: see Aumann (1964a, 1964b).

N3.2 Chicken, The War of Attrition, and Correlated Strategies

- The game of Chicken discussed in the text is simpler than the game acted out in the movie *Rebel Without a Cause*, in which the players race towards a cliff and the winner is the player who jumps out of his car last. The pure-strategy space in the movie game is continuous and the payoffs are discontinuous at the cliff’s edge, which makes the game more difficult

to analyze technically. (Looking ahead to Section 4.1, recall the importance of a “tremble” in the movie.)

- Technical difficulties arise in some models with a continuum of actions and mixed strategies. In The Welfare Game, the Government chose a single number, a probability, on the continuum from zero to one. If we allowed the Government to mix over a continuum of aid levels, it would choose a function, a probability density, over the continuum. The original game has a finite number of elements in its strategy set, so its mixed extension still has a strategy space in $\mathbf{R}^n$. But with a continuous strategy set extended by a continuum of mixed strategies for each pure strategy, the mathematics become difficult. A finite number of mixed strategies can be allowed without much problem, but usually that is not satisfactory.

Games in continuous time frequently run into this problem. Sometimes it can be avoided by clever modelling, as in Fudenberg & Tirole’s (1986b) continuous-time war of attrition with asymmetric information. They specify as strategies the length of time firms would proceed to *Continue* given their beliefs about the type of the other player, in which case there is a pure strategy equilibrium.

- **Differential games** are played in continuous time. The action is a function describing the value of a state variable at each instant, so the strategy maps the game’s past history to such a function. Differential games are solved using dynamic optimization. A book-length treatment is Bagchi (1984).
- Fudenberg & Levine (1986) show circumstances under which the equilibria of games with infinite strategy spaces can be found as the limits of equilibria of games with finite strategy spaces.

N3.4 Randomizing Versus Mixing: The Auditing Game

- Auditing Game I is similar to a game called The Police Game. Care must be taken in such games that one does not use a simultaneous- move game when a sequential game is appropriate. Also, discrete strategy spaces can be misleading. In general, economic analysis assumes that costs rise convexly in the amount of an activity and benefits rise concavely. Modelling a situation with a 2-by-2 game uses just two discrete levels of the activity, so the concavity or convexity is lost in the simplification. If the true functions are linear, as in auditing costs which rise linearly with the probability of auditing, this is no great loss. If the true costs rise convexly, as in the case where the hours a policeman must stay on the street each day are increased, then a 2-by-2 model can be misleading. Be especially careful not to press the idea of a mixed-strategy equilibrium too hard if a pure-strategy equilibrium would exist when intermediate strategies are allowed. See Tsebelis (1989) and the criticism of it in J. Hirshleifer & Rasmusen (1992).
- Douglas Diamond (1984) shows the implications of monitoring costs for the structure of financial markets. A fixed cost to monitoring investments motivates the creation of a financial intermediary to avoid repetitive monitoring by many investors.
- Baron & Besanko (1984) study auditing in the context of a government agency which can at some cost collect information on the true production costs of a regulated firm.
- Mookherjee & Png (1989) and Border & Sobel (1987) have examined random auditing in the context of taxation. They find that if a taxpayer is audited he ought to be more than

compensated for his trouble if it turns out he was telling the truth. Under the optimal contract, the truth-telling taxpayer should be delighted to hear that he is being audited. The reason is that a reward for truthfulness widens the differential between the agent's payoff when he tells the truth and when he lies.

Why is such a scheme not used? It is certainly practical, and one would think it would be popular with the voters. One reason might be the possibility of corruption; if being audited leads to a lucrative reward, the government might purposely choose to audit its friends. The current danger seems even worse, though, since the government can audit its enemies and burden them with the trouble of an audit even if they have paid their taxes properly.

- Government action strongly affects what information is available as well as what is contractible. In 1988, for example, the United States passed a law sharply restricting the use of lie detectors for testing or monitoring. Previous to the restriction, about two million workers had been tested each year. (“Law Limiting Use of Lie Detectors is Seen Having Widespread Effect” *Wall Street Journal*, p. 13, 1 July 1988), “American Polygraph Association,” <http://www.polygraph.org/betasite/menu8.html> (Viewed August 31, 2003), Eric Rasmusen, “Bans on Lie Detector Tests,” <http://mypage.iu.edu/~erasmuse/archives1.htm#august10a> (Viewed August 31, 2003).)
- Section 3.4 shows how random actions come up in auditing and in mixed strategies. Another use for randomness is to reduce transactions costs. In 1983, for example, Chrysler was bargaining over how much to pay Volkswagen for a Detroit factory. The two negotiators locked themselves into a hotel room and agreed not to leave till they had an agreement. When they narrowed the price gap from \$100 million to \$5 million, they agreed to flip a coin. (Chrysler won.) How would you model that? “Chrysler Hits Brakes, Starts Saving Money After Shopping Spree,” *Wall Street Journal*, p. 1, 12 January 1988. See also David Friedman’s ingenious idea in Chapter 15 of *Law’s Order* of using a 10% probability of death to replace a 6-year prison term (http://www.daviddfriedman.com/Academic/Course/Pages/L_and_E_LS_98/Why_Is_Law/Why_Is_Law_Chapter_15/Why_Is_Law_Chapter_15.html [viewed August 31, 2003])

N3.5 Continuous Strategies: The Cournot Game

- An interesting class of simple continuous payoff games are the **Colonel Blotto games** (Tukey [1949], McDonald & Tukey [1949]). In these games, two military commanders allocate their forces to m different battlefields, and a battlefield contributes more to the payoff of the commander with the greater forces there. A distinguishing characteristic is that player i ’s payoff increases with the value of player i ’s particular action relative to player j ’s, and i ’s actions are subject to a budget constraint. Except for the budget constraint, this is similar to the tournaments of Section 8.2.
- Considerable work has been done characterizing the Cournot model. A representative article is Gaudet & Salant (1991) on conditions which ensure a unique equilibrium.
- “Stability” is a word used in many different ways in game theory and economics. The natural meaning of a stable equilibrium is that it has dynamics which cause the system to return to that point after being perturbed slightly, and the discussion of the stability of Cournot equilibrium was in that spirit. The uses of the term by von Neumann & Morgenstern (1944) and Kohlberg & Mertens (1986) are entirely different.

- The term “Stackelberg equilibrium” is not clearly defined in the literature. It is sometimes used to denote equilibria in which players take actions in a given order, but since that is just the perfect equilibrium (see Section 4.1) of a well-specified extensive form, I prefer to reserve the term for the Nash equilibrium of the duopoly quantity game in which one player moves first, which is the context of Chapter 3 of Stackelberg (1934).

An alternative definition is that a Stackelberg equilibrium is a strategy combination in which players select strategies in a given order and in which each player’s strategy is a best response to the fixed strategies of the players preceding him and the yet-to-be-chosen strategies of players succeeding him, i.e., a situation in which players precommit to strategies in a given order. Such an equilibrium would not generally be either Nash or perfect.

- Stackelberg (1934) suggested that sometimes the players are confused about which of them is the leader and which the follower, resulting in the disequilibrium outcome called **Stackelberg warfare**.
- With linear costs and demand, total output is greater in Stackelberg equilibrium than in Cournot. The slope of the reaction curve is less than one, so Apex’s output expands more than Brydcox’s contracts. Total output being greater, the price is less than in the Cournot equilibrium.
- A useful application of Stackelberg equilibrium is to an industry with a dominant firm and a **competitive fringe** of smaller firms that sell at capacity if the price exceeds their marginal cost. These smaller firms act as Stackelberg leaders (not followers), since each is small enough to ignore its effect on the behavior of the dominant firm. The oil market could be modelled this way with OPEC as the dominant firm and producers such as Britain on the fringe.

N3.6 Continuous Strategies: The Bertrand Game, Strategic Complements, and Strategic Substitutes (formerly 14.2, new)

- We can also work out the Cournot equilibrium for demand functions (38) and (39), but product differentiation does not affect it much. Start by expressing the price in the demand curve in terms of quantities alone, obtaining

$$p_a = 12 - \frac{1}{2}q_a + \frac{1}{2}p_b \quad (45)$$

and

$$p_b = 12 - \frac{1}{2}q_b + \frac{1}{2}p_a. \quad (46)$$

After substituting from (46) into (45) and solving for p_a , we obtain

$$p_a = 24 - \frac{2}{3}q_a - \frac{1}{3}q_b. \quad (47)$$

The first-order condition for Apex’s maximization problem is

$$\frac{d\pi_a}{dq_a} = 24 - 3 - \frac{4}{3}q_a - \frac{1}{3}q_b = 0, \quad (48)$$

which gives rise to the reaction function

$$q_a = 15.75 - \frac{1}{4}q_b. \quad (49)$$

We can guess that $q_a = q_b$. It follows from (49) that $q_a = 12.6$ and the market price is 11.4. On checking, you would find this to indeed be a Nash equilibrium. But reaction function (49) has much the same shape as if there were no product differentiation, unlike when we moved from undifferentiated Bertrand to differentiated Bertrand competition.

- For more on the technicalities of strategic complements and strategic substitutes, see Bulow, Geanakoplos & Klemperer (1985) and Milgrom & Roberts (1990). If the strategies are strategic complements, Milgrom & Roberts (1990) and Vives (1990) show that pure-strategy equilibria exist. These models often explain peculiar economic phenomenon nicely, as in Peter Diamond (1982) on search and business cycles and Douglas Diamond and P. Dybvig (1983) on bank runs. If the strategies are strategic substitutes, existence of pure-strategy equilibria is more troublesome; see Pradeep Dubey, Ori Haimanko & Andriy Zapechelnuk (2002).

Problems

3.1. Presidential Primaries

Smith and Jones are fighting it out for the Democratic nomination for President of the United States. The more months they keep fighting, the more money they spend, because a candidate must spend one million dollars a month in order to stay in the race. If one of them drops out, the other one wins the nomination, which is worth 11 million dollars. The discount rate is r per month. To simplify the problem, you may assume that this battle could go on forever if neither of them drops out. Let θ denote the probability that an individual player will drop out each month in the mixed-strategy equilibrium.

- (a) In the mixed-strategy equilibrium, what is the probability θ each month that Smith will drop out? What happens if r changes from 0.1 to 0.15?
- (b) What are the two pure-strategy equilibria?
- (c) If the game only lasts one period, and the Republican wins the general election (for Democrat payoffs of zero) if both Democrats refuse to exit, what is the probability γ with which each candidate exits in a symmetric equilibrium?

3.2. Running from the Gestapo

Two risk-neutral men, Schmidt and Braun, are walking south along a street in Nazi Germany when they see a single Gestapo agent coming to check their papers. Only Braun has his papers (unknown to the Gestapo, of course). The Gestapo agent will catch both men if both or neither of them run north, but if just one runs, he must choose which one to stop—the walker or the runner. The penalty for being without papers is 24 months in prison. The penalty for running away from an agent of the State is 24 months in prison, on top of the sentences for any other charges, but the conviction rate for this offense is only 25 percent. The two friends want to maximize their joint welfare, which the Gestapo man wants to minimize. Braun moves first, then Schmidt, then the Gestapo.

- (a) What is the outcome matrix for outcomes that might be observed in equilibrium? (Use θ for the probability that the Gestapo chases the runner and γ for the probability that Braun runs.)
- (b) What is the probability that the Gestapo agent chases the runner, (call it θ^*)?
- (c) What is the probability that Braun runs, (call it γ^*)?
- (d) Since Schmidt and Braun share the same objectives, is this a cooperative game?

3.3. Uniqueness in Matching Pennies

In the game Matching Pennies, Smith and Jones each show a penny with either heads or tails up. If they choose the same side of the penny, Smith gets both pennies; otherwise, Jones gets them.

- (a) Draw the outcome matrix for Matching Pennies.
- (b) Show that there is no Nash equilibrium in pure strategies.

- (c) Find the mixed-strategy equilibrium, denoting Smith's probability of *Heads* by γ and Jones's by θ .
- (d) Prove that there is only one mixed-strategy equilibrium.

3.4. Mixed Strategies in The Battle of the Sexes

Refer back to The Battle of the Sexes and Ranked Coordination in Section 1.4. Denote the probabilities that the man and woman pick *Prize Fight* by γ and θ .

- (a) Find an expression for the man's expected payoff.
- (b) What are the equilibrium values of γ and θ , and the expected payoffs?
- (c) Find the most likely outcome and its probability.
- (d) What is the equilibrium payoff in the mixed-strategy equilibrium for Ranked Coordination?
- (e) Why is the mixed-strategy equilibrium a better focal point in the Battle of the Sexes than in Ranked Coordination?

3.5. A Voting Paradox

Adam, Karl, and Vladimir are the only three voters in Podunk. Only Adam owns property. There is a proposition on the ballot to tax property-holders 120 dollars and distribute the proceeds equally among all citizens who do not own property. Each citizen dislikes having to go to the polling place and vote (despite the short lines), and would pay 20 dollars to avoid voting. They all must decide whether to vote before going to work. The proposition fails if the vote is tied. Assume that in equilibrium Adam votes with probability θ and Karl and Vladimir each vote with the same probability γ , but they decide to vote independently of each other.

- (a) What is the probability that the proposition will pass, as a function of θ and γ ?
- (b) What are the two possible equilibrium probabilities γ_1 and γ_2 with which Karl might vote? Why, intuitively, are there two symmetric equilibria?
- (c) What is the probability θ that Adam will vote in each of the two symmetric equilibria?
- (d) What is the probability that the proposition will pass?

3.6. Industry Output

Industry output is

- (a) lowest with monopoly, highest with a Cournot equilibrium
- (b) lowest with monopoly, highest with a Stackelberg equilibrium.
- (c) lowest with a Cournot, highest with a Stackelberg equilibrium.
- (d) lowest with a Stackelberg, highest with a Cournot equilibrium.

3.7. Duopoly Output

Three firms producing an identical product face the demand curve $P = 240 - \alpha Q$, and produce at

marginal cost β . Each firm picks its quantity simultaneously. If $\alpha = 1$ and $\beta = 40$, the equilibrium output of the industry is in the interval

- (a) $[0, 20]$
- (b) $[20, 100]$
- (c) $[100, 130]$
- (d) $[130, 200]$
- (e) $[200, \infty]$

3.8. Mixed Strategies

If a player uses mixed strategies in equilibrium,

- (a) All players are indifferent among all their strategies.
- (b) That player is indifferent among all his strategies.
- (c) That player is indifferent among the strategies he has a positive probability of choosing in equilibrium.
- (d) That player is indifferent among all his strategies except the ones that are weakly dominated.
- (e) None of the above.

3.9. Nash Equilibrium

Find the unique Nash equilibrium of the game in Table 9.

Table 9: A Game for the 1996 Midterm

		Column		
		<i>Left</i>	<i>Middle</i>	<i>Right</i>
	<i>Up</i>	1, 0	10, -1	0, 1
Row:	<i>Sideways</i>	-1, 0	-2, -2	-12, 4
	<i>Down</i>	0, 2	823, -1	2, 0

Payoffs to: (Row, Column).

3.10. Triopoly

Three companies provide tires to the Australian market. The total cost curve for a firm making Q tires is $TC = 5 + 20Q$, and the demand equation is $P = 100 - N$, where N is the total number of tires on the market.

According to the Cournot model, in which the firms's simultaneously choose quantities, what will the total industry output be?

3.11 (hard). Cournot with Heterogeneous Costs

On his job visit, Professor Schaffer of Michigan told me that in a Cournot model with a linear demand curve $P = \alpha - \beta Q$ and constant marginal cost C_i for firm i , the equilibrium industry output Q depends on $\Sigma_i C_i$, but not on the individual levels of C_i . I may have misremembered. Prove or disprove this assertion. Would your conclusion be altered if we made some other assumption on demand? Discuss.

3.12. Alba and Rome: Asymmetric Information and Mixed Strategies

A Roman, Horatius, unwounded, is fighting the three Curiatius brothers from Alba, each of whom is wounded. If Horatius continues fighting, he wins with probability 0.1, and the payoffs are (10,-10) for (Horatius, Curiatii) if he wins, and (-10,10) if he loses. With probability $\alpha = 0.5$, Horatius is panic-stricken and runs away. If he runs and the Curiatii do not chase him, the payoffs are (-20, 10). If he runs and the Curiatius brothers chase and kill him, the payoffs are (-21, 20). If, however, he is not panic-stricken, but he runs anyway and the Curiatii give chase, he is able to kill the fastest brother first and then dispose of the other two, for payoffs of (10,-10). Horatius is, in fact, not panic-stricken.

- (a) With what probability θ would the Curiatii give chase if Horatius were to run?
- (b) With what probability γ does Horatius run?
- (c) How would θ and γ be affected if the Curiatii falsely believed that the probability of Horatius being panic-stricken was 1? What if they believed it was 0.9?

3.13. Finding Nash Equilibria

Find all of the Nash equilibria for the game of Table 10.

Table 10: A Takeover Game

		Target		
		<i>Hard</i>	<i>Medium</i>	<i>Soft</i>
Raider:	<i>Hard</i>	-3, -3	-1, 0	4, 0
	<i>Medium</i>	0, 0	2, 2	3, 1
	<i>Soft</i>	0, 0	2, 4	3, 3

Payoffs to: (Raider, Target).

3.14. Risky Skating

Elena and Mary are the two leading figure skaters in the world. Each must choose during her training what her routine is going to look like. They cannot change their minds later and try to alter any details of their routines. Elena goes first in the Olympics, and Mary goes next. Each has five minutes for her performance. The judges will rate the routines on three dimensions, beauty, how high they jump, and whether they stumble after they jump. A skater who stumbles is sure to lose, and if both Elena and Mary stumble, one of the ten lesser skaters will win, though those ten skaters have no chance otherwise.

Elena and Mary are exactly equal in the beauty of their routines, and both of them know this, but they are not equal in their jumping ability. Whoever jumps higher without stumbling will definitely win. Elena's probability of stumbling is $P(h)$, where h is the height of the jump, and P is increasing smoothly and continuously in h . (In calculus terms, let P' and P'' both exist, and P' be positive) Mary's probability is $P(h) - .1$ — that is, it is 10 percent less for equal heights.

Let us define as $h=0$ the maximum height that the lesser skaters can achieve, and assume that $P(0) = 0$.

- (a) Show that it cannot be an equilibrium for both Mary and Elena to choose the same value for h (Call them M and E).
- (b) Show for any pair of values (M, E) that it cannot be an equilibrium for Mary and Elena to choose those values.
- (c) Describe the optimal strategies to the best of your ability. (Do not get hung up on trying to answer this question; my expectations are not high here.)
- (d) What is a business analogy? Find some situation in business or economics that could use this same model.

3.15. The Kosovo War

Senator Robert Smith of New Hampshire said of the US policy in Serbia of bombing but promising not to use ground forces, “It’s like saying we’ll pass on you but we won’t run the football.” (*Human Events*, p. 1, April 16, 1999.) Explain what he meant, and why this is a strong criticism of U.S. policy, using the concept of a mixed strategy equilibrium. (Foreign students: in American football, a team can choose to throw the football (to pass it) or to hold it and run with it to move towards the goal.) Construct a numerical example to compare the U.S. expected payoff in (a) a mixed strategy equilibrium in which it ends up not using ground forces, and (b) a pure strategy equilibrium in which the U.S. has committed not to use ground forces.

3.16. IMF Aid

Consider the game of Table 11.

		Debtor	
		Reform	Waste
IMF	Aid	3,2	-1,3
	No Aid	-1,1	0,0

Payoffs to: (IMF, Debtor).

- (a) What is the exact form of every Nash equilibrium?
- (b) For what story would this matrix be a good model?

3.17. Coupon Competition

Two marketing executives are arguing. Smith says that reducing our use of coupons will make us a less aggressive competitor, and that will hurt our sales. Jones says that reducing our use of coupons will make us a less aggressive competitor, but that will end up helping our sales.

Discuss, using the effect of reduced coupon use on your firm’s reaction curve, under what circumstance each executive could be correct.

3.18. Rent Seeking

I mentioned that Rogerson (1982) uses a game very similar to “Patent Race for a New Market” on p. 374 to analyze competition for a government monopoly franchise.⁹ See if you can do this too. What can you predict about the welfare results of such competition?

3.19. Does not exist. xxx

3.20. Entry for Buyout

Find the equilibrium in “Entry for Buyout” if all the parameters of the numerical example on page 388¹⁰ are the same except that the marginal cost of output is $c = 20$ instead of $c = 10$.

⁹xxx Fix page reference.

¹⁰xxx Fix this.

4 Dynamic Games with Symmetric Information

4.1 Subgame Perfectness

In this chapter we will make heavy use of the extensive form to study games with moves that occur in sequence. We start in Section 4.1 with a refinement of the Nash equilibrium concept called perfectness that incorporates sensible implications of the order of moves. Perfectness is illustrated in Section 4.2 with a game of entry deterrence. Section 4.3 expands on the idea of perfectness using the example of nuisance suits, meritless lawsuits brought in the hopes of obtaining a settlement out of court. Nuisance suits show the importance of a threat being made credible and how sinking costs early or having certain nonmonetary payoffs can benefit a player. This example will also be used to discuss the open-set problem of weak equilibria in games with continuous strategy spaces, in which a player offering a contract chooses its terms to make the other player indifferent about accepting or rejecting. The last perfectness topic will be renegotiation: the idea that when there are multiple perfect equilibria, the players will coordinate on equilibria that are Pareto optimal in subgames but not in the game as a whole.

The Perfect Equilibrium of Follow the Leader I


Subgame perfectness is an equilibrium concept based on the ordering of moves and the distinction between an equilibrium path and an equilibrium. The **equilibrium path** is the path through the game tree that is followed in equilibrium, but the equilibrium itself is a strategy combination, which includes the players' responses to other players' deviations from the equilibrium path. These off-equilibrium responses are crucial to decisions on the equilibrium path. A threat, for example, is a promise to carry out a certain action if another player deviates from his equilibrium actions, and it has an influence even if it is never used.

Perfectness is best introduced with an example. In Section 2.1, a flaw of Nash equilibrium was revealed in the game Follow the Leader I, which has three pure strategy Nash equilibria of which only one is reasonable. The players are Smith and Jones, who choose disk sizes. Both their payoffs are greater if they choose the same size and greatest if they coordinate on *Large*. Smith moves first, so his strategy set is $\{Small, Large\}$. Jones' strategy is more complicated, because it must specify an action for each information set, and Jones's information set depends on what Smith chose. A typical element of Jones's strategy set is $(Large, Small)$, which specifies that he chooses *Large* if Smith chose *Large*, and *Small* if Smith chose *Small*. From the strategic form we found the following three Nash equilibria.

Equilibrium	Strategies	Outcome
E_1	$\{Large, (Large, Large)\}$	Both pick <i>Large</i> .
E_2	$\{Large, (Large, Small)\}$	Both pick <i>Large</i> .
E_3	$\{Small, (Small, Small)\}$	Both pick <i>Small</i> .

Only Equilibrium E_2 is reasonable, because the order of the moves should matter to the decisions players make. The problem with the strategic form, and thus with simple Nash equilibrium, is that it ignores who moves first. Smith moves first, and it seems reasonable that Jones should be allowed—in fact should be required—to rethink his strategy after Smith moves.

Figure 1: *Follow the Leader I*



Consider Jones’s strategy of $(Small, Small)$ in equilibrium E_3 . If Smith deviated from equilibrium by choosing *Large*, it would be unreasonable for Jones to stick to the response *Small*. Instead, he should also choose *Large*. But if Smith expected a response of *Large*, he would have chosen *Large* in the first place, and E_3 would not be an equilibrium. A similar argument shows that it would be irrational for Jones to choose $(Large, Large)$, and we are left with E_2 as the unique equilibrium.

We say that equilibria E_1 and E_3 are Nash equilibria but not “perfect” Nash equilibria. A strategy combination is a perfect equilibrium if it remains an equilibrium on all possible paths, including not only the equilibrium path but all the other paths, which branch off into different “subgames.”

A **subgame** is a game consisting of a node which is a singleton in every player’s information partition, that node’s successors, and the payoffs at the associated end nodes.¹

¹Technically, this is a *proper* subgame because of the information qualifier, but no economist is so ill-bred as to use any other kind of subgame.

A strategy combination is a **subgame perfect Nash equilibrium** if (a) it is a Nash equilibrium for the entire game; and (b) its relevant action rules are a Nash equilibrium for every subgame.

The extensive form of Follow the Leader I in Figure 1 (a reprise of Figure 1 from Chapter 2) has three subgames: (1) the entire game, (2) the subgame starting at node J_1 , and (3) the subgame starting at node J_2 . Strategy combination E_1 is not a subgame perfect equilibrium because it is only Nash in subgames (1) and (3), not in subgame (2). Strategy combination E_3 is not a subgame perfect equilibrium because it is only Nash in subgames (1) and (2), not in subgame (3). Strategy combination E_2 is perfect because it is Nash in all three subgames.

The term **sequential rationality** is often used to denote the idea that a player should maximize his payoffs at each point in the game, re-optimizing his decisions at each point and taking into account the fact that he will re-optimize in the future. This is a blend of the economic ideas of ignoring sunk costs and rational expectations. Sequential rationality is so standard a criterion for equilibrium now that often I will speak of “equilibrium” without the qualifier when I wish to refer to an equilibrium that satisfies sequential rationality in the sense of being a “subgame perfect equilibrium” or, in a game of asymmetric information, a “perfect Bayesian equilibrium.”

One reason why perfectness (the word “subgame” is usually left off) is a good equilibrium concept is because it represents the idea of sequential rationality. A second reason is that a weak Nash equilibrium is not robust to small changes in the game. So long as he is certain that Smith will not choose *Large*, Jones is indifferent between the never-to-be-used responses (*Small* if *Large*) and (*Large* if *Large*). Equilibria E_1 , E_2 , and E_3 are all weak Nash equilibria because of this. But if there is even a small probability that Smith will choose *Large*—perhaps by mistake—then Jones would prefer the response (*Large* if *Large*), and equilibria E_1 and E_3 are no longer valid. Perfectness is a way to eliminate some of these less robust weak equilibria. The small probability of a mistake is called a **tremble**, and Section 6.1 returns to this **trembling hand** approach as one way to extend the notion of perfectness to games of asymmetric information.

For the moment, however, the reader should note that the tremble approach is distinct from sequential rationality. Consider the Tremble Game in Figure 2. This game has three Nash equilibria, all weak: (*Exit*, *Down*), (*Exit*, *Up*), and (*Remain*, *Up*). Only (*Exit*, *Up*) and (*Remain*, *Up*) are subgame perfect, because although *Down* is weakly Jones’s best response to Smith’s *Exit*, it is inferior if Smith chooses *Remain*. In the subgame starting with Jones’s move, the only subgame perfect equilibrium is for Jones to choose *Up*. The possibility of trembles, however, rules out (*Remain*, *Up*) as an equilibrium. If Jones has even an infinitesimal chance of trembling and choosing *Down*, Smith will choose *Exit* instead of *Remain*. Also, Jones will choose *Up*, not *Down*, because if Smith trembles and chooses *Remain*, Jones prefers *Up* to *Down*. This leaves only (*Exit*, *Up*) as an equilibrium, despite the fact that it is weakly Pareto dominated by (*Remain*, *Up*).




Figure 2: The Tremble Game: Trembling Hand Versus Subgame Perfectness

4.2 An Example of Perfectness: Entry Deterrence I

We turn now to a game in which perfectness plays a role just as important as in Follow the Leader I but in which the players are in conflict. An old question in industrial organization is whether an incumbent monopolist can maintain his position by threatening to wage a price war against any new firm that enters the market. This idea was heavily attacked by Chicago School economists such as McGee (1958) on the grounds that a price war would hurt the incumbent more than collusion with the entrant. Game theory can present this reasoning very cleanly. Let us consider a single episode of possible entry and price warfare, which nobody expects to be repeated. We will assume that even if the incumbent chooses to collude with the entrant, maintaining a duopoly is difficult enough that market revenue drops considerably from the monopoly level.

Entry Deterrence I

Players

Two firms, the entrant and the incumbent.

The Order of Play

- 1 The entrant decides whether to *Enter* or *Stay Out*.
- 2 If the entrant enters, the incumbent can *Collude* with him, or *Fight* by cutting the price drastically.

Payoffs

Market profits are 300 at the monopoly price and 0 at the fighting price. Entry costs are 10. Duopoly competition reduces market revenue to 100, which is split evenly.

Table 1 Entry Deterrence I

		Incumbent	
		<i>Collude</i>	<i>Fight</i>
Entrant:	<i>Enter</i>	40,50 ←	−10, 0
	<i>Stay Out</i>	0, 300 ↔	0,300

Payoffs to: (Entrant, Incumbent).

The strategy sets can be discovered from the order of play. They are $\{Enter, Stay Out\}$ for the entrant, and $\{Collude \text{ if entry occurs}, Fight \text{ if entry occurs}\}$ for the incumbent. The game has the two Nash equilibria indicated in boldface in Table 1, $(Enter, Collude)$ and $(Stay Out, Fight)$. The equilibrium $(Stay Out, Fight)$ is weak, because the incumbent would just as soon *Collude* given that the entrant is staying out.

Figure 3 Entry Deterrence I

A piece of information has been lost by condensing from the extensive form, Figure 3, to the strategic form, Table 1, i.e., the fact that the entrant gets to move first. Once he has chosen *Enter*, the incumbent's best response is *Collude*. The threat to fight is not credible and would be employed only if the incumbent could bind himself to fight, in which case he never does fight, because the entrant chooses to stay out. The equilibrium $(Stay Out, Fight)$ is Nash but not subgame perfect, because if the game is started after the entrant has already entered, the incumbent's best response is *Collude*. This does not prove that collusion is inevitable in duopoly, but it is the equilibrium for Entry Deterrence I.

The trembling hand interpretation of perfect equilibrium can be used here. So long as it is certain that the entrant will not enter, the incumbent is indifferent between *Fight* and

Collude, but if there were even a small probability of entry— perhaps because of a lapse of good judgement by the entrant— the incumbent would prefer *Collude* and the Nash equilibrium would be broken.

Perfectness rules out threats that are not credible. Entry Deterrence I is a good example because if a communication move were added to the game tree, the incumbent might tell the entrant that entry would be followed by fighting, but the entrant would ignore this noncredible threat. If, however, some means existed by which the incumbent could precommit himself to fight entry, the threat would become credible. The next section will look at one context, nuisance lawsuits, in which such precommitment might be possible

Should the Modeller Ever Use Nonperfect Equilibria?

A game in which a player can commit himself to a strategy can be modelled in two ways:

- 1 As a game in which nonperfect equilibria are acceptable, or
- 2 By changing the game to replace the action *Do X* with *Commit to Do X* at an earlier node.

An example of (2) in Entry Deterrence I is to reformulate the game so the incumbent moves first, deciding in advance whether or not to choose *Fight* before the entrant moves. Approach (2) is better than (1) because if the modeller wants to let players commit to some actions and not to others, he can do this by carefully specifying the order of play. Allowing equilibria to be nonperfect forbids such discrimination and multiplies the number of equilibria. Indeed, the problem with subgame perfectness is not that it is too restrictive but that it still allows too many strategy combinations to be equilibria in games of asymmetric information. A subgame must start at a single node and not cut across any player's information set, so often the only subgame will be the whole game and subgame perfectness does not restrict equilibrium at all. Section 6.1 discusses perfect Bayesian equilibrium and other ways to extend the perfectness concept to games of asymmetric information.

4.3 Credible Threats, Sunk Costs, and the Open-set Problem in the Game of Nuisance Suits

Like the related concepts of sunk costs and rational expectations, sequential rationality is a simple idea with tremendous power. This section will show that power in another simple game, one which models nuisance suits. We have already come across one application of game theory to law, in the Png (1983) model of Section 2.5. In some ways, law is particularly well suited to analysis by game theory because the legal process is so concerned with conflict and the provision of definite rules to regulate that conflict. In what other field could an article be titled: "An Economic Analysis of Rule 68," as Miller (1986) does in his discussion of the federal rule of procedure that penalizes a losing litigant who had refused to accept a settlement offer. The growth in the area can be seen by comparing the overview in the Ayres's (1990) review of the first edition of the present book with the entire book by Baird, Gertner & Picker (1994). In law, even more clearly than in business, a major objective is to avoid inefficient outcomes by restructuring the rules, and nuisance suits are one of the inefficiencies that a good policy maker hopes to eliminate.

Nuisance suits are lawsuits with little chance of success, whose only possible purpose seems to be the hope of a settlement out of court. In the context of entry deterrence people commonly think large size is an advantage and a large incumbent will threaten a small entrant, but in the context of nuisance suits people commonly think large size is a disadvantage and a wealthy corporation is vulnerable to extortionary litigation. Nuisance Suits I models the essentials of the situation: bringing suit is costly and has little chance of success, but because defending the suit is also costly the defendant might pay a generous amount to settle it out of court. The model is similar to the Png Settlement Game of Chapter 2 in many respects, but here the model will be one of symmetric information and we will make explicit the sequential rationality requirement that was implicit the discussion in Chapter 2.

Nuisance Suits I: Simple Extortion

Players

A plaintiff and a defendant.

The Order of Play

- 1 The plaintiff decides whether to bring suit against the defendant at cost c .
- 2 The plaintiff makes a take-it-or-leave-it settlement offer of $s > 0$.
- 3 The defendant accepts or rejects the settlement offer.
- 4 If the defendant rejects the offer, the plaintiff decides whether to give up or go to trial at a cost p to himself and d to the defendant.
- 5 If the case goes to trial, the plaintiff wins amount x with probability γ and otherwise wins nothing.

Payoffs

Figure 4 shows the payoffs. Let $\gamma x < p$, so the plaintiff's expected winnings are less than his marginal cost of going to trial.




Figure 4: The Extensive Form for *Nuisance Suits*

The perfect equilibrium is

Plaintiff: *Do nothing*, *Offer s* , *Give up*

Defendant: *Reject*

Outcome: The plaintiff does not bring a suit.

The equilibrium settlement offer s can be any positive amount. Note that the equilibrium specifies actions at all four nodes of the game, even though only the first is reached in equilibrium.

To find a perfect equilibrium the modeller starts at the end of the game tree, following the advice of Dixit & Nalebuff (1991, p. 34) to “Look ahead and reason back.” At node P_3 , the plaintiff will choose *Give up*, since by assumption $\gamma x - c - p < -c$. This is because the suit is brought only in the hope of settlement, not in the hope of winning at trial. At node D_1 , the defendant, foreseeing that the plaintiff will give up, rejects any positive settlement offer. This makes the plaintiff’s offer at P_2 irrelevant, and, looking ahead to a payoff of $-c$ from choosing *Sue* at P_1 , the plaintiff chooses *Do nothing*.

Thus, if nuisance suits are brought, it must be for some reason other than the obvious one, the plaintiff’s hope of extracting a settlement offer from a defendant who wants to avoid trial costs. This is fallacious because the plaintiff himself bears trial costs and hence cannot credibly make the threat. It is fallacious even if the defendant’s legal costs would be much higher than the plaintiff’s (d much bigger than p), because the relative size of the costs does not enter into the argument.

One might wonder how risk aversion affects this conclusion. Might not the defendant settle because he is more risk averse than the plaintiff? That is a good question, but Nuisance Suits I can be adapted to risk-averse players with very little change. Risk would enter at the trial stage, as a final move by Nature to decide who wins. In Nuisance Suits I, γx represented the expected value of the award. If both the defendant and the plaintiff are equally risk averse, γx can still represent the expected payoff from the award—one simply interprets x and 0 as the utility of the cash award and the utility of an award of 0, rather than as the actual cash amounts. If the players have different degrees of risk aversion, the expected loss to the defendant is not the same as the expected gain to the plaintiff, and the payoffs must be adjusted. If the defendant is more risk averse, the payoffs from *Go to trial* would change to $(-c - p + \gamma x, -\gamma x - y - d)$, where y represents the extra disutility of risk to the defendant. This, however, makes no difference to the equilibrium. The crux of the game is that the plaintiff is unwilling to go to trial because of the cost to himself, and the cost to the defendant, including the cost of bearing risk, is irrelevant.

If nuisance suits are brought, it must therefore be for some more complicated reason. Already, in chapter 2, we looked at one reason for litigation to reach trial in the *Png Settlement Game*: incomplete information. That is probably the most important explanation and it has been much studied, as can be seen from the surveys by Cooter & Rubinfeld (1989) and Kennan and R. Wilson (1993). In this section, though, let us confine ourselves to explanations where the probability of the suit’s success is common knowledge. Even

then, costly threats might be credible because of sinking costs strategically (*Nuisance Suits II*), or because of the nonmonetary payoffs resulting from going to trial (*Nuisance Suits III*) .

Nuisance Suits II : Using Sunk Costs Strategically ²

Let us now modify the game so that the plaintiff can pay his lawyer the amount p in advance, with no refund if the case settles. This inability to obtain a refund actually helps the plaintiff, by changing the payoffs from the game so his payoff from *Give up* is $-c - p$, compared to $-c - p + \gamma x$ from *Go to trial*. Having sunk the legal costs, he will go to trial if $\gamma x > 0$ — that is, if he has any chance of success at all.³

This, in turn, means that the plaintiff would only prefer settlement to trial if $s > \gamma x$. The defendant would prefer settlement to trial if $s < \gamma x + d$, so there is a positive **settlement range** of $[\gamma x, \gamma x + d]$ within which both players are willing to settle. The exact amount of the settlement depends on the bargaining power of the parties, something to be examined in chapter 11. Here, allowing the plaintiff to make a take-it-or-leave-it offer means that $s = \gamma x + d$ in equilibrium, and if $\gamma x + d > p + c$, the nuisance suit will be brought even though $\gamma x < p + c$. Thus, the plaintiff is bringing the suit only because he can extort d , the amount of the defendant's legal costs.

Even though the plaintiff can now extort a settlement, he does it at some cost to himself, so an equilibrium with nuisance suits will require that

$$-c - p + \gamma x + d \geq 0 \quad (1)$$

If inequality (1) is false, then, even if the plaintiff could extract the maximum possible settlement of $s = \gamma x + d$, he would not do so, because he would then have to pay $c + p$ before reaching the settlement stage. This implies that a totally meritless suit (with $\gamma = 0$), would not be brought unless the defendant had higher legal costs than the plaintiff ($d > p$). If inequality (1) is satisfied, however, the following strategy combination is a perfect equilibrium:

Plaintiff: *Sue, Offer $s = \gamma x + d$, Go to trial*

Defendant: *Accept $s \leq \gamma x + d$*

Outcome: Plaintiff sues and offers to settle, to which the defendant agrees.

An obvious counter to the plaintiff's ploy would be for the defendant to also sink his costs, by paying d before the settlement negotiations, or even before the plaintiff decides to file suit. Perhaps this is one reason why large corporations use in house counsel, who are paid a salary regardless of how many hours they work, as well as outside counsel, hired by the hour. If so, nuisance suits cause a social loss—the wasted time of the lawyers, d —even if nuisance suits are never brought, just as aggressor nations cause social loss in the form of world military expenditure even if they never start a war.⁴

²The inspiration for this model is Rosenberg & Shavell (1985).

³Have a figure with the previous NS I Extensive Form and the new one side by side, with the differences highlighted.

⁴Nonrefundable lawyer's fees, paid in advance, have traditionally been acceptable, but a New York

Two problems, however, face the defendant who tries to sink the cost d . First, although it saves him γx if it deters the plaintiff from filing suit, it also means the defendant must pay the full amount d . This is worthwhile if the plaintiff has all the bargaining power, as in Nuisance Suits II, but it might not be if s lay in the middle of the settlement range because the plaintiff was not able to make a take-it-or-leave-it offer. If settlement negotiations resulted in s lying exactly in the middle of the settlement range, so $s = \gamma x + \frac{d}{2}$, then it might not be worthwhile for the defendant to sink d to deter nuisance suits that would settle for $\gamma x + \frac{d}{2}$.

Second, there is an asymmetry in litigation: the plaintiff has the choice of whether to bring suit or not. Since it is the plaintiff who has the initiative, he can sink p and make the settlement offer before the defendant has the chance to sink d . The only way for the defendant to avoid this is to pay d well in advance, in which case the expenditure is wasted if no possible suits arise. What the defendant would like best would be to buy legal insurance which, for a small premium, would pay all defense costs in future suits that might occur. As we will see in Chapters 7 and 9, however, insurance of any kind faces problems arising from asymmetric information. In this context, there is the “moral hazard” problem, in that once the defendant is insured he has less incentive to avoid causing harm to the plaintiff and provoking a lawsuit.

The Open-set Problem in Nuisance Suits II

Nuisance Suits II illustrates a technical point that arises in a great many games with continuous strategy spaces and causes great distress to novices in game theory. The equilibrium in Nuisance Suits II is only a weak Nash equilibrium. The plaintiff proposes $s = \gamma x + d$, and the defendant has the same payoff from accepting or rejecting, but in equilibrium the defendant accepts the offer with probability one, despite his indifference. This seems arbitrary, or even silly. Should not the plaintiff propose a slightly lower settlement to give the defendant a strong incentive to accept it and avoid the risk of having to go to trial? If the parameters are such that $s = \gamma x + d = 60$, for example, why does the plaintiff risk holding out for 60 when he might be rejected and most likely receive 0 at trial, when he could offer 59 and give the defendant a strong incentive to accept?

One answer is that no other equilibrium exists besides $s = 60$. Offering 59 cannot be part of an equilibrium because it is dominated by offering 59.9; offering 59.9 is dominated by offering 59.99, and so forth. This is known as the **open-set problem**, because the set of offers that the defendant strongly wishes to accept is open and has no maximum—it is bounded at 60, but a set must be bounded *and closed* to guarantee that a maximum exists.

A second answer is that under the assumptions of rationality and Nash equilibrium the objection’s premise is false because the plaintiff bears no risk whatsoever in offering $s = 60$. It is fundamental to Nash equilibrium that each player believe that the others will follow equilibrium behavior. Thus, if the equilibrium strategy combination says that

court recently ruled they were unethical. The court thought that such fees unfairly restricted the client’s ability to fire his lawyer, an example of how ignorance of game theory can lead to confused rule-making. See “Nonrefundable Lawyers’ Fees, Paid in Advance, are Unethical, Court Rules,” *Wall Street Journal*, January 29, 1993, p. B3, citing *In the matter of Edward M. Cooperman*, Appellate Division of the Supreme Court, Second Judicial Department, Brooklyn, 90-00429.

the defendant will accept $s \leq 60$, the plaintiff can offer 60 and believe it will be accepted. This is really just to say that a weak Nash equilibrium is still a Nash equilibrium, a point emphasized in chapter 3 in connection with mixed strategies.

A third answer is that the problem is an artifact of using a model with a continuous strategy space, and it disappears if the strategy space is made discrete. Assume that s can only take values in multiples of 0.01, so it could be 59.0, 59.01, 59.02, and so forth, but not 59.001 or 59.002. The settlement part of the game will now have two perfect equilibria. In the strong equilibrium E1, $s = 59.99$ and the defendant accepts any offer $s < 60$. In the weak equilibrium E2, $s = 60$ and the defendant accepts any offer $s \leq 60$. The difference is trivial, so the discrete strategy space has made the model more complicated without any extra insight.⁵

One can also specify a more complicated bargaining game to avoid the issue of how exactly the settlement is determined. Here one could say that the settlement is not proposed by the plaintiff, but simply emerges with a value halfway through the settlement range, so $s = \gamma x + \frac{d}{2}$. This seems reasonable enough, and it adds a little extra realism to the model at the cost of a little extra complexity. It avoids the open-set problem, but only by avoiding being clear about how s is determined. I call this kind of modelling **blackboxing**, because it is as if at some point in the game, variables with certain values go into a black box and come out the other side with values determined by an exogenous process. Blackboxing is perfectly acceptable as long as it neither drives nor obscures the point the model is making. Nuisance Suits III will illustrate this method.

Fundamentally, however, the point to keep in mind is that games are models, not reality. They are meant to clear away the unimportant details of a real situation and simplify it down to the essentials. Since a model is trying to answer a question, it should focus on what answers that question. Here, the question is why nuisance suits might be brought, so it is proper to exclude details of the bargaining if they are irrelevant to the answer. Whether a plaintiff offers 59.99 or 60, and whether a rational person accepts an offer with probability 0.99 or 1.00, is part of the unimportant detail, and whatever approach is simplest should be used. If the modeller really thinks that these are important matters, they can indeed be modelled, but they are not important in this context.

One source of concern over the open-set problem, I think, is that perhaps that the payoffs are not quite realistic, because the players should derive utility from hurting “unfair” players. If the plaintiff makes a settlement offer of 60, keeping the entire savings from avoiding the trial for himself, everyday experience tells us that the defendant will indignantly refuse the offer. Guth *et al.* (1982) have found in experiments that people turn down bargaining offers they perceive as unfair, as one might expect. If indignation is truly important, it can be explicitly incorporated into the payoffs, and if that is done, the open-set problem returns. Indignation is not boundless, whatever people may say. Suppose that accepting a settlement offer that benefits the plaintiff more than the defendant gives a disutility of x to the defendant because of his indignation at his unjust treatment. The

⁵A good example of the ideas of discrete money values and sequential rationality is in Robert Louis Stevenson’s story, “The Bottle Imp” (Stevenson [1987]). The imp grants the wishes of the bottle’s owner but will seize his soul if he dies in possession of it. Although the bottle cannot be given away, it can be sold, but only at a price less than that for which it was purchased.

plaintiff will then offer to settle for exactly $60 - x$, so the equilibrium is still weak and the defendant is still indifferent between accepting and rejecting the offer. The open-set problem persists, even after realistic emotions are added to the model.

I have spent so much time on the open-set problem not because it is important but because it arises so often and is a sticking point for people unfamiliar with modelling. It is not a problem that disturbs experienced modellers, unlike other basic issues we have already encountered—for example, the issue of how a Nash equilibrium comes to be common knowledge among the players—but it is important to understand why it is not important.

Nuisance Suits III: Malice

One of the most common misconceptions about game theory, as about economics in general, is that it ignores non-rational and non-monetary motivations. Game theory does take the basic motivations of the players to be exogenous to the model, but those motivations are crucial to the outcome and they often are not monetary, although payoffs are always given numerical values. Game theory does not call somebody irrational who prefers leisure to money or who is motivated by the desire to be world dictator. It does require the players' emotions to be carefully gauged to determine exactly how the actions and outcomes affect the players' utility.

Emotions are often important to lawsuits, and law professors tell their students that when the cases they study seem to involve disputes too trivial to be worth taking to court, they can guess that the real motivations are emotional. Emotions could enter in a variety of distinct ways. The plaintiff might simply like going to trial, which can be expressed as a value of $p < 0$. This would be true of many criminal cases, because prosecutors like news coverage and want credit with the public for prosecuting certain kinds of crime. The Rodney King trials of 1992 and 1993 were of this variety; regardless of the merits of the cases against the policemen who beat Rodney King, the prosecutors wanted to go to trial to satisfy the public outrage, and when the state prosecutors failed in the first trial, the federal government was happy to accept the cost of bringing suit in the second trial. A different motivation is that the plaintiff might derive utility from the fact of winning the case quite separately from the monetary award, because he wants a public statement that he is in the right. This is a motivation in bringing libel suits, or for a criminal defendant who wants to clear his good name.

A different emotional motivation for going to trial is the desire to inflict losses on the defendant, a motivation we will call “malice,” although it might as inaccurately be called “righteous anger.” In this case, d enters as a positive argument in the plaintiff's utility function. We will construct a model of this kind, called Nuisance Suits III, and assume that $\gamma = 0.1$, $c = 3$, $p = 14$, $d = 50$, and $x = 100$, and that the plaintiff receives additional utility of 0.1 times the defendant's disutility. Let us also adopt the blackboxing technique discussed earlier and assume that the settlement s is in the middle of the settlement range. The payoffs conditional on suit being brought are

$$\pi_{plaintiff}(Defendant\ accepts) = s - c + 0.1s = 1.1s - 3 \quad (2)$$

and

$$\begin{aligned}\pi_{\text{plaintiff}}(Go\ to\ trial) &= \gamma x - c - p + 0.1(d + \gamma x) \\ &= 10 - 3 - 14 + 6 = -1.\end{aligned}\tag{3}$$

Now, working back from the end in accordance with sequential rationality, note that since the plaintiff's payoff from *Give Up* is -3 , he will go to trial if the defendant rejects the settlement offer. The overall payoff from bringing a suit that eventually goes to trial is still -1 , which is worse than the payoff of 0 from not bringing suit in the first place, but if s is high enough, the payoff from bringing suit and settling is higher still. If s is greater than 1.82 ($= \frac{-1+3}{1.1}$, rounded), the plaintiff prefers settlement to trial, and if s is greater than about 2.73 ($= \frac{0+3}{1.1}$, rounded), he prefers settlement to not bringing the suit at all.

In determining the settlement range, the relevant payoff is the expected incremental payoff since the suit was brought. The plaintiff will settle for any $s \geq 1.82$, and the defendant will settle for any $s \leq \gamma x + d = 60$, as before. The settlement range is $[1.82, 60]$, and $s = 30.91$. The settlement offer is no longer the maximizing choice of a player, and hence is moved to the outcome in the equilibrium description below.

Plaintiff: *Sue, Go to Trial*

Defendant: *Accept any $s \leq 60$*

Outcome: The plaintiff sues and offers $s = 30.91$, and the defendant accepts the settlement.

Perfectness is important here because the defendant would like to threaten never to settle and be believed. The plaintiff would not bring suit given his expected payoff of -1 from bringing a suit that goes to trial, so a believable threat would be effective. But such a threat is not believable. Once the plaintiff does bring suit, the only Nash equilibrium in the remaining subgame is for the defendant to accept his settlement offer. This is interesting because the plaintiff, despite his willingness to go to trial, ends up settling out of court. When information is symmetric, as it is here, there is a tendency for equilibria to be efficient. Although the plaintiff wants to hurt the defendant, he also wants to keep his expenses low. Thus, he is willing to hurt the defendant less if it enables him to save on his own legal costs.

One final point before leaving these models is that much of the value of modelling comes simply from setting up the rules of the game, which helps to show what is important in a situation. One problem that arises in setting up a model of nuisance suits is deciding what a “nuisance suit” really is. In the game of Nuisance Suits, it has been defined as a suit whose expected damages do not repay the plaintiff's costs of going to trial. But having to formulate a definition brings to mind another problem that might be called the problem of nuisance suits: that the plaintiff brings suits he knows will not win unless the court makes a mistake. Since the court might make a mistake with very high probability, the games above would not be appropriate models— γ would be high, and the problem is not that the plaintiff's expected gain from trial is low, but that it is high. This, too, is an important problem, but having to construct a model shows that it is different.

4.4 Recoordination to Pareto-dominant Equilibria in Subgames: Pareto Perfection

One simple refinement of equilibrium that was mentioned in chapter 1 is to rule out any strategy combinations that are Pareto dominated by Nash equilibria. Thus, in the game of Ranked Coordination, the inferior Nash equilibrium would be ruled out as an acceptable equilibrium. The idea behind this is that in some unmodelled way the players discuss their situation and coordinate to avoid the bad equilibria. Since only Nash equilibria are discussed, the players' agreements are self-enforcing and this is a more limited suggestion than the approach in cooperative game theory according to which the players make binding agreements.

The coordination idea can be taken further in various ways. One is to think about coalitions of players coordinating on favorable equilibria, so that two players might coordinate on an equilibrium even if a third player dislikes it. Bernheim, Peleg, & Whinston (1987) and Bernheim & Whinston (1987) define a Nash strategy combination as a **coalition-proof Nash equilibrium** if no coalition of players could form a self-enforcing agreement to deviate from it. They take the idea further by subordinating it to the idea of sequential rationality. The natural way to do this is to require that no coalition would deviate in future subgames, a notion called by various names, including **renegotiation proofness**, **recoordination** (e.g., Laffont & Tirole [1993], p. 460), and **Pareto perfection** (e.g., Fudenberg & Tirole (1991a), p. 175). The idea has been used extensively in the analysis of infinitely repeated games, which are particularly subject to the problem of multiple equilibria; Abreu, Pearce & Stachetti (1986) is an example of this literature. Whichever name is used, the idea is distinct from the renegotiation problem in the principal-agent models to be studied in Chapter 8, which involves the rewriting of earlier binding contracts to make new binding contracts.

The best way to demonstrate the idea of Pareto perfection is by an illustration, the Pareto Perfection Puzzle, whose extensive form is shown in Figure 5. In this game Smith chooses *In* or *Outside Option 1*, which yields payoffs of 10 to each player. Jones then chooses *Outside Option 2*, which yields 20 to each player, or initiates either a coordination game or a prisoner's dilemma. Rather than draw the full subgames in extensive form, Figure 5 inserts the payoff matrix for the subgames.




Figure 5: The Pareto Perfection Puzzle

The Pareto Perfection Puzzle illustrates the complicated interplay between perfectness and Pareto dominance. The pareto-dominant strategy combination is (*In*, *Prisoner's Dilemma*|*In*, any actions in the coordination subgame, the actions yielding (50,50) in the *Prisoner's Dilemma* subgame). Nobody expects this strategy combination to be an equilibrium, since it is neither perfect nor Nash. Perfectness tells us that if the *Prisoner's Dilemma* subgame is reached, the payoffs will be (0,0), and if the coordination subgame is reached they will be either (1,1) or (2,30). In light of this, the perfect equilibria of the Pareto Perfection Puzzle are:

E1: (*In*, outside option 2|*In*, the actions yielding (1,1) in the coordination subgame, the actions yielding (0,0) in the *Prisoner's Dilemma* subgame). The payoffs are (20,20).

E2: (outside option 1, coordination game|*In*, the actions yielding (2,30) in the coordination subgame, the actions yielding (0,0) in the *Prisoner's Dilemma* subgame). The payoffs are (10,10).

If one applies Pareto dominance without perfection, E1 will be the equilibrium, since both players prefer it. If the players can recoordinate at any point and change their expectations, however, then if play of the game reaches the coordination subgame, the players will recoordinate on the actions yielding (2,30). Pareto perfection thus knocks out E1 as an equilibrium. Not only does it rule out the Pareto-dominant strategy combination that yields (50,50) as an equilibrium, it also rules out the Pareto-dominant perfect strategy combination that yields (20,20) as an equilibrium. Rather, the payoff is (10,10). Thus, Pareto perfection is not the same thing as simply picking the Pareto-dominant perfect strategy combination.

It is difficult to say which equilibrium is best here, since this is an abstract game and we cannot call upon details from the real world to refine the model. The approach of applying an equilibrium refinement is not as likely to yield results as using the intuition behind the refinement. The intuition here is that the players will somehow coordinate on Pareto-dominant equilibria, perhaps finding open discussion helpful. If we ran an experiment on student players using the Pareto Perfection Puzzle, I would expect to reach different equilibria depending on what communication is allowed. If the players are allowed to talk only before the game starts, it seems more likely that E1 would be the equilibrium, since players could agree to play it and would have no chance to explicitly recoordinate later. If the players could talk at any time as the game proceeded, E2 becomes more plausible. Real-world situations arise with many different communications technologies, so there is no one right answer.

Notes

N4.1 Subgame perfectness

- The terms “perfectness” and “perfection” are used synonymously. Selten (1965) proposed the equilibrium concept in an article written in German. “Perfectness” is used in Selten (1975) and conveys an impression of completeness more appropriate to the concept than the goodness implied by “perfection.” “Perfection,” however, is more common.
- It is debatable whether the definition of subgame ought to include the original game. Gibbon (1992, p. 122) does not, for example, and modellers usually do not in their conversation.
- Perfectness is not the only way to eliminate weak Nash equilibria like (*Stay Out, Collude*). In Entry Deterrence I, (*Enter, Collude*) is the only iterated dominance equilibrium, because *Fight* is weakly dominated for the incumbent.
- The distinction between perfect and non-perfect Nash equilibria is like the distinction between **closed loop** and **open loop** trajectories in dynamic programming. Closed loop (or **feedback**) trajectories can be revised after they start, like perfect equilibrium strategies, while open loop trajectories are completely prespecified (though they may depend on state variables). In dynamic programming the distinction is not so important, because prespecified strategies do not change the behavior of other players. No threat, for example, is going to alter the pull of the moon’s gravity on a rocket.
- A subgame can be infinite in length, and infinite games can have non-perfect equilibria. The infinitely repeated *Prisoner’s Dilemma* is an example; here every subgame looks exactly like the original game, but begins at a different point in time.
- **Sequential rationality in macroeconomics.** In macroeconomics the requirement of **dynamic consistency** or **time consistency** is similar to perfectness. These terms are less precisely defined than perfectness, but they usually require that strategies need only be best responses in subgames starting from nodes on the equilibrium path, instead of all subgames. Under this interpretation, time consistency is a less stringent condition than perfectness.

The Federal Reserve, for example, might like to induce inflation to stimulate the economy, but the economy is stimulated only if the inflation is unexpected. If the inflation is expected, its effects are purely bad. Since members of the public know that the Fed would like to fool them, they disbelieve its claims that it will not generate inflation (see Kydland & Prescott [1977]). Likewise, the government would like to issue nominal debt, and promises lenders that it will keep inflation low, but once the debt is issued, the government has an incentive to inflate its real value to zero. One reason the US Federal Reserve Board was established to be independent of Congress in the United States was to diminish this problem.

The amount of game theory used in macroeconomics has been increasing at a fast rate. For references see Canzoneri & Henderson’s 1991 book, which focusses on international coordination and pays particular attention to trigger strategies.

- Often, irrationality— behavior that is automatic rather than strategic— is an advantage. The Doomsday Machine in the movie *Dr Strangelove* is one example. The Soviet Union decides that it cannot win a rational arms race against the richer United States, so it creates a bomb which automatically blows up the entire world if anyone explodes a nuclear bomb. The movie also illustrates a crucial detail without which such irrationality is worse than useless: you have to tell the other side that you have the Doomsday Machine.

President Nixon reportedly told his aide H.R. Haldeman that he followed a more complicated version of this strategy: “I call it the Madman Theory, Bob. I want the North Vietnamese to believe that I’ve reached the point where I might do *anything* to stop the war. We’ll just slip the word to them that ‘for God’s sake, you know Nixon is obsessed about Communism. We can’t restrain him when he’s angry— and he has his hand on the nuclear button’— and Ho Chi Minh himself will be in Paris in two days begging for peace”(Haldeman & DiMona [1978] p. 83). The Gang of Four model in section 6.4 tries to model a situation like this.

- The “lock-up agreement” is an example of a credible threat: in a takeover defense, the threat to destroy the firm is made legally binding. See Macey & McChesney (1985) p. 33.

N4.3 An example of perfectness: Entry Deterrence I

- The Stackelberg equilibrium of a duopoly game (section 3.4) can be viewed as the perfect equilibrium of a Cournot game modified so that one player moves first, a game similar to Entry Deterrence I. The player moving first is the Stackelberg leader and the player moving second is the Stackelberg follower. The follower could threaten to produce a high output, but he will not carry out his threat if the leader produces a high output first.
- Perfectness is not so desirable a property of equilibrium in biological games. The reason the order of moves matters is because the rational best reply depends on the node at which the game has arrived. In many biological games the players act by instinct and unthinking behavior is not unrealistic.
- Reinganum & Stokey (1985) is a clear presentation of the implications of perfectness and commitment illustrated with the example of natural resource extraction.

Problems

4.1. Repeated Entry Deterrence

Consider two repetitions without discounting of the game Entry Deterrence I from Section 4.2. Assume that there is one entrant, who sequentially decides whether to enter two markets that have the same incumbent.

- a Draw the extensive form of this game.
- b What are the 16 elements of the strategy sets of the entrant?
- c What is the subgame perfect equilibrium?
- d What is one of the nonperfect Nash equilibria?

4.2. The Three-Way Duel (after Shubik (1954))

Three gangsters armed with pistols, Al, Bob, and Curly, are in a room with a suitcase containing 120 thousand dollars. Al is the least accurate, with a 20 percent chance of killing his target. Bob has a 40 percent probability. Curly is slow but sure; he kills his target with 70 percent probability. For each, the value of his own life outweighs the value of any amount of money. Survivors split the money.

- (a) Suppose each gangster has one bullet and the order of shooting is first Al, then Bob, then Curly. Assume also that each gangster must try to kill another gangster when his turn comes. What is an equilibrium strategy combination and what is the probability that each of them dies in that equilibrium? Hint: Do not try to draw a game tree.
- (b) Suppose now that each gangster has the additional option of shooting his gun at the ceiling, which may kill somebody upstairs but has no direct effect on his payoff. Does the strategy combination that you found was an equilibrium in part (a) remain an equilibrium?
- (c) Replace the three gangsters with three companies, Apex, Brydox, and Costco, which are competing with slightly different products. What story can you tell about their advertising strategies?
- (d) In the United States, before the general election a candidate must win the nomination of his party. It is often noted that candidates are reluctant to be seen as the frontrunner in the race for the nomination of their party, Democrat or Republican. In the general election, however, no candidate ever minds being seen to be ahead of his rival from the other party. Why?
- (e) In the 1920's, several men vied for power in the Soviet Union after Lenin died. First Stalin and Zinoviev combined against Trotsky. Then Stalin and Bukharin combined against Zinoviev. Then Stalin turned on Bukharin. Relate this to Curly, Bob, and Al.

4.3. Heresthetics in Pliny and the freedmens' trial (Pliny, 1963, pp. 221-4, Riker, 1986, pp. 78-88)

Afranius Dexter died mysteriously, perhaps dead by his own hand, perhaps killed by his freedmen (servants a step above slaves), or perhaps killed by his freedmen by his own orders. The freedmen

went on trial before the Roman Senate. Assume that 45 percent of the senators favor acquittal, 35 percent favor banishment, and 20 percent favor execution, and that the preference rankings in the three groups are $A \succ B \succ E$, $B \succ A \succ E$, and $E \succ B \succ A$. Also assume that each group has a leader and votes as a bloc.

- a Modern legal procedure requires the court to decide guilt first and then assign a penalty if the accused is found guilty. Draw a tree to represent the sequence of events (this will not be a game tree, since it will represent the actions of groups of players, not of individuals). What is the outcome in a perfect equilibrium?
- b Suppose that the acquittal bloc can pre-commit to how they will vote in the second round if guilt wins in the first round. What will they do, and what will happen? What would the execution bloc do if they could control the second-period vote of the acquittal bloc?
- c The normal Roman procedure began with a vote on execution versus no execution, and then voted on the alternatives in a second round if execution failed to gain a majority. Draw a tree to represent this. What would happen in this case?
- d Pliny proposed that the Senators divide into three groups, depending on whether they supported acquittal, banishment, or execution, and that the outcome with the most votes should win. This proposal caused a roar of protest. Why did he propose it?
- e Pliny did not get the result he wanted with his voting procedure. Why not?
- (f Suppose that personal considerations made it most important to a senator that he show his stand by his vote, even if he had to sacrifice his preference for a particular outcome. If there were a vote over whether to use the traditional Roman procedure or Pliny's procedure, who would vote with Pliny, and what would happen to the freedmen?

4.4. DROPPED

4.5. Garbage Entry

Mr. Turner is thinking of entering the garbage collection business in a certain large city. Currently, Cutright Enterprises has a monopoly, earning 40 million dollars from the 40 routes the city offers up for bids. Turner thinks he can take away as many routes as he wants from Cutright, at a profit of 1.5 million per route for him. He is worried, however, that Cutright might resort to assassination, killing him to regain their lost routes. He would be willing to be assassinated for profit of 80 million dollars, and assassination would cost Cutright 6 million dollars in expected legal costs and possible prison sentences.

How many routes should Turner try to take away from Cutright?

5 Reputation and Repeated Games with Symmetric Information

September 11, 1999. February 3, 2000. February 6, 2000. November 29, 2003. Eric Rasmusen, Erasmuse@indiana.edu. [Http://mypage.iu.edu/~erasmuse](http://mypage.iu.edu/~erasmuse). Footnotes starting with xxx are the author's notes to himself. Comments welcomed.

5.1 Finitely Repeated Games and the Chainstore Paradox

Chapter 4 showed how to refine the concept of Nash equilibrium to find sensible equilibria in games with moves in sequence over time, so-called dynamic games. An important class of dynamic games is repeated games, in which players repeatedly make the same decision in the same environment. Chapter 5 will look at such games, in which the rules of the game remain unchanged with each repetition and all that changes is the “history” which grows as time passes, and, if the number of repetitions is finite, the approach of the end of the game. It is also possible for asymmetry of information to change over time in a repeated game since players' moves may convey their private information, but Chapter 5 will confine itself to games of symmetric information.

Section 5.1 will show the perverse unimportance of repetition for the games of Entry Deterrence and the Prisoner's Dilemma, a phenomenon known as the Chainstore Paradox. Neither discounting, probabilistic end dates, infinite repetitions, nor precommitment are satisfactory escapes from the Chainstore Paradox. This is summarized in the Folk Theorem of Section 5.2. Section 5.2 will also discuss strategies which punish players who fail to cooperate in a repeated game— strategies such as the Grim Strategy, Tit-for-Tat, and Minimax. Section 5.3 builds a framework for reputation models based on the Prisoner's Dilemma, and Section 5.4 presents one particular reputation model, the Klein-Leffler model of product quality. Section 5.5 concludes the chapter with an overlapping generations model of consumer switching costs which uses the idea of Markov strategies to narrow down the number of equilibria.

The Chainstore Paradox

Suppose that we repeat Entry Deterrence I 20 times in the context of a chainstore that is trying to deter entry into 20 markets where it has outlets. We have seen that entry into just one market would not be deterred, but perhaps with 20 markets the outcome is different because the chainstore would fight the first entrant to deter the next 19.

The repeated game is much more complicated than the **one-shot game**, as the unrepeated version is called. A player's action is still to *Enter* or *Stay Out*, to *Fight* or *Collude*, but his strategy is a potentially very complicated rule telling him what action to choose depending on what actions both players took in each of the previous periods. Even the five-round repeated Prisoner's Dilemma has a strategy set for each player with over two billion strategies, and the number of strategy combinations is even greater (Sugden [1986], p. 108).

The obvious way to solve the game is from the beginning, where there is the least

past history on which to condition a strategy, but that is not the easy way. We have to follow Kierkegaard, who said, “Life can only be understood backwards, but it must be lived forwards” (Kierkegaard 1938, p. 465). In picking his first action, a player looks ahead to its implications for all the future periods, so it is easiest to start by understanding the end of a multi-period game, where the future is shortest.

Consider the situation in which 19 markets have already been invaded (and maybe the chainstore fought, or maybe not). In the last market, the subgame in which the two players find themselves is identical to the one-shot Entry Deterrence I, so the entrant will *Enter* and the chainstore will *Collude*, regardless of the past history of the game. Next, consider the next-to-last market. The chainstore can gain nothing from building a reputation for ferocity, because it is common knowledge that he will *Collude* with the last entrant anyway. So he might as well *Collude* in the 19th market. But we can say the same of the 18th market and— by continuing backward induction— of every market, including the first. This result is called the **Chainstore Paradox** after Selten (1978).

Backward induction ensures that the strategy combination is a subgame perfect equilibrium. There are other Nash equilibria— (*Always Fight, Never Enter*), for example— but because of the Chainstore Paradox they are not perfect.

The Repeated Prisoner’s Dilemma

The Prisoner’s Dilemma is similar to Entry Deterrence I. Here the prisoners would like to commit themselves to *Deny*, but, in the absence of commitment, they *Confess*. The Chainstore Paradox can be applied to show that repetition does not induce cooperative behavior. Both prisoners know that in the last repetition, both will *Confess*. After 18 repetitions, they know that no matter what happens in the 19th, both will *Confess* in the 20th, so they might as well *Confess* in the 19th too. Building a reputation is pointless, because in the 20th period it is not going to matter. Proceeding inductively, both players *Confess* in every period, the unique perfect equilibrium outcome.

In fact, as a consequence of the fact that the one-shot Prisoner’s Dilemma has a dominant strategy equilibrium, confessing is the only Nash outcome for the repeated Prisoner’s Dilemma, not just the only perfect outcome. The argument of the previous paragraph did not show that confessing was the unique Nash outcome. To show subgame perfectness, we worked back from the end using longer and longer subgames. To show that confessing is the only Nash outcome, we do not look at subgames, but instead rule out successive classes of strategies from being Nash. Consider the portions of the strategy which apply to the equilibrium path (that is, the portions directly relevant to the payoffs). No strategy in the class that calls for *Deny* in the last period can be a Nash strategy, because the same strategy with *Confess* replacing *Deny* would dominate it. But if both players have strategies calling for confessing in the last period, then no strategy that does not call for confessing in the next-to-last period is Nash, because a player should deviate by replacing *Deny* with *Confess* in the next-to-last period. The argument can be carried back to the first period, ruling out any class of strategies that does not call for confessing everywhere along the equilibrium path.

The strategy of always confessing is not a dominant strategy, as it is in the one-

shot game, because it is not the best response to various suboptimal strategies such as (*Deny until the other player Confesses, then Deny for the rest of the game*). Moreover, the uniqueness is only on the equilibrium path. Nonperfect Nash strategies could call for cooperation at nodes far away from the equilibrium path, since that action would never have to be taken. If Row has chosen (*Always Confess*), one of Column's best responses is (*Always Confess unless Row has chosen Deny ten times; then always Deny*).

5.2 Infinitely Repeated Games, Minimax Punishments, and the Folk Theorem

The contradiction between the Chainstore Paradox and what many people think of as real world behavior has been most successfully resolved by adding incomplete information to the model, as will be seen in Section 6.4. Before we turn to incomplete information, however, we will explore certain other modifications. One idea is to repeat the Prisoner's Dilemma an infinite number of times instead of a finite number (after all, few economies have a known end date). Without a last period, the inductive argument in the Chainstore Paradox fails.

In fact, we can find a simple perfect equilibrium for the infinitely repeated Prisoner's Dilemma in which both players cooperate—a game in which both players adopt the Grim Strategy.

Grim Strategy

1 *Start by choosing Deny.*

2 *Continue to choose Deny unless some player has chosen Confess, in which case choose Confess forever.*

Notice that the Grim Strategy says that even if a player is the first to deviate and choose *Confess*, he continues to choose *Confess* thereafter.

If Column uses the Grim Strategy, the Grim Strategy is weakly Row's best response. If Row cooperates, he will continue to receive the high (*Deny, Deny*) payoff forever. If he confesses, he will receive the higher (*Confess, Deny*) payoff once, but the best he can hope for thereafter is the (*Confess, Confess*) payoff.

Even in the infinitely repeated game, cooperation is not immediate, and not every strategy that punishes confessing is perfect. A notable example is the strategy of Tit-for-Tat.

Tit-for-Tat

1 *Start by choosing Deny.*

2 *Thereafter, in period n choose the action that the other player chose in period $(n - 1)$.*

If Column uses Tit-for-Tat, Row does not have an incentive to *Confess* first, because if Row cooperates he will continue to receive the high (*Deny, Deny*) payoff, but if

he confesses and then returns to Tit-for-Tat, the players alternate *(Confess, Deny)* with *(Deny, Confess)* forever. Row's average payoff from this alternation would be lower than if he had stuck to *(Deny, Deny)*, and would swamp the one-time gain. But Tit-for-Tat is almost never perfect in the infinitely repeated Prisoner's Dilemma without discounting, because it is not rational for Column to punish Row's initial *Confess*. Adhering to Tit-for-Tat's punishments results in a miserable alternation of *Confess* and *Deny*, so Column would rather ignore Row's first *Confess*. The deviation is not from the equilibrium path action of *Deny*, but from the off-equilibrium action rule of *Confess in response to a Confess*. Tit-for-Tat, unlike the Grim Strategy, cannot enforce cooperation.¹

Unfortunately, although eternal cooperation is a perfect equilibrium outcome in the infinite game under at least one strategy, so is practically anything else, including eternal confessing. The multiplicity of equilibria is summarized by the Folk Theorem, so called because its origins are hazy.

Theorem 1 (the Folk Theorem)

In an infinitely repeated n -person game with finite action sets at each repetition, any combination of actions observed in any finite number of repetitions is the unique outcome of some subgame perfect equilibrium given

Condition 1: *The rate of time preference is zero, or positive and sufficiently small;*

Condition 2: *The probability that the game ends at any repetition is zero, or positive and sufficiently small; and*

Condition 3: *The set of payoff combinations that strictly Pareto dominate the minimax payoff combinations in the mixed extension of the one-shot game is n -dimensional.*

What the Folk Theorem tells us is that claiming that particular behavior arises in a perfect equilibrium is meaningless in an infinitely repeated game. This applies to any game that meets conditions 1 to 3, not just to the Prisoner's Dilemma. If an infinite amount of time always remains in the game, a way can always be found to make one player willing to punish some other player for the sake of a better future, even if the punishment currently hurts the punisher as well as the punished. Any finite interval of time is insignificant compared to eternity, so the threat of future reprisal makes the players willing to carry out the punishments needed.

We will next discuss conditions 1 to 3.

Condition 1: Discounting

The Folk Theorem helps answer the question of whether discounting future payments lessens the influence of the troublesome Last Period. Quite to the contrary, with discounting, the present gain from confessing is weighted more heavily and future gains from cooperation more lightly. If the discount rate is very high the game almost returns to being one-shot. When the real interest rate is 1,000 percent, a payment next year is little better than a payment a hundred years hence, so next year is practically irrelevant. Any model that relies on a large number of repetitions also assumes that the discount rate is not too high.

Allowing a little discounting is none the less important to show there is no discontinuity

¹See Kalai, Samet & Stanford (1988) and Problem 5.5 for elaboration of this point.

at the discount rate of zero. If we come across an undiscounted, infinitely repeated game with many equilibria, the Folk Theorem tells us that adding a low discount rate will not reduce the number of equilibria. This contrasts with the effect of changing the model by having a large but finite number of repetitions, a change which often eliminates all but one outcome by inducing the Chainstore Paradox.

A discount rate of zero supports many perfect equilibria, but if the rate is high enough, the only equilibrium outcome is eternal confessing. We can calculate the critical value for given parameters. The Grim Strategy imposes the heaviest possible punishment for deviant behavior. Using the payoffs for the Prisoner's Dilemma from Table 2a in the next section, the equilibrium payoff from the Grim Strategy is the current payoff of 5 plus the value of the rest of the game, which from Table 2 of Chapter 4 is $\frac{5}{r}$. If Row deviated by confessing, he would receive a current payoff of 10, but the value of the rest of the game would fall to 0. The critical value of the discount rate is found by solving the equation $5 + \frac{5}{r} = 10 + 0$, which yields $r = 1$, a discount rate of 100 percent or a discount factor of $\delta = 0.5$. Unless the players are extremely impatient, confessing is not much of a temptation.

Condition 2: A probability of the game ending

Time preference is fairly straightforward, but what is surprising is that assuming that the game ends in each period with probability θ does not make a drastic difference. In fact, we could even allow θ to vary over time, so long as it never became too large. If $\theta > 0$, the game ends in finite time with probability one; or, put less dramatically, the expected number of repetitions is finite, but it still behaves like a discounted infinite game, because the expected number of future repetitions is always large, no matter how many have already occurred. The game still has no Last Period, and it is still true that imposing one, no matter how far beyond the expected number of repetitions, would radically change the results.

The following two situations are different from each other.

“1 The game will end at some uncertain date before T .”

“2 There is a constant probability of the game ending.”

In situation (1), the game is like a finite game, because, as time passes, the maximum length of time still to run shrinks to zero. In situation (2), even if the game will end by T with high probability, if it actually lasts until T the game looks exactly the same as at time zero. The fourth verse from the hymn “Amazing grace” puts this stationarity very nicely (though I expect it is supposed to apply to a game with $\theta = 0$).

*When we've been there ten thousand years,
Bright shining as the sun,
We've no less days to sing God's praise
Than when we'd first begun.*

Condition 3: Dimensionality

The “minimax payoff” mentioned in theorem 5.1 is the payoff that results if all the other players pick strategies solely to punish player i , and he protects himself as best he can.


The set of strategies s_{-i}^{i*} is a set of $(n - 1)$ **minimax strategies** chosen by all the players except i to keep i 's payoff as low as possible, no matter how he responds. s_{-i}^{i*} solves

$$\underset{s_{-i}}{\text{Minimize}} \quad \underset{s_i}{\text{Maximum}} \quad \pi_i(s_i, s_{-i}). \quad (1)$$

Player i 's **minimax payoff**, **minimax value**, or **security value** is his payoff from the solution of (1).

The dimensionality condition is needed only for games with three or more players. It is satisfied if there is some payoff combination for each player in which his payoff is greater than his minimax payoff but still different from the payoff of every other player. Figure 1 shows how this condition is satisfied for the two-person Prisoner's Dilemma of Table 2a a few pages beyond this paragraph, but not for the two-person Ranked Coordination game. It is also satisfied by the n -person Prisoner's Dilemma in which a solitary confessor gets a higher payoff than his cooperating fellow-prisoners, but not by the n -person Ranked Coordination game, in which all the players have the same payoff. The condition is necessary because establishing the desired behavior requires some way for the other players to punish a deviator without punishing themselves.

Figure 1: The Dimensionality Condition



An alternative to the dimensionality condition in the Folk Theorem is

Condition 3': The repeated game has a “desirable” subgame-perfect equilibrium in which the strategy combination $\bar{s}$ played each period gives player i a payoff that exceeds his payoff from some other “punishment” subgame-perfect equilibrium in which the strategy combination $\underline{s}^i$ is played each period:

$$\exists \bar{s} : \forall i, \exists \underline{s}^i : \pi_i(\underline{s}^i) < \pi_i(\bar{s}).$$

Condition 3' is useful because sometimes it is easy to find a few perfect equilibria. To enforce the desired pattern of behavior, use the “desirable” equilibrium as a carrot and the “punishment” equilibrium as a self-enforcing stick (see Rasmusen [1992a]).

Minimax and Maximin

In discussions of strategies which enforce cooperation, the question of deciding on the maximum severity of punishment strategies frequently arises. The idea of the minimax strategy is useful for this in that the minimax strategy is defined as the most severe sanction possible if the offender does not cooperate in his own punishment. The corresponding strategy for the offender, trying to protect himself from punishment, is the maximin strategy:

*The strategy s_i^* is a **maximin strategy** for player i if, given that the other players pick strategies to make i 's payoff as low as possible, s_i^* gives i the highest possible payoff. In our notation, s_i^* solves*

$$\underset{s_i}{\text{Maximize}} \quad \underset{s_{-i}}{\text{Minimum}} \quad \pi_i(s_i, s_{-i}). \quad (2)$$

The following formulae show how to calculate the minimax and maximin strategies for a two-player game with Player 1 as i .

$$\begin{array}{llll} \text{Maximin:} & \underset{s_1}{\text{Maximum}} & \underset{s_2}{\text{Minimum}} & \pi_1. \\ \text{Minimax:} & \underset{s_2}{\text{Minimum}} & \underset{s_1}{\text{Maximum}} & \pi_1. \end{array}$$

In the Prisoner's Dilemma, the minimax and maximin strategies are both *Confess*. Although the Welfare Game (table 3.1) has only a mixed strategy Nash equilibrium, if we restrict ourselves to the pure strategies² the Pauper's maximin strategy is *Try to Work*, which guarantees him at least 1, and his strategy for minimaxing the Government is *Be Idle*, which prevents the Government from getting more than zero.

Under minimax, Player 2 is purely malicious but must move first (at least in choosing a mixing probability) in his attempt to cause player 1 the maximum pain. Under maximin, Player 1 moves first, in the belief that Player 2 is out to get him. In variable-sum games, minimax is for sadists and maximin for paranoids. In zero-sum games, the players are merely neurotic: minimax is for optimists and maximin for pessimists.

The maximin strategy need not be unique, and it can be in mixed strategies. Since maximin behavior can also be viewed as minimizing the maximum loss that might be suffered, decision theorists refer to such a policy as a **minimax criterion**, a catchier phrase (Luce & Raiffa [1957], p. 279).

It is tempting to use maximin strategies as the basis of an equilibrium concept. A **maximin equilibrium** is made up of a maximin strategy for each player. Such a strategy

²xxx what if we don't?

might seem reasonable because each player then has protected himself from the worst harm possible. Maximin strategies have very little justification, however, for a rational player. They are not simply the optimal strategies for risk-averse players, because risk aversion is accounted for in the utility payoffs. The players' implicit beliefs can be inconsistent in a maximin equilibrium, and a player must believe that his opponent would choose the most harmful strategy out of spite rather than self-interest if maximin behavior is to be rational.

The usefulness of minimax and maximin strategies is not in directly predicting the best strategies of the players, but in setting the bounds of how their strategies affect their payoffs, as in condition 3 of Theorem 1.

It is important to remember that minimax and maximin strategies are not always pure strategies. In the Minimax Illustration Game of Table 1, which I take from page 150 of Row can guarantee himself a payoff of 0 by choosing *Down*, so that is his maximin strategy. Column cannot hold Row's payoff down to 0, however, by using a pure minimax strategy. If Column chooses *Left*, Row can choose *Middle* and get a payoff of 1; if Column chooses *Right*, Row can choose *Up* and get a payoff of 1. Column can, however, hold Row's payoff down to 0 by choosing a mixed minimax strategy of (*Probability 0.5 of Left, Probability 0.5 of Right*). Row would then respond with *Down*, for a minimax payoff of 0, since either *Up*, *Middle*, or a mixture of the two would give him a payoff of $-0.5 (= 0.5(-2) + 0.5(1))$.³

Table 1 The Minimax Illustration Game

		Column	
		<i>Left</i>	<i>Right</i>
Row:	<i>Up</i>	$-2, \boxed{2}$	$\boxed{1}, -2$
	<i>Middle</i>	$\boxed{1}, -2$	$-2, \boxed{2}$
	<i>Down</i>	$0, \boxed{1}$	$0, \boxed{1}$
<i>Payoffs to: (Row, Column).</i>			

In two-person zero-sum games, minimax and maximin strategies are more directly useful, because when player 1 reduces player 2's payoff, he increases his own payoff. Punishing the other player is equivalent to rewarding yourself. This is the origin of the celebrated **Minimax Theorem** (von Neumann [1928]), which says that a minimax equilibrium exists in pure or mixed strategies for every two-person zero-sum game and is identical to the maximin equilibrium. Unfortunately, the games that come up in applications are almost never zero-sum games, so the Minimax Theorem is of limited applicability.

Precommitment

What if we use metastrategies, abandoning the idea of perfectness by allowing players to commit at the start to a strategy for the rest of the game? We would still want to keep the game noncooperative by disallowing binding promises, but we could model it as a game with simultaneous choices by both players, or with one move each in sequence.

³Column's maximin and minimax strategies can also be computed. The strategy for minimizing Column is (*Probability 0.5 of Up, Probability 0.5 of Middle*), his maximin strategy is (*Probability 0.5 of Left, Probability 0.5 of Right*), and his minimax payoff is 0.

If precommitted strategies are chosen simultaneously, the equilibrium outcome of the finitely repeated Prisoner's Dilemma calls for always confessing, because allowing commitment is the same as allowing equilibria to be nonperfect, in which case, as was shown earlier, the unique Nash outcome is always confessing.

A different result is achieved if the players precommit to strategies in sequence. The outcome depends on the particular values of the parameters, but one possible equilibrium is the following: Row moves first and chooses the strategy (*Deny* until Column *Confess*; thereafter always *Confess*), and Column chooses (*Deny* until the last period; then *Confess*). The observed outcome would be for both players to deny until the last period, and then for Row to again deny, but for Column to confess. Row would submit to this because if he chose a strategy that initiated confessing earlier, Column would choose a strategy of starting to confess earlier too. The game has a second-mover advantage.

5.3 Reputation: The One-sided Prisoner's Dilemma

Part II of this book will analyze moral hazard and adverse selection. Under moral hazard, a player wants to commit to high effort, but he cannot credibly do so. Under adverse selection, a player wants to prove he is high ability, but he cannot. In both, the problem is that the penalties for lying are insufficient. Reputation seems to offer a way out of the problem. If the relationship is repeated, perhaps a player is willing to be honest in early periods in order to establish a reputation for honesty which will be valuable to himself later.

Reputation seems to play a similar role in making threats to punish credible. Usually punishment is costly to the punisher as well as the punished, and it is not clear why the punisher should not let bygones be bygones. Yet in 1988 the Soviet Union paid off 70-year-old debt to dissuade the Swiss authorities from blocking a mutually beneficial new bond issue ("Soviets Agree to Pay Off Czarist Debt to Switzerland," *Wall Street Journal*, January 19, 1988, p. 60). Why were the Swiss so vindictive towards Lenin?

The questions of why players do punish and do not cheat are really the same questions that arise in the repeated Prisoner's Dilemma, where the fact of an infinite number of repetitions allows cooperation. That is the great problem of reputation. Since everyone knows that a player will *Confess*, choose low effort, or default on debt in the last period, why do they suppose he will bother to build up a reputation in the present? Why should past behavior be any guide to future behavior?

Not all reputation problems are quite the same as the Prisoner's Dilemma, but they have much the same flavor. Some games, like duopoly or the original Prisoner's Dilemma, are **two-sided** in the sense that each player has the same strategy set and the payoffs are symmetric. Others, such as the game of Product Quality (see below), are what we might call **one-sided Prisoner's Dilemmas**, which have properties similar to the Prisoner's Dilemma, but do not fit the usual definition because they are asymmetric. Table 2 shows the normal forms for both the original Prisoner's Dilemma and the one-sided version.⁴ The important difference is that in the one-sided Prisoner's Dilemma at least one player

⁴The exact numbers are different from the Prisoner's Dilemma in Table 1 in Chapter 1, but the ordinal rankings are the same. Numbers such as those in Table 2 of the present chapter are more commonly used,

really does prefer the outcome equivalent to $(Deny, Deny)$, which is $(High\ Quality, Buy)$ in Table 2b, to anything else. He confesses defensively, rather than both offensively and defensively. The payoff $(0,0)$ can often be interpreted as the refusal of one player to interact with the other, for example, the motorist who refuses to buy cars from Chrysler because he knows they once falsified odometers. Table 3 lists examples of both one-sided and two-sided games. Versions of the Prisoner's Dilemma with three or more players can also be classified as one-sided or two-sided, depending on whether or not all players find *Confess* a dominant strategy.

Table 2 Prisoner's Dilemmas

(a) Two-Sided (conventional)

		Column	
		<i>Deny</i>	<i>Confess</i>
Row:	<i>Deny</i>	5,5	→ -5,10
	<i>Confess</i>	↓ 10,-5	↓ 0,0
<i>Payoffs to: (Row, Column)</i>			

(b) One-Sided

		Consumer (Column)	
		<i>Buy</i>	<i>Boycott</i>
Seller (Row):	<i>High Quality</i>	5,5 ←	0,0
	<i>Low Quality</i>	↓ 10, -5 →	↓ 0,0
<i>Payoffs to: (Seller, Consumer)</i>			

Table 3 Some repeated games in which reputation is important

because it is convenient to normalize the $(Confess, Confess)$ payoffs to $(0,0)$ and to make most of the numbers positive rather than negative.

Application	Sidedness	Players	Actions
Prisoner's Dilemma	two-sided	Row Column	<i>Deny/Confess</i> <i>Deny/Confess</i>
Duopoly	two-sided	Firm Firm	<i>High price/Low price</i> <i>High price/Low price</i>
Employment	two-sided	Employer Employee	<i>Bonus/No bonus</i> <i>Work/Shirk</i>
Product Quality	one-sided	Consumer Seller	<i>Buy/Boycott</i> <i>High quality/low quality</i>
Entry Deterrence	one-sided	Incumbent Entrant	<i>Low price/High price</i> <i>Enter/Stay out</i>
Financial Disclosure	one-sided	Corporation Investor	<i>Truth/Lies</i> <i>Invest/Refrain</i>
Borrowing	one-sided	Lender Borrower	<i>Lend/Refuse</i> <i>Repay/Default</i>

The Nash and iterated dominance equilibria in the one-sided Prisoner's Dilemma are still (*Confess*, *Confess*), but it is not a dominant-strategy equilibrium. Column does not have a dominant strategy, because if Row were to choose *Deny*, Column would also choose *Deny*, to obtain the payoff of 5; but if Row chooses *Confess*, Column would choose *Confess*, for a payoff of zero. *Confess* is however, weakly dominant for Row, which makes (*Confess*, *Confess*) the iterated dominant strategy equilibrium. In both games, the players would like to persuade each other that they will cooperate, and devices that induce cooperation in the one-sided game will usually obtain the same result in the two-sided game.

5.4 Product Quality in an Infinitely Repeated Game

The Folk Theorem tells us that some perfect equilibrium of an infinitely repeated game—sometimes called an **infinite horizon model**—can generate any pattern of behavior observed over a finite number of periods. But since the Folk Theorem is no more than a mathematical result, the strategies that generate particular patterns of behavior may be unreasonable. The theorem's value is in provoking close scrutiny of infinite horizon models so that the modeller must show why his equilibrium is better than a host of others. He must go beyond satisfaction of the technical criterion of perfectness and justify the strategies on other grounds.

In the simplest model of product quality, a seller can choose between producing costly high quality or costless low quality, and the buyer cannot determine quality before he

purchases. If the seller would produce high quality under symmetric information, we have a one-sided Prisoner's Dilemma, as in Table 2b. Both players are better off when the seller produces high quality and the buyer purchases the product, but the seller's weakly dominant strategy is to produce low quality, so the buyer will not purchase. This is also an example of moral hazard, the topic of chapter 7.

A potential solution is to repeat the game, allowing the firm to choose quality at each repetition. If the number of repetitions is finite, however, the outcome stays the same because of the Chainstore Paradox. In the last repetition, the subgame is identical to the one-shot game, so the firm chooses low quality. In the next-to-last repetition, it is foreseen that the last period's outcome is independent of current actions, so the firm also chooses low quality, an argument that can be carried back to the first repetition.

If the game is repeated an infinite number of times, the Chainstore Paradox is inapplicable and the Folk Theorem says that a wide range of outcomes can be observed in equilibrium. Klein & Leffler (1981) construct a plausible equilibrium for an infinite period model. Their original article, in the traditional verbal style of UCLA, does not phrase the result in terms of game theory, but we will recast it here, as I did in Rasmusen (1989b). In equilibrium, the firm is willing to produce a high quality product because it can sell at a high price for many periods, but consumers refuse to ever buy again from a firm that has once produced low quality. The equilibrium price is high enough that the firm is unwilling to sacrifice its future profits for a one-time windfall from deceitfully producing low quality and selling it at a high price. Although this is only one of a large number of subgame perfect equilibria, the consumers' behavior is simple and rational: no consumer can benefit by deviating from the equilibrium.

Product Quality

Players

An infinite number of potential firms and a continuum of consumers.

The Order of Play

- 1 An endogenous number n of firms decide to enter the market at cost F .
- 2 A firm that has entered chooses its quality to be *High* or *Low*, incurring the constant marginal cost c if it picks *High* and zero if it picks *Low*. The choice is unobserved by consumers. The firm also picks a price p .
- 3 Consumers decide which firms (if any) to buy from, choosing firms randomly if they are indifferent. The amount bought from firm i is denoted q_i .
- 4 All consumers observe the quality of all goods purchased in that period.
- 5 The game returns to (2) and repeats.

Payoffs

The consumer benefit from a product of low quality is zero, but consumers are willing to buy quantity $q(p) = \sum_{i=1}^n q_i$ for a product believed to be high quality, where $\frac{dq}{dp} < 0$.

If a firm stays out of the market, its payoff is zero.

If firm i enters, it receives $-F$ immediately. Its current end-of-period payoff is $q_i p$ if it produces *Low* quality and $q_i(p - c)$ if it produces *High* quality. The discount rate is $r \geq 0$.

That the firm can produce low quality items at zero marginal cost is unrealistic, but it is only a simplifying assumption. By normalizing the cost of producing low quality to zero, we avoid having to carry an extra variable through the analysis without affecting the result.

The Folk Theorem tells us that this game has a wide range of perfect outcomes, including a large number with erratic quality patterns like (*High, High, Low, High, Low, Low...*). If we confine ourselves to pure-strategy equilibria with the stationary outcome of constant quality and identical behavior by all firms in the market, then the two outcomes are low quality and high quality. Low quality is always an equilibrium outcome, since it is an equilibrium of the one-shot game. If the discount rate is low enough, high quality is also an equilibrium outcome, and this will be the focus of our attention. Consider the following strategy combination:

Firms. $\tilde{n}$ firms enter. Each produces high quality and sells at price $\tilde{p}$. If a firm ever deviates from this, it thereafter produces low quality (and sells at the same price $\tilde{p}$). The values of $\tilde{p}$ and $\tilde{n}$ are given by equations (4) and (8) below.

Buyers. Buyers start by choosing randomly among the firms charging $\tilde{p}$. Thereafter, they remain with their initial firm unless it changes its price or quality, in which case they switch randomly to a firm that has not changed its price or quality.

This strategy combination is a perfect equilibrium. Each firm is willing to produce high quality and refrain from price-cutting because otherwise it would lose all its customers. If it has deviated, it is willing to produce low quality because the quality is unimportant, given the absence of customers. Buyers stay away from a firm that has produced low quality because they know it will continue to do so, and they stay away from a firm that has cut the price because they know it will produce low quality. For this story to work, however, the equilibrium must satisfy three constraints that will be explained in more depth in Section 7.3: incentive compatibility, competition, and market clearing.

The **incentive compatibility** constraint says that the individual firm must be willing to produce high quality. Given the buyers' strategy, if the firm ever produces low quality it receives a one-time windfall profit, but loses its future profits. The tradeoff is represented by constraint (3), which is satisfied if the discount rate is low enough.

$$\frac{q_i p}{1+r} \leq \frac{q_i(p-c)}{r} \quad (\text{incentive compatibility}). \quad (3)$$

Inequality (3) determines a lower bound for the price, which must satisfy

$$\tilde{p} \geq (1+r)c. \quad (4)$$

Condition (4) will be satisfied as an equality, because any firm trying to charge a price higher than the quality-guaranteeing $\tilde{p}$ would lose all its customers.

The second constraint is that competition drives profits to zero, so firms are indifferent between entering and staying out of the market.

$$\frac{q_i(p-c)}{r} = F. \quad (\text{competition}) \quad (5)$$

Treating (3) as an equation and using it to replace p in equation (5) gives

$$q_i = \frac{F}{c}. \quad (6)$$

We have now determined p and q_i , and only n remains, which is determined by the equality of supply and demand. The market does not always clear in models of asymmetric information (see Stiglitz [1987]), and in this model each firm would like to sell more than its equilibrium output at the equilibrium price, but the market output must equal the quantity demanded by the market.

$$nq_i = q(p). \quad (\text{market clearing}) \quad (7)$$

Combining equations (3), (6), and (7) yields

$$\tilde{n} = \frac{cq([1+r]c)}{F}. \quad (8)$$

We have now determined the equilibrium values, the only difficulty being the standard existence problem caused by the requirement that the number of firms be an integer (see note N5.4).

The equilibrium price is fixed because F is exogenous and demand is not perfectly inelastic, which pins down the size of firms. If there were no entry cost, but demand were still elastic, then the equilibrium price would still be the unique p that satisfied constraint (3), and the market quantity would be determined by $q(p)$, but F and q_i would be undetermined. If consumers believed that any firm which might possibly produce high quality paid an exogenous dissipation cost F , the result would be a continuum of equilibria. The firms' best response would be for $\tilde{n}$ of them to pay F and produce high quality at price $\tilde{p}$, where $\tilde{n}$ is determined by the zero profit condition as a function of F . Klein & Leffler note this indeterminacy and suggest that the profits might be dissipated by some sort of brand-specific capital. The history of the industry may also explain the number of firms. Schmalensee (1982) shows how a pioneering brand can retain a large market share because consumers are unwilling to investigate the quality of new brands.

This idea can be extended to labor contracts also, and is very similar to the notion of the "efficiency wage" in Section 8.1.

***5.5 Markov Equilibria and Overlapping Generations in the Game of Customer Switching Costs**

The next model demonstrates a general modelling technique, the **overlapping generations model**, in which different cohorts of otherwise identical players enter and leave the game with overlapping "lifetimes," and a new equilibrium concept, "Markov equilibrium." The best-known example of an overlapping-generations model is the original consumption-loans model of Samuelson (1958). The models are most often used in macroeconomics, but they can also be useful in microeconomics. Klemperer (1987) has stimulated considerable interest in customers who incur costs in moving from one seller to another. The model used here will be that of Farrell & C. Shapiro (1988).

Customer Switching Costs

Players

Firms Apex and Brydax, and a series of customers, each of whom is first called a youngster and then an oldster.

The Order of Play

- 1a Brydax, the initial incumbent, picks the incumbent price p_1^i .
- 1b Apex, the initial entrant, picks the entrant price p_1^e .
- 1c The oldster picks a firm.
- 1d The youngster picks a firm.
- 1e Whichever firm attracted the youngster becomes the incumbent.
- 1f The oldster dies and the youngster becomes an oldster.
- 2a Return to (1a), possibly with new identities for entrant and incumbent.

Payoffs

The discount factor is δ . The customer reservation price is R and the switching cost is c . The per period payoffs in period t are, for $j = (i, e)$,

$$\pi_{firm\ j} = \begin{cases} 0 & \text{if no customers are attracted.} \\ p_t^j & \text{if just oldsters or just youngsters are attracted.} \\ 2p_t^j & \text{if both oldsters and youngsters are attracted.} \end{cases}$$

$$\pi_{oldster} = \begin{cases} R - p_t^i & \text{if he buys from the incumbent.} \\ R - p_t^e - c & \text{if he switches to the entrant.} \end{cases}$$

$$\pi_{youngster} = \begin{cases} R - p_t^i & \text{if he buys from the incumbent.} \\ R - p_t^e & \text{if he buys from the entrant.} \end{cases}$$

Finding all the perfect equilibria of an infinite game like this one is difficult, so we will follow Farrell and Shapiro in limiting ourselves to the much easier task of finding the perfect Markov equilibrium, which is unique.

*A **Markov strategy** is a strategy that, at each node, chooses the action independently of the history of the game except for the immediately preceding action (or actions, if they were simultaneous).*

Here, a firm's Markov strategy is its price as a function of whether the particular is the incumbent or the entrant, and not a function of the entire past history of the game.

There are two ways to use Markov strategies: (1) just look for equilibria that use Markov strategies, and (2) disallow nonMarkov strategies and then look for equilibria. Because the first way does not disallow non-Markov strategies, the equilibrium must be such that no player wants to deviate by using any other strategy, whether Markov or not. This is just a way of eliminating possible multiple equilibria by discarding ones that use non-Markov strategies. The second way is much more dubious, because it requires the players not to use non-Markov strategies, even if they are best responses. A **perfect Markov equilibrium** uses the first approach: it is a perfect equilibrium that happens to use only Markov strategies.

Brydoux, the initial incumbent, moves first and chooses p^i low enough that Apex is not tempted to choose $p^e < p^i - c$ and steal away the oldsters. Apex's profit is p^i if it chooses $p^e = p^i$ and serves just youngsters, and $2(p^i - c)$ if it chooses $p^e = p^i - c$ and serves both oldsters and youngsters. Brydoux chooses p^i to make Apex indifferent between these alternatives, so

$$p^i = 2(p^i - c), \quad (9)$$

and

$$p^i = p^e = 2c. \quad (10)$$

In equilibrium, Apex and Brydoux take turns being the incumbent and charge the same price.

Because the game lasts forever and the equilibrium strategies are Markov, we can use a trick from dynamic programming to calculate the payoffs from being the entrant versus being the incumbent. The equilibrium payoff of the current entrant is the immediate payment of p^e plus the discounted value of being the incumbent in the next period:

$$\pi_e^* = p^e + \delta \pi_i^*. \quad (11)$$

The incumbent's payoff can be similarly stated as the immediate payment of p^i plus the discounted value of being the entrant next period:

$$\pi_i^* = p^i + \delta \pi_e^*. \quad (12)$$

We could use equation (10) to substitute for p^e and p^i , which would leave us with the two equations (11) and (12) for the two unknowns π_i^* and π_e^* , but an easier way to compute the payoff is to realize that in equilibrium the incumbent and the entrant sell the same amount at the same price, so $\pi_i^* = \pi_e^*$ and equation (12) becomes

$$\pi_i^* = 2c + \delta \pi_i^*. \quad (13)$$

It follows that

$$\pi_i^* = \pi_e^* = \frac{2c}{1 - \delta}. \quad (14)$$

Prices and total payoffs are increasing in the switching cost c , because that is what gives the incumbent market power and prevents ordinary competition of the ordinary Bertrand kind to be analyzed in section 13.2. The total payoffs are increasing in δ for the usual reason that future payments increase in value as δ approaches one.

*5.6 Evolutionary Equilibrium: Hawk-Dove

For most of this book we have been using the Nash equilibrium concept or refinements of it based on information and sequentiality, but in biology such concepts are often inappropriate. The lower animals are less likely than humans to think about the strategies of their opponents at each stage of a game. Their strategies are more likely to be preprogrammed and their strategy sets more restricted than the businessman's, if perhaps not more so than his customer's. In addition, behavior evolves, and any equilibrium must take account of the possibility of odd behavior caused by the occasional mutation. That the equilibrium is common knowledge, or that players cannot precommit to strategies, are not compelling

assumptions. Thus, the ideas of Nash equilibrium and sequential rationality are much less useful than when game theory is modelling rational players.

Game theory has grown to some importance in biology, but the style is different than in economics. The goal is not to explain how players would rationally pick actions in a given situation, but to explain how behavior evolves or persists over time under exogenous shocks. Both approaches end up defining equilibria to be strategy combinations that are best responses in some sense, but biologists care much more about the stability of the equilibrium and how strategies interact over time. In section 3.5, we touched briefly on the stability of the Cournot equilibrium, but economists view stability as a pleasing by-product of the equilibrium rather than its justification. For biologists, stability is the point of the analysis.

Consider a game with identical players who engage in pairwise contests. In this special context, it is useful to think of an equilibrium as a strategy combination such that no player with a new strategy can enter the environment (**invade**) and receive a higher expected payoff than the old players. Moreover, the invading strategy should continue to do well even if it plays itself with finite probability, or its invasion could never grow to significance. In the commonest model in biology, all the players adopt the same strategy in equilibrium, called an evolutionarily stable strategy. John Maynard Smith originated this idea, which is somewhat confusing because it really aims at an equilibrium concept, which involves a strategy combination, not just one player's strategy. For games with pairwise interactions and identical players, however, the evolutionarily stable strategy can be used to define an equilibrium concept.

*A strategy s^* is an **evolutionarily stable strategy**, or **ESS**, if, using the notation $\pi(s_i, s_{-i})$ for player i 's payoff when his opponent uses strategy s_{-i} , for every other strategy s' either*

$$\pi(s^*, s^*) > \pi(s', s^*) \tag{15}$$

or

$$\begin{aligned} (a) \quad & \pi(s^*, s^*) = \pi(s', s^*) \\ \text{and} \\ (b) \quad & \pi(s^*, s') > \pi(s', s'). \end{aligned} \tag{16}$$

If condition (15) holds, then a population of players using s^* cannot be invaded by a deviant using s' . If condition (16) holds, then s' does well against s^* , but badly against itself, so that if more than one player tried to use s' to invade a population using s^* , the invaders would fail.

We can interpret ESS in terms of Nash equilibrium. Condition (15) says that s^* is a strong Nash equilibrium (although not every strong Nash strategy is an ESS). Condition (16) says that if s^* is only a weak Nash strategy, the weak alternative s' is not a best response to itself. ESS is a refinement of Nash, narrowed by the requirement that ESS not only be a best response, but that (a) it have the highest payoff of any strategy used in equilibrium (which rules out equilibria with asymmetric payoffs), and (b) it be a strictly best response to itself.

The motivations behind the two equilibrium concepts are quite different, but the sim-

ilarities are useful because even if the modeller prefers ESS to Nash, he can start with the Nash strategies in his efforts to find an ESS.

As an example of (a), consider the Battle of the Sexes. In it, the mixed strategy equilibrium is an ESS, because a player using it has as high a payoff as any other player. The two pure strategy equilibria are not made up of ESS's, though, because in each of them one player's payoff is higher than the other's. Compare with Ranked Coordination, in which the two pure strategy equilibria and the mixed strategy equilibrium are all made up of ESS's. (The dominated equilibrium strategy is nonetheless an ESS, because given that the other players are using it, no player could do as well by deviating.)

As an example of (b), consider the Utopian Exchange Economy game in Table 4, adapted from problem 7.5 of Gintis (forthcoming). In Utopia, each citizen can produce either one or two units of individualized output. He will then go into the marketplace and meet another citizen. If either of them produced only one unit, trade cannot increase their payoffs. If both of them produced two, however, they can trade one unit for one unit, and both end up happier with their increased variety of consumption.

Table 4 The Utopian Exchange Economy Game

		Jones	
		<i>Low Output</i>	<i>High Output</i>
Smith:	<i>Low Output</i>	1, 1	$\leftrightarrow$ 1, 1
	<i>High Output</i>	1, 1	$\rightarrow$ 2, 2

Payoffs to: (Smith, Jones).

This game has three Nash equilibria, one of which is in mixed strategies. Since all strategies but *High Output* are weakly dominated, that alone is an ESS. *Low Output* fails to meet condition (16b), because it is not the strictly best response to itself. If the economy began with all citizens choosing *Low Output*, then if Smith deviated to *High Output* he would not do any better, but if *two* people deviated to *High Output*, they would do better in expectation because they might meet each other and receive the payoff of (2,2).

An Example of ESS: Hawk-Dove

The best-known illustration of the ESS is the game of Hawk-Dove. Imagine that we have a population of birds, each of whom can behave as an aggressive Hawk or a pacific Dove. We will focus on two randomly chosen birds, Bird One and Bird Two. Each bird has a choice of what behavior to choose on meeting another bird. A resource worth $V = 2$ "fitness units" is at stake when the two birds meet. If they both fight, the loser incurs a cost of $C = 4$, which means that the expected payoff when two Hawks meet is -1 ($= 0.5[2] + 0.5[-4]$) for each of them. When two Doves meet, they split the resource, for a payoff of 1 apiece. When a Hawk meets a Dove, the Dove flees for a payoff of 0, leaving the Hawk with a payoff of 2. Table 5 summarizes this.

Table 5 Hawk-Dove: economics notation

		Bird Two	
		<i>Hawk</i>	<i>Dove</i>
Bird One:	<i>Hawk</i>	-1,-1 →	2,0
	<i>Dove</i>	0, 2 ←	1,1
<i>Payoffs to: (Bird One, Bird Two)</i>			

These payoffs are often depicted differently in biology games. Since the two players are identical, one can depict the payoffs by using a table showing the payoffs only of the row player. Applying this to Hawk-Dove generates Table 6.

Table 6 Hawk-Dove: Biology Notation

		Bird Two	
		<i>Hawk</i>	<i>Dove</i>
Bird One:	<i>Hawk</i>	-1	2
	<i>Dove</i>	0	1
<i>Payoffs to: (Bird One)</i>			

Hawk-Dove is Chicken with new feathers. The two games have the same ordinal ranking of payoffs, as can be seen by comparing Table 5 with Chapter 3's Table 2, and their equilibria are the same except for the mixing parameters. Hawk-Dove has no symmetric pure-strategy Nash equilibrium, and hence no pure-strategy ESS, since in the two asymmetric Nash equilibria, *Hawk* gives a bigger payoff than *Dove*, and the doves would disappear from the population. In the ESS for this game, neither hawks nor doves completely take over the environment. If the population consisted entirely of hawks, a dove could invade and obtain a one-round payoff of 0 against a hawk, compared to the -1 that a hawk obtains against itself. If the population consisted entirely of doves, a hawk could invade and obtain a one-round payoff of 2 against a dove, compared to the 1 that a dove obtains against a dove.

In the mixed-strategy ESS, the equilibrium strategy is to be a hawk with probability 0.5 and a dove with probability 0.5, which can be interpreted as a population 50 percent hawks and 50 percent doves. As in the mixed-strategy equilibria in chapter 3, the players are indifferent as to their strategies. The expected payoff from being a hawk is the $0.5(2)$ from meeting a dove plus the $0.5(-1)$ from meeting another hawk, a sum of 0.5. The expected payoff from being a dove is the $0.5(1)$ from meeting another dove plus the $0.5(0)$ from meeting a hawk, also a sum of 0.5. Moreover, the equilibrium is stable in a sense similar to the Cournot equilibrium. If 60 percent of the population were hawks, a bird would have a higher fitness level as a dove. If "higher fitness" means being able to reproduce faster, the number of doves increases and the proportion returns to 50 percent over time.

The ESS depends on the strategy sets allowed the players. If two birds can base their behavior on commonly observed random events such as which bird arrives at the resource first, and $V < C$ (as specified above), then a strategy called the **bourgeois strategy** is an ESS. Under this strategy, the bird respects property rights like a good bourgeois; it

behaves as a hawk if it arrives first, and a dove if it arrives second, where we assume the order of arrival is random. The bourgeois strategy has an expected payoff of 1 from meeting itself, and behaves exactly like a 50:50 randomizer when it meets a strategy that ignores the order of arrival, so it can successfully invade a population of 50:50 randomizers. But the bourgeois strategy is a correlated strategy (see section 3.3), and requires something like the order of arrival to decide which of two identical players will play *Hawk*.

The ESS is suited to games in which all the players are identical and interacting in pairs. It does not apply to games with non-identical players—wolves who can be wily or big and deer who can be fast or strong—although other equilibrium concepts of the same flavor can be constructed. The approach follows three steps, specifying (1) the initial population proportions and the probabilities of interactions, (2) the pairwise interactions, and (3) the dynamics by which players with higher payoffs increase in number in the population. Economics games generally use only the second step, which describes the strategies and payoffs from a single interaction.

The third step, the evolutionary dynamics, is especially foreign to economics. In specifying dynamics, the modeller must specify a difference equation (for discrete time) or differential equation (for continuous time) that describes how the strategies employed change over iterations, whether because players differ in the number of their descendants or because they learn to change their strategies over time. In economics games, the adjustment process is usually degenerate: the players jump instantly to the equilibrium. In biology games, the adjustment process is slower and cannot be derived from theory. How quickly the population of hawks increases to relative to doves depends on the metabolism of the bird and the length of a generation.

Slow dynamics also makes the starting point of the game important, unlike the case when adjustment is instantaneous. Figure 2, taken from D. Friedman (1991), shows a way to graphically depict evolution in a game in which all three strategies of *Hawk*, *Dove*, and *Bourgeois* are used. A point in the triangle represents a proportion of the three strategies in the population. At point E_3 , for example, half the birds play *Hawk*, half play *Dove*, and none play *Bourgeois*, while at E_4 all the birds play *Bourgeois*.

Figure 2: Evolutionary Dynamics in the Hawk-Dove- Bourgeois Game




Figure 2 shows the result of dynamics based on a function specified by Friedman that gives the rate of change of a strategy's proportion based on its payoff relative to the other two strategies. Points E_1 , E_2 , E_3 , and E_4 are all fixed points in the sense that the proportions do not change no matter which of these points the game starts from. Only point E_4 represents an evolutionarily stable equilibrium, however, and if the game starts with any positive proportion of birds playing *Bourgeois*, the proportions tend towards E_4 . The original Hawk-Dove which excluded the bourgeois strategy can be viewed as the HD line at the bottom of the triangle, and E_3 is evolutionarily stable in that restricted game.

Figure 2 also shows the importance of mutation in biological games. If the population of birds is 100 percent dove, as at E_2 , it stays that way in the absence of mutation, since if there are no hawks to begin with, the fact that they would reproduce at a faster rate than doves becomes irrelevant. If, however, a bird could mutate to play *Hawk* and then pass this behavior on to his offspring, then eventually some bird would do so and the mutant strategy would be successful. The technology of mutations can be important to the ultimate equilibrium. In more complicated games than Hawk-Dove, it can matter whether mutations happen to be small, accidental shifts to strategies similar to those that are currently being played, or can be of arbitrary size, so that a superior strategy quite different from the existing strategies might be reached.

The idea of mutation is distinct from the idea of evolutionary dynamics, and it is possible to use one without the other. In economics models, a mutation would correspond to the appearance of a new action in the action set of one of the players in a game. This is one way to model innovation: not as research followed by stochastic discoveries, but as accidental learning. The modeller might specify that the discovered action becomes available to players slowly through evolutionary dynamics, or instantly, in the usual style of economics. This style of research has promise for economics, but since the technologies of dynamics and mutation are important there is a danger of simply multiplying models without reliable results unless the modeller limits himself to a narrow context and bases his technology on empirical measurements.

N5.1 Finitely repeated games and the Chainstore Paradox

- The Chainstore Paradox does not apply to all games as neatly as to Entry Deterrence and the Prisoner's Dilemma. If the one-shot game has only one Nash equilibrium, the perfect equilibrium of the finitely repeated game is unique and has that same outcome. But if the one-shot game has multiple Nash equilibria, the perfect equilibrium of the finitely repeated game can have not only the one-shot outcomes, but others besides. See Benoit & Krishna (1985), Harrington (1987), and Moreaux (1985).
- John Heywood is Bartlett's source for the term "tit-for-tat," from the French "tant pour tant."
- A realistic expansion of a game's strategy space may eliminate the Chainstore Paradox. D. Hirshleifer & Rasmusen (1989), for example, show that allowing the players in a multi-person finitely repeated Prisoner's Dilemma to ostracize offenders can enforce cooperation even if there are economies of scale in the number of players who cooperate and are not ostracized.
- The peculiarity of the unique Nash equilibrium for the repeated Prisoner's Dilemma was noticed long before Selten (1978) (see Luce & Raiffa [1957] p. 99), but the term Chainstore Paradox is now generally used for all unravelling games of this kind.
- An **epsilon-equilibrium** is a strategy combination s^* such that no player has more than an ϵ incentive to deviate from his strategy given that the other players do not deviate. Formally,

$$\forall i, \pi_i(s_i^*, s_{-i}^*) \geq \pi_i(s_i', s_{-i}^*) - \epsilon, \forall s_i' \in S_i. \quad (17)$$

Radner (1980) has shown that cooperation can arise as an ϵ -equilibrium of the finitely repeated Prisoner's Dilemma. Fudenberg & Levine (1986) compare the ϵ -equilibria of finite games with the Nash equilibria of infinite games. Other concepts besides Nash can also use the ϵ -equilibrium idea.

- A general way to decide whether a mathematical result is a trick of infinity is to see if the same result is obtained as the limit of results for longer and longer finite models. Applied to games, a good criterion for picking among equilibria of an infinite game is to select one which is the limit of the equilibria for finite games as the number of periods gets longer. Fudenberg & Levine (1986) show under what conditions one can find the equilibria of infinite-horizon games by this process. For the Prisoner's Dilemma, (*Always Confess*) is the only equilibrium in all finite games, so it uniquely satisfies the criterion.
- Defining payoffs in games that last an infinite number of periods presents the problem that the total payoff is infinite for any positive payment per period. Ways to distinguish one infinite amount from another include the following.

1 Use an **overtaking criterion**. Payoff stream π is preferred to $\tilde{\pi}$ if there is some time T^* such that for every $T \geq T^*$,

$$\sum_{t=1}^T \delta^t \pi_t > \sum_{t=1}^T \delta^t \tilde{\pi}_t.$$

2 Specify that the discount rate is strictly positive, and use the present value. Since payments in distant periods count for less, the discounted value is finite unless the payments are growing faster than the discount rate.

3 Use the average payment per period, a tricky method since some sort of limit needs to be taken as the number of periods averaged goes to infinity.

Whatever the approach, game theorists assume that the payoff function is **additively separable** over time, which means that the total payoff is based on the sum or average, possibly discounted, of the one-shot payoffs. Macroeconomists worry about this assumption, which rules out, for example, a player whose payoff is very low if any of his one-shot payoffs dips below a certain subsistence level. The issue of separability will arise again in section 13.5 when we discuss durable monopoly.

- Ending in finite time with probability one means that the limit of the probability the game has ended by date t approaches one as t tends to infinity; the probability that the game lasts till infinity is zero. Equivalently, the expectation of the end date is finite, which it could not be were there a positive probability of an infinite length.

N5.2 Infinitely Repeated Games, Minimax Punishments, and the Folk Theorem

- References on the Folk Theorem include Aumann (1981), Fudenberg & Maskin (1986), Fudenberg & Tirole (1991a, pp. 152-62), and Rasmusen (1992a). The most commonly cited version of the Folk Theorem says that if conditions 1 to 3 are satisfied, then:

Any payoff combination that strictly Pareto dominates the minimax payoff combinations in the mixed extension of an n -person one-shot game with finite action sets is the average payoff in some perfect equilibrium of the infinitely repeated game.

- The evolutionary approach can also be applied to the repeated Prisoner's Dilemma. Boyd & Lorberbaum (1987) show that no pure strategy, including Tit-for-Tat, is evolutionarily stable in a population-interaction version of the Prisoner's Dilemma. J. Hirshleifer & Martinez-Coll (1988) have found that Tit-for-Tat is no longer part of an ESS in an evolutionary Prisoner's Dilemma if (1) more complicated strategies have higher computation costs; or (2) sometimes a *Deny* is observed to be a *Confess* by the other player.
- **Trigger strategies** of *trigger-price strategies* are an important kind of strategies for repeated games. Consider the oligopolist facing uncertain demand (as in Stigler [1964]). He cannot tell whether the low demand he observes facing him is due to Nature or to price cutting by his fellow oligopolists. Two things that could trigger him to cut his own price in retaliation are a series of periods with low demand or one period of especially low demand. Finding an optimal trigger strategy is a difficult problem (see Porter [1983a]). Trigger strategies are usually not subgame perfect unless the game is infinitely repeated, in which case they are a subset of the equilibrium strategies. Recent work has looked carefully at what trigger strategies are possible and optimal for players in infinitely repeated games; see Abreu, Pearce & Stacchetti (1990).

Empirical work on trigger strategies includes Porter (1983b), who examines price wars between railroads in the 19th century, and Slade (1987), who concluded that price wars among gas stations in Vancouver used small punishments for small deviations rather than big punishments for big deviations.

- A macroeconomist's technical note related to the similarity of infinite games and games with a constant probability of ending is Blanchard (1979), which discusses speculative bubbles.

- In the repeated Prisoner's Dilemma, if the end date is infinite with positive probability and only one player knows it, cooperation is possible by reasoning similar to that of the Gang of Four theorem in Section 6.4.
- Any Nash equilibrium of the one-shot game is also a perfect equilibrium of the finitely or infinitely repeated game.

N5.3 Reputation: The One-sided Prisoner's Dilemma

- *A game that is repeated an infinite number of times without discounting is called a **supergame**.*

There is no connection between the terms "supergame" and "subgame."

- The terms, "one-sided" and "two-sided" Prisoner's Dilemma, are my inventions. Only the two-sided version is a true Prisoner's Dilemma according to the definition of note N1.2.
- Empirical work on reputation is scarce. One worthwhile effort is Jarrell & Peltzman (1985), which finds that product recalls inflict costs greatly in excess of the measurable direct costs of the operations. The investigations into actual business practice of Macaulay (1963) is much cited and little imitated. He notes that reputation seems to be more important than the written details of business contracts.
- **Vengeance and Gratitude.** Most models have excluded these feelings (although see J. Hirshleifer [1987]), which can be modelled in two ways.

1 A player's current utility from *Confess* or *Deny* depends on what the other player has played in the past; or

2 A player's current utility depends on current actions and the other players' current utility in a way that changes with past actions of the other player.

The two approaches are subtly different in interpretation. In (1), the joy of revenge is in the action of confessing. In (2), the joy of revenge is in the discomfiture of the other player. Especially if the players have different payoff functions, these two approaches can lead to different results.

N5.4 Product Quality in an infinitely repeated game

- The game of Product Quality game may also be viewed as a principal agent model of moral hazard (see chapter 7). The seller (an agent), takes the action of choosing quality that is unobserved by the buyer (the principal), but which affects the principal's payoff, an interpretation used in much of the Stiglitz (1987) survey of the links between quality and price.

The intuition behind the Klein & Leffler model is similar to the explanation for high wages in the Shapiro & Stiglitz (1984) model of involuntary unemployment (section 8.1). Consumers, seeing a low price, realize that with a price that low the firm cannot resist lowering quality to make short-term profits. A large margin of profit is needed for the firm to decide on continuing to produce high quality.

- A paper related to Klein & Leffler (1981) is Shapiro (1983), which reconciles a high price with free entry by requiring that firms price under cost during the early periods to build up a reputation. If consumers believe, for example, that any firm charging a high price for any of the first five periods has produced a low quality product, but any firm charging a high price thereafter has produced high quality, then firms behave accordingly and the beliefs are confirmed. That the beliefs are self-confirming does not make them irrational; it only means that many different beliefs are rational in the many different equilibria.
- An equilibrium exists in the Product Quality model only if the entry cost F is just the right size to make n an integer in equation (8). Any of the usual assumptions to get around the integer problem could be used: allowing potential sellers to randomize between entering and staying out; assuming that for historical reasons, n firms have already entered; or assuming that firms lie on a continuum and the fixed cost is a uniform density across firms that have entered.

N5.5 Markov equilibria and overlapping generations in the game of Customer Switching Costs

- We assumed that the incumbent chooses its price first, but the alternation of incumbency remains even if we make the opposite assumption. The natural assumption is that prices are chosen simultaneously, but because of the discontinuity in the payoff function, that subgame has no equilibrium in pure strategies.

N5.6 Evolutionary equilibrium: the Hawk-Dove Game

- Dugatkin & Reeve (1998) is an edited volume of survey articles on different applications of game theory to biology. Dawkins (1989) is a good verbal introduction to evolutionary conflict. See also Axelrod & Hamilton (1981) for a short article on biological applications of the Prisoner's Dilemma, Hines (1987) for a survey, and Maynard Smith (1982) for a book. Weibull (1995) is a more recent treatment. J. Hirshleifer (1982) compares the approaches of economists and biologists. Boyd & Richerson (1985) uses evolutionary game theory to examine cultural transmission, which has important differences from purely genetic transmission.

Problems

5.1: Overlapping Generations (see Samuelson [1958])

There is a long sequence of players. One player is born in each period t , and he lives for periods t and $t + 1$. Thus, two players are alive in any one period, a youngster and an oldster. Each player is born with one unit of chocolate, which cannot be stored. Utility is increasing in chocolate consumption, and a player is very unhappy if he consumes less than 0.3 units of chocolate in a period: the per-period utility functions are $U(C) = -1$ for $C < 0.3$ and $U(C) = C$ for $C \geq 0.3$, where C is consumption. Players can give away their chocolate, but, since chocolate is the only good, they cannot sell it. A player's action is to consume X units of chocolate as a youngster and give away $1 - X$ to some oldster. Every person's actions in the previous period are common knowledge, and so can be used to condition strategies upon.

- 5.1a If there is finite number of generations, what is the unique Nash equilibrium?
- 5.1b If there are an infinite number of generations, what are two Pareto-ranked perfect equilibria?
- 5.1c If there is a probability θ at the end of each period (after consumption takes place) that barbarians will invade and steal all the chocolate (leaving the civilized people with payoffs of -1 for any X), what is the highest value of θ that still allows for an equilibrium with $X = 0.5$?

5.2. Product Quality with Lawsuits

Modify the Product Quality game of section 5.4 by assuming that if the seller misrepresents his quality he must, as a result of a class-action suit, pay damages of x per unit sold, where $x \in (0, c]$ and the seller becomes liable for x at the time of sale.

- 5.2a What is $\tilde{p}$ as a function of x, F, c , and r ? Is $\tilde{p}$ greater than when $x = 0$?
- 5.2b What is the equilibrium output per firm? Is it greater than when $x = 0$?
- 5.2c What is the equilibrium number of firms? Is it greater than when $x = 0$?
- 5.2d If, instead of x per unit, the seller pays X to a law firm to successfully defend him, what is the incentive compatibility constraint?

5.3. Repeated Games (see Benoit & Krishna [1985])

Players Benoit and Krishna repeat the game in Table 7 three times, with discounting:

Table 7: A Benoit-Krishna Game

		Krishna		
		<i>Deny</i>	<i>Waffle</i>	<i>Confess</i>
Benoit:	<i>Deny</i>	10,10	-1, -12	-1, 15
	<i>Waffle</i>	-12, -1	8,8	-1, -1
	<i>Confess</i>	15, -1	8, -1	0, 0

Payoffs to: (Benoit, Krishna).

- (a) Why is there no equilibrium in which the players play *Deny* in all three periods?
- 5.3b Describe a perfect equilibrium in which both players pick *Deny* in the first two periods.
- 5.3c Adapt your equilibrium to the twice-repeated game.
- 5.3d Adapt your equilibrium to the T -repeated game.
- 5.3e What is the greatest discount rate for which your equilibrium still works in the three-period game?

5.4. Repeated Entry Deterrence

Assume that Entry Deterrence I is repeated an infinite number of times, with a tiny discount rate and with payoffs received at the start of each period. In each period, the entrant chooses *Enter* or *Stay out*, even if he entered previously.

- 5.4a What is a perfect equilibrium in which the entrant enters each period?
- 5.4b Why is (*Stay out*, *Fight*) not a perfect equilibrium?
- 5.4c What is a perfect equilibrium in which the entrant never enters?
- 5.4d What is the maximum discount rate for which your strategy combination in part (c) is still an equilibrium?

5.5. The Repeated Prisoner's Dilemma

Set $P = 0$ in the general Prisoner's Dilemma in table 1.9, and assume that $2R > S + T$.

- 5.5a Show that the Grim Strategy, when played by both players, is a perfect equilibrium for the infinitely repeated game. What is the maximum discount rate for which the Grim Strategy remains an equilibrium?
- 5.5b Show that Tit-for-Tat is not a perfect equilibrium in the infinitely repeated Prisoner's Dilemma with no discounting.

5.6. Evolutionarily Stable Strategies

A population of scholars are playing the following coordination game over their two possible conversation topics over lunch, football and economics. Let $N_t(F)$ and $N_t(E)$ be the numbers who talk football and economics in period t , and let θ be the percentage who talk football, so $\theta = \frac{N(\text{football})}{N(\text{football}) + N(\text{economics})}$. Government regulations requiring lunchtime attendance and stipulating the topics of conversation have maintained the values $\theta = 0.5$, $N_t(F) = 50,000$ and $N_t(E) = 50,000$ up to this year's deregulatory reform. In the future, some people may decide to go home for lunch instead, or change their conversation. Table 8 shows the payoffs.

Table 8 Evolutionarily stable strategies

		Scholar 2	
		Football (θ)	Economics ($1 - \theta$)
Scholar 1	Football (θ)	1,1	0,0
	Economics ($1 - \theta$)	0,0	5,5
Payoffs to: (Scholar 1, Scholar 2)			

- 5.6a) There are three Nash equilibria: (*Football, Football*), (*Economics, Economics*), and a mixed-strategy equilibrium. What are the evolutionarily stable strategies?
- 5.6b Let $N_t(s)$ be the number of scholars playing a particular strategy in period t and let $\pi_t(s)$ be the payoff. Devise a Markov difference equation to express the population dynamics from period to period: $N_{t+1}(s) = f(N_t(s), \pi_t(s))$. Start the system with a population of 100,000, half the scholars talking football and half talking economics. Use your dynamics to finish Table 9.

Table 9: Conversation dynamics

t	$N_t(F)$	$N_t(E)$	θ	$\pi_t(F)$	$\pi_t(E)$
-1	50,000	50,000	0.5	0.5	2.5
0					
1					
2					

- 5.6c Repeat part (b), but specifying non-Markov dynamics, in which $N_{t+1}(s) = f(N_t(s), \pi_t(s), \pi_{t-1}(s))$.

5.7. Grab the Dollar

Table 10 shows the payoffs for the simultaneous-move game of Grab the Dollar. A silver dollar is put on the table between Smith and Jones. If one grabs it, he keeps the dollar, for a payoff of 4 utils. If both grab, then neither gets the dollar, and both feel bitter. If neither grabs, each gets to keep something.

Table 10: Grab the Dollar

		Jones	
		Grab (θ)	Wait ($1 - \theta$)
Smith:	Grab (θ)	-1, -1	4, 0
	Wait ($1 - \theta$)	0, 4	1, 1
Payoffs to: (Smith, Jones)			

- (5.7a) What are the evolutionarily stable strategies?
- 5.7b Suppose each player in the population is a point on a continuum, and that the initial amount of players is 1, evenly divided between *Grab* and *Wait*. Let $N_t(s)$ be the amount of players playing a particular strategy in period t and let $\pi_t(s)$ be the payoff. Let the population dynamics be $N_{t+1}(i) = (2N_t(i)) \left(\frac{\pi_t(i)}{\sum_j \pi_t(j)} \right)$. Find the missing entries in Table 11.

Table 11: Grab the Dollar: dynamics

t	$N_t(G)$	$N_t(W)$	$N_t(total)$	θ	$\pi_t(G)$	$\pi_t(w)$
0	0.5	0.5	1	0.5	1.5	0.5
1						
2						

5.7c Repeat part (b), but with the dynamics $N_{t+t}(s) = [1 + \frac{\pi_t(s)}{\sum_j \pi_t(j)}][2N_t(s)]$.

5.7d Which three games that have appeared so far in the book resemble Grab the Dollar?

6 Dynamic Games with Incomplete Information

6.1 Perfect Bayesian Equilibrium: Entry Deterrence II and III

Asymmetric information, and, in particular, incomplete information, is enormously important in game theory. This is particularly true for dynamic games, since when the players have several moves in sequence, their earlier moves may convey private information that is relevant to the decisions of players moving later on. Revealing and concealing information are the basis of much of strategic behavior and are especially useful as ways of explaining actions that would be irrational in a nonstrategic world.

Chapter 4 showed that even if there is symmetric information in a dynamic game, Nash equilibrium may need to be refined using subgame perfectness if the modeller is to make sensible predictions. Asymmetric information requires a somewhat different refinement to capture the idea of sunk costs and credible threats, and Section 6.1 sets out the standard refinement of perfect Bayesian equilibrium. Section 6.2 shows that even this may not be enough refinement to guarantee uniqueness and discusses further refinements based on out-of-equilibrium beliefs. Section 6.3 uses the idea to show that a player's ignorance may work to his advantage, and to explain how even when all players know something, lack of common knowledge still affects the game. Section 6.4 introduces incomplete information into the repeated Prisoner's Dilemma and shows the Gang of Four solution to the Chain-store Paradox of Chapter 5. Section 6.5 describes the celebrated Axelrod tournament, an experimental approach to the same paradox.¹

Subgame Perfectness is Not Enough

¹xxx Somewhere put in the table for the Prisoner dilemma payoffs.

Table 2: The Prisoner's Dilemma


		Column	
		<i>Deny</i>	<i>Confess</i>
Row	<i>Deny</i>	-1,-1 → -10, 0	
	<i>Confess</i>	0,-10 → - 8,-8	

Payoffs to: (Row,Column)

In games of asymmetric information, we will still require that an equilibrium be subgame perfect, but the mere forking of the game tree might not be relevant to a player's decision, because with asymmetric information he does not know which fork the game has taken. Smith might know he is at one of two different nodes depending on whether Jones has high or low production costs, but if he does not know the exact node, the “subgames” starting at each node are irrelevant to his decisions. In fact, they are not even subgames as we have defined them, because they cut across Smith's information sets. This can be seen in an asymmetric information version of Entry Deterrence I (Section 4.2). In Entry Deterrence I, the incumbent colluded with the entrant because fighting him was more costly than colluding once the entrant had entered. Now, let us set up the game to allow some entrants to be *Strong* and some *Weak* in the sense that it is more costly for the incumbent to choose *Fight* against a *Strong* entrant than a *Weak* one. The incumbent's payoff from *Fight|Strong* will be 0, as before, but his payoff from *Fight|Weak* will be X , where X will take values ranging from 0 (Entry Deterrence I) to 300 (Entry Deterrence IV and V) in different versions of the game.

Entry Deterrence II, III, and IV will all have the extensive form shown in Figure 1. With 50 percent probability, the incumbent's payoff from *Fight* is X rather than the 0 in Entry Deterrence I, but the incumbent does not know which payoff is the correct one in the particular realization of the game. This is modelled as an initial move by Nature, who chooses between the entrant being *Weak* or *Strong*, unobserved by the incumbent.

Figure 1 Entry Deterrence II, III, and IV



Entry Deterrence II: Fighting is Never Profitable

In Entry Deterrence II, $X = 1$, so information is not very asymmetric. It is common knowledge that the incumbent never benefits from *Fight*, even though his exact payoff might be zero or might be one. Unlike in the case of Entry Deterrence I, however, subgame perfectness does not rule out any Nash equilibria, because the only subgame is the subgame

starting at node N , which is the entire game. A subgame cannot start at nodes E_1 or E_2 , because neither of those nodes are singletons in the information partitions. Thus, the implausible Nash equilibrium, (*Stay Out*, *Fight*), escapes elimination by a technicality.

The equilibrium concept needs to be refined in order to eliminate the implausible equilibrium. Two general approaches can be taken: either introduce small “trembles” into the game, or require that strategies be best responses given rational beliefs. The first approach takes us to the “trembling hand-perfect” equilibrium, while the second takes us to the “perfect Bayesian” and “sequential” equilibrium. The results are similar whichever approach is taken.

Trembling-hand Perfectness

Trembling-hand perfectness is an equilibrium concept introduced by Selten (1975) according to which a strategy that is to be part of an equilibrium must continue to be optimal for the player even if there is a small chance that the other player will pick an out-of-equilibrium action (i.e., that the other player’s hand will “tremble”).

Trembling-hand perfectness is defined for games with finite action sets as follows.

*The strategy profile s^* is a **trembling-hand perfect** equilibrium if for any ϵ there is a vector of positive numbers $\delta_1, \dots, \delta_n \in [0, 1]$ and a vector of completely mixed strategies $\sigma_1, \dots, \sigma_n$ such that the perturbed game where every strategy is replaced by $(1 - \delta_i)s_i + \delta_i\sigma_i$ has a Nash equilibrium in which every strategy is within distance ϵ of s^* .*

Every trembling-hand perfect equilibrium is subgame perfect; indeed, Section 4.1 justified subgame perfectness using a tremble argument. Unfortunately, it is often hard to tell whether a strategy profile is trembling-hand perfect, and the concept is undefined for games with continuous strategy spaces because it is hard to work with mixtures of a continuum (see note N3.1). Moreover, the equilibrium depends on which trembles are chosen, and deciding why one tremble should be more common than another may be difficult.

Perfect Bayesian Equilibrium and Sequential Equilibrium

The second approach to asymmetric information, introduced by Kreps & Wilson (1982b) in the spirit of Harsanyi (1967), is to start with prior beliefs, common to all players, that specify the probabilities with which Nature chooses the types of the players at the beginning of the game. Some of the players observe Nature’s move and update their beliefs, while other players can update their beliefs only by deductions they make from observing the actions of the informed players.

The deductions used to update beliefs are based on the actions specified by the equilibrium. When players update their beliefs, they assume that the other players are following the equilibrium strategies, but since the strategies themselves depend on the beliefs, an equilibrium can no longer be defined based on strategies alone. Under asymmetric information, an equilibrium is a strategy profile and a set of beliefs such that the strategies are best responses. The combination of beliefs and strategies is called an **assessment** by Kreps and Wilson.

On the equilibrium path, all that the players need to update their beliefs are their priors and Bayes' s Rule, but off the equilibrium path this is not enough. Suppose that in equilibrium, the entrant always enters. If for whatever reason the impossible happens and the entrant stays out, what is the incumbent to think about the probability that the entrant is weak? Bayes' s Rule does not help, because when $Prob(data) = 0$, which is the case for data such as *Stay Out* which is never observed in equilibrium, the posterior belief cannot be calculated using Bayes' s Rule. From section 2.4,

$$Prob(Weak|Stay Out) = \frac{Prob(Stay Out|Weak)Prob(Weak)}{Prob(Stay Out)}. \quad (1)$$

The posterior $Prob(Weak|Stay Out)$ is undefined, because (6.1) requires dividing by zero.

A natural way to define equilibrium is as a strategy profile consisting of best responses given that equilibrium beliefs follow Bayes' s Rule and out-of-equilibrium beliefs follow a specified pattern that does not contradict Bayes' s Rule.

A perfect Bayesian equilibrium is a strategy profile s and a set of beliefs μ such that at each node of the game:

- (1) The strategies for the remainder of the game are Nash given the beliefs and strategies of the other players.
- (2) The beliefs at each information set are rational given the evidence appearing thus far in the game (meaning that they are based, if possible, on priors updated by Bayes' s Rule, given the observed actions of the other players under the hypothesis that they are in equilibrium).

Kreps & Wilson (1982b) use this idea to form their equilibrium concept of sequential equilibrium, but they impose a third condition, defined only for games with discrete strategies, to restrict beliefs a little further:

- (3) The beliefs are the limit of a sequence of rational beliefs, i.e., if (μ^*, s^*) is the equilibrium assessment, then some sequence of rational beliefs and completely mixed strategies converges to it:

$$(\mu^*, s^*) = \lim_{n \rightarrow \infty} (\mu^n, s^n) \text{ for some sequence } (\mu^n, s^n) \text{ in } \{\mu, s\}.$$

Condition (3) is quite reasonable and makes sequential equilibrium close to trembling-hand perfect equilibrium, but it adds more to the concept's difficulty than to its usefulness. If players are using the sequence of completely mixed strategies s^n , then every action is taken with some positive probability, so Bayes'Rule can be applied to form the beliefs μ^n after any action is observed. Condition (3) says that the equilibrium assessment has to be the limit of some such sequence (though not of every such sequence). For the rest of the book we will use perfect Bayesian equilibrium and dispense with condition (3), although it usually can be satisfied.

Sequential equilibria are always subgame perfect (condition (1) takes care of that). Every trembling-hand perfect equilibrium is a sequential equilibrium, and "almost every" sequential equilibrium is trembling hand perfect. Every sequential equilibrium is perfect Bayesian, but not every perfect Bayesian equilibrium is sequential.

Back to Entry Deterrence II

Armed with the concept of the perfect Bayesian equilibrium, we can find a sensible equilibrium for Entry Deterrence II .

Entrant: $Enter|Weak, Enter|Strong$

Incumbent: $Collude$

Beliefs: $Prob(Strong| Stay Out) = 0.4$

In this equilibrium the entrant enters whether he is *Weak* or *Strong*. The incumbent's strategy is *Collude*, which is not conditioned on Nature's move, since he does not observe it. Because the entrant enters regardless of Nature's move, an out-of-equilibrium belief for the incumbent if he should observe *Stay Out* must be specified, and this belief is arbitrarily chosen to be that the incumbent's subjective probability that the entrant is *Strong* is 0.4 given his observation that the entrant deviated by choosing *Stay Out*. Given this strategy profile and out-of-equilibrium belief, neither player has incentive to change his strategy.

There is no perfect Bayesian equilibrium in which the entrant chooses *Stay Out*. *Fight* is a bad response even under the most optimistic possible belief, that the entrant is *Weak* with probability 1. Notice that perfect Bayesian equilibrium is not defined structurally, like subgame perfectness, but rather in terms of optimal responses. This enables it to come closer to the economic intuition which we wish to capture by an equilibrium refinement.

Finding the perfect Bayesian equilibrium of a game, like finding the Nash equilibrium, requires intelligence. Algorithms are not useful. To find a Nash equilibrium, the modeller thinks about his game, picks a plausible strategy profile, and tests whether the strategies are best responses to each other. To make it a perfect Bayesian equilibrium, he notes which actions are never taken in equilibrium and specifies the beliefs that players use to interpret those actions. He then tests whether each player's strategies are best responses given his beliefs at each node, checking in particular whether any player would like to take an out-of-equilibrium action in order to set in motion the other players' out-of-equilibrium beliefs and strategies. This process does not involve testing whether a player's beliefs are beneficial to the player, because players do not choose their own beliefs; the priors and out-of-equilibrium beliefs are exogenously specified by the modeller.

One might wonder why the beliefs have to be specified in Entry Deterrence II. Does not the game tree specify the probability that the entrant is *Weak*? What difference does it make if the entrant stays out? Admittedly, Nature does choose each type with probability 0.5, so if the incumbent had no other information than this prior, that would be his belief. But the entrant's action might convey additional information. The concept of perfect Bayesian equilibrium leaves the modeller free to specify how the players form beliefs from that additional information, so long as the beliefs do not violate Bayes' Rule. (A technically valid choice of beliefs by the modeller might still be met with scorn, though, as with any silly assumption.) Here, the equilibrium says that if the entrant stays out, the incumbent believes he is *Strong* with probability 0.4 and *Weak* with probability 0.6, beliefs that are arbitrary but do not contradict Bayes' s Rule.

In Entry Deterrence II the out-of-equilibrium beliefs do not and should not matter.

If the entrant chooses *Stay Out*, the game ends, so the incumbent's beliefs are irrelevant. Perfect Bayesian equilibrium was only introduced as a way out of a technical problem. In the next section, however, the precise out-of-equilibrium beliefs will be crucial to which strategy profiles are equilibria.

6.2 Refining Perfect Bayesian Equilibrium: the PhD Admissions Game

Entry Deterrence III: Fighting is Sometimes Profitable

In Entry Deterrence III, assume that $X = 60$, not $X = 1$. This means that fighting is more profitable for the incumbent than collusion if the entrant is *Weak*. As before, the entrant knows if he is *Weak*, but the incumbent does not. Retaining the prior after observing out-of-equilibrium actions, which in this game is $Prob(Strong) = 0.5$, is a convenient way to form beliefs that is called **passive conjectures**. The following is a perfect Bayesian equilibrium which uses passive conjectures.

A plausible pooling equilibrium for Entry Deterrence III

Entrant: $Enter|Weak, Enter|Strong$

Incumbent:] $Collude$

Beliefs: $Prob(Strong|Stay Out) = 0.5$

In choosing whether to enter, the entrant must predict the incumbent's behavior. If the probability that the entrant is *Weak* is 0.5, the expected payoff to the incumbent from choosing *Fight* is 30 ($= 0.5[0] + 0.5[60]$), which is less than the payoff of 50 from *Collude*. The incumbent will collude, so the entrant enters. The entrant may know that the incumbent's payoff is actually 60, but that is irrelevant to the incumbent's behavior.

The out-of-equilibrium belief does not matter to this first equilibrium, although it will in other equilibria of the same game. Although beliefs in a perfect Bayesian equilibrium must follow Bayes' s Rule, that puts very little restriction on how players interpret out-of-equilibrium behavior. Out-of-equilibrium behavior is "impossible," so when it does occur there is no obvious way the player should react. Some beliefs may seem more reasonable than others, however, and Entry Deterrence III has another equilibrium that requires less plausible beliefs off the equilibrium path.

An implausible equilibrium for Entry Deterrence III

Entrant: $Stay Out|Weak, Stay Out|Strong$

Incumbent: $Fight$

Beliefs: $Prob(Strong|Enter) = 0.1$

This is an equilibrium because if the entrant were to deviate and enter, the incumbent would calculate his payoff from fighting to be 54 ($= 0.1[0] + 0.9[60]$), which is greater than the *Collude* payoff of 50. The entrant would therefore stay out.

The beliefs in the implausible equilibrium are different and less reasonable than in the plausible equilibrium. Why should the incumbent believe that weak entrants would enter

mistakenly nine times as often as strong entrants? The beliefs do not violate Bayes' s Rule, but they have no justification.

The reasonableness of the beliefs is important because if the incumbent uses passive conjectures, the implausible equilibrium breaks down. With passive conjectures, the incumbent would want to change his strategy to *Collude*, because the expected payoff from *Fight* would be less than 50. The implausible equilibrium is less robust with respect to beliefs than the plausible equilibrium, and it requires beliefs that are harder to justify.

Even though dubious outcomes may be perfect Bayesian equilibria, the concept does have some bite, ruling out other dubious outcomes. There does not, for example, exist an equilibrium in which the entrant enters only if he is *Strong* and stays out if he is *Weak* (called a “separating equilibrium” because it separates out different types of players). Such an equilibrium would have to look like this:

A conjectured separating equilibrium for Entry Deterrence III

Entrant: *Stay Out*|*Weak*, *Enter*|*Strong*

Incumbent: *Collude*

No out-of-equilibrium beliefs are specified for the conjectures in the separating equilibrium because there is no out-of-equilibrium behavior about which to specify them. Since the incumbent might observe either *Stay out* or *Enter* in equilibrium, the incumbent will always use Bayes' s Rule to form his beliefs. He will believe that an entrant who stays out must be weak and an entrant who enters must be strong. This conforms to the idea behind Nash equilibrium that each player assumes that the other follows the equilibrium strategy, and then decides how to reply. Here, the incumbent's best response, given his beliefs, is *Collude*|*Enter*, so that is the second part of the proposed equilibrium. But this cannot be an equilibrium, because the entrant would want to deviate. Knowing that entry would be followed by collusion, even the weak entrant would enter. So there cannot be an equilibrium in which the entrant enters only when strong.

The PhD Admissions Game

Passive conjectures may not always be the most satisfactory belief, as the next example shows. Suppose that a university knows that 90 percent of the population hate economics and would be unhappy in its PhD program, and 10 percent love economics and would do well. In addition, it cannot observe the applicant's type. If the university rejects an application, its payoff is 0 and the applicant's is -1 because of the trouble needed to apply. If the university accepts the application of someone who hates economics, the payoffs of both university and student are -10 , but if the applicant loves economics, the payoffs are $+20$ for each player. Figure 2 shows this game in extensive form. The population proportions are represented by a node at which Nature chooses the student to be a *Lover* or *Hater* of economics.

Figure 2 PhD Admissions

The PhD Admissions Game is a signalling game of the kind we will look at in Chapter 10. It has various perfect Bayesian equilibria that differ in their out-of-equilibrium beliefs, but the equilibria can be divided into two distinct categories, depending on the outcome: the **separating equilibrium**, in which the lovers of economics apply and the haters do not, and the **pooling equilibrium**, in which neither type of student applies.

A separating equilibrium for PhD Admissions

Student: *Apply* | *Lover*, *Do Not Apply* | *Hater*

University: *Admit*

The separating equilibrium does not need to specify out-of-equilibrium beliefs, because Bayes' s Rule can always be applied whenever both of the two possible actions *Apply* and *Do Not Apply* can occur in equilibrium.

A pooling equilibrium for PhD Admissions

Student: *Do Not Apply* | *Lover*, *Do Not Apply* | *Hater*

University: *Reject*

Beliefs: $Prob(Hater|Apply) = 0.9$ (passive conjectures)

The pooling equilibrium is supported by passive conjectures. Both types of students refrain from applying because they believe correctly that they would be rejected and receive a payoff of -1 ; and the university is willing to reject any student who foolishly applied, believing that he is a *Hater* with 90 percent probability.

Because the perfect Bayesian equilibrium concept imposes no restrictions on out-of-equilibrium beliefs, researchers starting with McLennan (1985) have come up with a variety of exotic refinements of the equilibrium concept. Let us consider whether various alternatives to passive conjectures would support the pooling equilibrium in PhD Admissions.

Passive Conjectures. $\text{Prob}(\text{Hater}|\text{Apply}) = 0.9$

This is the belief specified above, under which out-of-equilibrium behavior leaves beliefs unchanged from the prior. The argument for passive conjectures is that the student's application is a mistake, and that both types are equally likely to make mistakes, although *Haters* are more common in the population. This supports the pooling equilibrium.

The Intuitive Criterion. $\text{Prob}(\text{Hater}|\text{Apply}) = 0$

Under the Intuitive Criterion of Cho & Kreps (1987), if there is a type of informed player who could not benefit from the out-of-equilibrium action no matter what beliefs were held by the uninformed player, the uninformed player's belief must put zero probability on that type. Here, the *Hater* could not benefit from applying under any possible beliefs of the university, so the university puts zero probability on an applicant being a *Hater*. This argument will not support the pooling equilibrium, because if the university holds this belief, it will want to admit anyone who applies.

Complete Robustness. $\text{Prob}(\text{Hater}|\text{Apply}) = m, 0 \leq m \leq 1$

Under this approach, the equilibrium strategy profile must consist of responses that are best, given any and all out-of-equilibrium beliefs. Our equilibrium for Entry Deterrence II satisfied this requirement. Complete robustness rules out a pooling equilibrium in PhD Admissions, because a belief like $m = 0$ makes accepting applicants a best response, in which case only the *Lover* will apply. A useful first step in analyzing conjectured pooling equilibria is to test whether they can be supported by extreme beliefs such as $m = 0$ and $m = 1$.

An ad hoc specification. $\text{Prob}(\text{Hater}|\text{Apply}) = 1$

Sometimes the modeller can justify beliefs by the circumstances of the particular game. Here, one could argue that anyone so foolish as to apply knowing that the university would reject them could not possibly have the good taste to love economics. This supports the pooling equilibrium also.

An alternative approach to the problem of out-of-equilibrium beliefs is to remove its origin by building a model in which every outcome is possible in equilibrium because different types of players take different equilibrium actions. In PhD Admissions, we could assume that there are a few students who both love economics and actually enjoy writing applications. Those students would always apply in equilibrium, so there would never be a pure pooling equilibrium in which nobody applied, and Bayes' s Rule could always be used. In equilibrium, the university would always accept someone who applied, because applying is never out-of-equilibrium behavior and it always indicates that the applicant is a *Lover*. This approach is especially attractive if the modeller takes the possibility of trembles literally, instead of just using it as a technical tool.

The arguments for different kinds of beliefs can also be applied to Entry Deterrence III, which had two different pooling equilibria and no separating equilibrium. We used passive conjectures in the "plausible" equilibrium. The intuitive criterion would not restrict beliefs

at all, because both types would enter if the incumbent's beliefs were such as to make him collude, and both would stay out if they made him fight. Complete robustness would rule out as an equilibrium the strategy profile in which the entrant stays out regardless of type, because the optimality of staying out depends on the beliefs. It would support the strategy profile in which the entrant enters and out-of-equilibrium beliefs do not matter.

The Importance of Common Knowledge: Entry Deterrence IV and V

To demonstrate the importance of common knowledge, let us consider two more versions of Entry Deterrence. We will use passive conjectures in both. In Entry Deterrence III, the incumbent was hurt by his ignorance. Entry Deterrence IV will show how he can benefit from it, and Entry Deterrence V will show what can happen when the incumbent has the same information as the entrant but the information is not common knowledge.

Entry Deterrence IV: the incumbent benefits from ignorance

To construct Entry Deterrence IV, let $X = 300$ in Figure 1, so fighting is even more profitable than in Entry Deterrence III but the game is otherwise the same: the entrant knows his type, but the incumbent does not. The following is the unique perfect Bayesian equilibrium in pure strategies.²

Equilibrium for Entry Deterrence IV

Entrant: *Stay Out* | *Weak*, *Stay Out* | *Strong*

Incumbent: *Fight*

Beliefs: $Prob(Strong|Enter) = 0.5$ (passive conjectures)

This equilibrium can be supported by other out-of-equilibrium beliefs, but no equilibrium is possible in which the entrant enters. There is no pooling equilibrium in which both types of entrant enter, because then the incumbent's expected payoff from *Fight* would be 150 ($= 0.5[0] + 0.5[300]$), which is greater than the *Collude* payoff of 50. There is no separating equilibrium, because if only the strong entrant entered and the incumbent always colluded, the weak entrant would be tempted to imitate him and enter as well.

In Entry Deterrence IV, unlike Entry Deterrence III, the incumbent benefits from his own ignorance, because he would always fight entry, even if the payoff were (unknown to himself) just zero. The entrant would very much like to communicate the costliness of fighting, but the incumbent would not believe him, so entry never occurs.

Entry Deterrence V: Lack of Common Knowledge of Ignorance

In Entry Deterrence V, it may happen that both the entrant and the incumbent know the payoff from (*Enter*, *Fight*), but the entrant does not know whether the incumbent knows. The information is known to both players, but is not common knowledge.

²There exists a plausible mixed-strategy equilibrium too: Entrant: *Enter* if *Strong*, *Enter* with probability $m = .2$ if *Weak*; Incumbent: *Collude* with probability $n = .2$. The payoff from this is only 150, so if the equilibrium were one in mixed strategies, ignorance would *not* help.

Figure 3 depicts this somewhat complicated situation. The game begins with Nature assigning the entrant a type, *Strong* or *Weak* as before. This is observed by the entrant but not by the incumbent. Next, Nature moves again and either tells the incumbent the entrant's type or remains silent. This is observed by the incumbent, but not by the entrant. The four games starting at nodes G_1 to G_4 represent different combinations of payoffs from (*Enter*, *Fight*) and knowledge of the incumbent. The entrant does not know how well informed the incumbent is, so the entrant's information partition is $(\{G_1, G_2\}, \{G_3, G_4\})$.

Figure 3 Entry Deterrence V

Equilibrium for Entry Deterrence V

Entrant: *Stay Out*|*Weak*, *Stay Out*|*Strong*

Incumbent: *Fight*|*Nature said "Weak"*, *Collude* |*Nature said "Strong"*, *Fight* |*Nature said nothing*

Beliefs: $Prob(\text{Strong}|\text{Enter, Nature said nothing}) = 0.5$ (passive conjectures)

Since the entrant puts a high probability on the incumbent not knowing, the entrant should stay out, because the incumbent will fight for either of two reasons. With probability 0.9, Nature has said nothing and the incumbent calculates his expected payoff from *Fight* to be 150, and with probability 0.05 ($= 0.1[0.5]$) Nature has told the incumbent that the entrant is weak and the payoff from *Fight* is 300. Even if the entrant is strong and Nature tells this to the incumbent, the entrant would choose *Stay Out*, because he does not know that the incumbent knows, and his expected payoff from *Enter* would be -5 ($= [0.9][-10] + 0.1[40]$).

If it were common knowledge that the entrant was strong, the entrant would enter and the incumbent would collude. If it is known by both players, but not common knowledge, the entrant stays out, even though the incumbent would collude if he entered. Such is the importance of common knowledge.

6.4 Incomplete Information in the Repeated Prisoner's Dilemma: The Gang of Four Model

Chapter 5 explored various ways to steer between the Scylla of the Chainstore Paradox and the Charybdis of the Folk Theorem to find a resolution to the problem of repeated games. In the end, uncertainty turned out to make little difference to the problem, but incomplete information was left unexamined in chapter 5. One might imagine that if the players did not know each others' types, the resulting confusion might allow cooperation. Let us investigate this by adding incomplete information to the finitely repeated Prisoner's Dilemma and finding the perfect Bayesian equilibria.

One way to incorporate incomplete information would be to assume that a large number of players are irrational, but that a given player does not know whether any other player is of the irrational type or not. In this vein, one might assume that with high probability Row is a player who blindly follows the strategy of Tit-for-Tat. If Column thinks he is playing against a Tit-for-Tat player, his optimal strategy is to *Deny* until near the last period (how near depending on the parameters), and then *Confess*. If he were not certain of this, but the probability were high that he faced a Tit-for-Tat player, Row would choose that same strategy. Such a model begs the question, because it is not the incompleteness of the information that drives the model, but the high probability that one player blindly uses Tit-for-Tat. Tit-for-Tat is not a rational strategy, and to assume that many players use it is to assume away the problem. A more surprising result is that a small amount of incomplete information can make a big difference to the outcome.³

The Gang of Four Model

One of the most important explanations of reputation is that of Kreps, Milgrom, Roberts & Wilson (1982), hereafter referred to as the Gang of Four. In their model, a few players are genuinely unable to play any strategy but Tit-for-Tat, and many players pretend to be of that type. The beauty of the model is that it requires only a small amount of incomplete information, and a low probability γ that player Row is a Tit-for-Tat player. It is not unreasonable to suppose that a world which contains Neo-Ricardians and McGovernites contains a few mildly irrational tit-for-tat players, and such behavior is especially plausible among consumers, who are subject to less evolutionary pressure than firms.

It may even be misleading to call the Tit-for-Tat "irrational", because they may just have unusual payoffs, particularly since we will assume that they are rare. The unusual players have a small direct influence, but they matter because other players imitate them. Even if Column knows that with high probability Row is just pretending to be Tit-for-Tat, Column does not care what the truth is so long as Row keeps on pretending. Hypocrisy is not only the tribute vice pays to virtue; it can be just as good for deterring misbehavior.

Theorem 6.1: The Gang of Four theorem

Consider a T -stage, repeated prisoner's dilemma, without discounting but with a probability

³Begging the question is not as illegitimate in modelling as in rhetoric, however, because it may indicate that the question is a vacuous one in the first place. If the payoffs of the Prisoner's Dilemma are not those of most of the people one is trying to model, the Chainstore Paradox becomes irrelevant.

γ of a *Tit-for-Tat* player. In any perfect Bayesian equilibrium, the number of stages in which either player chooses *Confess* is less than some number M that depends on γ but not on T .

The significance of the Gang of Four theorem is that while the players do resort to *Confess* as the last period approaches, the number of periods during which they *Confess* is independent of the total number of periods. Suppose $M = 2,500$. If $T = 2,500$, there might be a *Confess* every period. But if $T = 10,000$, there are 7,500 periods without a *Confess*. For reasonable probabilities of the unusual type, the number of periods of cooperation can be much larger. Wilson (unpublished) has set up an entry deterrence model in which the incumbent fights entry (the equivalent of *Deny* above) up to seven periods from the end, although the probability the entrant is of the unusual type is only 0.008.

The Gang of Four Theorem characterizes the equilibrium outcome rather than the equilibrium. Finding perfect Bayesian equilibria is difficult and tedious, since the modeller must check all the out-of-equilibrium subgames, as well as the equilibrium path. Modellers usually content themselves with describing important characteristics of the equilibrium strategies and payoffs. Section 14.3 contains a somewhat more detailed description of what happens in a model of repeated entry deterrence with incomplete information.

To get a feeling for why Theorem 6.1 is correct, consider what would happen in a 10,001 period game with a probability of 0.01 that Row is playing the Grim Strategy of *Deny* until the first *Confess*, and *Confess* every period thereafter. If the payoffs are as in table 5.2a, a best response for Column to a known grim player is (*Confess* only in the last period, unless Row chooses *Confess* first, in which case respond with *Confess*). Both players will choose *Deny* until the last period, and Column's payoff will be 50,010 ($= (10,000)(5) + 10$). Suppose for the moment that if Row is not grim, he is highly aggressive, and will choose *Confess* every period. If Column follows the strategy just described, the outcome will be (*Confess*, *Deny*) in the first period and (*Confess*, *Confess*) thereafter, for a payoff to Column of $-5 (= -5 + (10,000)(0))$. If the probabilities of the two outcomes are 0.01 and 0.99, Column's expected payoff from the strategy described is 495.15. If instead he follows a strategy of (*Confess* every period), his expected payoff is just 0.1 ($= 0.01(10) + 0.99(0)$). It is clearly in Column's advantage to take a chance by cooperating with Row, even if Row has a 0.99 probability of following a very aggressive strategy.

The aggressive strategy, however, is not Row's best response to Column's strategy. A better response is for Row to choose *Deny* until the second-to-last period, and then to choose *Confess*. Given that Column is cooperating in the early periods, Row will cooperate also. This argument has not described what the Nash equilibrium actually is, since the iteration back and forth between Row and Column can be continued, but it does show why Column chooses *Deny* in the first period, which is the leverage the argument needs: the payoff is so great if Row is actually the grim player that it is worthwhile for Column to risk a low payoff for one period.

The Gang of Four Theorem provides a way out of the Chainstore Paradox, but it creates a problem of multiple equilibria in much the same way as the infinitely repeated game. For one thing, if the asymmetry is two-sided, so both players might be unusual

types, it becomes much less clear what happens in threat games such as Entry Deterrence. Also, what happens depends on which unusual behaviors have positive, if small, probability. Theorem 6.2 says that the modeller can make the average payoffs take any particular values by making the game last long enough and choosing the form of the irrationality carefully.

Theorem 6.2: The Incomplete Information Folk Theorem(Fudenberg & Maskin [1986] p. 547)

For any two-person repeated game without discounting, the modeller can choose a form of irrationality so that for any probability $\epsilon > 0$ there is some finite number of repetitions such that with probability $(1 - \epsilon)$ a player is rational and the average payoffs in some sequential equilibrium are closer than ϵ to any desired payoffs greater than the minimax payoffs.

6.5 The Axelrod Tournament

Another way to approach the repeated Prisoner's Dilemma is through experiments, such as the round robin tournament described by political scientist Robert Axelrod in his 1984 book. Contestants submitted strategies for a 200-repetition Prisoner's Dilemma. Since the strategies could not be updated during play, players could precommit, but the strategies could be as complicated as they wished. If a player wanted to specify a strategy which simulated subgame perfectness by adapting to past history just as a noncommitted player would, he was free to do so, but he could also submit a non-perfect strategy such as Tit-for-Tat or the slightly more forgiving Tit-for-Two-Tats. Strategies were submitted in the form of computer programs that were matched with each other and played automatically. In Axelrod's first tournament, 14 programs were submitted as entries. Every program played every other program, and the winner was the one with the greatest sum of payoffs over all the plays. The winner was Anatol Rapoport, whose strategy was Tit-for-Tat.

The tournament helps to show which strategies are robust against a variety of other strategies in a game with given parameters. It is quite different from trying to find a Nash equilibrium, because it is not common knowledge what the equilibrium is in such a tournament. The situation could be viewed as a game of incomplete information in which Nature chooses the number and cognitive abilities of the players and their priors regarding each other.

After the results of the first tournament were announced, Axelrod ran a second tournament, adding a probability $\theta = 0.00346$ that the game would end each round so as to avoid the Chainstore Paradox. The winner among the 62 entrants was again Anatol Rapoport, and again he used Tit-for-Tat.

Before choosing his tournament strategy, Rapoport had written an entire book on the Prisoner's Dilemma in analysis, experiment, and simulation (Rapoport & Chammah [1965]). Why did he choose such a simple strategy as Tit-for-Tat? Axelrod points out that Tit-for-Tat has three strong points.

1. It never initiates confessing (**niceness**);
2. It retaliates instantly against confessing (**provokability**);
3. It forgives a confesser who goes back to cooperating (it is **forgiving**).

Despite these advantages, care must be taken in interpreting the results of the tournament. It does not follow that Tit-for-Tat is the best strategy, or that cooperative behavior should always be expected in repeated games.

First, Tit-for-Tat never beats any other strategy in a one-on-one contest. It won the tournament by piling up points through cooperation, having lots of high score plays and very few low score plays. In an elimination tournament, Tit-for-Tat would be eliminated very early, because it scores *high* payoffs but never the *highest* payoff.

Second, the other players' strategies matter to the success of Tit-for-Tat. In neither tournament were the strategies submitted a Nash equilibrium. If a player knew what strategies he was facing, he would want to revise his own. Some of the strategies submitted in the second tournament would have won the first, but they did poorly because the environment had changed. Other programs, designed to try to probe the strategies of their opposition, wasted too many (*Confess, Confess*) episodes on the learning process, but if the games had lasted a thousand repetitions they would have done better.

Third, in a game in which players occasionally confessed because of trembles, two Tit-for-Tat players facing each other would do very badly. The strategy instantly punishes a confessing player, and it has no provision for ending the punishment phase.

Optimality depends on the environment. When information is complete and the payoffs are all common knowledge, confessing is the only equilibrium outcome, but in practically any imaginable situation, information is slightly incomplete, so cooperation becomes more plausible. Tit-for-Tat is suboptimal for any given environment, but it is robust across environments, and that is its advantage.

6.6 Credit and the Age of the Firm: the Diamond Model

An example of another way to look at reputation is Diamond's model of credit terms, which seeks to explain why older firms get cheaper credit using a game similar to the Gang of Four model. Telser (1966) suggested that predatory pricing would be a credible threat if the incumbent had access to cheaper credit than the entrant, and so could hold out for more periods of losses before going bankrupt. While one might wonder whether this is effective protection against entry—what if the entrant is a large old firm from another industry?—we shall focus on how better-established firms might get cheaper credit.

D. Diamond (1989) aims to explain why old firms are less likely than young firms to default on debt. His model has both adverse selection, because firms differ in type, and

moral hazard, because they take hidden actions. The three types of firms, R, S, and RS, are “born” at time zero and borrow to finance projects at the start of each of T periods. We must imagine that there are overlapping generations of firms, so that at any point in time a variety of ages are coexisting, but the model looks at the lifecycle of only one generation. All the players are risk neutral. Type RS firms can choose independently risky projects with negative expected values or safe projects with low but positive expected values. Although the risky projects are worse in expectation, if they are successful the return is much higher than from safe projects. Type R firms can only choose risky projects, and type S firms only safe projects. At the end of each period the projects bring in their profits and loans are repaid, after which new loans and projects are chosen for the next period. Lenders cannot tell which project is chosen or what a firm’s current profits are, but they can seize the firm’s assets if a loan is not repaid, which always happens if the risky project was chosen and turned out unsuccessfully.

This game foreshadows two other models of credit that will be described in this book, the Repossession Game of section 8.4 and the Stiglitz-Weiss model of section 9.6. Both will be one-shot games in which the bank worried about not being repaid; in the Repossession Game because the borrower did not exert enough effort, and in the Stiglitz-Weiss model because he was of an undesirable type that could not repay. The Diamond model is a mixture of adverse selection and moral hazard: the borrowers differ in type, but some borrowers have a choice of action.

The equilibrium path has three parts. The RS firms start by choosing risky projects. Their downside risk is limited by bankruptcy, but if the project is successful the firm keeps large residual profits after repaying the loan. Over time, the number of firms with access to the risky project (the RS’s and R’s) diminishes through bankruptcy, while the number of S’s remains unchanged. Lenders can therefore maintain zero profits while lowering their interest rates. When the interest rate falls, the value of a stream of safe investment profits minus interest payments rises relative to the expected value of the few periods of risky returns minus interest payments before bankruptcy. After the interest rate has fallen enough, the second phase of the game begins when the RS firms switch to safe projects at a period we will call t_1 . Only the tiny and diminishing group of type R firms continue to choose risky projects. Since the lenders know that the RS firms switch, the interest rate can fall sharply at t_1 . A firm that is older is less likely to be a type R, so it is charged a lower interest rate. Figure 4 shows the path of the interest rate over time.

Figure 4 The interest rate over time

Towards period T , the value of future profits from safe projects declines and even with a low interest rate the RS's are again tempted to choose risky projects. They do not all switch at once, however, unlike in period t_1 . In period t_1 , if a few RS's had decided to switch to safe projects, the lenders would have been willing to lower the interest rate, which would have made switching even more attractive. If a few firms switch to risky projects at some time t_2 , on the other hand, the interest rate rises and switching to risky projects becomes more attractive—a result that will also be seen in the Lemons model in Chapter 9. Between t_2 and t_3 , the RS's follow a mixed strategy, an increasing number of them choosing risky projects as time passes. The increasing proportion of risky projects causes the interest rate to rise. At t_3 , the interest rate is high enough and the end of the game is close enough that the RS's revert to the pure strategy of choosing risky projects. The interest rate declines during this last phase as the number of RS's diminishes because of failed risky projects.

One might ask, in the spirit of modelling by example, why the model contains three types of firms rather than two. Types S and RS are clearly needed, but why type R? The little extra detail in the game description allows simplification of the equilibrium, because with three types bankruptcy is never out-of-equilibrium behaviour, since the failing firm might be a type R. Bayes's Rule can therefore always be applied, eliminating the problem of ruling out peculiar beliefs and absurd perfect bayesian equilibria.

This is a Gang of Four model but differs from previous examples in an important respect: the Diamond model is not stationary, and as time progresses, some firms of types R and RS go bankrupt, which changes the lenders' payoff functions. Thus, it is not, strictly speaking, a repeated game.

N6.1 Perfect Bayesian equilibrium: Entry Deterrence I and II

- Section 4.1 showed that even in games of perfect information, not every subgame perfect equilibrium is trembling-hand perfect. In games of perfect information, however, every subgame perfect equilibrium is a perfect Bayesian equilibrium, since no out-of-equilibrium beliefs need to be specified.

N6.2 Refining Perfect Bayesian equilibrium: the PhD Admissions Game

- Fudenberg & Tirole (1991b) is a careful analysis of the issues involved in defining perfect Bayesian equilibrium.
- Section 6.2 is about debatable ways of restricting beliefs such as passive conjectures or equilibrium dominance, but less controversial restrictions are sometimes useful. In a three-player game, consider what happens when Smith and Jones have incomplete information about Brown, and then Jones deviates. If it was Brown himself who had deviated, one might think that the other players might deduce something about Brown's type. But should they update their priors on Brown because Jones has deviated? Especially, should Jones update his beliefs, just because he himself deviated? Passive conjectures seems much more reasonable.


If, to take a second possibility, Brown himself does deviate, is it reasonable for the out-of-equilibrium beliefs to specify that Smith and Jones update their beliefs about Brown in different ways? This seems dubious in light of the Harsanyi doctrine that everyone begins with the same priors.

On the other hand, consider a tremble interpretation of out-of-equilibrium moves. Maybe if Jones trembles and picks the wrong strategy, that really does say something about Brown's type. Jones might tremble more often, for example, if Brown's type is strong than if it is weak. Jones himself might learn from his own trembles. Once we are in the realm of non-Bayesian beliefs, it is hard to know what to do without a real-world context.

- For discussions of the appropriateness of different equilibrium concepts in actual economic models see Rubinstein (1985b) on bargaining, Shleifer & Vishny (1986) on greenmail and D. Hirshleifer & Titman (1990) on tender offers.
- **Exotic refinements.** Binmore (1990) and Kreps (1990b) are booklength treatments of rationality and equilibrium concepts.
- **The Beer-Quiche Game** of Cho & Kreps (1987). To illustrate their “intuitive criterion”, Cho and Kreps use the Beer-Quiche Game. In this game, Player I might be either weak or strong in his duelling ability, but he wishes to avoid a duel even if he thinks he can win. Player II wishes to fight a duel only if player I is weak, which has a probability of 0.1. Player II does not know player I's type, but he observes what player I has for breakfast. He knows that weak players prefer quiche for breakfast, while strong players prefer beer. The payoffs are shown in Figure 5.

Figure 5 illustrates a few twists on how to draw an extensive form. It begins with Nature's choice of *Strong* or *Weak* in the middle of the diagram. Player I then chooses whether to breakfast on *beer* or *quiche*. Player II's nodes are connected by a dotted line if they are in the same information set. Player II chooses *Duel* or *Don't*, and payoffs are then received.

Figure 5 The Beer-Quiche Game



This game has two perfect Bayesian equilibrium outcomes, both of which are pooling. In E_1 , player I has beer for breakfast regardless of type, and Player II chooses not to duel. This is supported by the out-of-equilibrium belief that a quiche-eating player I is weak with probability over 0.5, in which case player II would choose to duel on observing quiche. In E_2 , player I has quiche for breakfast regardless of type, and player II chooses not to duel. This is supported by the out-of-equilibrium belief that a beer-drinking player I is weak with probability greater than 0.5, in which case player II would choose to duel on observing beer. Passive conjectures and the intuitive criterion both rule out equilibrium E_2 . According to the reasoning of the intuitive criterion, player I could deviate without fear of a duel by giving the following convincing speech,

I am having beer for breakfast, which ought to convince you I am strong. The only conceivable benefit to me of breakfasting on beer comes if I am strong. I would never wish to have beer for breakfast if I were weak, but if I am strong and this message is convincing, then I benefit from having beer for breakfast.

N6.5 The Axelrod tournament

- Hofstadter (1983) is a nice discussion of the Prisoner's Dilemma and the Axelrod tournament by an intelligent computer scientist who came to the subject untouched by the preconceptions or training of economics. It is useful for elementary economics classes. Axelrod's 1984 book provides a fuller treatment.

Problems

6.1. Cournot Duopoly under Incomplete Information about Costs

This problem introduces incomplete information into the Cournot model of Chapter 3 and allows for a continuum of player types.

- 6.1a Modify the Cournot Game of Chapter 3 by specifying that Apex's average cost of production be c per unit, while Brydoux's remains zero. What are the outputs of each firm if the costs are common knowledge? What are the numerical values if $c = 10$?
- 6.1b Let Apex' cost c be c_{max} with probability θ and 0 with probability $1 - \theta$, so Apex is one of two types. Brydoux does not know Apex's type. What are the outputs of each firm?
- 6.1c Let Apex's cost c be drawn from the interval $[0, c_{max}]$ using the uniform distribution, so there is a continuum of types. Brydoux does not know Apex's type. What are the outputs of each firm?
- 6.1d Outputs were 40 for each firm in the zero-cost game in chapter 3. Check your answers in parts (b) and (c) by seeing what happens if $c_{max} = 0$.
- 6.1e Let $c_{max} = 20$ and $\theta = 0.5$, so the expectation of Apex's average cost is 10 in parts (a), (b), and (c). What are the average outputs for Apex in each case?
- 6.1f Modify the model of part (b) so that $c_{max} = 20$ and $\theta = 0.5$, but somehow $c = 30$. What outputs do your formulas from part (b) generate? Is there anything this could sensibly model?

Problem 6.2. Limit Pricing (see Milgrom and Roberts [1982a])

An incumbent firm operates in the local computer market, which is a natural monopoly in which only one firm can survive. The incumbent can price *Low*, losing 40 in profits, or *High*, losing nothing. It knows its own operating cost C , which is 20 with probability 0.75 and 30 with probability 0.25. A potential entrant knows those probabilities, but not the incumbent's exact cost. In the second period, The entrant can enter at a cost of 100, and its operating cost of 25 is common knowledge. If there are two firms in the market, each incurs a loss of 50, but one then drops out and the survivor earns the monopoly revenue of 200. There is no discounting; $r=0$

- 6.2a Draw the extensive form for this game.
- 6.2b Why is there no perfect equilibrium in which the incumbent prices *Low* only when its costs are 20? (No separating equilibrium)
- 6.2c In a perfect Bayesian equilibrium in which the incumbent prices *Low* regardless of its costs (a pooling equilibrium), what do out-of-equilibrium beliefs have to be specified?
- 6.2d What are two different perfect Bayesian equilibria for this game?
- 6.2e What is a set of out-of-equilibrium beliefs that do not support a pooling equilibrium at a *Low* price?

6.3. Symmetric Information and Prior Beliefs

In the Expensive-Talk Game of Table 1, the Battle of the Sexes is preceded by a communication move in which the man chooses *Silence* or *Talk*. *Talk* costs 1 payoff unit, and consists of a declaration by the man that he is going to the prize fight. This declaration is just talk; it is not binding on him.

Table 1: Subgame payoffs in the Expensive-Talk Game

		Woman	
		<i>Fight</i>	<i>Ballet</i>
Man:	<i>Fight</i>	3,1	0,0
	<i>Ballet</i>	0,0	1,3

Payoffs to: (Man, Woman)

- (6.3a) Draw the extensive form for this game, putting the man's move first in the simultaneous-move subgame.
- (6.3b) What are the strategy sets for the game? (Start with the woman's.)
- 6.3c What are the three perfect pure-strategy equilibrium outcomes in terms of observed actions? (Remember: strategies are not the same thing as outcomes.)
- 6.3d Describe the equilibrium strategies for a perfect equilibrium in which the man chooses to talk.
- 6.3e The idea of "forward induction" says that an equilibrium should remain an equilibrium even if strategies dominated in that equilibrium are removed from the game and the procedure is iterated. Show that this procedure rules out SBB as an equilibrium outcome. (See Van Damme [1989]. In fact, this procedure rules out TFF (*Talk, Fight, Fight*) also.)

6.4. Lack of common knowledge

This problem looks at what happens if the parameter values in Entry Deterrence V are changed.

- 6.4a If the value for the belief, $Pr(Strong|Enter, Nature\ said\ nothing)$, were .05 or .95, would such beliefs support the equilibrium in section 6.3?
- 6.4b Why is the equilibrium in section 6.3 not an equilibrium if 0.7 is the probability that Nature tells the incumbent?
- 6.4c Describe the equilibrium if 0.7 is the probability that Nature tells the incumbent. For what out-of-equilibrium beliefs does this remain the equilibrium?

September 4, 1999. February 3, 2000. February 6, 2000. November 29, 2003. Eric Rasmusen, Erasmuse@indiana.edu. [Http: //mypage.iu.edu/~erasmuse](http://mypage.iu.edu/~erasmuse). Footnotes starting with xxx are the author's notes to himself. Comments welcomed.

Part II Asymmetric Information

7 Moral Hazard: Hidden Actions

7.1 Categories of Asymmetric Information Models

It used to be that the economist's generic answer to someone who brought up peculiar behavior which seemed to contradict basic theory was "It must be some kind of price discrimination." Today, we have a new answer: "It must be some kind of asymmetric information." In a game of asymmetric information, player Smith knows something that player Jones does not. This covers a broad range of models (including price discrimination nowadays), so perhaps it is not surprising that so many situations come under its rubric. We will divide games of asymmetric information into five categories, to be studied in four chapters.

1 Moral hazard with hidden actions (Chapters 7 and 8).

Smith and Jones begin with symmetric information and agree to a contract, but then Smith takes an action unobserved by Jones. Information is complete.

2 Adverse selection (Chapter 9).

Nature begins the game by choosing Smith's type (his payoff and strategies), unobserved by Jones. Smith and Jones then agree to a contract. Information is incomplete.

3 Mechanism Design in Adverse Selection and in Moral Hazard with Hidden Information (Chapter 10).

Jones is designing a contract for Smith designed to elicit Smith's private information. This may happen under adverse selection— in which case Smith knows the information prior to contracting— or moral hazard with hidden information— in which case Smith will learn it after contracting.

4,5 Signalling and Screening (Chapter 11).

Nature begins the game by choosing Smith's type, unobserved by Jones. To demonstrate his type, Smith takes actions that Jones can observe. If Smith takes the action before they agree to a contract, he is signalling; if he takes it afterwards, he is being screened. Information is incomplete.

Signalling and screening are special cases of adverse selection, which is itself a situation of hidden knowledge. Information is complete in either kind of moral hazard, and incomplete in adverse selection, signalling, and screening.

Note that some people may say that information *becomes* incomplete in a model of moral hazard with hidden knowledge, even though it is complete at the start of the game. That statement runs contrary to the definition of complete information in Chapter 2, however. The most important distinctions to keep in mind are whether or not the players agree to a contract before or after information becomes asymmetric and whether their own actions are common knowledge.

We will make heavy use of the principal-agent model to analyze asymmetric information. Usually this term is applied to moral hazard models, since the problems studied in the law of agency usually involve an employee who disobeys orders by choosing the wrong actions, but the paradigm will be useful in all of these contexts. The two players are the

principal and the agent, who are usually representative individuals. The principal hires an agent to perform a task, and the agent acquires an informational advantage about his type, his actions, or the outside world at some point in the game. It is usually assumed that the players can make a binding **contract** at some point in the game, which is to say that the principal can commit to paying the agent an agreed sum if he observes a certain outcome. In the implicit background of such models are courts which will punish any player who breaks a contract in a way that can be proven with public information.

*The **principal** (or **uninformed player**) is the player who has the coarser information partition.*

*The **agent** (or **informed player**) is the player who has the finer information partition.*

Figure 1: Categories of Asymmetric Information Models

Figure 1 shows the game trees for five principal-agent models. In each model, the principal (P) offers the agent (A) a contract, which he accepts or rejects. In some, Nature (N) makes a move or the agent chooses an effort level, message, or signal. The moral hazard models are games of complete information with uncertainty. The principal offers a contract, and after the agent accepts, Nature adds noise to the task being performed. In moral hazard with hidden actions, (a) in Figure 1, the agent moves before Nature and in moral hazard with hidden knowledge, (b) in Figure 1, the agent moves after Nature and conveys a “message” to the principal about Nature’s move.

Adverse selection models have incomplete information, so Nature moves first and picks the type of the agent, generally on the basis of his ability to perform the task. In the simplest model, Figure 1(c), the agent simply accepts or rejects the contract. If the agent can send a “signal” to the principal, as in Figures 1(d) and 1(e), the model is signalling if he sends the signal before the principal offers a contract, and is screening otherwise. A “signal” is different from a “message” because it is not a costless statement, but a costly action. Some adverse selection models contain uncertainty and some do not.

A problem we will consider in detail is that of an employer (the principal) hiring a worker (the agent). If the employer knows the worker's ability but not his effort level, the problem is moral hazard with hidden actions. If neither player knows the worker's ability at first, but the worker discovers it once he starts working, the problem is moral hazard with hidden knowledge. If the worker knows his ability from the start, but the employer does not, the problem is adverse selection. If, in addition to the worker knowing his ability from the start, he can acquire credentials before he makes a contract with the employer, the problem is signalling. If the worker acquires his credentials in response to a wage offer made by the employer, the problem is screening.

The five categories are only gradually rising from the swirl of the literature on agency models, and the definitions are not well established. In particular, some would argue that what I have called moral hazard with hidden knowledge and screening are essentially the same as adverse selection. Myerson (1991, p. 263), for example, suggests calling the problem of players taking the wrong action "moral hazard" and the problem of misreporting information "adverse selection." Many economists do not realize that screening and signalling are different and use the terms interchangeably. "Signal" is such a useful word that it is often used simply to indicate any variable conveying information. Most people have not thought very hard about any of the definitions, but the importance of the distinctions will become clear as we explore the properties of the models. For readers whose minds are more synthetic than analytic, Table 1 may be as helpful as anything in clarifying the categories.

Table 1: Applications of the Principal-Agent Model

	Principal	Agent	Effort or type and signal
Moral hazard with hidden actions	Insurance company Insurance company Plantation owner Bondholders Tenant Landlord Society	Policyholder Policyholder Sharecropper Stockholders Landlord Tenant Criminal	Care to avoid theft Drinking and smoking Farming effort Riskiness of corporate projects Upkeep of the building Upkeep of the building Number of robberies
Moral hazard with hidden knowledge	Shareholders FDIC	Company president Bank	Investment decision Safety of loans
Adverse selection	Insurance company Employer	Policyholder Worker	Infection with HIV virus Skill
Signalling and screening	Employer Buyer Investor	Worker Seller Stock issuer	Skill and education Durability and warranty Stock value and percentage retained

Section 7.2 discusses the roles of uncertainty and asymmetric information in a principal-

agent model of moral hazard with hidden actions, called the Production Game, and Section 7.3 shows how various constraints are satisfied in equilibrium. Section 7.4 collects several unusual contracts produced under moral hazard and discusses the properties of optimal contracts using the example of the Broadway Game.

7.2 A Principal-Agent Model: The Production Game

In the archetypal principal-agent model, the principal is a manager and the agent a worker. In this section we will devise a series of these games, the last of which will be the standard principal-agent model.

Denote the monetary value of output by $q(e)$, which is increasing in effort, e . The agent's utility function $U(e, w)$ is decreasing in effort and increasing in the wage, w , while the principal's utility $V(q - w)$ is increasing in the difference between output and the wage.

The Production Game

Players

The principal and the agent.

The order of play

- 1 The principal offers the agent a wage w .
- 2 The agent decides whether to accept or reject the contract.
- 3 If the agent accepts, he exerts effort e .
- 4 Output equals $q(e)$, where $q' > 0$.

Payoffs

If the agent rejects the contract, then $\pi_{agent} = \bar{U}$ and $\pi_{principal} = 0$.

If the agent accepts the contract, then $\pi_{agent} = U(e, w)$ and $\pi_{principal} = V(q - w)$.

An assumption common to most principal-agent models is that either the principal or the agent is one of many perfect competitors. In the background, other principals compete to employ the agent, so the principal's equilibrium profit equals zero; or many agents compete to work for the principal, so the agent's equilibrium utility equals the minimum for which he will accept the job, called the **reservation utility**, $\bar{U}$. There is some reservation utility level even if the principal is a monopolist, however, because the agent has the option of remaining unemployed if the wage is too low.

One way of viewing the assumption in the Production Game that the principal moves first is that many agents compete for one principal. The order of moves allows the principal to make a take-it-or-leave-it offer, leaving the agent with as little bargaining room as if he had to compete with a multitude of other agents. This is really just a modelling convenience,

however, since the agent's reservation utility, $\bar{U}$, can be set at the level a principal would have to pay the agent in competition with other principals. This level of $\bar{U}$ can even be calculated, since it is the level at which the principal's payoff from profit maximization using the optimal contract is driven down to the principal's reservation utility by competition with other principals. Here the principal's reservation utility is zero, but that too can be chosen to fit the situation being modelled. As in the game of Nuisance Suits in Section 4.3, the main concern in choosing who makes the offer is to avoid getting caught up in a bargaining subgame.

Refinements of the equilibrium concept will not be important in this chapter. Nash equilibrium will be sufficient, because information is complete and the concerns of perfect Bayesian equilibrium will not arise. Subgame perfectness will be required, since otherwise the agent might commit to reject any contract that does not give him all of the gains from trade, but it will not drive the important results.

We will go through a series of five versions of the Production Game in this chapter.

Production Game I: Full Information. In the first version of the game, every move is common knowledge and the contract is a function $w(e)$.

Finding the equilibrium involves finding the best possible contract from the point of view of the principal, given that he must make the contract acceptable to the agent and that he foresees how the agent will react to the contract's incentives. The principal must decide what he wants the agent to do and what incentive to give him to do it.

The agent must be paid some amount $\tilde{w}(e)$ to exert effort e , where $\tilde{w}(e)$ is defined to be the w that solves the participation constraint

$$U(e, w(e)) = \bar{U}. \quad (1)$$

Thus, the principal's problem is

$$\underset{e}{\text{Maximize}} \quad V(q(e) - \tilde{w}(e)) \quad (2)$$

The first-order condition for this problem is

$$V'(q(e) - \tilde{w}(e)) \left(\frac{\partial q}{\partial e} - \frac{\partial \tilde{w}}{\partial e} \right) = 0, \quad (3)$$

which implies that

$$\frac{\partial q}{\partial e} = \frac{\partial \tilde{w}}{\partial e}. \quad (4)$$

From the implicit function theorem (see section 13.4) and the participation constraint,

$$\frac{\partial \tilde{w}}{\partial e} = - \left(\frac{\frac{\partial U}{\partial e}}{\frac{\partial U}{\partial w}} \right). \quad (5)$$


Combining equations (4) and (5) yields

$$\frac{\partial U}{\partial \tilde{w}} \frac{\partial q}{\partial e} = - \frac{\partial U}{\partial e}. \quad (6)$$

Equation (6) says that at the optimal effort level, e^* , the marginal utility to the agent which would result if he kept all the marginal output from extra effort equals the marginal disutility to him of that effort.

Figure 2 shows this graphically. The agent has indifference curves in effort-wage space that slope upwards, since if his effort rises his wage must increase also to keep his utility the same. The principal's indifference curves also slope upwards, because although he does not care about effort directly, he does care about output, which rises with effort. The principal might be either risk averse or risk neutral; his indifference curve is concave rather than linear in either case because Figure 2 shows a technology with diminishing returns to effort. If effort starts out being higher, extra effort yields less additional output so the wage cannot rise as much without reducing profits.

Figure 2: The Efficient Effort Level in Production Game I



Under perfect competition among the principals the profits are zero, so the reservation utility, $\bar{U}$, will be at the level such that at the profit-maximizing effort e^* , $\tilde{w}(e^*) = q(e^*)$, or

$$U(e^*, q(e^*)) = \bar{U}. \quad (7)$$

The principal selects the point on the $U = \bar{U}$ indifference curve that maximizes his profits, at effort e^* and wage w^* . He must then design a contract that will induce the agent to choose this effort level. The following three contracts are equally effective under full information.

1 The **forcing contract** sets $w(e^*) = w^*$ and $w(e \neq e^*) = 0$. This is certainly a strong incentive for the agent to choose exactly $e = e^*$.

2 The **threshold contract** sets $w(e \geq e^*) = w^*$ and $w(e < e^*) = 0$. This can be viewed as a flat wage for low effort levels, equal to 0 in this contract, plus a bonus if effort reaches e^* . Since the agent dislikes effort, the agent will choose exactly $e = e^*$.

3 The **linear contract** shown in Figure 2 sets $w(e) = \alpha + \beta e$, where α and β are chosen so that $w^* = \alpha + \beta e^*$ and the contract line is tangent to the indifference curve $U = \bar{U}$ at e^* . The most northwesterly of the agent's indifference curves that touch this contract line touches it at e^* .

Let's now fit out Production Game I with specific functional forms. Suppose the agent exerts effort $e \in [0, \infty]$, and output equals $q(e) = 100 * \log(1 + e)$. If the agent rejects the contract, let $\pi_{agent} = \bar{U} = 3$ and $\pi_{principal} = 0$, whereas if the agent accepts the contract, let $\pi_{agent} = U(e, w) = \log(w) - e^2$ and $\pi_{principal} = q(e) - w(e)$.

The agent must be paid some amount $\tilde{w}(e)$ to exert effort e , where $\tilde{w}(e)$ is defined to solve the participation constraint,

$$U(e, w(e)) = \bar{U}, \quad \text{so } \log(\tilde{w}(e)) - e^2 = 3. \quad (8)$$

Knowing the particular functional form as we do, we can solve (8) for the wage function:

$$\tilde{w}(e) = \text{Exp}(3 + e^2). \quad (9)$$

This makes sense. As effort rises, the wage must rise to compensate, and rise more than exponentially if utility is to be kept equal to 3.

The principal's problem is

$$\underset{e}{\text{Maximize}} \quad V(q(e) - \tilde{w}(e)) = 100 * \log(1 + e) - \text{Exp}(3 + e^2) \quad (10)$$

The first order condition for this problem is

$$V'(q(e) - \tilde{w}(e)) \left(\frac{\partial q}{\partial e} - \frac{\partial \tilde{w}}{\partial e} \right) = 0, \quad (11)$$

or, for our problem, since the firm is risk-neutral and $V' = 1$,

$$\frac{100}{1 + e} - 2e(\text{Exp}(3 + e^2)) = 0, \quad (12)$$

We can simplify the first order condition a little to get


$$(2e + 2e^2)\text{Exp}(3 + e^2) = 100, \quad (13)$$

but this cannot be solved analytically. Using the computer program Mathematica, I found that $e^* \approx .77$, from which, using the formulas above, we get $q^* \approx 100 * \log(1 + .77) \approx 57.26$ and $w^* \approx 36.50$.

Here, the implicit function theorem was not needed, because specifying the functional forms allowed us to find the solution using algebra instead.

Note that if $\bar{U}$ were high enough, the principal's payoff would be zero. If the market for agents were competitive, this is what would happen, since the agent's reservation payoff would be from working for another principal.

Figure 3: Three contracts that induce effort e^* for wage w^*



To implement the contract, a number of types of contracts could be used, as shown in Figure 3.

1 The **forcing contract** sets $w(e^*) = w^*$ and $w(e \neq .77) = 0$. Here, $w(.77) = 37$ (rounding up) and $w(e \neq e^*) = 0$.

2 The **threshold contract** sets $w(e \geq e^*) = w^*$ and $w(e < e^*) = 0$. Here, $w(e \geq .77) = 37$ and $w(e < .77) = 0$.

3 The **linear contract** sets $w(e) = \alpha + \beta e$, where α and β are chosen so that $w^* = \alpha + \beta e^*$ and the contract line is tangent to the indifference curve $U = \bar{U}$ at e^* . The slope of that indifference curve is the derivative of the $\tilde{w}(e)$ function, which is

$$\frac{\partial \tilde{w}(e)}{\partial e} = 2e * \text{Exp}(3 + e^2). \quad (14)$$

At $e^* = .77$, this takes the value 56. That is the β for the linear contract. The α must solve $w(e^*) = 37 = \alpha + 56(.77)$, so $\alpha = -7$.

You ought to be a little concerned as to whether the linear contract satisfies the incentive compatibility constraint. We constructed it so that it satisfied the participation constraint, because if the agent chooses $e = 0.77$, his utility will be 3. But might he prefer to choose some larger or smaller e and get even more utility? No, because his utility is concave. That makes the indifference curve convex, so its slope is always increasing and no preferable indifference curve touches the equilibrium contract line.

Before going on to versions of the game with asymmetric information, it will be useful to look at one other version of the game with full information, in which the agent, not the principal, proposes the contract. This will be called Production Game II.

Production Game II: Full Information. Agent Moves First.

In this version, every move is common knowledge and the contract is a function $w(e)$. The order of play, however, is now as follows

The order of play

- 1 The agent offers the principal a contract $w(e)$.
- 2 The principal decides whether to accept or reject the contract.
- 3 If the principal accepts, the agent exerts effort e .
- 4 Output equals $q(e)$, where $q' > 0$.

In this game, the agent has all the bargaining power, not the principal. The participation constraint is now that the principal must earn zero profits, so $q(e) - w(e) \geq 0$. The agent will maximize his own payoff by driving the principal to exactly zero profits, so $w(e) = q(e)$. Substituting $q(e)$ for $w(e)$ to account for the participation constraint, the maximization problem for the agent in proposing an effort level e at a wage $w(e)$ can therefore be written as

$$\underset{e}{\text{Maximize}} \quad U(e, q(e)) \quad (15)$$

The first-order condition is

$$\frac{\partial U}{\partial e} + \left(\frac{\partial U}{\partial q} \right) \left(\frac{\partial q}{\partial e} \right) = 0. \quad (16)$$

Since $\frac{\partial U}{\partial q} = \frac{\partial U}{\partial w}$ when the wages equals output, equation (16) implies that

$$\frac{\partial U}{\partial w} \frac{\partial q}{\partial e} = - \frac{\partial U}{\partial e}. \quad (17)$$

Comparing this with equation (6), the equation when the principal had the bargaining power, it is clear that e^* is identical in Production Games I and II. It does not matter who has the bargaining power; the efficient effort level stays the same.

Figure 2 (a few pages back) can be used to illustrate this game. Suppose that $V_1 = 0$. The agent must choose a point on the $V = 0$ indifference curve that maximizes his own utility, and then provide himself with contract incentives to choose that point. The agent's payoff is highest at effort e^* given that he must make $V = 0$, and all three contracts described in Production Game I provide him with the correct incentives.

The efficient-effort level is independent of which side has the bargaining power because the gains from efficient production are independent of how those gains are distributed so long as each party has no incentive to abandon the relationship. This is the same lesson as that of the Coase theorem, which says that under general conditions the activities undertaken will be efficient and independent of the distribution of property rights (Coase [1960]). This property of the efficient-effort level means that the modeller is free to make the assumptions on bargaining power that help to focus attention on the information problems he is studying.

Production Game III: A Flat Wage Under Certainty

In this version of the game, the principal can condition the wage neither on effort nor on output. This is modelled as a principal who observes neither effort nor output, so information is asymmetric.

It is easy to imagine a principal who cannot observe effort, but it seems very strange that he cannot observe output, especially since he can deduce the output from the value of his payoff. It is not ridiculous that he cannot base wages on output, however, because a contract must be enforceable by some third party such as a court. Law professors complain about economists who speak of “unenforceable contracts.” In law school, a contract is defined as an enforceable agreement, and most of a contracts class is devoted to discovering which agreements are contracts. A simple promise to give someone money without any obligation on his part, for example, is not something a court will enforce, nor is it possible in practice for a court to enforce a contract in which someone agrees to pay a barber \$50 “if the haircut is especially good” but \$10 otherwise. A court can only enforce contingencies it can observe. In the extreme, Production Game III is appropriate. Output is not **contractible** (the court will not enforce a contract) or **verifiable** (the court cannot observe output), which usually leads to the same outcome as when output is unobservable to the two parties to the agreement.

The outcome of Production Game III is simple and inefficient. If the wage is non-negative, the agent accepts the job and exerts zero effort, so the principal offers a wage of zero.

If there is nothing on which to condition the wage, the agency problem cannot be solved by designing the contract carefully. If it is to be solved at all, it will be by some other means such as reputation or repetition of the game, the solutions of Chapter 5. Typically, however, there is some contractible variable such as output upon which the principal can condition the wage. Such is the case in Production Game IV.

Production Game IV: an output-based wage under certainty

In this version, the principal cannot observe effort but can observe output and specify the contract to be $w(q)$.

Now the principal picks not a number w but a function $w(q)$. His problem is not quite so straightforward as in Production Game I, where he picked the function $w(e)$, but here, too, it is possible to achieve the efficient effort level e^* despite the unobservability of effort. The principal starts by finding the optimal effort level e^* , as in Production Game I. That effort yields the efficient output level $q^* = q(e^*)$. To give the agent the proper incentives, the contract must reward him when output is q^* . Again, a variety of contracts could be used. The forcing contract, for example, would be any wage function such that $U(e^*, w(q^*)) = \bar{U}$ and $U(e, w(q)) < \bar{U}$ for $e \neq e^*$.

Production Game IV shows that the unobservability of effort is not a problem in itself, if the contract can be conditioned on something which is observable and perfectly correlated with effort. The true agency problem occurs when that perfect correlation breaks down, as in Production Game V.

Production Game V: Output-Based Wage under Uncertainty.

In this version, the principal cannot observe effort but can observe output and specify the contract to be $w(q)$. Output, however, is a function $q(e, \theta)$ both of effort and the state

of the world $\theta \in \mathbf{R}$, which is chosen by Nature according to the probability density $f(\theta)$ as the new move (5) of the game. Move (5) comes just after the agent chooses effort, so the agent cannot choose a low effort knowing that Nature will take up the slack. (If the agent can observe Nature's move before his own, the game becomes moral hazard with hidden knowledge and hidden actions).

Because of the uncertainty about the state of the world, effort does not map cleanly onto the observed output in Production Game V. A given output might have been produced by any of several different effort levels, so a forcing contract will not necessarily achieve the desired effort. Unlike the case in Production Game IV, here the principal cannot deduce that $e = e^*$ from the fact that $q = q^*$. Moreover, even if the contract does induce the agent to choose e^* , if it does so by penalizing him heavily when $q \neq q^*$ it will be expensive for the principal. The agent's expected utility must be kept equal to $\bar{U}$ because of the participation constraint, and if the agent is sometimes paid a low wage because output happens not to equal q^* , he must be paid more when output does equal q^* to make up for it. If the agent is risk averse, this variability in his wage requires that his expected wage be higher than the w^* found earlier, because he must be compensated for the extra risk. There is a tradeoff between incentives and insurance against risk.

Moral hazard becomes a problem when $q(e)$ is not a one-to-one function because a single value of e might result in any of a number of values of q , depending on the value of θ . In this case the output function is not invertible; knowing q , the principal cannot deduce the value of e perfectly without assuming equilibrium behavior on the part of the agent.

The combination of unobservable effort and lack of invertibility in Production Game V means that no contract can induce the agent to put forth the efficient effort level without incurring extra costs, which usually take the form of an extra risk imposed on the agent. We will still try to find a contract that is efficient in the sense of maximizing welfare given the informational constraints. The terms "first-best" and "second-best" are used to distinguish these two kinds of optimality.

*A **first-best contract** achieves the same allocation as the contract that is optimal when the principal and the agent have the same information set and all variables are contractible.*

*A **second-best contract** is Pareto optimal given information asymmetry and constraints on writing contracts.*

The difference in welfare between the first-best world and the second-best world is the cost of the agency problem.

The first four production games were easier because the principal could find a first-best contract without searching very far. But even defining the strategy space in a game like Production Game V is tricky, because the principal may wish to choose a very complicated function $w(q)$. Finding the optimal contract when a forcing contract cannot be used becomes a difficult problem without general answers, because of the tremendous variety of possible contracts. The rest of the chapter will show how the problem may at least be approached, if not actually solved.

7.3 The Incentive Compatibility, Participation, and Competition Constraints

The principal's objective in Production Game V is to maximize his utility knowing that the agent is free to reject the contract entirely and that the contract must give the agent an incentive to choose the desired effort. These two constraints arise in every moral hazard problem, and they are named the **participation constraint** and the **incentive compatibility constraint**. Mathematically, the principal's problem is

$$\begin{aligned} \text{Maximize} \quad & EV(q(\tilde{e}, \theta) - w(q(\tilde{e}, \theta))) \\ & w(\cdot) \end{aligned} \tag{18}$$

subject to

$$\tilde{e} = \underset{e}{\operatorname{argmax}} EU(e, w(q(e, \theta))) \quad (\text{incentive compatibility constraint}) \tag{18a}$$

$$EU(\tilde{e}, w(q(\tilde{e}, \theta))) \geq \bar{U} \quad (\text{participation constraint}) \tag{18b}$$

The incentive-compatibility constraint takes account of the fact that the agent moves second, so the contract must induce him to voluntarily pick the desired effort. The participation constraint, also called the **reservation utility** or **individual rationality** constraint, requires that the worker prefer the contract to leisure, home production, or alternative jobs.

Expression (18) is the way an economist instinctively sets up the problem, but setting it up is often as far as he can get with the **first-order condition approach**. The difficulty is not just that the maximizer is choosing a wage function instead of a number, because control theory or the calculus of variations can solve such problems. Rather, it is that the constraints are nonconvex—they do not rule out a nice convex set of points in the space of wage functions such as the constraint “ $w \geq 4$ ” would, but rather rule out a very complicated set of possible wage functions.

A different approach, developed by Grossman & Hart (1983) and called the **three-step procedure** by Fudenberg & Tirole (1991a), is to focus on contracts that induce the agent to pick a particular action rather than to directly attack the problem of maximizing profits. The first step is to find for each possible effort level the set of wage contracts that induce the agent to choose that effort level. The second step is to find the contract which supports that effort level at the lowest cost to the principal. The third step is to choose the effort level that maximizes profits, given the necessity to support that effort with the costly wage contract from the second step.

To support the effort level e , the wage contract $w(\cdot)$ must satisfy the incentive compatibility and participation constraints. Mathematically, the problem of finding the least cost $C(\tilde{e})$ of supporting the effort level $\tilde{e}$ combines steps one and two.

$$\begin{aligned} C(\tilde{e}) = \quad & \text{Minimum} \quad Ew(q(\tilde{e}, \theta)) \\ & w(\cdot) \end{aligned} \tag{19}$$

subject to constraints (18a) and (18b).

Step three takes the principal's problem of maximizing his payoff, expression (18), and restates it as

$$\underset{\tilde{e}}{\text{Maximize}} \quad EV(q(\tilde{e}, \theta) - C(\tilde{e})). \quad (20)$$

After finding which contract most cheaply induces each effort, the principal discovers the optimal effort by solving problem (20).

Breaking the problem into parts makes it easier to solve. Perhaps the most important lesson of the three-step procedure, however, is to reinforce the points that the goal of the contract is to induce the agent to choose a particular effort level and that asymmetric information increases the cost of the inducements.

7.4 Optimal Contracts: The Broadway Game

The next game, inspired by Mel Brooks's offbeat film *The Producers*, illustrates a peculiarity of optimal contracts: sometimes the agent's reward should not increase with his output. Investors advance funds to the producer of a Broadway show that might succeed or might fail. The producer has the choice of embezzling or not embezzling the funds advanced to him, with a direct gain to himself of 50 if he embezzles. If the show is a success, the revenue is 500 if he did not embezzle and 100 if he did. If the show is a failure, revenue is -100 in either case, because extra expenditure on a fundamentally flawed show is useless.

Broadway Game I

Players

Producer and investors.

The order of play

- 1 The investors offer a wage contract $w(q)$ as a function of revenue q .
- 2 The producer accepts or rejects the contract.
- 3 The producer chooses to *Embezzle* or *Do not embezzle*.
- 4 Nature picks the state of the world to be *Success* or *Failure* with equal probability. Table 2 shows the resulting revenue q .

Payoffs

The producer is risk averse and the investors are risk neutral. The producer's payoff is $U(100)$ if he rejects the contract, where $U' > 0$ and $U'' < 0$, and the investors' payoff is 0. Otherwise,

$$\pi_{\text{producer}} = \begin{cases} U(w(q) + 50) & \text{if he embezzles} \\ U(w(q)) & \text{if he is honest} \end{cases}$$

$$\pi_{\text{investors}} = q - w(q)$$

Table 2: Profits in Broadway Game I

		State of the World	
		<i>Failure</i> (0.5)	<i>Success</i> (0.5)
Effort	<i>Embezzle</i>	−100	+100
	<i>Do not embezzle</i>	−100	+500

Another way to tabulate outputs, shown in Table 3, is to put the probabilities of outcomes in the boxes, with effort in the rows and output in the columns.

Table 3: Probabilities of Profits in Broadway Game I

		Profit			
		−100	+100	+500	Total
Effort	<i>Embezzle</i>	0.5	0.5	0	1
	<i>Do not embezzle</i>	0.5	0	0.5	1

The investors will observe q to equal either -100 , $+100$, or $+500$, so the producer's contract will specify at most three different wages: $w(-100)$, $w(+100)$, and $w(+500)$. The producer's expected payoffs from his two possible actions are

$$\pi(\textit{Do not embezzle}) = 0.5U(w(-100)) + 0.5U(w(+500)) \quad (21)$$

and

$$\pi(\textit{Embezzle}) = 0.5U(w(-100) + 50) + 0.5U(w(+100) + 50). \quad (22)$$

The incentive compatibility constraint is $\pi(\textit{Do not embezzle}) \geq \pi(\textit{Embezzle})$, so

$$0.5U(w(-100)) + 0.5U(w(+500)) \geq 0.5U(w(-100) + 50) + 0.5U(w(+100) + 50), \quad (23)$$

and the participation constraint is

$$\pi(\textit{Do not embezzle}) = 0.5U(w(-100)) + 0.5U(w(+500)) \geq U(100). \quad (24)$$

The investors want the participation constraint (24) to be satisfied at as low a dollar cost as possible. This means they want to impose as little risk on the producer as possible, since he requires a higher expected value for his wage if the risk is higher. Ideally, $w(-100) = w(+500)$, which provides full insurance. The usual agency tradeoff is between smoothing out the agent's wage and providing him with incentives. Here, no tradeoff is required, because of a special feature of the problem: there exists an outcome that could not occur unless the producer chooses the undesirable action. That outcome is $q = +100$, and it means that the following **boiling-in-oil contract** provides both riskless wages and effective incentives.

$$\begin{aligned} w(+500) &= 100 \\ w(-100) &= 100 \\ w(+100) &= -\infty \end{aligned}$$

Under this contract, the producer's wage is a flat 100 when he does not embezzle, so the participation constraint is satisfied. It is also binding, because it is satisfied as an equality, and the investors would have a higher payoff if the constraint were relaxed. If the producer does embezzle, he faces the payoff of $-\infty$ with probability 0.5, so the incentive compatibility constraint is satisfied. It is nonbinding, because it is satisfied as a strong inequality and the investors' equilibrium payoff does not fall if the constraint is tightened a little by making the producer's earnings from embezzlement slightly higher. Note that the cost of the contract to the investors is 100 in equilibrium, so that their overall expected payoff is $0.5(-100) + 0.5(+500) - 100 = 100$, which is greater than zero and thus gives the investors enough return to be willing to back the show.

The boiling-in-oil contract is an application of the **sufficient statistic condition**, which states that for incentive purposes, if the agent's utility function is separable in effort and money, wages should be based on whatever evidence best indicates effort, and only incidentally on output (see Holmstrom [1979] and note N7.2). In the spirit of the three-step procedure, what the principal wants is to induce the agent to choose the appropriate effort, *Do not embezzle*, and his data on what the agent chose is the output. In equilibrium (though not out of it), the datum $q = +500$ contains exactly the same information as the datum $q = -100$. Both lead to the same posterior probability that the agent chose *Do not embezzle*, so the wages conditioned on each datum should be the same. We need to insert the qualifier "in equilibrium," because to form the posterior probabilities the principal needs to have some beliefs as to the agent's behavior. Otherwise, the principal could not interpret $q = -100$ at all.

Milder contracts than this would also be effective. Two wages will be used in equilibrium, a low wage $\underline{w}$ for an output of $q = 100$ and a high wage $\bar{w}$ for any other output. The participation and incentive compatibility constraints provide two equations to solve for these two unknowns. To find the mildest possible contract, the modeller must also specify a function for $U(w)$ which, interestingly enough, was unnecessary for finding the first boiling-in-oil contract. Let us specify that

$$U(w) = 100w - 0.1w^2. \quad (25)$$

A quadratic utility function like this is only increasing if its argument is not too large, but since the wage will not exceed $w = 1000$, it is a reasonable utility function for this model. Substituting (25) into the participation constraint (24) and solving for the full-insurance high wage $\bar{w} = w(-100) = w(+500)$ yields $\bar{w} = 100$ and a reservation utility of 9000. Substituting into the incentive compatibility constraint, (23), yields

$$9000 \geq 0.5U(100 + 50) + 0.5U(\underline{w} + 50). \quad (26)$$

When (26) is solved using the quadratic equation, it yields (with rounding error), $\underline{w} \leq 5.6$. A low wage of $-\infty$ is far more severe than what is needed.

If both the producer and the investors were risk averse, risk sharing would change the part of the contract that applied in equilibrium. The optimal contract would then provide for $w(-100) < w(+500)$ to share the risk. The principal would have a lower marginal utility of wealth when output was +500, so he would be better able to pay an extra dollar of wages in that state than when output was -100.

One of the oddities of Broadway Game I is that the wage is higher for an output of -100 than for an output of $+100$. This illustrates the idea that the principal's aim is to reward not output, but input. If the principal pays more simply because output is higher, he is rewarding Nature, not the agent. People usually believe that higher pay for higher output is "fair," but Broadway Game I shows that this ethical view is too simple. Higher effort usually leads to higher output, so higher pay is usually a good incentive, but this is not invariably true.

The decoupling of reward and result has broad applications. Becker (1968) in criminal law and Polinsky & Che (1991) in tort law note that if society's objective is to keep the amount of enforcement costs and harmful behavior low, the penalty applied should not simply be matched to the harm. Very high penalties that are seldom inflicted will provide the proper incentives and keep enforcement costs low, even though a few unlucky offenders will receive penalties all out of proportion to the harm they have caused.

A less gaudy name for a boiling-in-oil contract is the alliterative "**shifting support scheme**," so named because the contract depends on the support of the output distribution being different when effort is optimal than when effort is other than optimal. Put more simply, the set of possible outcomes under optimal effort must be different from the set of possible outcomes under any other effort level. As a result, certain outputs show without doubt that the producer embezzled. Very heavy punishments inflicted only for those outputs achieve the first-best because a non-embezzling producer has nothing to fear.

Figure 4: Shifting Supports in an Agency Model




Figure 4 shows shifting supports in a model where output can take not three, but a continuum of values. If the agent shirks instead of working, certain low outputs become possible and certain high outputs become impossible. In a case like this, where the support of the output shifts when behavior changes, boiling-in-oil contracts are useful: the wage is $-\infty$ for the low outputs possible only under shirking. In Figure 4b, on the other hand, the support just shrinks under shirking, so boiling in oil is appropriate. When there is a limit to the amount the agent can be punished, or the support is the same under all actions, the threat of boiling-in-oil might not achieve the first-best contract, but similar contracts can still be used. The conditions favoring boiling-in-oil contracts are

1 The agent is not very risk averse.

2 There are outcomes with high probability under shirking that have low probability under optimal effort.

3 The agent can be severely punished.

4 It is credible that the principal will carry out the severe punishment.

Selling the Store

Another first-best contract that can sometimes be used is **selling the store**. Under this arrangement, the agent buys the entire output for a flat fee paid to the principal, becoming the **residual claimant**, since he keeps every additional dollar of output that his extra effort produces. This is equivalent to fully insuring the principal, since his payoff becomes independent of the moves of the agent and of Nature.

In Broadway Game I, selling the store takes the form of the producer paying the investors 100 ($= 0.5[-100] + 0.5[+500] - 100$) and keeping all the profits for himself. The drawbacks are that (1) the producer might not be able to afford to pay the investors the flat price of 100; and (2) the producer might be risk averse and incur a heavy utility cost in bearing the entire risk. These two drawbacks are why producers go to investors in the first place.

Public Information that Hurts the Principal and the Agent

We can modify Broadway Game I to show how having more public information available can hurt both players. This will also provide a little practice in using information sets. Let us split *Success* into two states of the world, *Minor Success* and *Major Success*, which have probabilities 0.3 and 0.2 as shown in Table 4.

Table 4: Profits in Broadway Game II

		State of the World		
		<i>Failure</i> (0.5)	<i>Minor Success</i> (0.3)	<i>Major Success</i> (0.2)
Effort	<i>Embezzle</i>	−100	−100	+400
	<i>Do not embezzle</i>	−100	+450	+575

Under the optimal contract,

$$w(-100) = w(+450) = w(+575) > w(+400) + 50. \quad (27)$$

This is so because the producer is risk averse and only the datum $q = +400$ is proof that the producer embezzled. The optimal contract must do two things: deter embezzlement and pay the producer as predictable a wage as possible. For predictability, the wage is made constant unless $q = +400$. To deter embezzlement, the producer must be punished

if $q = +400$. As in Broadway Game I, the punishment would not have to be infinitely severe, and the minimum effective punishment could be calculated in the same way as in that game. The investors would pay the producer a wage of 100 in equilibrium and their expected payoff would be 100 ($= 0.5(-100) + 0.3(450) + 0.2(575) - 100$). Thus, a contract can be found for Broadway Game II such that the agent would not embezzle.

But consider what happens when the information set is refined so that before the agent takes his action both he and the principal can tell whether the show will be a major success or not. Let us call this Broadway Game III. Under the refinement, each player's initial information partition is

$$(\{Failure, Minor Success\}, \{Major Success\}),$$

instead of the original coarse partition

$$(\{Failure, Minor Success, Major Success\}).$$

If the information sets were refined all the way to singletons, this would be very useful to the investors because they could abstain from investing in a failure and they could easily determine whether the producer embezzled or not. As it is, however, the refinement does not help the investors decide when to finance the show. If they could still hire the producer and prevent him from embezzling at a cost of 100, the payoff from investing in a major success would be 475 ($= 575 - 100$). But the payoff from investing in a show given the information set $\{Failure, Minor Success\}$ would be about 6.25, which is still positive ($(\frac{0.5}{0.5+0.3})(-100) + (\frac{0.3}{0.5+0.3})(450) - 100$). So the improvement in information is no help with respect to the decision of when to invest.

Although the refinement has no direct effect on the efficiency of investment, it ruins the producer's incentives. If he observes $\{Failure, Minor Success\}$, he is free to embezzle without fear of the oil-boiling output of +400. He would still refrain from embezzling if he observed $\{Major Success\}$, but no contract that does not impose risk on a nonembezzling producer can stop him from embezzling if he observes $\{Failure, Minor Success\}$. Whether a risky contract can be found that would prevent the producer from embezzling at a cost of less than 6.25 to the investors depends on the producer's risk aversion. If he is very risk averse, the cost of the incentive is more than 6.25, and the investors will give up investing in shows that might be minor successes. Better information reduces welfare, because it increases the producer's temptation to misbehave.

N7.1 Categories of asymmetric information models

- The separation of asymmetric information into hidden actions and hidden knowledge is suggested in Arrow (1985) and commented upon in Hart & Holmstrom (1987). The term “hidden knowledge” seems to have become more popular than “hidden information,” which I used in the first edition.
- Empirical work on agency problems includes Joskow (1985, 1987) on coal mining, Masten & Crocker (1985) on natural gas contracts, Monteverde & Teece (1982) on auto components, Murphy (1986) on executive compensation, Rasmusen (1988b) on the mutual organization in banking, Staten & Umbeck (1982) on air traffic controllers and disability payments, and Wolfson (1985) on the reputation of partners in oil-drilling.
- A large literature of nonmathematical theoretical papers looks at organizational structure in the light of the agency problem. See Alchian & Demsetz (1972), Fama (1980), and Klein, Crawford, & Alchian (1978). Milgrom & Roberts (1992) have written a book on organization theory that describes what has been learned about the principal-agent problem at a technical level that MBA students can understand. There may be much to be learned from intelligent economists of the past also; note that part III, chapter 8, section 12 of Pigou’s *Economics of Welfare* (1932/1920) has an interesting discussion of the advantage of piece-rate work, which can more easily induce each worker to choose the correct effort when abilities differ as well as efforts.
- For examples of agency problems, see “Many Companies Now Base Workers’ Raises on Their Productivity,” *Wall Street Journal*, November 15, 1985, pp. 1, 15; “Big Executive Bonuses Now Come with a Catch: Lots of Criticism,” *Wall Street Journal*, May 15, 1985, p. 33; “Bribery of Retail Buyers is Called Pervasive,” *Wall Street Journal*, April 1, 1985, p. 6; “Some Employers Get Tough on Use of Air-Travel Prizes,” *Wall Street Journal*, March 22, 1985, p. 27.
- We have lots of “prinsipuls” in economics. I find this paradigm helpful for remembering spelling: “The principal’s principal principle was to preserve his principal.”
- “Principal” and “Agent” are legal terms, and agency is an important area of the law. Economists have focussed on quite different questions than lawyers. Economists focus on effort: how the principal induces the agent to do things. Lawyers focus on malfeasance and third parties: how the principal stops the agent from doing the wrong things and who bears the burden if he fails. If, for example, the manager of a tavern enters into a supply contract against the express command of the owner, who must be disappointed—the owner or the third party supplier?
- *Double-sided moral hazard.* The text described one-sided moral hazard. Moral hazard can also be double-sided, as when each player takes actions unobservable by the other that affect the payoffs of both of them. An example is tort negligence by both plaintiff and defendant: if a careless auto driver hits a careless pedestrian, and they go to law, the court must try to allocate blame, and the legislature must try to set up laws to induce the proper amount of care. Landlords and tenants also face double moral hazard, as implied in Table 1.
- A common convention in principal-agent models is to make one player male and the other female, so that “his” and “her” can be used to distinguish between them. I find this

distracting, since gender is irrelevant to most models and adds one more detail for the reader to keep track of. If readers naturally thought “male” when they saw “principal,” this would not be a problem— but they do not.

N7.2 A principal-agent model: the Production Game

- In Production Game III, we could make the agent’s utility depend on the state of the world as well as on effort and wages. Little would change from the simpler model.
- The model in the text uses “effort” as the action taken by the agent, but effort is used to represent a variety of real-world actions. The cost of pilferage by employees is an estimated \$8 billion a year in the USA. Employers have offered rewards for detection, one even offering the option of a year of twice-weekly lottery tickets instead of a lump sum. The Chicago department store Marshall Field’s, with 14,000 workers, in one year gave out 170 rewards of \$500 each, catching almost 500 dishonest employees. (“Hotlines and Hefty Rewards: Retailers Step Up Efforts to Curb Employee Theft,” *Wall Street Journal*, September 17, 1987, p. 37.)

For an illustration of the variety of kinds of “low effort,” see “Hermann Hospital Estate, Founded for the Poor, has Benefited the Wealthy, Investigators Allege,” *Wall Street Journal*, March 13, 1985, p. 4, which describes such forms of misbehavior as pleasure trips on company funds, high salaries, contracts for redecorating awarded to girlfriends, phony checks, kicking back real estate commissions, and investing in friendly companies. Nonprofit enterprises, often lacking both principles and principals, are especially vulnerable, as are governments, for the same reason.

- The Production Game assumes that the agent dislikes effort. Is this realistic? People differ. My father tells of his experience in the navy when the sailors were kept busy by being ordered to scrape loose paint. My father found it a way to pass the time but says that other sailors would stop chipping when they were not watched, preferring to stare into space. *De gustibus non est disputandum* (“About tastes there can be no arguing”). But even if effort has positive marginal utility at low levels, it has negative marginal utility at high enough levels— including, perhaps, at the efficient level. This is as true for professors as for sailors.
- Suppose that the principal does not observe the variable θ (which might be effort), but he does observe t and x (which might be output and profits). From Holmstrom (1979) and Shavell (1979) we have, restated in my words,

The Sufficient Statistic Condition. *If t is a sufficient statistic for θ relative to x , then the optimal contract needs to be based only on t if both principal and agent have separable utility functions.*

*The variable t is a **sufficient statistic** for θ relative to x if, for all t and x ,*

$$Prob(\theta|t, x) = Prob(\theta|t). \quad (28)$$

This implies, from Bayes’ Rule, that $Prob(t, x|\theta) = Prob(x|t)Prob(t|\theta)$; that is, x depends on θ only because x depends on t and t depends on θ .

The sufficient statistic condition is closely related to the Rao-Blackwell Theorem (see Cox & Hinkley [1974] p. 258), which says that the decision rule for nonstrategic decisions ought not to be random.

Gjesdal (1982) notes that if the utility functions are not separable, the theorem does not apply and randomized contracts may be optimal. Suppose there are two actions the agent might take. The principal prefers action X , which reduces the agent's risk aversion, to action Y , which increases it. The principal could offer a randomized wage contract, so the agent would choose action X and make himself less risk averse. This randomization is not a mixed strategy. The principal is not indifferent to high and low wages; he prefers to pay a low wage, but we allow him to commit to a random wage earlier in the game.

N7.3 The Incentive Compatibility, Participation, and Competition Constraints

- Discussions of the first-order condition approach can be found in Grossman & Hart (1983) and Hart & Holmstrom (1987).
- The term, “individual rationality constraint,” is perhaps more common, but “participation constraint” is more sensible. Since in modern modelling every constraint requires individuals to be rational, the name is ill-chosen.
- **Paying the agent more than his reservation wage.** If agents compete to work for principals, the participation constraint is binding whenever there are only two possible outcomes or whenever the agent's utility function is separable in effort and wages. Otherwise, it might happen that the principal picks a contract giving the agent more expected utility than is necessary to keep him from quitting. The reason is that the principal not only wants to keep the agent working, but to choose a high effort.
- If the distribution of output satisfies the **monotone likelihood ratio property** (MLRP), the optimal contract specifies higher pay for higher output. Let $f(q|e)$ be the probability density of output. The MLRP is satisfied if

$$\forall e' > e, \text{ and } q' > q, \quad f(q'|e')f(q|e) - f(q'|e)f(q|e') > 0, \quad (29)$$

or, in other words, if when $e' > e$, the ratio $f(q|e')/f(q|e)$ is increasing in q . Alternatively, f satisfies the MLRP if $q' > q$ implies that q' is a more favorable message than q in the sense of Milgrom (1981b). Less formally, the MLRP is satisfied if the ratio of the likelihood of a high effort to a low effort rises with observed output. The distributions in the Broadway Game of Section 7.4 violate the MLRP, but the normal, exponential, Poisson, uniform, and chi-square distributions all satisfy it. Stochastic dominance does not imply the MLRP. If effort of 0 produces outputs of 10 or 12 with equal probability, and effort of 1 produces outputs of 11 or 13 also with equal probability, the second distribution is stochastically dominant, but the MLRP is not satisfied.

- Finding general conditions that allow the modeller to characterize optimal contracts is difficult. Much of Grossman & Hart (1983) is devoted to the rather obscure Spanning Condition or Linear Distribution Function Condition (LDFC), under which the first order condition approach is valid. The survey by Hart & Holmstrom (1987) makes a valiant attempt at explaining the LDFC.

N7.4 Optimal Contracts: The Broadway Game

- Daniel Asquith suggested the idea behind Broadway Game II.
- Franchising is one compromise between selling the store and paying a flat wage. See Mathewson & Winter (1985), Rubin (1978), and Klein & Saft (1985).
- Mirrlees (1974) is an early reference on the idea of the boiling-in-oil contract.
- Broadway Game II shows that improved information could reduce welfare by increasing a player's incentive to misbehave. This is distinct from the nonstrategic insurance reason why improved information can be harmful. Suppose that Smith is insuring Jones against hail ruining Jones' wheat crop during the next year, increasing Jones' expected utility and giving a profit to Smith. If someone comes up with a way to forecast the weather before the insurance contract is agreed upon, both players will be hurt. Insurance will break down, because if it is known that hail will ruin the crop, Smith will not agree to share the loss, and if it is known there will be no hail, Jones will not pay a premium for insurance. Both players prefer not knowing the outcome in advance.

Problems

7.1: First-Best Solutions in a Principal-Agent Model

Suppose an agent has the utility function of $U = \sqrt{w} - e$, where e can assume the levels 0 or 1. Let the reservation utility level be $\bar{U} = 3$. The principal is risk neutral. Denote the agent's wage, conditioned on output, as $\underline{w}$ if output is 0 and $\bar{w}$ if output is 100. Table 5 shows the outputs.

Table 5: A Moral Hazard Game

Effort	Probability of Output of		Total
	0	100	
<i>Low</i> ($e = 0$)	0.3	0.7	1
<i>High</i> ($e = 1$)	0.1	0.9	1

- What would the agent's effort choice and utility be if he owned the firm?
- If agents are scarce and principals compete for them, what will the agent's contract be under full information? His utility?
- If principals are scarce and agents compete to work for them, what would the contract be under full information? What will the agent's utility and the principal's profit be in this situation?
- Suppose that $U = w - e$. If principals are the scarce factor and agents compete to work for principals, what would the contract be when the principal cannot observe effort? (Negative wages are allowed.) What will be the agent's utility and the principal's profit be in this situation?

7.2: The Principal-Agent Problem

Suppose the agent has a utility function of $U = \sqrt{w} - e$, where e can assume the levels 0 or 7, and a reservation utility of $\bar{U} = 4$. The principal is risk neutral. Denote the agent's wage, conditioned on output, as $\underline{w}$ if output is 0 and $\bar{w}$ if output is 1,000. Only the agent observes his effort. Principals compete for agents. Table 6 shows the output.

Table 6: Output from Low and High Effort

Effort	Probability of output of		Total
	0	1,000	
<i>Low</i> ($e = 0$)	0.9	0.1	1
<i>High</i> ($e = 7$)	0.2	0.8	1

- What are the incentive compatibility, participation, and zero-profit constraints for obtaining high effort?

- b What would utility be if the wage were fixed and could not depend on output or effort?
- c What is the optimal contract? What is the agent's utility?
- d What would the agent's utility be under full information? Under asymmetric information, what is the agency cost (the lost utility) as a percentage of the utility the agent receives?

7.3. Why Entrepreneurs Sell Out

Suppose an agent has a utility function of $U = \sqrt{w} - e$, where e can assume the levels 0 or 2.4, and his reservation utility is $\bar{U} = 7$. The principal is risk neutral. Denote the agent's wage, conditioned on output, as $w(0)$, $w(49)$, $w(100)$, or $w(225)$. Table 7 shows the output.

Table 7: Entrepreneurs Selling Out

Method	Probability of output of				Total
	0	49	100	225	
<i>Safe</i> ($e = 0$)	0.1	0.1	0.8	0	1
<i>Risky</i> ($e = 2.4$)	0	0.5	0	0.5	1

- a What would the agent's effort choice and utility be if he owned the firm?
- b If agents are scarce and principals compete for them, what will the agent's contract be under full information? His utility?
- c If principals are scarce and agents compete to work for principals, what will the contract be under full information? What will the agent's utility and the principal's profit be in this situation?
- d If agents are the scarce factor, and principals compete for them, what will the contract be when the principal cannot observe effort? What will the agent's utility and the principal's profit be in this situation?

7.4. Authority

A salesman must decide how hard to work on his own time on getting to know a potential customer. If he exerts effort X , he incurs a utility cost $X^2/2$. With probability X , he can then go to customer X and add V to his own earnings. With probability $(1-X)$, he offends the customer, and on going to him would subtract L from his earnings. The boss will receive benefit B from the sale in either case. The ranking of these numbers is $V > L > B > 0$. The boss and the salesman have equal bargaining power, and are free to make side payments to each other.

- a What is the first-best level of effort, X_a ?
- b If the boss has the authority to block the salesman from selling to this customer, but cannot force him to sell, what value will X take?

- c If the salesman has the authority over the decision on whether to sell to this customer, and can bargain for higher pay, what will his effort be?
- d Rank the effort levels X_a , X_b , and X_c in the previous three sections.

7.5. Worker Effort

A worker can be *Careful* or *Careless*, efforts which generate mistakes with probabilities 0.25 and 0.75. His utility function is $U = 100 - 10/w - x$, where w is his wage and x takes the value 2 if he is careful, and 0 otherwise. Whether a mistake is made is contractible, but effort is not. Risk-neutral employers compete for the worker, and his output is worth 0 if a mistake is made and 20 otherwise. No computation is needed for any part of this problem.

- a Will the worker be paid anything if he makes a mistake?
- b Will the worker be paid more if he does not make a mistake?
- c How would the contract be affected if employers were also risk averse?
- d What would the contract look like if a third category, “slight mistake,” with an output of 19, occurs with probability 0.1 after *Careless* effort and with probability zero after *Careful* effort?

7.6. The Source of Inefficiency

In the hidden actions problem facing an employer, inefficiency arises because

- (a) The worker is risk averse.
- (b) The worker is risk neutral.
- (c) No contract can induce high effort.
- (d) The type of the worker is unknown.
- (e) The level of risk aversion of the worker is unknown.

7.7. Optimal Compensation

An agent’s utility function is $U = (\log(\text{wage}) - \text{effort})$. What should his compensation scheme be if different (output,effort) pairs have the probabilities in Table 8?

- (a) The agent should be paid exactly his output.
- (b) The same wage should be paid for outputs of 1 and 100.
- (c) The agent should receive more for an output of 100 than of 1, but should receive still lower pay if output is 2.
- (d) None of the above.

Table 8: Output Probabilities

		Output		
		1	2	100
Effort	High	0.5	0	0.5
	Low	0.1	0.8	0.1

7.8. Effort and Output, Multiple Choices

The utility function of the agents whose situation is depicted in Table 9 is $U = w + \sqrt{w} - \alpha e$, and his reservation utility is 0. Principals compete for agents, and have reservation profits of zero. Principals are risk neutral.

Table 9: Output Probabilities

		Effort	
		Low ($e = 0$)	High ($e = 5$)
Output	$y = 0$	0.9	0.5
	$y = 100$	0.1	0.5

- a If $\alpha = 2$, then if the agent's action can be observed by the principal, his equilibrium utility is in the interval
- $[-\infty, 0.5]$
 - $[0.5, 5]$
 - $[5, 10]$
 - $[10, 40]$
 - $[40, \infty]$
- b If $\alpha = 10$, then if the agent's action can be observed by the principal, his equilibrium utility is in the interval
- $[-\infty, 0.5]$
 - $[0.5, 5]$
 - $[5, 10]$
 - $[10, 40]$
 - $[40, \infty]$
- c If $\alpha = 5$, then if the agent's action can be observed by the principal, his equilibrium effort level is
- Low
 - High
 - A mixed strategy effort, sometimes low and sometimes high
- d If $\alpha = 2$, then if the agent's action cannot be observed by the principal, and he must be paid a flat wage, his wage will be in the interval
- $[-\infty, 2]$
 - $[2, 5]$
 - $[5, 8]$
 - $[8, 12]$
 - $[12, \infty]$
- e If the agent owns the firm, and $\alpha = 2$, will his utility be higher or lower than in the case where he works for the principal and his action can be observed?
- Higher
 - Lower
 - Exactly the same.
- f If the agent owns the firm, and $\alpha = 2$, his equilibrium utility is in the interval
- $[-\infty, 0.5]$

- (b) $[0.5, 5]$
- (c) $[5, 10]$
- (d) $[10, 40]$
- (e) $[40, \infty]$

g If the agent owns the firm, and $\alpha = 8$, his equilibrium utility is in the interval

- (a) $[-\infty, 0.5]$
- (b) $[0.5, 5]$
- (c) $[5, 10]$
- (d) $[10, 40]$
- (e) $[40, \infty]$

7.9. Hiring a Lawyer

A one-man firm with concave utility function $U(X)$ hires a lawyer to sue a customer for breach of contract. The lawyer is risk-neutral and effort averse, with a convex disutility of effort. What can you say about the optimal contract? What would be the practical problem with such a contract, if it were legal?

7.10. Constraints

An agent has the utility function $U = \log(w) - e$, where e can take the levels 0 and 4, and his reservation utility is $\bar{U} = 4$. His principal is risk-neutral. Denote the agent's wage conditioned on output as $\underline{w}$ if output is 0 and $\bar{w}$ if output is 10. Only the agent observes his effort. Principals compete for agents. Output is as shown in Table 10.

Table 10: Effort and Outputs

Effort	Probability of Outputs		
	0	10	Total
<i>Low</i> ($e = 0$)	0.9	0.1	1
<i>High</i> ($e = 4$)	0.2	0.8	1

What are the incentive compatibility and participation constraints for obtaining high effort?

7.11. Constraints Again

Suppose an agent has the utility function $U = \log(w) - e$, where e can take the levels 1 or 3, and a reservation utility of $\bar{U}$. The principal is risk-neutral. Denote the agent's wage conditioned on output as $\underline{w}$ if output is 0 and $\bar{w}$ if output is 100. Only the agent observes his effort. Principals compete for agents, and outputs occur according to Table 11.

Table 11: Efforts and Outputs

Effort	Probability of Outputs	
	0	100
<i>Low</i> ($e = 1$)	0.9	0.1
<i>High</i> ($e = 3$)	0.5	0.5

What conditions must the optimal contract satisfy, given that the principal can only observe output, not effort? You do not need to solve out for the optimal contract— just provide the

equations which would have to be true. Do not just provide inequalities— if the condition is a binding constraint, state it as an equation.

7.12. Bankruptcy Constraints

A risk-neutral principal hires an agent with utility function $U = w - e$ and reservation utility $\bar{U} = 5$. Effort is either 0 or 10. There is a bankruptcy constraint: $w \geq 0$. Output is given by Table 12.

Table 12: Bankruptcy

Effort	Probability of Outputs		Total
	0	400	
<i>Low</i> ($e = 0$)	0.5	0.5	1
<i>High</i> ($e = 10$)	0.1	0.9	1

- What would be the agent's effort choice and utility if he owned the firm?
- If agents are scarce and principals compete for them what will be the agent's contract under full information? His utility?
- If principals are scarce and agents compete to work for them, what will the contract be under full information? What will the agent's utility be?
- If principals are scarce and agents compete to work for them, what will the contract be when the principal cannot observe effort? What will the payoffs be for each player?
- Suppose there is no bankruptcy constraint. If principals are the scarce factor and agents compete to work for them, what will the contract be when the principal cannot observe effort? What will the payoffs be for principal and agent?

7.13. The Game Wizard

A high-tech firm is trying to develop the game Wizard 1.0. It will have revenues of 200,000 *if it succeeds, and 0* if it fails. Success depends on the programmer. If he exerts high effort, the probability of success is .8. If he exerts low effort, it is .6. The programmer requires wages of at least 50,000 *if he can exert low effort, but 70,000* if he must exert high effort. (Let's just use payoffs in thousands of dollars, so 70,000 dollars will be written as 70.)

- Prove that high effort is first-best efficient.
- Explain why high effort would be inefficient if the probability of success when effort is low were .75.
- Let the probability of success with low effort go back to .6 for the remainder of the problem. If you cannot monitor the programmer and cannot pay him a wage contingent on success, what should you do?
- Now suppose you can make the wage contingent on success. Let the wage be S if Wizard is successful, and F if it fails. S and F will have to satisfy two conditions: a participation constraint and an incentive compatibility constraint. What are they?

- (e) What is a contract that will achieve the first best?
- (f) What is the optimal contract if you cannot pay a programmer a negative wage?

7.14. The Supercomputer Salesman

If a salesman exerts high effort, he will sell a supercomputer this year with probability .9. If he exerts low effort, he will succeed with probability .5. The company will make a profit of 2 million dollars if the sale is made. The salesman would require a wage of \$50,000 if he had to exert low effort, but \$70,000 if he had to exert high effort, he is risk neutral, and his utility is separable in effort and money. (Let's just use payoffs in thousands of dollars, so 70,000 dollars will be written as 70, and 2 million dollars will be 2000)

- (a) Prove that high effort is first-best efficient.
- (b) How high would the probability of success with low effort have to be for high effort to be inefficient?
- (c) If you cannot monitor the programmer and cannot pay him a wage contingent on success, what should you do?
- (d) Now suppose you can make the wage contingent on success. Let the wage be S if he makes a sale and F if he does not. S and F will have to satisfy two conditions: a participation constraint and an incentive compatibility constraint. What are they?
- (e) What is a contract that will achieve the first best?
- (f) Now suppose the salesman is risk averse, and his utility from money is $\log(w)$. Set up the participation and incentive compatibility constraints again.
- (g) You do not need to solve for the optimal contract. Using the $\log(w)$ utility function assumption, however, will the expected payment by the firm in the optimal contract rise, fall, or stay the same, compared with what it was in part (e) for the risk neutral salesman?
- (h) You do not need to solve for the optimal contract. Using the $\log(w)$ utility function assumption, however, will the gap between S and F in the optimal contract rise, fall, or stay the same, compared with what it was in part (e) for the risk neutral salesman?

8 Further Topics in Moral Hazard

Moral hazard deserves two chapters of this book. As we will see, adverse selection will sneak in with two also, since signalling is really just an elaboration of the adverse selection model, but moral hazard is perhaps even more important. It is really just the study of incentives, one of the central concepts of economics. And so in the chapter we will be going through a hodge-podge of special situations of moral hazard in which chapter 7's paradigm of providing the right incentives for effort by satisfying a participation constraint and an incentive compatibility constraint do not apply so straightforwardly.

The chapter begins with efficiency wages—high wages provided in situations where the it is so important to provide incentive compatibility that the principal is willing to abandon a tight participation constraint. Section 8.2 will be about tournaments—situations where competition between two agents can be used to simplify the optimal contract. After an excursion into various institutions, we will go to a big problem for incentive contracts: how does the principal restrain himself from being too merciful to a wayward agent, when mercy is not only kind but profitable? Section 8.5 then abandons the algebraic paradigm altogether to pursue a diagrammatic approach to the classic problem of moral hazard in insurance, and section 8.6 concludes with another special case: the teams problem, in which the unobservable efforts of many agents produce one observable output.

8.1 Efficiency Wages

One's first thought is that the basic idea of an incentive contract is to punish the agent if he chooses the wrong action. That is not quite right. Rather, the basic idea of an incentive contract is to create a difference between the agent's expected payoff from right and wrong actions. That can be done either with the stick of punishment or the carrot of reward.

It is important to keep this in mind, because sometimes punishments are simply not available. Consider the following game.

The Lucky Executive Game

Players

A corporation and an executive.

The Order of play

- 1 The corporation offers the executive a contract which pays $w(q) \geq 0$ depending on profit, q .
- 2 The executive accepts the contract, or rejects it and receives his reservation utility of $\bar{U} = 5$
- 3 The executive exerts effort e of either 0 or 10.
- 4 Nature chooses profit according to Table 1.

Payoffs

Both players are risk neutral. The corporation's payoff is $q - w$. The executive's payoff is $w - e$ if he accepts the contract.

Table 1: Output in the Lucky Executive Game

Effort	Probability of Outputs		Total
	0	400	
<i>Low</i> ($e = 0$)	0.5	0.5	1
<i>High</i> ($e = 10$)	0.1	0.9	1

Since both players are risk neutral, you might think that the first-best can be achieved by selling the store, putting the entire risk on the agent. The participation constraint if the executive exerts high effort is

$$0.1[w(0) - 10] + 0.9[w(400) - 10] \geq 5, \quad (1)$$

so his expected wage must equal 15. The incentive compatibility constraint is

$$0.5w(0) + 0.5w(400) \leq 0.1w(0) + 0.9w(400) - 10, \quad (2)$$

which can be rewritten as $w(400) - w(0) \geq 25$, so the gap between the executive's wage for high output and low output must equal at least 25.

A contract that satisfies both constraints is $\{w(0) = -345, w(400) = 55\}$. But this contract is not feasible, because the game requires $w(q) \geq 0$. This is an example of the common and realistic **bankruptcy constraint**; the principal cannot punish the agent by taking away more than the agent owns in the first place. The worst the boss can do is fire the worker. (In fact, the same problem would arise in a slavery regime that allowed the owner to kill his slaves—there, the worst the boss can do is kill the worker.) So what can be done?

What can be done is to use the carrot instead of the stick and abandon satisfying the participation constraint as an equality. All that is needed from constraint (2) is a gap between the high wage and the low wage of 25. Setting the low wage as low as is feasible, the corporation can use the contract $\{w(0) = 0, w(400) = 25\}$, and this will induce high effort. Notice, however, that the executive's expected utility will be $.1(0) + .9(25) - 10 = 12.5$, more than double his reservation utility of 5. He is very happy in this equilibrium— but the corporation is reasonably happy, too. The corporation's payoff is $337.5 (= 0.1(0 - 0) +$

$0.9(400 - 25)$, compared with the $195 (= 0.5(0 - 5) + 0.5(400 - 5))$ it would get if it paid a lower expected wage. Since high enough punishments are infeasible, the corporation has to use higher rewards.

Executives, of course, will now be lining up to work for this corporation, since they can get an expected utility of 12.5 there and only 5 elsewhere. If, in fact, there was some chance of the current executive dying and his job opening up, potential successors would be willing to pass up alternative jobs in order to be in position to get this unusually attractive job. Thus, the model generates unemployment. These are the two parts of the idea of the **efficiency wage**: the employer pays a wage higher than that needed to attract workers, and workers are willing to be unemployed in order to get a chance at the efficiency-wage job.

Shapiro & Stiglitz (1984) showed in more detail how involuntary unemployment can be explained by a principal-agent model. When all workers are employed at the market wage, a worker who is caught shirking and fired can immediately find another job just as good. Firing is ineffective and effective penalties like boiling-in-oil are excluded from the strategy spaces of legal businesses. Becker & Stigler (1974) suggested that workers post performance bonds, but if workers are poor this is impractical. Without bonds or boiling-in-oil, the worker chooses low effort and receives a low wage.

To induce a worker not to shirk, the firm can offer to pay him a premium over the market-clearing wage, which he loses if he is caught shirking and fired. If one firm finds it profitable to raise the wage, however, so do all firms. One might think that after the wages equalized, the incentive not to shirk would disappear. But when a firm raises its wage, its demand for labor falls, and when all firms raise their wages, the market demand for labor falls, creating unemployment. Even if all firms pay the same wage, a worker has an incentive not to shirk, because if he were fired he would stay unemployed, and even if there is a random chance of leaving the unemployment pool, the unemployment rate rises sufficiently high that workers choose not to risk being caught shirking. The equilibrium is not first-best efficient, because even though the marginal revenue of labor equals the wage, it exceeds the marginal disutility of effort, but it is efficient in a second-best sense. By deterring shirking, the hungry workers hanging around the factory gates are performing a socially valuable function (but they mustn't be paid for it!).

The idea of paying high wages to increase the threat of dismissal is old, and can even be found in *The Wealth of Nations* (Smith [1776] p. 207). What is new in Shapiro & Stiglitz (1984) is the observation that unemployment is generated by these “efficiency wages.” These firms behave paradoxically. They pay workers more than necessary to attract them, and outsiders who offer to work for less are turned away. Can this explain why “overqualified” jobseekers are unsuccessful and mediocre managers are retained? Employers are unwilling to hire someone talented, because he could find another job after being fired for shirking, and trustworthiness matters more than talent in some jobs.

This discussion should remind you of the game of Product Quality of section 5.4. There too, purchasers paid more than the reservation price in order to give the seller an incentive to behave properly, because a seller who misbehaved could be punished by termination of the relationship. The key characteristics of such models are a constraint on the amount of

contractual punishment for misbehavior and a participation constraint that is not binding in equilibrium. In addition, although the Lucky Executive Game works even with just one period, many versions, including the Product Quality Game, rely on there being a repeated game (infinitely repeated, or otherwise avoiding the Chainstore Paradox). Repetition allows for a situation in which the agent could considerably increase his payoff in one period by misbehavior such as stealing or low quality, but refrains because he would lose his position and lose all the future efficiency wage payments.

8.2 Tournaments

Games in which relative performance is important are called **tournaments**. Tournaments are similar to auctions, the difference being that the actions of the losers matter directly, unlike in auctions. Like auctions, they are especially useful when the principal wants to elicit information from the agents. A principal-designed tournament is sometimes called a **yardstick competition** because the agents provide the measure for their wages.

Farrell (2001) uses a tournament to explain how “slack” might be the major source of welfare loss from monopoly, an old idea usually prompted by faulty reasoning. The usual claim is that monopolists are inefficient because, unlike competitive firms, they do not have to maximize profits to survive. This relies on the dubious assumption that firms care about survival, not profits. Farrell makes a subtler point: although the shareholders of a monopoly maximize profit, the managers maximize their own utility, and moral hazard is severe without the benchmark of other firms’ performances.

Let firm Apex have two possible production techniques, *Fast* and *Careful*. Independently for each technique, Nature chooses production cost $c = 1$ with probability θ and $c = 2$ with probability $1 - \theta$. The manager can either choose a technique at random or investigate the costs of both techniques at a utility cost to himself of α . The shareholders can observe the resulting production cost, but not whether the manager investigates. If they see the manager pick *Fast* and a cost of $c = 2$, they do not know whether he chose it without investigating, or investigated both techniques and found they were both costly. The wage contract is based on what the shareholders can observe, so it takes the form (w_1, w_2) , where w_1 is the wage if $c = 1$ and w_2 if $c = 2$. The manager’s utility is $\log w$ if he does not investigate, $\log w - \alpha$ if he does, and the reservation utility of $\log \bar{w}$ if he quits.

If the shareholders want the manager to investigate, the contract must satisfy the self-selection constraint

$$U(\text{not investigate}) \leq U(\text{investigate}). \quad (3)$$

If the manager investigates, he still fails to find a low-cost technique with probability $(1 - \theta)^2$, so (3) is equivalent to

$$\theta \log w_1 + (1 - \theta) \log w_2 \leq [1 - (1 - \theta)^2] \log w_1 + (1 - \theta)^2 \log w_2 - \alpha. \quad (4)$$

The self-selection constraint is binding, since the shareholders want to keep the manager's compensation to a minimum. Turning inequality (4) into an equality and simplifying yields

$$\theta(1 - \theta)\log \frac{w_1}{w_2} = \alpha. \quad (5)$$

The participation constraint, which is also binding, is $U(\bar{w}) = U(\text{investigate})$, or

$$\log \bar{w} = [1 - (1 - \theta)^2]\log w_1 + (1 - \theta)^2\log w_2 - \alpha. \quad (6)$$

Solving equations (5) and (6) together for w_1 and w_2 yields

$$\begin{aligned} w_1 &= \bar{w}e^{\alpha/\theta}. \\ w_2 &= \bar{w}e^{-\alpha/(1-\theta)}. \end{aligned} \quad (7)$$

The expected cost to the firm is

$$[1 - (1 - \theta)^2]\bar{w}e^{\alpha/\theta} + (1 - \theta)^2\bar{w}e^{-\alpha/(1-\theta)}. \quad (8)$$

If the parameters are $\theta = 0.1$, $\alpha = 1$, and $\bar{w} = 1$, the rounded values are $w_1 = 22,026$ and $w_2 = 0.33$, and the expected cost is 4,185. Quite possibly, the shareholders decide it is not worth making the manager investigate.

But suppose that Apex has a competitor, Brydax, in the same situation. The shareholders of Apex can threaten to boil their manager in oil if Brydax adopts a low-cost technology and Apex does not. If Brydax does the same, the two managers are in a prisoner's dilemma, both wishing not to investigate, but each investigating from fear of the other. The forcing contract for Apex specifies $w_1 = w_2$ to fully insure the manager, and boiling-in-oil if Brydax has lower costs than Apex. The contract need satisfy only the participation constraint that $\log w - \alpha = \log \bar{w}$, so $w = 2.72$ and the cost of learning to Apex is only 2.72, not 4,185. Competition raises efficiency, not through the threat of firms going bankrupt but through the threat of managers being fired.

8.3 Institutions and Agency Problems (formerly section 8.6)

Ways to Alleviate Agency Problems

Usually when agents are risk averse, the first-best cannot be achieved, because some tradeoff must be made between providing the agent with incentives and keeping his compensation from varying too much between states of the world, or because it is not possible to punish him sufficiently. We have looked at a number of different ways to solve the problem, and at this point a listing might be useful. Each method is illustrated by application to the particular problem of executive compensation, which is empirically important, and interesting both because explicit incentive contracts are used and because they are not used more often (see Baker, Jensen & Murphy [1988]).

1 **Reputation** (sections 5.3, 5.4, 6.4, 6.6).

Managers are promoted on the basis of past effort or truthfulness.

2 **Risk-sharing contracts** (sections 7.2, 7.3, 7.4).

The executive receives not only a salary, but call options on the firm's stock. If he reduces the stock value, his options fall in value.

3 **Boiling in oil** (section 7.4).

If the firm would only become unable to pay dividends if the executive shirked and was unlucky, the threat of firing him when the firm skips a dividend will keep him working hard.

4 **Selling the store** (section 7.4).

The managers buy the firm in a leveraged buyout.

5 **Efficiency wages** (section 8.1).

To make him fear losing his job, the executive is paid a higher salary than his ability warrants (cf. Rasmusen [1988b] on mutual banks).

6 **Tournaments** (section 8.2).

Several vice presidents compete and the winner succeeds the president.

7 **Monitoring** (section 3.4).

The directors hire a consultant to evaluate the executive's performance.

8 **Repetition.**

Managers are paid less than their marginal products for most of their career, but are rewarded later with higher salaries or generous pensions if their career record has been good.

9 **Changing the type of the agent**

Older executives encourage the younger by praising ambition and hard work.

We have talked about all but the last two solutions. Repetition enables the contract to come closer to the first-best if the discount rate is low (Radner [1985]). Production Game V failed to attain the first-best in section 7.2 because output depended on both the agent's effort and random noise. If the game were repeated 50 times with independent drawings of the noise, the randomness would average out and the principal could form an accurate estimate of the agent's effort. This is, in a sense, begging the question, by saying that in the long run effort can be deduced after all.

Changing the agent's type by increasing the direct utility from desirable or decreasing that from undesirable behavior is a solution that has received little attention from economists, who have focussed on changing the utility by changing monetary rewards. Akerlof (1983), one of the few papers on the subject of changing type, points out that the moral education of children, not just their intellectual education, affects their productivity and success. The attitude of economics, however, has been that while virtuous agents exist, the rules of an organization need to be designed with the unvirtuous agents in mind. As the Chinese thinker Han Fei Tzu said some two thousand years ago,

Hardly ten men of true integrity and good faith can be found today, and yet the offices of the state number in the hundreds. If they must be filled by men of integrity and good faith, then there will never be enough men to go around; and if the offices are left unfilled, then those whose business it is to govern will dwindle in numbers while disorderly men increase. Therefore the way of the enlightened ruler is to unify the laws instead of seeking for wise men, to lay down firm policies instead of longing for men of good faith. (Han Fei Tzu [1964], p. 109 from his chapter, “The Five Vermin”)

The number of men of true integrity has probably not increased as fast as the size of government, so Han Fei Tzu’s observation remains valid, but it should be kept in mind that honest men do exist and honesty can enter into rational models. There are tradeoffs between spending to foster honesty and spending for other purposes, and there may be tradeoffs between using the second-best contracts designed for agents indifferent about the truth and using the simpler contracts appropriate for honest agents.

Government Institutions and Agency Problems

The field of law is well suited to analysis by principal-agent models. Even in the nineteenth century, Holmes (1881, p. 31) conjectured in *The Common Law* that the reason why sailors at one time received no wages if their ship was wrecked was to discourage them from taking to the lifeboats too early instead of trying to save it. The reason why such a legal rule may have been suboptimal is not that it was unfair—presumably sailors knew the risk before they set out—but because incentive compatibility and insurance work in opposite directions. If sailors are more risk averse than ship owners, and pecuniary advantage would not add much to their effort during storms, then the owner ought to provide insurance to the sailors by guaranteeing them wages whether the voyage succeeds or not.

Another legal question is who should bear the cost of an accident: the victim (for example, a pedestrian hit by a car) or the person who caused it (the driver). The economist’s answer is that it depends on who has the most severe moral hazard. If the pedestrian could have prevented the accident at the lowest cost, he should pay; otherwise, the driver. This idea of the **least-cost avoider** is extremely useful in the economic analysis of law, and is a major theme of Posner’s classic treatise on law and economics (Posner [1992]). Insurance or wealth transfer may also enter as considerations. If pedestrians are more risk averse, drivers should bear the cost, and, according to some political views, if pedestrians are poorer, drivers should bear the cost. Note that this last consideration—wealth transfer—is not relevant to private contracts. If a principal earning zero profits is required to bear the cost of work accidents, for example, the agent’s wage will be lower than if he bore them instead.

Criminal law is also concerned with tradeoffs between incentives and insurance. Holmes (1881, p. 40) also notes, approvingly, that Macaulay’s draft of the Indian Penal Code made breach of contract for the carriage of passengers a criminal offense. The reason is that the palanquin-bearers were too poor to pay damages for abandoning their passengers in desolate regions, so the power of the State was needed to provide for heavier punishments than bankruptcy. In general, however, the legal rules actually used seem to diverge more

from optimality in criminal law than civil law. If, for example, there is no chance that an innocent man can be convicted of embezzlement, boiling embezzlers in oil might be good policy, but most countries would not allow this. Taking the example a step further, if the evidence for murder is usually less convincing than for embezzling, our analysis could easily indicate that the penalty for murder should be less, but such reasoning offends the common notion that the severity of punishment should be matched with harm from the crime.

Private Institutions and Agency Problems

While agency theory can be used to explain and perhaps improve government policy, it also helps explain the development of many curious private institutions. Agency problems are an important hindrance to economic development, and may explain a number of apparently irrational practices. Popkin (1979, pp. 66, 73, 157) notes a variety of these. In Vietnam, for example, absentee landlords were more lenient than local landlords, but improved the land less, as one would expect of principals who suffer from informational disadvantages *vis-à-vis* their agents. Along the pathways in the fields, farmers would plant early-harvesting rice that the farmer's family could harvest by itself in advance of the regular crop, so that hired labor could not grab handfuls as they travelled. In thirteenth century England, beans were seldom grown, despite their nutritional advantages, because they were too easy to steal. Some villages tried to solve the problem by prohibiting anyone from entering the beanfields except during certain hours marked by the priest's ringing the church bell, so everyone could tend and watch their beans at the same official time.

In less exotic settings, moral hazard provides another reason besides tax benefits why employees take some of their wages in fringe benefits. Professors are granted some of their wages in university computer time because this induces them to do more research. Having a zero marginal cost of computer time is a way around the moral hazard of slacking on research, despite being a source of moral hazard in wasting computer time. A less typical but more imaginative example is that of the bank in Minnesota which, concerned about its image, gave each employee \$100 in credit at certain clothing stores to upgrade their style of dress. By compromising between paying cash and issuing uniforms the bank could hope to raise both its profits and the utility of its employees. ("The \$100 Sounds Good, but What do They Wear on the Second Day?" *Wall Street Journal*, October 16, 1987, p. 17.)

Longterm contracts are an important occasion for moral hazard, since so many variables are unforeseen, and hence noncontractible. The term **opportunism** has been used to describe the behavior of agents who take advantage of noncontractibility to increase their payoff at the expense of the principal (see Williamson [1975] and Tirole [1986]). Smith may be able to extract a greater payment from Jones than was agreed upon in their contract, because when a contract is incomplete, Smith can threaten to harm Jones in some way. This is called **hold-up potential** (Klein, Crawford, & Alchian [1978]). Hold-up potential can even make an agent introduce competing agents into the game, if competition is not so extreme as to drive rents to zero. Michael Granfield tells me that Fairchild once developed a new patent on a component of electronic fuel injection systems that it sought to sell to another firm, TRW. TRW offered a much higher price if Fairchild would license its patent to other producers, fearing the hold-up potential of buying from just one supplier. TRW could have tried writing a contract to prevent hold-up, but knew that it would be difficult

to prespecify all the ways that Fairchild could cause harm, including not only slow delivery, poor service, and low quality, but also sins of omission like failing to sufficiently guard the plant from shutdown due to accidents and strikes.

It should be clear from the variety of these examples that moral hazard is a common problem. Now that the first flurry of research on the principal-agent problem has finished, researchers are beginning to use the new theory to study institutions that were formerly relegated to descriptive “soft” scholarly work.

***8.4 Renegotiation: the Repossession Game**

Renegotiation comes up in two very different contexts in game theory. Chapter 4 looked at the situation where players can coordinate on Pareto-superior subgame equilibria that might be Pareto inferior for the entire game, an idea linked to the problem of selecting among multiple equilibria. This section looks at a completely different context, one in which the players have signed a binding contract, but in a subsequent subgame, both players might agree to scrap the old contract and write a new one using the old contract as a starting point in their negotiations. Here, the questions are not about equilibrium selection, but instead concern which strategies should be allowed in the game. This is an issue that frequently arises in principal-agent models, and it was first pointed out in the context of hidden knowledge by Dewatripont (1989). Here we will use a model of hidden actions to illustrate renegotiation, a model in which a bank wants to lend money to a consumer so that he can buy a car, and must worry whether the consumer will work hard enough to repay the loan.

The Repossession Game

Players

A bank and a consumer.

The Order of Play

- 1 The bank can do nothing or it can offer the consumer an auto loan which allows him to buy a car that costs 11, but requires him to pay back L or lose possession of the car to the bank.
- 2 The consumer accepts or rejects the loan.
- 3 The consumer chooses to *Work*, for an income of 15, or *Play*, for an income of 8. The disutility of work is 5.
- 4 The consumer repays the loan or defaults.
 - 4a In one version of the game, the bank offers to settle for an amount S and leave possession of the car to the consumer.
 - 4b The consumer accepts or rejects the settlement S .
- 5 If the bank has not been paid L or S , it repossesses the car.

Payoffs

If the bank does not make any loan or the consumer rejects it, both players' payoffs are zero. The value of the car is 12 to the consumer and 7 to the bank, so the bank's payoff if a loan is made is

$$\pi_{bank} = \begin{cases} L - 11 & \text{if the original loan is repaid} \\ S - 11 & \text{if a settlement is made} \\ 7 - 11 & \text{if the car is repossessed.} \end{cases}$$

If the consumer chooses *Work* his income W is 15 and his disutility of effort D is -5 . If he chooses *Play*, then $W = 8$ and $D = 0$. His payoff is

$$\pi_{consumer} = \begin{cases} W + 12 - L - D & \text{if the original loan is repaid} \\ W + 12 - S - D & \text{if a settlement is made} \\ W - D & \text{if the car is repossessed.} \end{cases}$$

We will consider two versions of the game, both of which allow commitment in the sense of legally binding agreements over transfers of money and wealth but do not allow the consumer to commit directly to *Work*. If the consumer does not repay the loan, the bank has the legal right to repossess the car, but the bank cannot have the consumer thrown into prison for breaking a promise to choose *Work*. Where the two versions of the game will differ is in whether they allow the renegotiation moves (4a) and (4b).

Repossession Game I

The first version of the game does not allow renegotiation, so moves (4a) and (4b) are dropped from the game. In equilibrium, the bank will make the loan at a rate of $L = 12$, and the consumer will choose *Work* and repay the loan. Working back from the end of the game in accordance with sequential rationality, the consumer is willing to repay because by repaying 12 he receives a car worth 12.¹ He will choose *Work* because he can then repay the loan and his payoff will be 10 ($= 15 + 12 - 12 - 5$), but if he chooses *Play* he will not be able to repay the loan and the bank will repossess the car, reducing his payoff to 8 ($= 8 - 0$). The bank will offer a loan at $L = 12$ because the consumer will repay it and that is the maximum repayment to which the consumer will agree. The bank's equilibrium payoff is 1 ($= 12 - 11$). This is an efficient outcome because the consumer does buy the car, which he values at more than its cost to the car dealer, although it is the bank rather than the consumer that gains the surplus, because of the bank's bargaining power over the terms of the loan.

Repossession Game II

The second version of the game does allow renegotiation, so moves (4a) and (4b) are added back into the game. Renegotiation turns out to be harmful, because it results in an equilibrium in which the bank refuses to make a loan, reducing the payoffs of bank and consumer to (0,10) instead of (1,10); the gains from trade are lost.

¹As usual, we could change the model slightly to make the consumer strongly desire to repay the loan, by substituting a bargaining subgame that splits the gains from trade between bank and consumer rather than specifying that the bank make a take-it-or-leave-it offer. See section 4.3.

The equilibrium in Repossession Game I breaks down because the consumer would deviate by choosing *Play*. In Repossession Game I, this would result in the bank repossessing the car, and in Repossession Game II, the bank still has the right to do this, for a payoff of $-4 (= 7 - 11)$. If the bank chooses to renegotiate and offer $S = 8$, however, this settlement will be accepted by the consumer, since in exchange he gets to keep a car worth 12, and the payoffs of bank and consumer are $-3 (= 8 - 11)$ and $12 (= 8 + 12 - 8)$. Thus, the bank will renegotiate, and the consumer will have increased his payoff from 10 to 12 by choosing *Play*. Looking ahead to this from move (1), however, the bank will see that it can do better by refusing to make the loan, resulting in the payoffs $(0, 10)$. The bank cannot even break even by raising the loan rate L . If $L = 30$, for instance, the consumer will still happily accept, knowing that when he chooses *Play* and defaults the ultimate amount he will pay will be just $S = 8$.

Renegotiation has a paradoxical effect. In the subgame starting with consumer default it increases efficiency, by allowing the players to make a Pareto improvement over an inefficient punishment. In the game as a whole, however, it reduces efficiency by preventing players from using punishments to deter inefficient actions. This is true of any situation in which punishment imposes a deadweight loss instead of being simply a transfer from the punished to the punisher. This may be why American judges are less willing than the general public to impose punishments on criminals. By the time a criminal reaches the courtroom, extra years in jail have no beneficial effect (incapacitation aside) and impose real costs on both criminal and society, and judges are unwilling to impose sentences which in each particular case are inefficient.

The renegotiation problem also comes up in principal-agent models because of risk bearing by a risk-averse agent when the principal is risk neutral. Optimal contracts impose risk on risk-averse agents to provide incentives for high effort or self selection. If at some point in the game it is common knowledge that the agent has chosen his action or report, but Nature has not yet moved, the agent bears needless risk. The principal knows the agent has already moved, so the two of them are willing to recontract to put the risk from Nature's move back on the principal. But the expected future recontracting makes a joke of the original contract and reduces the agent's incentives for effort or truthfulness.

The Repossession Game illustrates other ideas besides renegotiation. Note that it is a game of perfect information but has the feel of a game of moral hazard with hidden actions. This is because the game has an implicit bankruptcy constraint, so that the contract cannot sufficiently punish the consumer for an inefficient choice of effort. Restricting the strategy space has the same effect as restricting the information available to a player. It is another example of the distinction between observability and contractibility—the consumer's effort is observable, but it is not really contractible, because the bankruptcy constraint prevents him from being punished for his low effort.

This game also illustrates the difficulty of deciding what “bargaining power” means. This is a term that is very important to how many people think about law and public policy but which they define hazily. Chapter 11 will analyze bargaining in great detail, using the paradigm of splitting a pie. The natural way to think of bargaining power is to treat it as the ability to get a bigger share of the pie. Here, the pie to be split is the surplus of 1 from the consumer's purchase of a car at cost 11 which will yield him 12 in utility. Both

versions of the Repossession Game give all the bargaining power to the bank in the sense that where there is a surplus to be split, the bank gets 100 percent of it. But this does not help the bank in Repossession Game II, because the consumer can put himself in a position where the bank ends up a loser from the transaction despite its bargaining power.


*8.5 State-space Diagrams: Insurance Games I and II (formerly section 7.5)

The principal-agent models so far in the chapter have been presented in terms of algebraic equations or outcome matrices. Another approach, especially useful when the strategy space is continuous, is to use diagrams. The term “moral hazard” comes from the insurance industry. Suppose Mr Smith (the agent) is considering buying theft insurance for a car with a value of 12. Figure 1, which illustrates his situation, is an example of a **State-space diagram**, a diagram whose axes measure the values of one variable in two different states of the world. Before Smith buys insurance, his dollar wealth is 0 if there is a theft and 12 otherwise, depicted as his endowment, $\omega = (12, 0)$. The point $(12, 0)$ indicates a wealth of 12 in one state and 0 in the other, while the point $(6, 6)$ indicates a wealth of 6 in each state.

One cannot tell the probabilities of each state just by looking at the state-space diagram. Let us specify that if Smith is careful where he parks, the state *Theft* occurs with probability 0.5, but if he is careless the probability rises to 0.75. He is risk averse, and, other things equal, he has a mild preference to be careless, a preference worth only ϵ to him. Other things are not equal, however, and he would choose to be careful were he uninsured because of the high correlation of carelessness with theft.

The insurance company (the principal) is risk neutral, perhaps because it is owned by diversified shareholders. We assume that no transaction costs are incurred in providing insurance and that the market is competitive, a switch from Production Game V, where the principal collected all the gains from trade. If the insurance company can require Smith to park carefully, it offers him insurance at a premium of 6, with a payout of 12 if theft occurs, leaving him with an allocation of $C_1 = (6, 6)$. This satisfies the competition constraint because it is the most attractive contract any company can offer without making losses. Smith, whose allocation is 6 no matter what happens, is **fully insured**. In state-space diagrams, allocations like C_1 which fully insure one player are on the 45° line through the origin, the line along which his allocations in the two states are equal.

Figure 1 Insurance Game I



The game is described below in a specification that includes two insurance companies to simulate a competitive market. For Smith, who is risk averse, we must distinguish between dollar *allocations* such as (12,0) and utility *payoffs* such as $0.5U(12) + 0.5U(0)$. The curves in Figure 1 are labelled in units of utility for Smith and dollars for the insurance company.

Insurance Game I: observable care

Players

Smith and two insurance companies.

The Order of Play

- 1 Smith chooses to be either *Careful* or *Careless*, observed by the insurance company.
- 2 Insurance company 1 offers a contract (x, y) , in which Smith pays premium x and receives compensation y if there is a theft.
- 3 Insurance company 2 also offers a contract of the form (x, y) .
- 4 Smith picks a contract.
- 5 Nature chooses whether there is a theft, with probability 0.5 if Smith is *Careful* or 0.75 if Smith is *Careless*.

Payoffs

Smith is risk averse and the insurance companies are risk neutral. The insurance company not picked by Smith has a payoff of zero.

Smith's utility function U is such that $U' > 0$ and $U'' < 0$. If Smith picks contract (x, y) , the payoffs are:

If Smith chooses *Careful*,

$$\begin{aligned}\pi_{Smith} &= 0.5U(12 - x) + 0.5U(0 + y - x) \\ \pi_{company} &= 0.5x + 0.5(x - y), \text{ for his insurer.}\end{aligned}$$

If Smith chooses *Careless*,

$$\begin{aligned}\pi_{Smith} &= 0.25U(12 - x) + 0.75U(0 + y - x) + \epsilon \\ \pi_{company} &= 0.25x + 0.75(x - y), \text{ for his insurer.}\end{aligned}$$

In the equilibrium of Insurance Game I Smith chooses to be *Careful* because he foresees that otherwise his insurance will be more expensive. Figure 1 is the corner of an Edgeworth box which shows the indifference curves of Smith and his insurance company given that Smith's care keeps the probability of a theft down to 0.5. The company is risk neutral, so its indifference curve, $\pi_i = 0$, is a straight line with slope $-1/1$. Its payoffs are higher on indifference curves such as $\pi_i = 6$ that are closer to the origin and thus have smaller expected payouts to Smith. The insurance company is indifferent between ω and C_1 , at both of which its profits are zero. Smith is risk averse, so if he is *Careful* his indifference curves are closest to the origin on the 45 degree line, where his wealth in the two states is equal. Picking the numbers 66 and 83 for concreteness, I have labelled his original indifference curve $\pi_s = 66$ and drawn the preferred indifference curve $\pi_s = 83$ through the equilibrium contract C_1 . The equilibrium contract is C_1 , which satisfies the competition constraint by generating the highest expected utility for Smith that allows nonnegative profits to the company.

Insurance Game I is a game of symmetric information. Insurance Game II changes that. Suppose that

1. The company cannot observe Smith's action; or
2. The state insurance commission does not allow contracts to require Smith to be careful; or
3. A contract requiring Smith to be careful is impossible to enforce because of the cost of proving carelessness.

In each case Smith's action is a noncontractible variable, so we model all three the same way by putting Smith's move second. The new game is like Production Game V, with uncertainty, unobservability, and two levels of output, *Theft* and *No Theft*. The insurance company may not be able to directly observe Smith's action, but his dominant strategy is to be *Careless*, so the company knows the probability of a theft is 0.75. Insurance Game II is the same as Insurance Game I except for the following.

Insurance Game II: unobservable care

The Order of Play

1. Insurance company 1 offers a contract of form (x, y) , under which Smith pays premium x and receives compensation y if there is a theft.

2. Insurance company 2 offers a contract of form (x, y)
3. Smith picks a contract.
4. Smith chooses either *Careful* or *Careless*.
5. Nature chooses whether there is a theft, with probability 0.5 if Smith is *Careful* or 0.75 if Smith is *Careless*.

Smith's dominant strategy is *Careless*, so in contrast to Insurance Game I, the insurance company must offer a contract with a premium of 9 and a payout of 12 to prevent losses, which leaves Smith with an allocation $C_2 = (3, 3)$. Making thefts more probable reduces the slopes of both players' indifference curves, because it decreases the utility of points to the southeast of the 45 degree line and increases utility to the northwest. In Figure 2, the insurance company's isoprofit curve swivels from the solid line $\pi_i = 0$ to the dotted line $\tilde{\pi}_i = 0$. It swivels around ω because that is the point at which the company's profit is independent of how probable it is that Smith's car will be stolen, since the company is not insuring him at point ω . Smith's indifference curve also swivels, from the solid curve $\pi_s = 66$ to the dotted curve $\tilde{\pi}_s = 66 + \epsilon$. It swivels around the intersection of the $\pi_s = 66$ curve with the 45 degree line, because on that line the probability of theft does not affect Smith's payoff. The ϵ difference appears because Smith gets to choose the action *Careless*, which he slightly prefers.

Figure 2: Insurance Game II with full and partial insurance




Figure 2 shows that no full insurance contract will be offered. The contract C_1 is acceptable to Smith, but not to the insurance company, because it earns negative profits, and the contract C_2 is acceptable to the insurance company, but not to Smith, who prefers ω . Smith would like to commit himself to being careful, but he cannot make his commitment

credible. If the means existed to prove his honesty, he would use them even if they were costly. He might, for example, agree to buy off-street parking even though locking his car would be cheaper, if verifiable.


Although no full insurance contract such as C_1 or C_2 is mutually agreeable, other contracts can be used. Consider the partial insurance contract C_3 in Figure 2, which has a premium of 6 and a payout of 8. Smith would prefer C_3 to his endowment of $\omega = (12, 0)$ whether he chooses *Careless* or *Careful*. We can think of C_3 in two ways:

1. Full insurance except for a **deductible** of four. The insurance company pays for all losses in excess of four.
2. Insurance with a **coinsurance** rate of one-third. The insurance company pays two-thirds of all losses.

The outlook is bright, because Smith chooses *Careful* under a partial insurance contract like C_3 . The moral hazard is “small” in the sense that Smith barely prefers *Careless*. With even a small deductible, Smith would choose *Careful* and the probability of theft would fall to 0.5, allowing the company to provide much more generous insurance. The solution of full insurance is “almost” reached. In reality, we rarely observe truly full insurance, because insurance contracts repay only the price of the car and not the bother of replacing it, which is great enough to deter owners from leaving their cars unlocked.

Figure 3 illustrates effort choice under partial insurance. Smith has a choice between dashed indifference curves (*Careless*) and solid ones (*Careful*). To the southeast of the 45 degree line, the dashed indifference curve for a particular utility level is always above that utility’s solid indifference curve. Offered contract C_4 , Smith chooses *Careful*, remaining on the solid indifference curve, so C_4 yields zero profit to the insurance company. In fact, the competing insurance companies will offer contract C_5 in equilibrium, which is almost full insurance, but just almost, so that Smith will choose *Careful* to avoid the small amount of risk he still bears.

Figure 3: More on Partial Insurance in Insurance Game II



Thus, as in the principal-agent model there is a tradeoff between efficient effort and efficient risk allocation. Even when the ideal of full insurance and efficient effort cannot be reached, there exists some best choice like C_5 in the set of feasible contracts, a second-best insurance contract that recognizes the constraints of informational asymmetry.

*8.6 Joint Production by Many Agents: The Holmstrom Teams Model

To conclude this chapter, let us switch our focus from the individual agent to a group of agents. We have already looked at tournaments, which involve more than one agent, but a tournament still takes place in a situation where each agent's output is distinct. The tournament is a solution to the standard problem, and the principal could always fall back on other solutions such as individual risk-sharing contracts. In this section, the existence of a group of agents results in destroying the effectiveness of the individual risk-sharing contracts, because observed output is a joint function of the unobserved effort of many agents. Even though there is a group, a tournament is impossible, because only one output is observed. The situation has much of the flavor of the Civic Duty Game of chapter 3: the actions of a group of players produce a joint output, and each player wishes that the others would carry out the costly actions. A teams model is defined as follows.

*A **team** is a group of agents who independently choose effort levels that result in a single output for the entire group.*

We will look at teams using the following game.

Teams
(Holmstrom [1982])

Players

A principal and n agents.

The order of play

- 1 The principal offers a contract to each agent i of the form $w_i(q)$, where q is total output.
- 2 The agents decide whether or not to accept the contract.
- 3 The agents simultaneously pick effort levels e_i , ($i = 1, \dots, n$).
- 4 Output is $q(e_1, \dots, e_n)$.

Payoffs

If any agent rejects the contract, all payoffs equal zero. Otherwise,

$$\begin{aligned}\pi_{principal} &= q - \sum_{i=1}^n w_i; \\ \pi_i &= w_i - v_i(e_i), \text{ where } v'_i > 0 \text{ and } v''_i < 0.\end{aligned}$$

Despite the risk neutrality of the agents, “selling the store” fails to work here, because the team of agents still has the same problem as the employer had. The team’s problem is cooperation between agents, and the principal is peripheral.

Denote the efficient vector of actions by e^* . An efficient contract is

$$w_i(q) = \begin{cases} b_i & \text{if } q \geq q(e^*) \\ 0 & \text{if } q < q(e^*) \end{cases} \quad (9)$$

where $\sum_{i=1}^n b_i = q(e^*)$ and $b_i > v_i(e_i^*)$.

Contract (9) gives agent i the wage b_i if all agents pick the efficient effort, and nothing if any of them shirks, in which case the principal keeps the output. The teams model gives one reason to have a principal: he is the residual claimant who keeps the forfeited output. Without him, it is questionable whether the agents would carry out the threat to discard the output if, say, output were 99 instead of the efficient 100. There is a problem of dynamic consistency. The agents would like to commit in advance to throw away output, but only because they never have to do so in equilibrium. If the modeller wishes to disallow discarding output, he imposes the **budget-balancing constraint** that the sum of the wages exactly equal the output, no more and no less. But budget balancing creates a problem for the team that is summarized in Proposition 1.

Proposition 1. *If there is a budget-balancing constraint, no differentiable wage contract $w_i(q)$ generates an efficient Nash equilibrium.*

Agent i ’s problem is

$$\underset{e_i}{\text{Maximize}} \quad w_i(q(e)) - v_i(e_i). \quad (10)$$

His first-order condition is

$$\left(\frac{dw_i}{dq} \right) \left(\frac{dq}{de_i} \right) - \frac{dv_i}{de_i} = 0. \quad (11)$$

With budget balancing and a linear utility function, the Pareto optimum maximizes the sum of utilities (something not generally true), so the optimum solves


$$\underset{e_1, \dots, e_n}{\text{Maximize}} \quad q(e) - \sum_{i=1}^n v_i(e_i) \quad (12)$$

The first-order condition is that the marginal dollar contribution to output equal the marginal disutility of effort:

$$\frac{dq}{de_i} - \frac{dv_i}{de_i} = 0. \quad (13)$$

Equation (13) contradicts (11), the agent's first-order condition, because $\frac{dw_i}{dq}$ is not equal to one. If it were, agent i would be the residual claimant and receive the entire marginal increase in output— but under budget balancing, not every agent can do that. Because each agent bears the entire burden of his marginal effort and only part of the benefit, the contract does not achieve the first-best. Without budget balancing, on the other hand, if the agent shirked a little he would gain the entire leisure benefit from shirking, but he would lose his entire wage under the optimal contract.

Figure 4 (new): xxx Contracts in the Holmstrom Teams Model



Discontinuities in Public Good Payoffs

Ordinarily, there is a free rider problem if several players each pick a level of effort which increases the level of some public good whose benefits they share. Noncooperatively, they choose effort levels lower than if they could make binding promises. Mathematically, let identical risk-neutral players indexed by i choose effort levels e_i to produce amount $q(e_1, \dots, e_n)$ of the public good, where q is a continuous function. Player i 's problem is

$$\underset{e_i}{\text{Maximize}} \quad q(e_1, \dots, e_n) - e_i, \quad (14)$$

which has first order condition

$$\frac{\partial q}{\partial e_i} - 1 = 0, \quad (15)$$

whereas the greater, first-best effort n -vector e^* is characterized by

$$\sum_{i=1}^n \frac{\partial q}{\partial e_i} - 1 = 0. \quad (16)$$

If the function q were discontinuous at e^* (for example, if $q = 0$ if $e_i < e_i^*$ for any i), the strategy profile e^* could be a Nash equilibrium. In the game of Teams, the same effect is at work. Although output is not discontinuous, contract (9) is constructed as if it were (as if $q = 0$ if $e_i \neq e_i^*$ for any i), in order to obtain the same incentives.

The first-best can be achieved because the discontinuity at e^* makes every player the marginal, decisive player. If he shirks a little, output falls drastically and with certainty. Either of the following two modifications restores the free rider problem and induces shirking:

- 1 Let q be a function not only of effort but of random noise—Nature moves after the players. Uncertainty makes the *expected* output a continuous function of effort.
- 2 Let players have incomplete information about the critical value—Nature moves before the players and chooses e^* . Incomplete information makes the estimated output a continuous function of effort.

The discontinuity phenomenon is common. Examples, not all of which note the problem, include:

- 1 Effort in teams (Holmstrom [1982], Rasmusen [1987]).
- 2 Entry deterrence by an oligopoly (Bernheim [1984b], Waldman [1987]).
- 3 Output in oligopolies with trigger strategies (Porter [1983a]).
- 4 Patent races (Section 14.1).
- 5 Tendering shares in a takeover (Grossman & Hart [1980], Section 14.2).
- 6 Preferences for levels of a public good.

N8.1 Efficiency wages

- Which is the better, carrot or stick, is an interesting question. Two misconceptions that might lead one to think sticks are more powerful should be cleared up. First, if the agent is risk averse, equal dollar punishments and rewards lead to the punishment disutility being greater than the reward utility. Second, regression to the mean can easily lead a principal to think sticks work better than carrots in practice. Suppose a teacher assigns equal utility rewards and punishments to a student depending on his performance on tests, and that the student's effort is, in fact, constant. If the student is lucky on a test, he will do well and be rewarded, but will probably do worse on the next test. If the student is unlucky, he will be punished, and will do better on the next test. The naive teacher will think that rewards hurt performance and punishments help it. See Robyn Dawes's 1988 book, *Rational Choice in an Uncertain World* for a good exposition of this and other pitfalls of reasoning (especially pages 84-87). Kahneman, Slovic & Tversky (1982) covers similar material.
- For surveys of the efficiency wage literature, see the article by L. Katz (1986), the book of articles edited by Akerlof & Yellen (1986), and the book-length survey by Weiss (1990).
- While the efficiency wage model does explain involuntary unemployment, it does not explain cyclical changes in unemployment. There is no reason for the unemployment needed to control moral hazard to fluctuate widely and create a business cycle.
- The efficiency wage idea is essentially the same idea as in the Klein & Leffler (1981) model of product quality formalized in section 5.3. If no punishment is available for player who is tempted to misbehave, a punishment can be created by giving him something to take away. This something can be a high-paying job or a loyal customer. It is also similar to the idea of **co-opting** opponents familiar in politics and university administration. To tame the radical student association, give them an office of their own which can be taken away if they seize the dean's office. Rasmusen (1988b) shows yet another context: when depositors do not know which investments are risky and which are safe, mutual bank managers can be highly paid to deter them from making risky investments that might cost them their jobs.
- Adverse selection can also drive an efficiency wage model. We will see in Chapter 9 that a customer might be willing to pay a high price to attract sellers of high-quality cars when he cannot detect quality directly.

N8.2 Tournaments

- An article which stimulated much interest in tournaments is Lazear & Rosen (1981), which discusses in detail the importance of risk aversion and adverse selection.
- One example of a tournament is the two-year, three-man contest for the new chairman of Citicorp. The company named three candidates as vice-chairmen: the head of consumer banking, the head of corporate banking, and the legal counsel. Earnings reports were even split into three components, two of which were the corporate and consumer banking (the third was the "investment" bank, irrelevant to the tournament). See "What Made Reed Wriston's Choice at Citicorp," *Business Week*, July 2, 1984, p. 25.

- General Motors has tried a tournament among its production workers. During a depressed year, management credibly threatened to close down the auto plant with the lowest productivity. Reportedly, this did raise productivity. Such a tournament is interesting because it helps explain why a firm's supply curve could be upward sloping even if all its plants are identical, and why it might hold excess capacity. Should information on a plant's current performance have been released to other plants? See "Unions Say Auto Firms Use Interplant Rivalry to Raise Work Quotas," *Wall Street Journal*, November 8, 1983, p. 1.
- Under adverse selection, tournaments must be used differently than under moral hazard because agents cannot control their effort. Instead, tournaments are used to deter agents from accepting contracts in which they must compete for a prize with other agents of higher ability.
- Interfirm management tournaments run into difficulties when shareholders want managers to cooperate in some arenas. If managers collude in setting prices, for example, they can also collude to make life easier for each other.
- Antle & Smith (1986) is an empirical study of tournaments in managers' compensation. Rosen (1986) is a theoretical model of a labor tournament in which the prize is promotion.
- Suppose a firm conducts a tournament in which the best-performing of its vice-presidents becomes the next president. Should the firm fire the most talented vice-president before it starts the tournament? The answer is not obvious. Maybe in the tournament's equilibrium, Mr Talent works less hard because of his initial advantage, so that all of the vice-presidents retain the incentive to work hard.
- A tournament can reward the winner, or shoot the loser. Which is better? Nalebuff & Stiglitz (1983) say to shoot the loser, and Rasmusen (1987) finds a similar result for teams, but for a different reason. Nalebuff & Stiglitz's result depends on uncertainty and a large number of agents in the tournament, while Rasmusen's depends on risk aversion. If a utility function is concave because the agent is risk averse, the agent is hurt more by losing a given sum than he would benefit by gaining it. Hence, for incentive purposes the carrot is inferior to the stick, a result unfortunate for efficiency since penalties are often bounded by bankruptcy or legal constraints.
- Using a tournament, the equilibrium effort might be greater in a second-best contract than in the first-best, even though the second-best is contrived to get around the problem of inducing sufficient effort. Also, a pure tournament, in which the prizes are distributed solely according to the ordinal ranking of output by the agents, is often inferior to a tournament in which an agent must achieve a significant margin of superiority over his fellows in order to win (Nalebuff & Stiglitz [1983]). Companies using sales tournaments sometimes have prizes for record yearly sales besides ordinary prizes, and some long distance athletic races have nonordinal prizes to avoid dull events in which the best racers run "tactical races."
- Organizational slack of the kind described in the Farrell model has important practical implications. In dealing with bureaucrats, one must keep in mind that they are usually less concerned with the organization's prosperity than with their own. In complaining about bureaucratic ineptitude, it may be much more useful to name particular bureaucrats and send them copies of the complaint than to stick to the abstract issues at hand. Private firms, at least, are well aware that customers help monitor agents.

N8.3 Institutions and agency problems

- Gaver & Zimmerman (1977) describes how a performance bond of 100 percent was required for contractors building the BART subway system in San Francisco. “Surety companies” generally bond a contractor for five to 20 times his net worth, at a charge of 0.6 percent of the bond per year, and absorption of their bonding capacity is a serious concern for contractors in accepting jobs.
- Even if a product’s quality need not meet government standards, the seller may wish to bind himself to them voluntarily. Stroh’s *Erlanger* beer proudly announces on every bottle that although it is American, “Erlanger is a special beer brewed to meet the stringent requirements of Reinheitsgebot, a German brewing purity law established in 1516.” Inspection of household electrical appliances by an independent lab to get the “ U_L ” listing is a similarly voluntary adherence to standards.
- The stock price is a way of using outside analysts to monitor an executive’s performance. When General Motors bought EDS, they created a special class of stock, GM-E, which varied with EDS performance and could be used to monitor it.

*N8.6 Joint production by many agents: the Holmstrom Teams Model

- **Team theory**, as developed by Marschak & Radner (1972) is an older mathematical approach to organization. In the old usage of “team” (different from the current, Holmstrom [1982] usage), several agents who have different information but cannot communicate it must pick decision rules. The payoff is the same for each agent, and their problem is coordination, not motivation.
- The efficient contract (9) supports the efficient Nash equilibrium, but it also supports a continuum of inefficient Nash equilibria. Suppose that in the efficient equilibrium all workers work equally hard. Another Nash equilibrium is for one worker to do no work and the others to work inefficiently hard to make up for him.
- **A teams contract with hidden knowledge.** In the 1920s, National City Co. assigned 20 percent of profits to compensate management as a group. A management committee decided how to share it, after each officer submitted an unsigned ballot suggesting the share of the fund that Chairman Mitchell should have, and a signed ballot giving his estimate of the worth of each of the other eligible officers, himself excluded. (Galbraith [1954] p. 157)
- **A First-best, budget-balancing contract when agents are risk averse.** Proposition 8.1 can be shown to hold for any contract, not just for differentiable sharing rules, but it does depend on risk neutrality and separability of the utility function. Consider the following contract from Rasmusen (1987):

$$w_i = \left\{ \begin{array}{ll} b_i & \text{if } q \geq q(e^*) \\ 0 \text{ with probability } (n-1)/n \\ q \text{ with probability } 1/n & \text{if } q < q(e^*) \end{array} \right\}$$

If the worker shirks, he enters a lottery. If his risk aversion is strong enough, he prefers the riskless return b_i , so he does not shirk. If agents’ wealth is unlimited, then for any positive risk aversion we could construct such a contract, by making the losers in the lottery accept negative pay.

- A teams contract like (9) is not a tournament. Only absolute performance matters, even though the level of absolute performance depends on what all the players do.

- **The budget-balancing constraint.** The legal doctrine of “consideration” makes it difficult to make binding, Pareto-suboptimal promises. An agreement is not a legal contract unless it is more than a promise: both parties have to receive something valuable for the courts to enforce the agreement.
- Adverse selection can be incorporated into a teams model. A team of workers who may differ in ability produce a joint output, and the principal tries to ensure that only high-ability workers join the team. (See Rasmusen & Zenger [1990]).

Problems

8.1. Monitoring with error

An agent has a utility function $U = \sqrt{w} - \alpha e$, where $\alpha = 1$ and e is either 0 or 5. His reservation utility level is $\bar{U} = 9$, and his output is 100 with low effort and 250 with high effort. Principals are risk neutral and scarce, and agents compete to work for them. The principal cannot condition the wage on effort or output, but he can, if he wishes, spend five minutes of his time, worth 10 dollars, to drop in and watch the agent. If he does that, he observes the agent *Daydreaming* or *Working*, with probabilities that differ depending on the agent's effort. He can condition the wage on those two things, so the contract will be $\{\underline{w}, \bar{w}\}$. The probabilities are given by Table 1.

Table 1: Monitoring with Error

Effort	Probability of	
	<i>Daydreaming</i>	<i>Working</i>
<i>Low</i> ($e = 0$)	0.6	0.4
<i>High</i> ($e = 5$)	0.1	0.9

- What are profits in the absence of monitoring, if the agent is paid enough to make him willing to work for the principal?
- Show that high effort is efficient under full information.
- If $\alpha = 1.2$, is high effort still efficient under full information?
- Under asymmetric information, with $\alpha = 1$, what are the participation and incentive compatibility constraints?
- Under asymmetric information, with $\alpha = 1$, what is the optimal contract?

8.2. Monitoring with Error: Second Offenses (see Rubinstein [1979]).

Individuals who are risk-neutral must decide whether to commit zero, one, or two robberies. The cost to society of robbery is 10, and the benefit to the robber is 5. No robber is ever convicted and jailed, but the police beat up any suspected robber they find. They beat up innocent people mistakenly sometimes, as shown by Table 2, which shows the probabilities of zero or more beatings for someone who commits zero, one, or two robberies.

Table 2: Crime

Robberies	Beatings		
	0	1	2
0	0.81	0.18	0.01
1	0.60	0.34	0.06
2	0.49	0.42	0.09

- How big should p^* , the disutility of a beating, be made to deter crime completely while inflicting a minimum of punishment on the innocent?

- b In equilibrium, what percentage of beatings are of innocent people? What is the payoff of an innocent man?
- c Now consider a more flexible policy, which inflicts heavier beatings on repeat offenders. If such flexibility is possible, what are the optimal severities for first- and second-time offenders? (call these p_1 and p_2). What is the expected utility of an innocent person under this policy?
- d Suppose that the probabilities are as given in Table 3. What is an optimal policy for first and second offenders?

Table 3: More Crime

Robberies	Beatings		
	0	1	2
0	0.9	0.1	0
1	0.6	0.3	0.1
2	0.5	0.3	0.2

8.3: Bankruptcy Constraints. A risk-neutral principal hires an agent with utility function $U = w - e$ and reservation utility $\bar{U} = 7$. Effort is either 0 or 20. There is a bankruptcy constraint: $w \geq 0$. Output is given by Table 4.

Table 4: Bankruptcy

Effort	Probability of output of		Total
	0	400	
<i>Low</i> ($e = 0$)	0.5	0.5	1
<i>High</i> ($e = 10$)	0.2	0.8	1

- a What would the agent's effort choice and utility be if he owned the firm?
- b If agents are scarce and principals compete for them, what will the agent's contract be under full information? His utility?
- c If principals are scarce and agents compete to work for them, what will the contract be under full information? What will the agent's utility be?
- d If principals are scarce and agents compete to work for them, what will the contract be when the principal cannot observe effort? What will the payoffs be for each player?
- e Suppose there is no bankruptcy constraint. If principals are the scarce factor and agents compete to work for them, what will the contract be when the principal cannot observe effort? What will the payoffs be for principal and agent?

8.4: Teams. A team of two workers produces and sells widgets for the principal. Each worker chooses high or low effort. An agent's utility is $U = w - 20$ if his effort is high, and $U = w$ if it is low, with a reservation utility of $\bar{U} = 0$. Nature chooses business conditions to be excellent, good, or bad, with probabilities θ_1 , θ_2 , and θ_3 . The principal observes output but not business conditions, as shown in Table 5.

Table 5: Team output

	Excellent (θ_1)	Good (θ_2)	Bad (θ_3)
<i>High, High</i>	100	100	60
<i>High, Low</i>	100	50	20
<i>Low, Low</i>	50	20	0

- Suppose $\theta_1 = \theta_2 = \theta_3$. Why is $\{(w(100) = 30, w(\text{not } 100) = 0), (High, High)\}$ not an equilibrium?
- Suppose $\theta_1 = \theta_2 = \theta_3$. Is it optimal to induce high effort? What is an optimal contract with nonnegative wages?
- Suppose $\theta_1 = 0.5$, $\theta_2 = 0.5$, and $\theta_3 = 0$. Is it optimal to induce high effort? What is an optimal contract (possibly with negative wages)?
- Should the principal stop the agents from talking to each other?

8.5: Efficiency Wages and Risk Aversion (see Rasmusen [1992c])

In each of two periods of work, a worker decides whether to steal amount v , and is detected with probability α and suffers legal penalty p if he, in fact, did steal. A worker who is caught stealing can also be fired, after which he earns the reservation wage w_0 . If the worker does not steal, his utility in the period is $U(w)$; if he steals, it is $U(w + v) - \alpha p$, where $U(w_0 + v) - \alpha p > U(w_0)$. The worker's marginal utility of income is diminishing: $U' > 0$, $U'' < 0$, and $\lim_{x \rightarrow \infty} U'(x) = 0$. There is no discounting. The firm definitely wants to deter stealing in each period, if at all possible.

- Show that the firm can indeed deter theft, even in the second period, and, in fact, do so with a second-period wage w_2^* that is higher than the reservation wage w_0 .
- Show that the equilibrium second-period wage w_2^* is higher than the first-period wage w_1^* .

8.6. The Revelation Principle

If you apply the Revelation Principle, that

- Increases the welfare of all the players in the model.
- Increases the welfare of just the player offering the contract.
- Increases the welfare of just the player accepting the contract.
- Makes the problem easier to model, but does not raise the welfare of the players.
- Makes the problem easier to model and raises the welfare of some players, but not all.

8.7 Machinery

Mr. Smith is thinking of buying a custom- designed machine from either Mr. Jones or Mr. Brown. This machine costs 5000 dollars to build, and it is useless to anyone but Smith. It is common knowledge that with 90 percent probability the machine will be worth 10,000 dollars to Smith at the time of delivery, one year from today, and with 10 percent probability it will only be worth 2,000 dollars. Smith owns assets of 1,000 dollars. At the time of contracting, Jones and Brown

believe there is there is a 20 percent chance that Smith is actually acting as an “undisclosed agent” for Anderson, who has assets of 50,000 dollars.

Find the price be under the following two legal regimes: (a) An undisclosed principal is not responsible for the debts of his agent; and (b) even an undisclosed principal is responsible for the debts of his agent. Also, explain (as part [c]) which rule a moral hazard model like this would tend to support.

9 Adverse Selection

9.1 Introduction: Production Game VI

In Chapter 7, games of asymmetric information were divided between games with moral hazard, in which agents are identical, and games with adverse selection, in which agents differ. In moral hazard with hidden knowledge and adverse selection, the principal tries to sort out agents of different types. In moral hazard with hidden knowledge, the emphasis is on the agent's action rather than his choice of contract, and agents accept contracts before acquiring information. Under adverse selection, the agent has private information about his type or the state of the world before he agrees to a contract, which means that the emphasis is on which contract he will accept.

For comparison with moral hazard, let us consider still another version of the Production Game of Chapters 7 and 8.

Production Game VI: Adverse Selection

Players

The principal and the agent.

The Order of Play

- (0) Nature chooses the agent's ability a , unobserved by the principal, according to distribution $F(a)$.
- (1) The principal offers the agent one or more wage contracts $w_1(q), w_2(q), \dots$
- (2) The agent accepts one contract or rejects them all.
- (3) Nature chooses a value for the state of the world, θ , according to distribution $G(\theta)$. Output is then $q = q(a, \theta)$.

Payoffs

If the agent rejects all contracts, then $\pi_{agent} = \bar{U}$ and $\pi_{principal} = 0$. Otherwise, $\pi_{agent} = U(w)$ and $\pi_{principal} = V(q - w)$.

Under adverse selection, it is not the worker's effort but his ability that is noncontractible. Without uncertainty (move (3)), the principal would provide a single contract specifying high wages for high output and low wages for low output, but unlike under moral hazard, either high or low output might be observed in equilibrium if both types of agent

accepted the contract. Also, in adverse selection, unlike moral hazard, offering multiple contracts can be an improvement over offering a single contract. The principal might, for example, provide a contract with a flat wage for the low-ability agents and an incentive contract for the high-ability agents. Production Game VIa puts specific functional forms into the game to illustrate equilibrium.

Production Game VIa: Adverse Selection, with Particular Parameters ’

Players

The principal and the agent.

The Order of Play

- (0) Nature chooses the agent’s ability a , unobserved by the principal, according to distribution $F(a)$, which puts probability 0.9 on low ability, $a = 0$, and probability 0.1 on high ability, $a = 10$.
- (1) The principal offers the agent one or more wage contracts $W_1 = \{w_1(q = 0), w_1(q = 10)\}$, $W_2 = \{w_2(q = 0), w_2(q = 10)\} \dots$
- (2) The agent accepts one contract or rejects them all.
- (3) Nature chooses a value for the state of the world, θ , according to distribution $G(\theta)$, which puts equal weight on 0 and 10. Output is then $q = \text{Min}(a + \theta, 10)$. (Thus, output is 0 or 10 for the low-ability agent, and always 10 for the high-ability.)

Payoffs

If the agent rejects all contracts, then depending on his type his reservation payoff is either $\pi_L = 3$ or $\pi_H = 4$ and the principal’s payoff is $\pi_{\text{principal}} = 0$.

Otherwise, $\pi_{\text{agent}} = w$ and $\pi_{\text{principal}} = q - w$.

An equilibrium is

Principal : Offer $W_1 = \{w_1(q = 0) = 3, w_1(q = 10) = 3\}$, $W_2 = \{w_2(q = 0) = 0, w_2(q = 10) = 4\}$

Low agent : Accept W_1 .

High agent : Accept W_2 .

As usual, this is a weak equilibrium. Both Low and High agents are indifferent about whether they accept or reject a contract. But the equilibrium indifference of the agents arises from the open-set problem; if the principal were to specify a wage of 2.01 for W_1 , for example, the low-ability agent would no longer be indifferent about accepting it.

This equilibrium can be obtained by what is a standard method for hidden-knowledge models. In hidden-action models, the principal tries to construct a contract which will induce the agent to take the single appropriate action. In hidden-knowledge models, the principal tries to make different actions attractive under different states of the world, so

the agent's choice depends on the hidden information. The principal's problem, as in Production Game V, is to maximize his profits subject to

- (1) **Incentive compatibility** (the agent picks the desired contract and actions).
- (2) **Participation** (the agent prefers the contract to his reservation utility).

In a model with hidden knowledge, the incentive compatibility constraint is customarily called the **self-selection constraint**, because it induces the different types of agents to pick different contracts. The big difference is that there will be an entire set of self-selection constraints, one for each type of agent or each state of the world, since the appropriate contract depends on the hidden information.

First, what action does the principal desire from each type of agent? The agents do not choose effort, but they do choose whether or not to work for the principal, and which contract to accept. The low-ability agent's expected output is $0.5(0) + 0.5(10) = 5$, compared to a reservation payoff of 3, so the principal will want to hire the low-ability agent if he can do it at an expected wage of 5 or less. The high ability agent's expected output is $0.5(10) + 0.5(10) = 10$, compared to a reservation payoff of 4, so the principal will want to hire the high-ability agent, if he can do it at an expected wage of 10 or less. The principal will want to induce the low-ability agent to choose a cheaper contract and not to choose the necessarily more expensive contract needed to attract the high-ability agent.

The participation constraints are

$$\begin{aligned} U_L(W_1) &\geq \bar{\pi}_L; \quad 0.5w_1(0) + 0.5w_1(10) \geq 3 \\ U_H(W_2) &\geq \bar{\pi}_H; \quad 0.5w_2(10) + 0.5w_2(10) \geq 4 \end{aligned} \tag{1}$$

Clearly the contracts $W_1 = \{3, 3\}$ and $W_2 = \{0, 10\}$ satisfy the participation constraints. The constraints show that both the low- output wage and the high-output wage matter to the low-ability agent, but only the high-output wage matters to the high-ability agent, so it makes sense to make W_2 as risky as possible.

The self selection constraints are

$$\begin{aligned} U_L(W_1) &\geq U_L(W_2); \quad 0.5w_1(0) + 0.5w_1(10) \geq 0.5w_2(0) + 0.5w_2(10) \\ U_H(W_2) &\geq U_H(W_1); \quad 0.5w_2(10) + 0.5w_2(10) \geq 0.5w_1(10) + 0.5w_1(10) \end{aligned} \tag{2}$$

The risky wage contract W_2 has to have a low enough expected return for the low-ability agent to deter him from accepting it; but the safe wage contract W_1 must be less attractive than W_1 to the high-ability agent. The contracts $W_1 = \{3, 3\}$ and $W_2 = \{0, 10\}$ do this, as can be seen by substituting their values into the constraints:

$$\begin{aligned} U_L(W_1) &\geq U_L(W_2); \quad 0.5(3) + 0.5(3) \geq 0.5(0) + 0.5(4) \\ U_H(W_2) &\geq U_H(W_1); \quad 0.5(4) + 0.5(4) \geq 0.5(3) + 0.5(3) \end{aligned} \tag{3}$$

Since the self selection and participation constraints are satisfied, the agents will not deviate from their equilibrium actions. All that remains to check is whether the principal

could increase his payoff. He cannot, because he makes a profit from either contract, and having driven the agents down to their reservation utilities, he cannot further reduce their pay.

As with hidden actions, if principals compete in offering contracts under hidden information, a **competition constraint** is added: the equilibrium contract must be as attractive as possible to the agent, since otherwise another principal could profitably lure him away. An equilibrium may also need to satisfy a part of the competition constraint not found in hidden actions models: either a **nonpooling constraint** or a **nonseparating constraint**. If one of several competing principals wishes to construct a pair of separating contracts C_1 and C_2 , he must construct it so that not only do agents choose C_1 and C_2 depending on the state of the world (to satisfy incentive compatibility), but also they prefer (C_1, C_2) to a pooling contract C_3 (to satisfy nonpooling). We only have one principal in Production Game VI, though, so competition constraints are irrelevant.

It is always true that the self selection and participation constraints must be satisfied for agents who accept the contracts, but it is not always the case that they accept different contracts.

*If all types of agents choose the same strategy in all states, the equilibrium is **pooling**. Otherwise, it is **separating**.*

The distinction between pooling and separating is different from the distinction between equilibrium concepts. A model might have multiple Nash equilibria, some pooling and some separating. Moreover, a single equilibrium— even a pooling one— can include several contracts, but if it is pooling the agent always uses the same strategy, regardless of type. If the agent’s equilibrium strategy is mixed, the equilibrium is pooling if the agent always picks the same mixed strategy, even though the messages and efforts would differ across realizations of the game.

These two terms came up in Section 6.2 in the game of PhD Admissions. Neither type of student applied in the pooling equilibrium, but one type did in the separating equilibrium. In a principal-agent model, the principal tries to design the contract to achieve separation unless the incentives turn out to be too costly. In Production Game VI, the equilibrium was separating, since the two types of agents choose different contracts.

A separating contract need not be fully separating. If agents who observe $\theta \leq 4$ accept contract C_1 but other agents accept C_2 , then the equilibrium is separating but it does not separate out every type. We say that the equilibrium is **fully revealing** if the agent’s choice of contract always conveys his private information to the principal. Between pooling and fully revealing equilibria are the **imperfectly separating** equilibria synonymously called **semi-separating**, **partially separating**, **partially revealing**, or **partially pooling** equilibria.

Production Game VI is a fairly complicated game, so let us start in Sections 9.2 and 9.3 with a certainty game, although we will return to uncertainty in Section 9.4. The first game will model a used car market in which the quality of the car is known to the seller but not the buyer, and the various versions of the game will differ in the types and numbers of the buyers and sellers. Section 9.4 will return to models with uncertainty, in a model

of adverse selection in insurance. One result there will be that a Nash equilibrium in pure strategies fails to exist for certain parameter values. Section 9.5 applies the idea of adverse selection to explain the magnitude of the bid-ask spread in financial markets, and Section 9.6 touches on a variety of other applications.

9.2 Adverse Selection under Certainty: Lemons I and II

Akerlof stimulated an entire field of research with his 1970 model of the market for shoddy used cars (“lemons”), in which adverse selection arises because car quality is better known to the seller than to the buyer. In agency terms, the principal contracts to buy from the agent a car whose quality, which might be high or low, is noncontractible despite the lack of uncertainty. Such a model may sound like moral hazard with hidden knowledge, but the difference is that in the used car market the seller has private information about his own type before making any kind of agreement. If, instead, the seller agreed to resell his car when he first bought it, the model would be moral hazard with hidden knowledge, because there would be no asymmetric information at the time of contracting, just an expectation of future asymmetry.

We will spend considerable time adding twists to a model of the market in used cars. The game will have one buyer and one seller, but this will simulate competition between buyers, as discussed in Section 7.2, because the seller moves first. If the model had symmetric information there would be no consumer surplus. It will often be convenient to discuss the game as if it had many sellers, interpreting a seller whom Nature randomly assigns a type as a population of sellers of different types, one of whom is drawn by Nature to participate in the game.

The Basic Lemons Model

Players

A buyer and a seller.

The Order of Play

- (0) Nature chooses quality type θ for the seller according to the distribution $F(\theta)$.
The seller knows θ , but while the buyer knows F , he does not know the θ of the particular seller he faces.
- (1) The buyer offers a price P .
- (2) The seller accepts or rejects.

Payoffs

If the buyer rejects the offer, both players receive payoffs of zero. Otherwise, $\pi_{buyer} = V(\theta) - P$ and $\pi_{seller} = P - U(\theta)$, where V and U will be defined later.

The payoffs of both players are normalized to zero if no transaction takes place. A normalization is part of the notation of the model rather than a substantive assumption.


Here, the model assigns the players' utility a base value of zero when no transaction takes place, and the payoff functions show changes from that base. The seller, for instance, gains P if the sale takes place but loses $U(\theta)$ from giving up the car.

There are various ways to specify $F(\theta)$, $U(\theta)$, and $V(\theta)$. We start with identical tastes and two types (Lemons I), and generalize to a continuum of types (Lemons II). Section 9.3 specifies first that the sellers are identical and value cars more than buyers (Lemons III), next that the sellers have heterogeneous tastes (Lemons IV). We will look less formally at other modifications involving risk aversion and the relative numbers of buyers and sellers.

Lemons I: Identical Tastes, Two Types of Sellers

Let good cars have quality 6,000 and bad cars (lemons) quality 2,000, so $\theta \in \{2,000, 6,000\}$, and suppose that half the cars in the world are of the first type and the other half of the second type. A payoff combination of (0,0) will represent the status quo, in which the buyer has \$50,000 and the seller has the car. Assume that both players are risk neutral and they value quality at one dollar per unit, so after a trade the payoffs are $\pi_{buyer} = \theta - P$ and $\pi_{seller} = P - \theta$. The extensive form is shown in Figure 1.

Figure 1: An Extensive Form for Lemons I



If he could observe quality at the time of his purchase, the buyer would be willing to accept a contract to pay \$6,000 for a good car and \$2,000 for a lemon. He cannot observe quality, and we assume that he cannot enforce a contract based on his discoveries once the purchase is made. Given these restrictions, if the seller offers \$4,000, a price equal to the average quality, the buyer will deduce that the seller does not have a good car. The very fact that the car is for sale demonstrates its low quality. Knowing that for \$4,000 he would be sold only lemons, the buyer would refuse to pay more than \$2,000. Let us assume that an indifferent seller sells his car, in which case half of the cars are traded in equilibrium, all of them lemons.

A friendly advisor might suggest to the owner of a good car that he wait until all the lemons have been sold and then sell his own car, since everyone knows that only good cars have remained unsold. But allowing for such behavior changes the model by adding a new action. If it were anticipated, the owners of lemons would also hold back and wait for the price to rise. Such a game could be formally analyzed as a War of Attrition (Section 3.2).

The outcome that half the cars are held off the market is interesting, though not startling, since half the cars do have genuinely higher quality. It is a formalization of Groucho Marx's wisecrack that he would refuse to join any club that would accept him as a member. Lemons II will have a more dramatic outcome.

Lemons II: Identical Tastes, a Continuum of Types of Sellers

One might wonder whether the outcome of Lemons I was an artifact of the assumption of just two types. Lemons II generalizes the game by allowing the seller to be any of a continuum of types. We will assume that the quality types are uniformly distributed between 2,000 and 6,000. The average quality is $\bar{\theta} = 4,000$, which is therefore the price the buyer would be willing to pay for a car of unknown quality if all cars were on the market. The probability density is zero except on the support $[2,000, 6,000]$, where it is $f(\theta) = 1/(6,000 - 2,000)$, and the cumulative density is

$$F(\theta) = \int_{2,000}^{\theta} f(x)dx. \quad (4)$$

After substituting the uniform density for $f(\theta)$ and integrating (1) we obtain


$$F(\theta) = \frac{\theta}{4,000} - 0.5. \quad (5)$$

The payoff functions are the same as in Lemons I.

The equilibrium price must be less than \$4,000 in Lemons II because, as in Lemons I, not all cars are put on the market at that price. Owners are willing to sell only if the quality of their cars is less than 4,000, so while the average quality of all used cars is 4,000, the average quality offered for sale is 3,000. The price cannot be \$4,000 when the average quality is 3,000, so the price must drop at least to \$3,000. If that happens, the owners of cars with values from 3,000 to 4,000 pull their cars off the market and the average of those remaining is 2,500. The acceptable price falls to \$2,500, and the unravelling continues until the price reaches its equilibrium level of \$2,000. But at $P = 2,000$ the number of cars on the market is infinitesimal. The market has completely collapsed!

Figure 2 puts the price of used cars on one axis and the average quality of cars offered for sale on the other. Each price leads to a different average quality, $\bar{\theta}(P)$, and the slope of $\bar{\theta}(P)$ is greater than one because average quality does not rise proportionately with price. If the price rises, the quality of the *marginal* car offered for sale equals the new price, but the quality of the *average* car offered for sale is much lower. In equilibrium, the average quality must equal the price, so the equilibrium lies on the 45° line through the origin. That line is a demand schedule of sorts, just as $\bar{\theta}(P)$ is a supply schedule. The only intersection is the point (\$2,000, 2,000).

Figure 2: Lemons II: Identical Tastes



9.3 Heterogeneous Tastes: Lemons III and IV


The outcome that no cars are traded is extreme, but there is no efficiency loss in either Lemons I or Lemons II. Since all the players have identical tastes, it does not matter who ends up owning the cars. But the players of this section, whose tastes differ, have real need of a market.

Lemons III : Buyers Value Cars More than Sellers

Assume that sellers value their cars at exactly their qualities θ , but that buyers have valuations 20 percent greater, and, moreover, outnumber the sellers. The payoffs if a trade occurs are $\pi_{buyer} = 1.2\theta - P$ and $\pi_{seller} = P - \theta$. In equilibrium, the sellers will capture the gains from trade.

In Figure 3, the curve $\bar{\theta}(P)$ is much the same as in Lemons II, but the equilibrium condition is no longer that price and average quality lie on the 45° line, but that they lie on the demand schedule $P(\bar{\theta})$, which has a slope of 1.2 instead of 1.0. The demand and supply schedules intersect only at $(P = \$3,000, \bar{\theta}(P) = 2,500)$. Because buyers are willing to pay a premium, we only see **partial adverse selection**; the equilibrium is partially pooling. The outcome is inefficient, because in a world of perfect information all the cars would be owned by the “buyers,” who value them more, but under adverse selection they only end up owning the low-quality cars.

**Figure 3: Adverse Selection When Buyers Value Cars More Than Sellers:
Lemons III**



Lemons IV : Sellers' Valuations Differ

In Lemons IV, we dig a little deeper to explain why trade occurs, and we model sellers as consumers whose valuations of quality have changed since they bought their cars. For a particular seller, the valuation of one unit of quality is $1 + \varepsilon$, where the random disturbance ε can be either positive or negative and has an expected value of zero. The disturbance could arise because of the seller's mistake—he did not realize how much he would enjoy driving when he bought the car—or because conditions changed—he switched to a job closer to home. Payoffs if a trade occurs are $\pi_{buyer} = \theta - P$ and $\pi_{seller} = P - (1 + \varepsilon)\theta$.

If $\varepsilon = -0.15$ and $\theta = 2,000$, then \$1,700 is the lowest price at which the player would resell his car. The average quality of cars offered for sale at price P is the expected quality of cars valued by their owners at less than P , i.e.,

$$\bar{\theta}(P) = E(\theta \mid (1 + \varepsilon)\theta \leq P). \quad (6)$$

Suppose that a large number of new buyers, greater in number than the sellers, appear in the market, and let their valuation of one unit of quality be \$1. The demand schedule, shown in Figure 4, is the 45° line through the origin. Figure 4 shows one possible shape for the supply schedule $\bar{\theta}(P)$, although to specify it precisely we would have to specify the distribution of the disturbances.




Figure 4: Lemons IV: Sellers' Valuations Differ

In contrast to Lemons I, II, and III, here if $P \geq \$6,000$ some car owners would be reluctant to sell, because they received positive disturbances to their valuations. The average quality of cars on the market is less than 4,000 even at $P = \$6,000$. On the other hand, even if $P = \$2,000$ some sellers with low quality cars *and* negative realizations of the disturbance still sell, so the average quality remains above 2,000. Under some distributions of ε , a few sellers hate their cars so much they would pay to have them taken away.

The equilibrium drawn in Figure 4 is $(P = \$2,600, \bar{\theta} = 2,600)$. Some used cars are sold, but the number is inefficiently low. Some of the sellers have high quality cars but negative disturbances, and although they would like to sell their cars to someone who values them more, they will not sell at a price of \$2,600.

A theme running through all four Lemons models is that when quality is unknown to the buyer, less trade occurs. Lemons I and II show how trade diminishes, while Lemons III and IV show that the disappearance can be inefficient because some sellers value cars less than some buyers. Next we will use Lemons III, the simplest model with gains from trade, to look at various markets with more sellers than buyers, excess supply, and risk-averse buyers.

More Sellers than Buyers

In analyzing *Lemons III*, we assumed that buyers outnumbered sellers. As a result, the sellers earned producer surplus. In the original equilibrium, all the sellers with quality less than 3,000 offered a price of \$3,000 and earned a surplus of up to \$1000. There were more buyers than sellers, so every seller who wished to sell was able to do so, but the price equalled the buyers' expected utility, so no buyer who failed to purchase was dissatisfied. The market cleared.

If, instead, sellers outnumber buyers, what price should a seller offer? At \$3,000, not all would-be sellers can find buyers. A seller who proposed a lower price would find willing buyers despite the somewhat lower expected quality. The buyer's tradeoff between lower price and lower quality is shown in Figure 3, in which the expected consumer surplus is the vertical distance between the price (the height of the supply schedule) and the demand schedule. When the price is \$3,000 and the average quality is 2,500, the buyer expects a consumer surplus of zero, which is $\$3,000 - \$1.2 \cdot 2,500$. The combination of price and quality that buyers like best is (\$2,000, 2,000), because if there were enough sellers with quality $\theta = 2,000$ to satisfy the demand, each buyer would pay $P = \$2,000$ for a car worth \$2,400 to him, acquiring a surplus of \$400. If there were fewer sellers, the equilibrium price would be higher and some sellers would receive producer surplus.

Heterogeneous Buyers: Excess Supply

If buyers have different valuations for quality, the market might not clear, as Charles Wilson (1980) points out. Assume that the number of buyers willing to pay \$1.2 per unit of quality exceeds the number of sellers, but that buyer Smith is an eccentric whose demand for high quality is unusually strong. He would pay \$100,000 for a car of quality 5,000 or greater, and \$0 for a car of any lower quality.

In Lemons III without Smith, the outcome is a price of \$3,000, an average market quality of 2,500, and a market quality range between 2,000 and 3,000. Smith would be unhappy with this, since he has zero probability of finding a car he likes. In fact, he would be willing to accept a price of \$6,000, so that all the cars, from quality 2,000 to 6,000, would be offered for sale and the probability that he buys a satisfactory car would rise from 0 to 0.25. But Smith would not want to buy all the cars offered to him, so the equilibrium has two prices, \$3,000 and \$6,000, with excess supply at the higher price.

Strangely enough, Smith's demand function is upward sloping. At a price of \$3,000, he is unwilling to buy; at a price of \$6,000, he is willing, because expected quality rises with price. This does not contradict basic price theory, for the standard assumption of *ceteris paribus* is violated. As the price increases, the quantity demanded would fall if all else stayed the same, but all else does not— quality rises.

Risk Aversion

We have implicitly assumed, by the choice of payoff functions, that the buyers and sellers are both risk neutral. What happens if they are risk averse— that is, if the marginal utilities of wealth and car quality are diminishing? Again we will use *Lemons III* and the assumption of many buyers.

On the seller's side, risk aversion changes nothing. The seller runs no risk because he knows exactly the price he receives and the quality he surrenders. But the buyer does bear risk, because he buys a car of uncertain quality. Although he would pay \$3,600 for a car he knows has quality 3,000, if he is risk averse he will not pay that much for a car with expected quality 3,000 but actual quality of possibly 2,500 or 3,500: he would obtain less utility from adding 500 quality units than from subtracting 500. The buyer would pay perhaps \$2,900 for a car whose expected quality is 3,000 where the demand schedule

is nonlinear, lying everywhere below the demand schedule of the risk-neutral buyer. As a result, the equilibrium has a lower price and average quality.

9.4 Adverse Selection under Uncertainty: *Insurance Game III*

The term “adverse selection,” like “moral hazard,” comes from insurance. Insurance pays more if there is an accident than otherwise, so it benefits accident-prone customers more than safe ones and a firm’s customers are “adversely selected” to be accident-prone. The classic article on adverse selection in insurance markets is Rothschild & Stiglitz (1976), which begins, “Economic theorists traditionally banish discussions of information to footnotes.” How things have changed! Within ten years, information problems came to dominate research in both microeconomics and macroeconomics.

We will follow Rothschild & Stiglitz in using state-space diagrams, and we will use a version of *The Insurance Game* of Section 8.5. Under moral hazard, Smith chose whether to be *Careful* or *Careless*. Under adverse selection, Smith cannot affect the probability of a theft, which is chosen by Nature. Rather, Smith is either *Safe* or *Unsafe*, and while he cannot affect the probability that his car will be stolen, he does know what the probability is.

Insurance Game III

Players

Smith and two insurance companies.

The Order of Play

- (0) Nature chooses Smith to be either *Safe*, with probability 0.6, or *Unsafe*, with probability 0.4. Smith knows his type, but the insurance companies do not.
- (1) Each insurance company offers its own contract (x, y) under which Smith pays premium x unconditionally and receives compensation y if there is a theft.
- (2) Smith picks a contract.
- (3) Nature chooses whether there is a theft, using probability 0.5 if Smith is *Safe* and 0.75 if he is *Unsafe*.

Payoffs.

Smith’s payoff depends on his type and the contract (x, y) that he accepts. Let $U' > 0$ and $U'' < 0$.

$$\pi_{\text{Smith}}(\text{Safe}) = 0.5U(12 - x) + 0.5U(0 + y - x).$$

$$\pi_{\text{Smith}}(\text{Unsafe}) = 0.25U(12 - x) + 0.75U(0 + y - x).$$

The companies’ payoffs depend on what types of customers accept their contracts, as shown in Table 1.

Table 1: Insurance Game III: Payoffs

Company payoff	Types of customers
0	no customers
$0.5x + 0.5(x - y)$	just <i>Safe</i>
$0.25x + 0.75(x - y)$	just <i>Unsafe</i>
$0.6[0.5x + 0.5(x - y)] + 0.4[0.25x + 0.75(x - y)]$	<i>Unsafe</i> and <i>Safe</i>

Smith is *Safe* with probability 0.6 and *Unsafe* with probability 0.4. Without insurance, Smith's dollar wealth is 12 if there is no theft and 0 if there is, depicted in Figure 5 as his endowment in state space, $\omega = (12, 0)$. If Smith is *Safe*, a theft occurs with probability 0.5, but if he is *Unsafe* the probability is 0.75. Smith is risk averse (because $U'' < 0$) and the insurance companies are risk neutral.




Figure 5: Insurance Game III: Nonexistence of a Pooling Equilibrium

If an insurance company knew that Smith was *Safe*, it could offer him insurance at a premium of 6 with a payout of 12 after a theft, leaving Smith with an allocation of (6, 6). This is the most attractive contract that is not unprofitable, because it fully insures Smith. Whatever the state, his allocation is 6.

Figure 5 shows the indifference curves of Smith and an insurance company. The insurance company is risk neutral, so its indifference curve is a straight line. If Smith will be a customer regardless of his type, the company's indifference curve based on its expected profits is ωF (although if the company knew that Smith was *Safe*, the indifference curve would be steeper, and if it knew he was *Unsafe*, the curve would be less steep). The insurance company is indifferent between ω and C_1 , at both of which its expected profits are zero. Smith is risk averse, so his indifference curves are convex, and closest to the origin along

the 45 degree if the probability of *Theft* is 0.5. He has two sets of indifference curves, solid if he is *Safe* and dotted if he is *Unsafe*.

Figure 5 shows why no Nash pooling equilibrium exists. To make zero profits, the equilibrium must lie on the line ωF . It is easiest to think about these problems by imagining an entire population of Smiths, whom we will call “customers.” Pick a contract C_1 anywhere on ωF and think about drawing the indifference curves for the *Unsafe* and *Safe* customers that pass through C_1 . *Safe* customers are always willing to trade *Theft* wealth for *No Theft* wealth at a higher rate than *Unsafe* customers. At any point, therefore, the slope of the solid (*Safe*) indifference curve is steeper than that of the dashed (*Unsafe*) curve. Since the slopes of the dashed and solid indifference curves differ, we can insert another contract, C_2 , between them and just barely to the right of ωF . The *Safe* customers prefer contract C_2 to C_1 , but the *Unsafe* customers stay with C_1 , so C_2 is profitable—since C_2 only attracts *Safes*, it need not be to the left of ωF to avoid losses. But then the original contract C_1 was not a Nash equilibrium, and since our argument holds for any pooling contract, no pooling equilibrium exists.

The attraction of the *Safe* customers away from pooling is referred to as **cream skimming**, although profits are still zero when there is competition for the cream. We next consider whether a separating equilibrium exists, using Figure 6. The zero profit condition requires that the *Safe* customers take contracts on ωC_4 and the *Unsafe*’s on ωC_3 .




Figure 6: A Separating Equilibrium for *Insurance Game III*

The *Unsafes* will be completely insured in any equilibrium, albeit at a high price. On the zero-profit line ωC_3 , the contract they like best is C_3 , which the *Safe*’s are not tempted to take. The *Safe*’s would prefer contract C_4 , but C_4 uniformly dominates C_3 , so it would attract *Unsafes* too, and generate losses. To avoid attracting *Unsafes*, the *Safe* contract must be below the *Unsafe* indifference curve. Contract C_5 is the fullest

insurance the *Safes* can get without attracting *Unsafes*: it satisfies the self-selection and competition constraints.

Contract C_5 , however, might not be an equilibrium either. Figure 7 is the same as Figure 6 with a few additional points marked. If one firm offered C_6 , it would attract both types, *Unsafe* and *Safe*, away from C_3 and C_5 , because it is to the right of the indifference curves passing through those points. Would C_6 be profitable? That depends on the proportions of the different types. The assumption on which the equilibrium of Figure 6 is based is that the proportion of *Safe*'s is 0.6, so that the zero-profit line for pooling contracts is ωF and C_6 would be unprofitable. In Figure 7 it is assumed that the proportion of *Safes* is higher, so the zero-profit line for pooling contracts would be $\omega F'$ and C_6 , lying to its left, is profitable. But we already showed that no pooling contract is Nash, so C_6 cannot be an equilibrium. Since neither a separating pair like (C_3, C_5) nor a pooling contract like C_6 is an equilibrium, no equilibrium whatsoever exists.




Figure 7: Curves for Which There is No Equilibrium in *Insurance Game III*

The essence of nonexistence here is that if separating contracts are offered, some company is willing to offer a superior pooling contract; but if a pooling contract is offered, some company is willing to offer a separating contract that makes it unprofitable. A monopoly would have a pure-strategy equilibrium, but in a competitive market only a mixed-strategy Nash equilibrium exists (see Dasgupta & Maskin [1986b]).

*9.5 Market Microstructure

The prices of securities such as stocks depend on what investors believe is the value of the

assets that underly them. The value are highly uncertain, and new information about them is constantly being generated. The market microstructure literature is concerned with how new information enters the market. In the paradigmatic situation, an informed trader has private information about the value which he hopes to use to make profitable trades, but other traders know that someone might have private information. This is a situation of adverse selection, because the informed trader has better information on the value of the stock, and no uninformed trader wants to trade with an informed trader. An institution that many markets have developed is that of the marketmaker or specialist, a trader in a particular stock who is always willing to buy or sell to keep the market going. Other traders feel safer in trading with the marketmaker than with a potentially informed trader, but this just transfers the adverse selection problem to the marketmaker.

The two models in this section will look at how a marketmaker deals with the problem of informed trading. Both are descendants of the verbal model in Bagehot (1971). (“Bagehot”, pronounced “badget”, is a pseudonym for Jack Treynor. See Glosten & Milgrom (1985) for a formalization.) In the Bagehot model, there may or may not be one or more informed traders, but the informed traders as a group have a trade of fixed size if they are present. The marketmaker must decide how big a bid-ask spread to charge. In the Kyle model, there is one informed trader, who decides how much to trade. On observing the imbalance of orders, the marketmaker decides what price to offer.

The Bagehot model is perhaps a better explanation of why marketmakers might charge a bid/ask spread even under competitive conditions and with zero transactions costs. Its assumption is that the marketmaker cannot change the price depending on volume, but must instead offer a price, and then accept whatever order comes along—a buy order, or a sell order.

The Bagehot Model

Players

The informed trader and two competing marketmakers.

The Order of Play

- (0) Nature chooses the asset value v to be either $p_0 - \delta$ or $p_0 + \delta$ with equal probability. The marketmakers never observe the asset value, nor do they observe whether anyone else observes it, but the “informed” trader observes v with probability θ .
- (1) The marketmakers choose their spreads s , offering prices $p_{bid} = p_0 - \frac{s}{2}$ at which they will buy the security and $p_{ask} = p_0 + \frac{s}{2}$ for which they will sell it.
- (2) The informed trader decides whether to buy one unit, sell one unit, or do nothing.
- (3) Noise traders buy n units and sell n units.

Payoffs

Everyone is risk neutral. The informed trader’s payoff is $v - p_{ask}$ if he buys, $p_{bid} - v$ if he

sells, and zero if he does nothing. The marketmaker who offers the highest p_{bid} trades with all the customers who wish to sell, and the marketmaker who offers the lowest p_{ask} trades with all the customers who wish to buy. If the marketmakers set equal prices, they split the market evenly. A marketmaker who sells x units gets a payoff of $x(p_{ask} - v)$, and a marketmaker who buys x units gets a payoff of $x(v - p_{bid})$.

This is a very simple game. Competition between the marketmakers will make their prices identical and their profits zero. The informed trader should buy if $v > p_{ask}$ and sell if $v < p_{bid}$. He has no incentive to trade if $[p_{bid}, p_{ask}]$.

A marketmaker will always lose money trading with the informed trader, but if $s > 0$, so $p_{ask} > p_0$ and $p_{bid} < p_0$, he will earn positive expected profits trading with the noise traders. Since a marketmaker could specialize in either sales or purchases, he must earn zero expected profits overall from either type of trade. Centering the bid-ask spread on the expected value of the stock, p_0 , ensures this. Marketmaker sales will be at the ask price of $(p_0 + s/2)$. With probability 0.5, this is above the true value of the stock, $(p_0 - \delta)$, in which case the informed trader will not buy but the marketmakers will earn a total profit of $n[(p_0 + s/2) - (p_0 - \delta)]$ from the noise traders. With probability 0.5, the ask price of $(p_0 + s/2)$ is below the true value of the stock, $(p_0 + \delta)$, in which case the informed trader will be informed with probability θ and buy one unit and the noise traders will buy n more in any case, so the marketmakers will earn a total expected profit of $(n + \theta)[(p_0 + s/2) - (p_0 + \delta)]$, a negative number. For marketmaker profits from sales at the ask price to be zero overall, this expected profit must be set to zero:

$$.5n[(p_0 + s/2) - (p_0 - \delta)] + .5(n + \theta)[(p_0 + s/2) - (p_0 + \delta)] = 0 \quad (7)$$

This equation implies that $n[s/2 + \delta] + (n + \theta)[s/2 - \delta] = 0$, so

$$s^* = \frac{2\delta\theta}{2n + \theta}. \quad (8)$$

The profit from marketmaker purchases must similarly equal zero, and will for the same spread s^* , though we will not go through the algebra here.

Equation (8) has a number of implications. First, the spread s^* is positive. Even though marketmakers compete and have zero transactions costs, they charge a different price to buy and to sell. They make money dealing with the noise traders but lose money with the informed trader, if he is present. The comparative statics reflect this. s^* rises in δ , the variance of the true value, because divergent true values increase losses from trading with the informed trader, and s^* falls in n , which reflects the number of noise traders relative to informed traders, because when there are more noise traders, the profits from trading with them are greater. The spread s^* rises in θ , the probability that the informed trader really has inside information, which is also intuitive but requires a little calculus to demonstrate starting from equation (8):

$$\frac{\partial s^*}{\partial \theta} = \frac{2\delta}{2n + \theta} - \frac{2\delta\theta}{(2n + \theta)^2} = \left(\frac{1}{(2n + \theta)^2} \right) (4\delta n + 2\delta\theta - 2\delta\theta) > 0. \quad (9)$$

The second model of market microstructure, important because it is commonly used as a foundation for more complicated models, is the Kyle model, which focuses on the

decision of the informed trader, not the marketmaker. The Kyle model is set up so that marketmaker observes the trade volume before he chooses the price.

The Kyle Model (Kyle [1985])

Players

The informed trader and two competing marketmakers.

The Order of Play

- (0) Nature chooses the asset value v from a normal distribution with mean p_0 and variance σ_v^2 , observed by the informed trader but not by the marketmakers.
- (1) The informed trader offers a trade of size $x(v)$, which is a purchase if positive and a sale if negative, unobserved by the marketmaker.
- (2) Nature chooses a trade of size u by noise traders, unobserved by the marketmaker, where u is distributed normally with mean zero and variance σ_u^2 .
- (3) The marketmakers observe the total market trade offer $y = x + u$, and choose prices $p(y)$.
- (4) Trades are executed. If y is positive (the market wants to purchase, in net), whichever marketmaker offers the lowest price executes the trades; if y is negative (the market wants to sell, in net), whichever marketmaker offers the highest price executes the trades. v is then revealed to everyone.

Payoffs

All players are risk neutral. The informed trader's payoff is $(v - p)x$. The marketmaker's payoff is zero if he does not trade and $(p - v)y$ if he does.

An equilibrium for this game is the strategy combination

$$x(v) = (v - p_0) \left(\frac{\sigma_u}{\sigma_v} \right) \quad (10)$$

and

$$p(y) = p_0 + \left(\frac{\sigma_v}{2\sigma_u} \right) y. \quad (11)$$

This is reasonable. It says that the informed trader will increase the size of his trade as v gets bigger relative to p_0 (and he will sell, not buy, if $v - p_0 < 0$), and the marketmaker will increase the price he charges for selling if y is bigger, meaning that more people want to sell, which is an indicator that the informed trader might be trading heavily. The variances of the asset value (σ_v^2) and the noise trading (σ_u^2) enter as one would expect, and they matter only in their relation to each other. If $\frac{\sigma_v^2}{\sigma_u^2}$ is large, then the asset value fluctuates more than the amount of noise trading, and it is difficult for the informed trader to conceal his trades under the noise. The informed trader will trade less, and a given amount of trading

will cause a greater response from the marketmaker. One might say that the market is less “liquid”: a trade of given size will have a greater impact on the price.

I will not (and cannot) prove uniqueness of the equilibrium, since it is very hard to check all possible combinations of nonlinear strategies, but I will show that $\{(10), (11)\}$ is Nash and is the unique linear equilibrium. To start, hypothesize that the informed trader uses a linear strategy, so

$$x(v) = \alpha + \beta v \quad (12)$$

for some constants α and β . Competition between the marketmakers means that their expected profits will be zero, which requires that the price they offer be the expected value of v . Thus, their equilibrium strategy $p(y)$ will be an unbiased estimate of v given their data y , where they know that y is normally distributed and that

$$\begin{aligned} y &= x + u \\ &= \alpha + \beta v + u. \end{aligned} \quad (13)$$

This means that their best estimate of v given the data y is, following the usual regression rule (which readers unfamiliar with statistics must accept on faith),

$$\begin{aligned} E(v|y) &= E(v) + \left(\frac{\text{cov}(v,y)}{\text{var}(y)} \right) y \\ &= p_0 + \left(\frac{\beta \sigma_v^2}{\beta^2 \sigma_v^2 + \sigma_u^2} \right) y \\ &= p_0 + \lambda y, \end{aligned} \quad (14)$$

where λ is a new shorthand variable to save writing out the term in parentheses in what follows.

The function $p(y)$ will be a linear function of y under our assumption that x is a linear function of v . Given that $p(y) = p_0 + \lambda y$, what must next be shown is that x will indeed be a linear function of v . Start by writing the informed trader’s expected payoff, which is

$$\begin{aligned} E\pi_i &= E([v - p(y)]x) \\ &= E([v - p_0 - \lambda(x + u)]x) \\ &= [v - p_0 - \lambda(x + 0)]x, \end{aligned} \quad (15)$$

since $E(u) = 0$. Maximizing the expected payoff with respect to x gives the first order condition

$$v - p_0 - 2\lambda x = 0, \quad (16)$$

which on rearranging becomes

$$x = -\frac{p_0}{2\lambda} + \left(\frac{1}{2\lambda} \right) v. \quad (17)$$

Equation (17) establishes that $x(v)$ is linear, given that $p(y)$ is linear. All that is left is to find the value of λ . See by comparing (17) and (12) that $\beta = \frac{1}{2\lambda}$. Substituting this β into the value of λ from (14) gives

$$\lambda = \frac{\beta \sigma_v^2}{\beta^2 \sigma_v^2 + \sigma_u^2} = \frac{\frac{\sigma_v^2}{2\lambda}}{\frac{\sigma_v^2}{(4\lambda^2)} + \sigma_u^2}, \quad (18)$$

which upon solving for λ yields $\lambda = \frac{\sigma_v}{2\sigma_u}$. Since $\beta = \frac{1}{2\lambda}$, it follows that $\beta = \frac{\sigma_u}{\sigma_v}$. These values of λ and β together with equation (17) give the strategies asserted at the start in equations (10) and (11).

The two main divisions of the field of finance are corporate finance and asset pricing. Corporate finance, the study of such things as project choice, capital structure, and mergers has the most obvious applications of game theory, but the Bagehot and Kyle models show that the same techniques are also important in asset pricing. For more information, I recommend Harris & Raviv (1995).

***9.6 A Variety of Applications**

Price Dispersion

Usually the best model for explaining price dispersion is a search model— Salop & Stiglitz (1977), for example, which is based on buyers whose search costs differ. But although we passed over it quickly in Section 9.3, the Lemons model with Smith, the quality-conscious consumer, generated not only excess supply, but price dispersion as well. Cars of the same average quality were sold for \$3,000 and \$6,000.

Similarly, while the most obvious explanation for why brands of stereo amplifiers sell at different prices is that customers are willing to pay more for higher quality, adverse selection contributes another explanation. Consumers might be willing to pay high prices because they know that high-priced brands could include both high-quality and low-quality amplifiers, whereas low-priced brands are invariably low quality. The low-quality amplifier ends up selling at two prices: a high price in competition with high-quality amplifiers, and, in different stores or under a different name, a low price aimed at customers less willing to trade dollars for quality.

This explanation does depend on sellers of amplifiers incurring a large enough fixed set-up or operating cost. Otherwise, too many low-quality brands would crowd into the market, and the proportion of high-quality brands would be too small for consumers to be willing to pay the high price. The low-quality brands would benefit as a group from entry restrictions: too many of them spoil the market, not through price competition but through degrading the average quality.

Health Insurance

Medical insurance is subject to adverse selection because some people are healthier than others. The variance in health is particularly high among old people, who sometimes have difficulty in obtaining insurance at all. Under basic economic theory this is a puzzle: the price should rise until supply equals demand. The problem is pooling: when the price of insurance is appropriate for the average old person, healthier ones stop buying. The price

must rise to keep profits nonnegative, and the market disappears, just as in Lemons II.

If the facts indeed fit this story, adverse selection is an argument for government-enforced pooling. If all old people are required to purchase government insurance, then while the healthier of them may be worse off, the vast majority could be helped.

Using adverse selection to justify medicare, however, points out how dangerous many of the models in this book can be. For policy questions, the best default opinion is that markets are efficient. On closer examination, we have found that many markets are inefficient because of strategic behavior or information asymmetry. It is dangerous, however, to immediately conclude that the government should intervene, because the same arguments applied to government show that the cure might be worse than the disease. The analyst of health care needs to take seriously the moral hazard and rent-seeking that arise from government insurance. Doctors and hospitals will increase the cost and amount of treatment if the government pays for it, and the transfer of wealth from young people to the elderly, which is likely to swamp the gains in efficiency, might distort the shape of the government program from the economist's ideal.

Henry Ford's Five-Dollar Day

In 1914 Henry Ford made a much-publicized decision to raise the wage of his auto workers to \$5 a day, considerably above the market wage. This pay hike occurred without pressure from the workers, who were non-unionized. Why did Ford do it?

The pay hike could be explained by either moral hazard or adverse selection. In accordance with the idea of efficiency wages (Section 8.1), Ford might have wanted workers who worried about losing their premium job at his factory, because they would work harder and refrain from shirking. Adverse selection could also explain the pay hike: by raising his wage Ford attracted a mixture of low -and high- quality workers, rather than low-quality alone (see Raff & Summers [1987]).

Bank Loans

Suppose that two people come to you for an unsecured loan of \$10,000. One offers to pay an interest rate of 10 percent and the other offers 200 percent. Who do you accept? Like the car buyer who chooses to buy at a high price, you may choose to lend at a low interest rate.

If a lender raises his interest rate, both his pool of loan applicants and their behavior change because adverse selection and moral hazard contribute to a rise in default rates. Borrowers who expect to default are less concerned about the high interest rate than dependable borrowers, so the number of loans shrinks and the default rate rises (see Stiglitz & Weiss [1981]). In addition, some borrowers shift to higher-risk projects with greater chance of default but higher yields when they are successful. In Section 6.6 we will go through the model of D. Diamond (1989) which looks at the evolution of this problem as firms age.

Whether because of moral hazard or adverse selection, asymmetric information can

also result in excess demand for bank loans. The savers who own the bank do not save enough at the equilibrium interest rate to provide loans to all the borrowers who want loans. Thus, the bank makes a loan to John Smith, while denying one to Joe, his observationally equivalent twin. Policymakers should carefully consider any laws that rule out arbitrary loan criteria or require banks to treat all customers equally. A bank might wish to restrict its loans to left-handed people, neither from prejudice nor because it is useful to ration loans according to some criterion arbitrary enough to avoid the moral hazard of favoritism by loan officers.

Bernanke (1983) suggests adverse selection in bank loans as an explanation for the Great Depression in the United States. The difficulty in explaining the Depression is not so much the initial stock market crash as the persistence of the unemployment that followed. Bernanke notes that the crash wiped out local banks and dispersed the expertise of the loan officers. After the loss of this expertise, the remaining banks were less willing to lend because of adverse selection, and it was difficult for the economy to recover.

Solutions to Adverse Selection

Even in markets where it apparently does not occur, the threat of adverse selection, like the threat of moral hazard, can be an important influence on market institutions. Adverse selection can be circumvented in a number of ways besides the contractual solutions we have been analyzing. I will mention some of them in the context of the used car market.

One set of solutions consists of ways to make car quality contractible. Buyers who find that their car is defective may have recourse to the legal system if the sellers were fraudulent, although in the United States the courts are too slow and costly to be fully effective. Other government bodies such as the Federal Trade Commission may do better by issuing regulations particular to the industry. Even without regulation, private warranties—promises to repair the car if it breaks down—may be easier to enforce than oral claims, by dispelling ambiguity about what level of quality is guaranteed.

Testing (the equivalent of moral hazard's monitoring) is always used to some extent. The prospective driver tries the car on the road, inspects the body, and otherwise tries to reduce information asymmetry. At a cost, he could even reverse the asymmetry by hiring mechanics, learning more about the car than the owner himself. The rule is not always *caveat emptor*; what should one's response be to an antique dealer who offers to pay \$500 for an apparently worthless old chair?

Reputation can solve adverse selection, just as it can solve moral hazard, but only if the transaction is repeated and the other conditions of the models in Chapters 5 and 6 are met. An almost opposite solution is to show that there are innocent motives for a sale; that the owner of the car has gone bankrupt, for example, and his creditor is selling the car cheaply to avoid the holding cost.

Penalties not strictly economic are also important. One example is the social ostracism inflicted by the friend to whom a lemon has been sold; the seller is no longer invited to dinner. Or, the seller might have moral principles that prevent him from defrauding buyers. Such principles, provided they are common knowledge, would help him obtain a higher

price in the used-car market. Akerlof himself has worked on the interaction between social custom and markets in his 1980 and 1983 articles. The second of these looks directly at the value of inculcating moral principles, using theoretical examples to show that parents might wish to teach their children principles, and that society might wish to give hiring preference to students from elite schools.

It is by violating the assumptions needed for perfect competition that asymmetric information enables government and social institutions to raise efficiency. This points to a major reason for studying asymmetric information: where it is important, noneconomic interference can be helpful instead of harmful. I find the social solutions particularly interesting since, as mentioned earlier in connection with health care, government solutions introduce agency problems as severe as the information problems they solve. Noneconomic behavior is important under adverse selection, in contrast to perfect competition, which allows an “Invisible Hand” to guide the market to efficiency, regardless of the moral beliefs of the traders. If everyone were honest, the lemons problem would disappear because the sellers would truthfully disclose quality. If some fraction of the sellers were honest, but buyers could not distinguish them from the dishonest sellers, the outcome would presumably be somewhere between the outcomes of complete honesty and complete dishonesty. The subject of market ethics is important, and would profit from investigation by scholars trained in economic analysis.

Notes

N9.1 Introduction: Production Game VI

- For an example of an adverse selection model in which workers also choose effort level, see Akerlof (1976) on the “rat race.” The model is not moral hazard, because while the employer observes effort, the worker’s types— their utility costs of hard work— are known only to themselves.
- In moral hazard with hidden knowledge, the contract must ordinarily satisfy only one participation constraint, whereas in adverse selection problems there is a different participation constraint for each type of agent. An exception is if there are constraints limiting how much an agent can be punished in different states of the world. If, for example, there are bankruptcy constraints, then, if the agent has different wealths across the N possible states of the world, there will be N constraints for how negative his wage can be, in addition to the single participation constraint. These can be looked at as **interim** participation constraints, since they represent the idea that the agent wants to get out of the contract once he observes the state of the world midway through the game.
- Gresham’s Law (“Bad money drives out good”) is a statement of adverse selection. Only debased money will be circulated if the payer knows the quality of his money better than the receiver. The same result occurs if quality is common knowledge, but for legal reasons the receiver is obligated to take the money, whatever its quality. An example of the first is Roman coins with low silver content; and of the second, Zambian currency with an overvalued exchange rate.
- Most adverse selection models have types that could be called “good” and “bad,” because one type of agent would like to pool with the other, who would rather be separate. It is also possible to have a model in which both types would rather separate— types of workers who prefer night shifts versus those who prefer day shifts, for example— or two types who both prefer pooling— male and female college students.
- Two curious features of labor markets is that workers of widely differing outputs seem to be paid identical wages and that tests are not used more in hiring decisions. Schmidt and Judiesch (as cited in Seligman [1992], p. 145) have found that in jobs requiring only unskilled and semi-skilled blue-collar workers, the top 1 percent of workers, as defined by performance on ability tests not directly related to output, were 50 percent more productive than the average. In jobs defined as “high complexity” the difference was 127 percent.

At about the same time as Akerlof (1970), another seminal paper appeared on adverse selection, Mirrlees (1971), although the relation only became clear later. Mirrlees looked at optimal taxation and the problem of how the government chooses a tax schedule given that it cannot observe the abilities of its citizens to earn income, and this began the literature on mechanism design. Used cars and income taxes do not appear similar, but in both situations an uninformed player must decide how to behave to another player whose type he does not know. Section 10.4 sets out a descendant of Mirrlees (1971) in a model of government procurement: much of government policy is motivated by the desire to create incentives for efficiency at minimum cost while eliciting information from individuals with superior information.

N9.2 Adverse Selection under Certainty: Lemons I and II

- Dealers in new cars and other durables have begun offering “extended-service contracts” in recent years. These contracts, offered either by the manufacturers or by independent companies, pay for repairs after the initial warranty expires. For reasons of moral hazard or adverse selection, the contracts usually do not cover damage from accidents. Oddly enough, they also do not cover items like oil changes despite their usefulness in prolonging engine life. Such contracts have their own problems, as shown by the fact that several of the independent companies went bankrupt in the late 1970s and early 1980s, making their contracts worthless. See “Extended-Service Contracts for New Cars Shed Bad Reputation as Repair Bills Grow,” *Wall Street Journal*, June 10, 1985, p. 25.
- Suppose that the cars of Lemons II lasted two periods and did not physically depreciate. A naive economist looking at the market would see new cars selling for \$6,000 (twice \$3,000) and old cars selling for \$2,000 and conclude that the service stream had depreciated by 33 percent. Depreciation and adverse selection are hard to untangle using market data.
- Lemons II uses a uniform distribution. For a general distribution F , the average quality $\bar{\theta}(P)$ of cars with quality P or less is

$$\bar{\theta}(P) = E(\theta | \theta \leq P) = \frac{\int_{-\infty}^P x F'(x) dx}{F(P)}. \quad (19)$$

Equation (19) also arises in physics (equation for a center of gravity) and nonlinear econometrics (the likelihood equation). Think of $\bar{\theta}(P)$ as a weighted average of the values of θ up to P , the weights being densities. Having multiplied by all these weights in the numerator, we have to divide by their “sum,” $F(P) = \int_{-\infty}^P F'(x) dx$, in the denominator, giving rise to equation (19).

N9.3 Heterogeneous Tastes: Lemons III and IV


- You might object to a model in which the buyers of used cars value quality more than the sellers, since the sellers are often richer people. Remember that quality here is “quality of used cars,” which is different from “quality of cars.” The utility functions could be made more complicated without abandoning the basic model. We could specify something like $\pi_{\text{buyer}} = \theta + k/\theta - P$, where $\theta^2 > k$. Such a specification implies that the lower is the quality of the car, the greater the difference between the valuations of buyer and seller.
- In Akerlof (1970) the quality of new cars is uniformly distributed between zero and 2, and the model is set up differently, with the demand and supply curves offered by different types of traders and net supply and gross supply presented rather confusingly. Usually the best way to model a situation in which traders sell some of their endowment and consume the rest is to use only gross supplies and demands. Each old owner supplies his car to the market, but in equilibrium he might buy it back, having better information about his car than the other consumers. Otherwise, it is easy to count a given unit of demand twice, once in the demand curve and once in the net supply curve.
- See Stiglitz (1987) for a good survey of the relation between price and quality. Leibenstein (1950) uses diagrams to analyze the implications of individual demand being linked to the market price of quantity in markets for “bandwagon,” “snob,” and “Veblen” goods. See also “Pricing of Products is Still an Art, Often Having Little Link to Costs,” *Wall Street Journal*, p. 29 (25 November 1981).

- Risk aversion is concerned only with variability of outcomes, not their level. If the quality of used cars ranges from 2,000 to 6,000, buying a used car is risky. If all used cars are of quality 2,000, buying a used car is riskless, because the buyer knows exactly what he is getting.

In Insurance Game III in Section 9.4, the separating contract for the *Unsafe* consumer fully insures him: he bears no risk. But in constructing the equilibrium, we had to be very careful to keep the *Unsafes* from being tempted by the risky contract designed for the *Safes*. Risk is a bad thing, but as with old age, the alternative is worse. If Smith were certain his car would be stolen, he would bear no risk, because he would be certain to have low utility.

- To the buyers in Lemons IV, the average quality of cars for a given price is stochastic because they do not know which values of ε were realized. To them, the curve $\bar{\theta}(P)$ is only the *expectation* of the average quality.
- **Lemons III': Minimum Quality of Zero.** If the minimum quality of car in Lemons III were 0, not 2,000, the resulting game (Lemons III') would be close to the original Akerlof (1970) specification. As Figure 8 shows, the supply schedule and the demand schedule intersect at the origin, so that the equilibrium price is zero and no cars are traded. The market has shut down entirely because of the unravelling effect described in Lemons II. Even though the buyers are willing to accept a quality lower than the dollar price, the price that buyers are willing to pay does not rise with quality as fast as the price needed to extract that average quality from the sellers, and a car of minimum quality is valued exactly the same by buyers and sellers. A 20 percent premium on zero is still zero. The efficiency implications are even stronger than before, because at the optimum all the old cars are sold to new buyers, but in equilibrium, none are.

Figure 8: Lemons III' When Buyers Value Cars More and the Minimum Quality is Zero



N9.4 Adverse Selection under Uncertainty: Insurance Game III

- Markets with two types of customers are very common in insurance, because it is easy to distinguish male from female, both those types are numerous, and the difference between them is important. Males under age 25 pay almost twice the auto insurance premiums of females, and females pay 10 to 30 percent less for life insurance. The difference goes both ways, however: Aetna charges a 35-year old woman 30 to 50 percent more than a man for medical insurance. One market in which rates do not differ much is disability insurance. Women do make more claims, but the rates are the same because relatively few women buy the product (*Wall Street Journal*, p. 21, 27 August 1987).

N9.6 A Variety of Applications

- Economics professors sometimes make use of self-selection for student exams. One of my colleagues put the following instructions on an MBA exam, after stating that either Question 5 or 6 must be answered.

“The value of Question 5 is less than that of Question 6. Question 5, however, is straightforward and the average student may expect to answer it correctly. Question 6 is more tricky: only those who have understood and absorbed the content of the course well will be able to answer it correctly... For a candidate to earn a final course grade of A or higher, it will be *necessary* for him to answer Question 6 successfully.”

Making the question even more self-referential, he asked the students for an explanation of its purpose.

Another of my colleagues tried asking who in his class would be willing to skip the exam and settle for an A–. Those students who were willing, received an A–. The others got A’s. But nobody had to take the exam (this method did upset a few people). More formally, Guasch & Weiss (1980) have looked at adverse selection and the willingness of workers with different abilities to take tests.

- Nalebuff & Scharfstein (1987) have written on testing, generalizing Mirrlees (1974), who showed how a forcing contract in which output is costlessly observed might attain efficiency by punishing only for very low output. In Nalebuff & Scharfstein, testing is costly and agents are risk averse. They develop an equilibrium in which the employer tests workers with small probability, using high-quality tests and heavy punishments to attain almost the first-best. Under a condition which implies that large expenditures on each test can eliminate false accusations, they show that the principal will test workers with small probability, but use expensive, accurate tests when he does test a worker, and impose a heavy punishment for lying.

Problems

9.1. Insurance with Equations and Diagrams

The text analyzes Insurance Game III using diagrams. Here, let us use equations too. Let $U(t) = \log(t)$.

- (a) Give the numeric values (x, y) for the full-information separating contracts C_3 and C_4 from Figure 6. What are the coordinates for C_3 and C_4 ?
- (b) Why is it not necessary to use the $U(t) = \log(t)$ function to find the values?
- (c) At the separating contract under incomplete information, C_5 , $x = 2.01$. What is y ? Justify the value 2.01 for x . What are the coordinates of C_5 ?
- (d) What is a contract C_6 that might be profitable and that would lure both types away from C_3 and C_5 ?

9.2: Testing and Commitment. Fraction β of workers are talented, with output $a_t = 5$, and fraction $(1 - \beta)$ are untalented, with output $a_u = 0$. Both types have a reservation wage of 1 and are risk neutral. At a cost of 2 to itself and 1 to the job applicant, employer Apex can test a job applicant and discover his true ability with probability θ , which takes a value of something over 0.5. There is just one period of work. Let $\beta = 0.001$. Suppose that Apex can commit itself to a wage schedule before the workers take the test, and that Apex must test all applicants and pay all the workers it hires the same wage, to avoid grumbling among workers and corruption in the personnel division.

- (a) What is the lowest wage, w_t , that will induce talented workers to apply? What is the lowest wage, w_u , that will induce untalented workers to apply? Which is greater?
- (b) What is the minimum accuracy value θ that will induce Apex to use the test? What are the firm's expected profits per worker who applies?
- (c) Now suppose that Apex can pay w_p to workers who pass the test and w_f to workers who flunk. What are w_p and w_f ? What is the minimum accuracy value θ that will induce Apex to use the test? What are the firm's expected profits per worker who applies?
- (d) What happens if Apex cannot commit to paying the advertised wage, and can decide each applicant's wage individually?
- (e) If Apex cannot commit to testing every applicant, why is there no equilibrium in which either untalented workers do not apply or the firm tests every applicant?

9.3. Finding the Mixed-Strategy Equilibrium in a Testing Game

Half of high school graduates are talented, producing output $a = x$, and half are untalented, producing output $a = 0$. Both types have a reservation wage of 1 and are risk neutral. At a cost of 2 to himself and 1 to the job applicant, an employer can test a graduate and discover his true ability. Employers compete with each other in offering wages but they cooperate in revealing test results, so an employer knows if an applicant has already been tested and failed. There is just one period of work. The employer cannot commit to testing every applicant or any fixed percentage of them.

- (a) Why is there no equilibrium in which either untalented workers do not apply or the employer tests every applicant?
- (b) In equilibrium, the employer tests workers with probability γ and pays those who pass the test w , the talented workers all present themselves for testing, and the untalented workers present themselves with probability α , where possibly $\gamma = 1$ or $\alpha = 1$. Find an expression for the equilibrium value of α in terms of w . Explain why α is not directly a function of x in this expression, even though the employer's main concern is that some workers have a productivity advantage of x .
- (c) If $x = 9$, what are the equilibrium values of α , γ , and w ?
- (d) If $x = 8$, what are the equilibrium values of α , γ , and w ?

9.4: Two-Time Losers. Some people are strictly principled and will commit no robberies, even if there is no penalty. Others are incorrigible criminals and will commit two robberies, regardless of the penalty. Society wishes to inflict a certain penalty on criminals as retribution. Retribution requires an expected penalty of 15 per crime (15 if detection is sure, 150 if it has probability 0.1, etc.). Innocent people are sometimes falsely convicted, as shown in Table 2.

Table 2: Two-Time Losers

Robberies	Convictions		
	0	1	2
0	0.81	0.18	0.01
2	0.49	0.42	0.09

Two systems are proposed: (i) a penalty of X for each conviction, and (ii) a penalty of 0 for the first conviction, and some amount P for the second conviction.

- (a) What must X and P be to achieve the desired amount of retribution?
- (b) Which system inflicts the smaller cost on innocent people? How much is the cost in each case?
- (c) Compare this with Problem 8.2. How are they different?

9.5. Insurance and State-Space Diagrams

Two types of risk-averse people, clean-living and dissolute, would like to buy health insurance. Clean-living people become sick with probability 0.3, and dissolute people with probability 0.9. In state-space diagrams with the person's wealth if he is healthy on the vertical axis and if he is sick on the horizontal, every person's initial endowment is (5,10), because his initial wealth is 10 and the cost of medical treatment is 5.

- (a) What is the expected wealth of each type of person?

- (b) Draw a state-space diagram with the indifference curves for a risk-neutral insurance company that insures each type of person separately. Draw in the post-insurance allocations C_1 for the dissolute and C_2 for the clean-living under the assumption that a person's type is contractible.
- (c) Draw a new state-space diagram with the initial endowment and the indifference curves for the two types of people that go through that point.
- (d) Explain why, under asymmetric information, no pooling contract C_3 can be part of a Nash equilibrium.
- (e) If the insurance company is a monopoly, can a pooling contract be part of a Nash equilibrium?

10 Mechanism Design ¹

10.1 The Revelation Principle and Moral Hazard with Hidden Knowledge

In Chapter 10 we will look at mechanism design. A mechanism is a set of rules that one player constructs and another freely accepts in order to convey information from the second player to the first. Thus, a mechanism consists of an information report by the second player and a mapping from each possible report to some action by the first.

Adverse selection models can be viewed as mechanism design. *Insurance Game III* was about an insurance company which wanted to know whether a customer was safe or not. In equilibrium it offers two contracts, an expensive full insurance contract preferred by the safe customers and a cheap partial insurance contract preferred by the unsafe. A mechanism design view is that the insurance company sets up a game in which a customer reports his type as Safe or Unsafe, whichever he prefers to report, and the company then assigns him either partial or full insurance as a consequence. The contract is a mechanism for getting the agents to truthfully report their types.

Mechanism design goes beyond simple adverse selection. It can be useful even when players begin a game with symmetric information or when both players have hidden information that they would like to exchange.

Section 10.1 introduces moral hazard with hidden knowledge and discusses a modelling simplification called the Revelation Principle and a paradox known as Unravelling. Section 10.2 uses diagrams to apply the model to sales quotas, and Section 10.3 uses a product quality game of Roger Myerson's to compare moral hazard with hidden information with adverse selection. Section 10.4 applies the principles of mechanism design to price discrimination. Section 10.4 contains a more complicated model of rate-of-return regulation by a government that constructs a mechanism to induce a regulated company to reveal how high its costs are. Section 10.4 introduces a multilateral mechanism, the Groves Mechanism, for use when the problem is to elicit truthful reports from not one but N agents who need to decide whether to invest in a public good.

Moral Hazard with Hidden Knowledge

¹xxx This chapter is unsatisfactory. The explanations are too difficult, and perhaps repetitive. Things need to be unified more. The order might be rearranged too. The notation should be unified.

Information is complete in moral hazard games, but in moral hazard with hidden knowledge the agent, but not the principal, observes a move of Nature after the game begins. Information is symmetric at the time of contracting but becomes asymmetric later. From the principal's point of view, agents are identical at the beginning of the game but develop private types midway through, depending on what they have seen. His chief concern is to give them incentives to disclose their types later, which gives games with hidden knowledge a flavor close to that of adverse selection. (In fact, an alternative name for this might be **post-contractual adverse selection**.) The agent might exert effort, but effort's contractibility is less important when the principal does not know which effort is appropriate because he is ignorant of the state of the world chosen by Nature. The main difference technically is that if information is symmetric at the start and only becomes asymmetric after a contract is signed, the participation constraint is based on the agent's expected payoffs across the different types of agent he might become. Thus, there is just one participation constraint even if there are eventually n possible types of agents in the model, rather than the n participation constraints that would be required in a standard adverse selection model.

There is more hope for obtaining efficient outcomes in moral hazard with hidden knowledge than in adverse selection or simple moral hazard. The advantage over adverse selection is that information is symmetric at the time of contracting, so neither player can use private information to extract surplus from the other by choosing inefficient contract terms. The advantage over simple moral hazard is that the post-contractual asymmetry is with respect to knowledge only, which is neutral in itself, rather than over whether the agent exerted high effort, which causes direct disutility to him.

For a comparison between the two types of moral hazard, let us modify *Production Game V* from Section 7.2 and turn it into a game of hidden knowledge.

Production Game VII: Hidden Knowledge

Players

The principal and the agent.

The Order of Play

- 1 The principal offers the agent a wage contract of the form $w(q, m)$, where q is output and m is a message to be sent by the agent.
- 2 The agent accepts or rejects the principal's offer.
- 3 Nature chooses the state of the world θ , according to probability distribution $F(\theta)$. The agent observes θ , but the principal does not.
- 4 If the agent accepts, he exerts effort e and sends a message m , both observed by the principal.
- 5 Output is $q(e, \theta)$.

Payoffs

If the agent rejects the contract, $\pi_{agent} = \bar{U}$ and $\pi_{principal} = 0$.

If the agent accepts the contract, $\pi_{agent} = U(e, w, \theta)$ and $\pi_{principal} = V(q - w)$.

The principal would like to know θ so he can tell which effort level is appropriate. In an ideal world he would employ an honest agent who always chose $m = \theta$, but in noncooperative games, talk is cheap. Since the agent's words are worthless, the principal must try to design a contract that either provides incentive for truth-telling or takes lying into account – he **implements** a **mechanism** to extract the agent's information.

Unravelling the Truth when Silence is the Only Alternative

Before going on to look at a self-selection contract, let us look at a special case in which hidden knowledge paradoxically makes no difference. The usual hidden knowledge model has no penalty for lying, but let us briefly consider what happens if the agent cannot lie, though he can be silent. Suppose that Nature uses the uniform distribution to assign the variable θ some value in the interval $[0, 10]$ and the agent's payoff is increasing in the principal's estimate of θ . Usually we assume that the agent can lie freely, sending a message m taking any value in $[0, 10]$, but let us assume instead that he cannot lie but he can conceal information. Thus, if $\theta = 2$, he can send the uninformative message $m \geq 0$ (equivalent to no message), or the message $m \geq 1$, or $m = 2$, but not the lie that $m \geq 4$.

When $\theta = 2$ the agent might as well send a message that is the exact truth: " $m = 2$." If he were to choose the message " $m \geq 1$ " instead, the principal's first thought might be to estimate θ as the average value of the interval $[1, 10]$, which is 5.5. But the principal would realize that no agent with a value of θ greater than 5.5 would want to send that message " $m \geq 1$ " if that was the resulting deduction. This realization restricts the possible interval to $[1, 5.5]$, which in turn has an average of 3.25. But then no agent with $\theta > 3.25$ would send the message " $m \geq 1$." The principal would continue this process of logical **unravelling** to conclude that $\theta = 1$. The message " $m \geq 0$ " would be even worse, making the principal believe that $\theta = 0$. In this model, "No news is bad news." The agent would therefore not send the message " $m \geq 1$ " and he would be indifferent between " $m = 2$ " and " $m \geq 2$ " because the principal would make the same deduction from either message.

Perfect unravelling is paradoxical, but that is because the assumptions behind the reasoning in the last paragraph are rarely satisfied in the real world. In particular, unpunishable lying and genuine ignorance allow information to be concealed. If the seller is free to lie without punishment then in the absence of other incentives he always pretends that his information is extremely favorable, so nothing he says conveys any information, good or bad. If he really is ignorant in some states of the world, then his silence could mean either that he has nothing to say or that he has nothing he wants to say. The unravelling argument fails because if he sends an uninformative message the buyers will attach some probability to "no news" instead of "bad news." Problem 10.1 explores unravelling further.

The Revelation Principle

A principal might choose to offer a contract that induces his agent to lie in equilibrium, since he can take lying into account when he designs the contract, but this complicates the analysis. Each state of the world has a single truth, but a continuum of lies. Generically speaking, almost everything is false. The following principle helps us simplify contract design.

The Revelation Principle. *For every contract $w(q, m)$ that leads to lying (that is, to $m \neq \theta$), there is a contract $w^*(q, m)$ with the same outcome for every θ but no incentive for the agent to lie.*

Many possible contracts make false messages profitable for the agent because when the state of the world is a he receives a reward of x_1 for the true report of a and $x_2 > x_1$ for the false report of b . A contract which gives the agent the same reward of x_2 regardless of whether he reports a or b would lead to exactly the same payoffs for each player while giving the agent no incentive to lie. The revelation principle notes that a truth-telling contract like this can always be found by imitating the relation between states of the world and payoffs in the equilibrium of a contract with lying. The idea can also be applied to games in which two players must make reports to each other.

Applied to concrete examples, the revelation principle is obvious. Suppose we are concerned with the effect on the moral climate of cheating on income taxes, but anyone who makes \$70,000 a year can claim he makes \$50,000 and we do not have the resources to catch him. The revelation principle says that we can rewrite the tax code to set the tax to be the same for taxpayers earning \$70,000 and for those earning \$50,000, and the same amount of taxes will be collected without anyone having incentive to lie. Applied to moral education, the principle says that the mother who agrees never to punish her daughter if she tells her all her escapades will never hear any untruths. Clearly, the principle's usefulness is not so much to improve outcomes as to simplify contracts. The principal (and the modeller) need only look at contracts which induce truth-telling, so the relevant strategy space is shrunk and we can add a third constraint to the incentive compatibility and participation constraints to help calculate the equilibrium:

(3) **Truth-telling.** The equilibrium contract makes the agent willing to choose $m = \theta$.

The revelation principle says that a truth-telling equilibrium exists, but not that it is unique. It may well happen that the equilibrium is a weak Nash equilibrium in which the optimal contract gives the agent no incentive to lie but also no incentive to tell the truth. This is similar to the open-set problem discussed in Section 4.3; the optimal contract may satisfy the agent's participation constraint but makes him indifferent between accepting and rejecting the contract. If agents derive the slightest utility from telling the truth, of course, then truth-telling becomes a strong equilibrium, but if their utility from telling the truth is really significant, it should be made an explicit part of the model. If the utility of truth-telling is strong enough, in fact, agency problems and the costs associated with them disappear. This is one reason why morality is useful to business.

10.2: An Example of Moral Hazard with Hidden Knowledge: *The Salesman Game*

Suppose the manager of a company has told his salesman to investigate a potential customer, who is either a *Pushover* or a *Bonanza*. If he is a *Pushover*, the efficient sales

effort is low and sales should be moderate. If he is a *Bonanza*, the effort and sales should be higher.

The Salesman Game

Players

A manager and a salesman.

The Order of Play

- 1 The manager offers the salesman a contract of the form $w(q, m)$, where q is sales and m is a message.
- 2 The salesman decides whether or not to accept the contract.
- 3 Nature chooses whether the customer is a *Bonanza* or a *Pushover* with probabilities 0.2 and 0.8. Denote the state variable “customer status” by θ . The salesman observes the state, but the manager does not.
- 4 If the salesman has accepted the contract, he chooses his sales level q , which implicitly measures his effort.

Payoffs

The manager is risk neutral and the salesman is risk averse. If the salesman rejects the contract, his payoff is $\bar{U} = 8$ and the manager’s is zero. If he accepts the contract, then

$$\begin{aligned}\pi_{manager} &= q - w \\ \pi_{salesman} &= U(q, w, \theta), \text{ where } \frac{\partial U}{\partial q} < 0, \frac{\partial^2 U}{\partial q^2} < 0, \frac{\partial U}{\partial w} > 0, \frac{\partial^2 U}{\partial w^2} < 0\end{aligned}$$

Figure 1 shows the indifference curves of manager and salesman, labelled with numerical values for exposition. The manager’s indifference curves are straight lines with slope 1 because he is acting on behalf of a risk-neutral company. If the wage and the quantity both rise by a dollar, profits are unchanged, and the profits do not depend directly on whether θ takes the value *Pushover* or *Bonanza*.

The salesman’s indifference curves also slope upwards, because he must receive a higher wage to compensate for the extra effort that makes q greater. They are convex because the marginal utility of dollars is decreasing and the marginal disutility of effort is increasing. As Figure 1 shows, the salesman has two sets of indifference curves, solid for *Pushovers* and dashed for *Bonanzas*, since the effort that secures a given level of sales depends on the state.

must satisfy the participation constraint,

$$0.8U(q_3, w_3, \text{Pushover}) + 0.2U(q_3, w_3, \text{Bonanza}) \geq 8. \quad (2)$$

The nature of the equilibrium depends on the shapes of the indifference curves. If they are shaped as in Figure 2, the equilibrium is separating, not pooling, and there does exist a first-best, fully revealing contract.




Figure 2: Indifference Curves for a Separating Equilibrium

$$\text{Separating Contract} \left\{ \begin{array}{l} \text{Agent announces } \textit{Pushover} : \quad w = \begin{cases} 0 & \text{if } q < q_1 \\ w_1 & \text{if } q \geq q_1 \end{cases} \\ \text{Agent announces } \textit{Bonanza} : \quad w = \begin{cases} 0 & \text{if } q < q_2 \\ w_2 & \text{if } q \geq q_2 \end{cases} \end{array} \right. \quad (3)$$

Again, we know from the revelation principle that we can narrow attention to contracts that induce the salesman to tell the truth. With the indifference curves of Figure 2, contract (3) induces the salesman to be truthful and the incentive compatibility constraint is satisfied. If the customer is a *Bonanza*, but the salesman claims to observe a *Pushover* and chooses q_1 , his utility is less than 8 because the point (q_1, w_1) lies below the $\tilde{U}_S = 8$ indifference curve. If the customer is a *Pushover* and the salesman claims to observe a *Bonanza*, then although (q_2, w_2) does yield the salesman a higher wage than (q_1, w_1) , the extra income is not worth the extra effort, because (q_2, w_2) is far below the indifference curve $U_S = 8$.

Another way to look at a separating equilibrium is to think of it as a choice of contracts rather than as one contract with different wages for different outputs. The salesman agrees

to work for the manager, and after he discovers what type the customer is he chooses either the contract (q_1, w_1) or the contract (q_2, w_2) , where each is a forcing contract that pays him 0 if after choosing the contract (q_i, w_i) he produces output of $q \neq q_i$. In this interpretation, the manager offers a **menu of contracts** and the salesman selects one of them after learning his type.

Sales contracts in the real world are often complicated because it is easy to measure sales and hard to measure efforts when workers who are out in the field away from direct supervision. The Salesman Game is a real problem. Gonik (1978) describes hidden knowledge contracts used by IBM's subsidiary in Brazil. Salesmen were first assigned quotas. They then announced their own sales forecast as a percentage of quota and chose from among a set of contracts, one for each possible forecast. Inventing some numbers for illustration, if Smith were assigned a quota of 400 and he announced 100 percent, he might get $w = 70$ if he sold 400 and $w = 80$ if he sold 450; but if he had announced 120 percent, he would have gotten $w = 60$ for 400 and $w = 90$ for 450. The contract encourages extra effort when the extra effort is worth the extra sales. The idea here, as in the Salesman Game, is to reward salesmen not just for high effort, but for appropriate effort.

The Salesman Game illustrates a number of ideas. It can have either a pooling or a separating equilibrium, depending on the utility function of the salesman. The revelation principle can be applied to avoid having to consider contracts in which the manager must interpret the salesman's lies. It also shows how to use diagrams when the algebraic functions are intractable or unspecified, a problem that does not arise in most of the two-valued numerical examples in this book.

10.3: Myerson Mechanism Design Example

Myerson (1991) uses a trading example in Sections 6.4 and 10.3 of his book to illustrate mechanism design. A seller has 100 units of a good. If it is high quality, he values it at 40 dollars per unit; if it is low quality, at 20 dollars. The buyer, who cannot observe quality before purchase, values high quality at 50 dollars per unit and low quality at 30 dollars. For efficiency, all of the good should be transferred from the seller to the buyer. The only way to get the seller to truthfully reveal the quality of the good, however, is for the buyer to say that if the seller admits the quality is bad, he will buy more units than if the seller claims it is good. Let us see how this works out.

Depending on who offers the contract and when it is offered, various games result. We will start with one in which the seller makes the offer, and does so before he knows whether his quality is high or low.

Myerson Trading Game I

Players

A buyer and a seller.

The Order of Play

- 1 The seller offers the buyer a contract $(Q_H, P_H, T_H, Q_L, P_L, T_L)$ under which the seller will later declare his quality to be high or low, and the buyer will first pay the lump sum T to the seller (perhaps with $T < 0$) and then buy Q units of the 100 the seller has available, at price P .
- 2 The buyer accepts or rejects the contract.
- 3 Nature chooses whether the seller's good is High quality (probability 0.2) or low quality (probability 0.8), unobserved by the buyer.
4. If the contract was accepted by both sides, the seller declares his type to be L or H and sells at the appropriate quantity and price as stated in the contract.

Payoffs

If the buyer rejects the contract, $\pi_{buyer} = 0$, $\pi_{seller H} = 40 * 100$, and $\pi_{seller L} = 20 * 100$.

If the buyer accepts the contract and the seller declares a type that has price P , quantity Q , and transfer T , then

$$\pi_{buyer|seller H} = -T + (50 - P)Q \quad \text{and} \quad \pi_{buyer|seller L} = -T + (30 - P)Q \quad (4)$$

and

$$\pi_{seller H} = T + 40(100 - Q) + PQ \quad \text{and} \quad \pi_{seller L} = T + 20(100 - Q) + PQ. \quad (5)$$

The seller wants to design a contract subject to two sets of constraints. First, the buyer must accept the contract. Thus, the participation constraint is²

$$\begin{aligned} 0.8\pi_{buyer|seller H}(Q_L, P_L, T_L) + 0.2\pi_{buyer|seller L}(Q_H, P_H, T_H) &\geq 0 \\ 0.8[-T_L + (30 - P_L)Q_L] + 0.2[-T_H + (30 - P_H)Q_H] &\geq 0 \end{aligned} \quad (6)$$

There might also be a participation constraint for the seller himself, because it might be that even when he designs the contract that maximizes his payoff, his payoff is no higher than when he refuses to offer a contract. He can always offer the acceptable if vacuous null contract $(Q_L = 0, P_L = 0, T_L = 0, Q_H = 0, P_H = 0, T_H = 0)$, however, so we do not need to write out the seller's participation constraint separately.

Second, the seller must design a contract that will induce himself to tell the truth later once he discovers his type. This is, of course a bit unusual— the seller is like a principal designing a contract for himself as agent. That is why things will be different in this chapter than in the chapters on basic moral hazard. What is happening is that the seller is trying to sell not just a good, but a contract, and so he must make the contract attractive to the

²Another kind of participation constraint would apply if the buyer had the option to reject purchasing anything, after accepting the contract and hearing the seller's type announcement. That would not make a difference here.

buyer. Thus, he faces incentive compatibility constraints: one for when he is low quality,

$$\begin{aligned}\pi_{seller\ L}(Q_L, P_L, T_L) &\geq \pi_{seller\ L}(Q_H, P_H, T_H) \\ 20(100 - Q_L) + P_L Q_L + T_L &\geq 20(100 - Q_H) + P_H Q_H + T_H,\end{aligned}\tag{7}$$

and one for when he has high quality,

$$\begin{aligned}\pi_{seller\ H}(Q_H, P_H, T_H) &\geq \pi_{seller\ H}(Q_L, P_L, T_L) \\ 40(100 - Q_H) + P_H Q_H + T_H &\geq 40(100 - Q_L) + P_L Q_L + T_L.\end{aligned}\tag{8}$$

To make the contract incentive compatible, the seller needs to set P_H greater than P_L , but if he does that it will be necessary to set Q_H less than Q_L . If he does that, then the low-quality seller will not be irresistably tempted to pretend his quality is high: he would be able to sell at a higher price, but not as great a quantity.

Since Q_H is being set below 100 only to make pretending to be high-quality unattractive, there is no reason to set Q_L below 100, so $Q_L = 100$. The buyer will accept the contract if $P_L \leq 30$, so the seller should set $P_L = 30$. The low-quality seller's incentive compatibility constraint, inequality (7), will be binding, and thus becomes

$$\begin{aligned}\pi_{seller\ L}(Q_L, P_L, T_L) &\geq \pi_{seller\ L}(Q_H, P_H, T_H) \\ 20(100 - 100) + 30 * 100 + 0 &= 20(100 - Q_H) + P_H Q_H + 0.\end{aligned}\tag{9}$$

Solving for Q_H gives us $Q_H = \frac{1000}{P_H - 20}$, which when substituted into the seller's payoff function yields

$$\begin{aligned}\pi_s &= 0.8\pi_{seller\ L}(Q_L, P_L, T_L) + 0.2\pi_{seller\ H}(Q_H, P_H, T_H) \\ &= 0.8[(20)(100 - Q_L) + P_L Q_L + T_L] + 0.2[(40)(100 - Q_H) + P_H Q_H + T_H] \\ &= 0.8[(20)(100 - 100) + 30 * 100 + 0] + 0.2[(40)(100 - \frac{1000}{P_H - 20}) + P_H(\frac{1000}{P_H - 20}) + 0]\end{aligned}\tag{10}$$

Maximizing with respect to P_H subject to the constraint that $P_H \leq 50$ (or else the buyer will turn down the contract) yields the corner solution of $P_H = 50$, which allows for $Q_H = 33\frac{1}{3}$.

The participation constraint for the buyer is already binding, so we do not need the transfers T_L and T_H to take away any remaining surplus, as we might in other situations.³ Thus, the equilibrium contract is

$$(Q_L = 100, P_L = 30, T_L = 0, Q_H = 33\frac{1}{3}, P_H = 50, T_H = 0).\tag{11}$$

Note that this mechanism will not work if further offers can be made after the end of the game. The mechanism is not first-best efficient; if the seller is high-quality, then he

³The transfers could be used to adjust the prices, too. We could have $Q_L = 20$ and $T_L = 1000$ in equation (10) without changing anything important.

only sells $33\frac{1}{3}$ units to the buyer instead of all 100, even though both realize that the buyer's value is 50 and the seller's is only 40. If they could agree to sell the remaining $66\frac{2}{3}$ units, then the mechanism would not be incentive compatible in the first place, though, because then the low-quality seller would pretend to be high-quality, first selling $33\frac{1}{3}$ units and then selling the rest. The importance of commitment is a general feature of mechanisms.

What if it is the buyer who makes the offer?

Myerson Trading Game II

The Order of Play

The same as in Myerson Trading Game I except that the buyer makes the contract offer in move (1) and the seller accepts or rejects in move (2).

Payoffs

The same as in Myerson Trading Game I.

The participation constraint in the buyer's mechanism design problem is

$$0.8\pi_{seller\ L}(Q_L, P_L, T_L) + 0.2\pi_{seller\ H}(Q_H, P_H, T_H) \geq 0. \quad (12)$$

The incentive compatibility constraints are just as they were before, since the buyer has to design a mechanism which makes the seller truthfully reveal his type.

As before, the mechanism will set $Q_L = 100$, but it will have to make $Q_H < 100$ to deter the low-quality seller from pretending he is high-quality. Also, $P_H \geq 40$, or the high-quality seller will pretend to be low-quality.

Suppose $P_H = 40$. The low-quality seller's incentive compatibility constraint, inequality (7), will be binding, and thus becomes

$$\begin{aligned} \pi_{seller\ L}(Q_L, P_L, T_L) &\geq \pi_{seller\ H}(Q_H, P_H, T_H) \\ 20(100 - 100) + P_L * 100 + 0 &= 20(100 - Q_H) + 40Q_H + 0. \end{aligned} \quad (13)$$

Solving for Q_H gives us $Q_H = 5P_L - 100$, which when substituted into the buyer's payoff function yields

$$\begin{aligned} \pi_b &= 0.8\pi_{b|L}(Q_L, P_L, T_L) + 0.2\pi_{b|H}(Q_H, P_H, T_H) \\ &= 0.8[(30 - P_L)Q_L] + 0.2[(50 - P_H)Q_H] \\ &= 0.8[(30 - P_L)100] + 0.2[(50 - 40)(5P_L - 100)] \\ &= 2400 - 80P_L + 10P_L - 200 = 2200 - 70P_L \end{aligned} \quad (14)$$

Maximizing with respect to P_L subject to the constraint that $P_L \geq 20$ (or else we would come out with $Q_H < 0$ to satisfy incentive compatibility constraint (9)) yields the corner solution of $P_L = 20$, which requires that $Q_H = 0$.

Would setting $P_H > 40$ help? No, because that just makes it harder to satisfy the low-quality seller's incentive compatibility constraint. We would continue to have $Q_H = 0$, and, of course, P_H does not matter if nothing is sold. And as before, we do not need to make use of transfers to make the participation constraint binding. Thus, the equilibrium contract has P_H take any possible value and

$$(Q_L = 100, P_L = 20, T_L = 0, Q_H = 0, T_H = 0). \quad (15)$$

In the next version of the game, we will continue to let the buyer make the offer, but he makes it at a time when the seller already knows his type. Thus, this will be an adverse selection model.

Myerson Trading Game III

The Order of Play

0. Nature chooses whether the seller's good is high quality (probability 0.2) or low quality (probability 0.8), unobserved by the buyer.
- 1 The buyer offers the seller a contract $(Q_H, P_H, T_H, Q_L, P_L, T_L)$ under which the seller will later declare his quality to be high or low, and the buyer will first pays the lump sum T to the seller (perhaps with $T < 0$) and then buy Q units of the 100 the seller has available, at price P .
- 2 The seller accepts or rejects the contract.
3. If the contract was accepted by both sides, the seller declares his type to be L or H and sells at the appropriate quantity and price as stated in the contract.

Payoffs

The same as in Myerson Trading Games I and II.

The incentive compatibility constraints are unchanged from the previous two versions of the game, but now the participation constraints are different for the two types of seller.

$$\pi_L(Q_L, P_L, T_L) \geq 0 \quad (16)$$

and

$$\pi_H(Q_H, P_H, T_H) \geq 0. \quad (17)$$

Any mechanism which satisfies these two constraints would also satisfy the single participation constraint in MTG II, since it says that a weighted average of the payoffs of the two sellers must be positive. Thus, any mechanism which maximized the buyer's payoff in MTG II would also maximize his payoff in MTG III, if it satisfied the tougher bifurcated

participation constraints. The mechanism we found for the game does satisfy the tougher constraints, so it is the optimal mechanism here too.

This is not a general feature of mechanisms. More generally the optimal mechanism will not have as high a payoff when one player starts the game with superior information, because of the extra constraints on the mechanism.

In the last of our versions of this game, the seller makes the offer, but after he knows his type.

Myerson Trading Game IV

The Order of Play

The same as in Myerson Trading Game III except that in (1) the seller makes the offer and in (2) the buyer accepts or rejects.

Payoffs

The same as in Myerson Trading Games I, II, and III.

The incentive compatibility constraints are the same as in the previous games, and the participation constraint is inequality (6), just as in *Myerson Trading Game I*. The big difference now is that unlike in the first three versions, MTG IV has an informed player making the contract offer. As a result, the form of the offer can convey information, and we have to consider out-of-equilibrium beliefs, as in the dynamic games of incomplete information in Chapter 6 (and we will see more of this in the signalling models of Chapter 11). Surprisingly, however, the importance of out-of-equilibrium beliefs does not lead to multiple equilibria. Instead, the equilibrium contract is

$$\text{M1: } (Q_L = 100, P_L = 30, T_L = 0, Q_H = 33\frac{1}{3}, P_H = 50, T_H = 0),$$

This is part of equilibrium under the out-of-equilibrium belief that if the seller offers any other contract, the buyer believes the quality is low.

This is the same equilibrium mechanism as in MTG I. It is interesting to compare it to two other mechanisms, M2 and M3, which satisfy the two incentive compatibility constraints and the participation constraint, but which are not equilibrium choices here:⁴

$$\text{M2: } (Q_L = 100, P_L = 28, T_L = 0, Q_H = 0, P_H = 40, T_H = 800).$$

$$\text{M3: } (Q_L = 100, P_L = 31\frac{3}{7}, T_L = 0, Q_H = 57\frac{1}{7}, P_H = 40, T_H = 0).$$

Mechanism M2 is interesting because the buyer expects a positive payoff of $(30 - 28)(100) = 200$ if the seller is low-quality and a negative payoff of 800 if the seller is high-quality, for an overall expected payoff of zero. The contract is incentive compatible because a low-quality seller could not increase his payoff of $28 \cdot 100$ by pretending to be high-quality (he would get $20 \cdot 100 + 800$ instead), and a high-quality seller would reduce his payoff

⁴xxx M2– hwo about an oo belief that a deviator is a HIGH type? Then no deviation is profitable.

of $(40 \cdot 100 + 800)$ if he pretended to have low quality. Here, for the first time, we see a positive value for the transfer T_H .

Under mechanism M3, the buyer expects a negative payoff of $(30 - 31 \frac{1}{7})(100) = -11 \frac{3}{7}$ if the seller is low-quality and a positive payoff of $(57 \frac{1}{7})(50 - 40) = 11 \frac{3}{7}$ if the seller is high-quality, for an overall expected payoff of zero. The contract is incentive compatible because a low-quality seller could not increase his payoff of $3,142 \frac{6}{7} (31 \frac{3}{7})(100)$ by pretending to be high-quality (he would get $(57 \frac{1}{7})(40) + (42 \frac{6}{7})(20)$ instead, which comes to the same figure), and a high-quality seller would reduce his payoff if he pretended to have low quality and sold something he valued at 40 at a price of $31 \frac{1}{7}$.

In MTG IV, unlike the previous versions of the game, the particular mechanism chosen in equilibrium is not necessarily the one that the player who offers the contract likes best. Instead, an informed offeror—here, the seller—must worry that his offer might make the uninformed receiver believe the offeror's type is undesirable.

Mechanism M1 maximizes the payoff of the average seller, as we found in MTG I, yielding the low-quality seller a payoff of 3,000 and the high-quality seller a payoff of 4,333 $(= (33 \frac{1}{3})(50) + 66 \frac{2}{3}(40))$, for an average payoff of 3,867. If the seller is high-quality, however, he would prefer mechanism M2, which has payoffs of 2800 and 4800 $(= 800 + 40(100))$, for an average payoff of 3200. If the seller is low-quality, he would prefer mechanism M3, which has payoffs of $3,142 \frac{6}{7}$ and 4000, for an average payoff of $3,314 \frac{6}{7}$.

Suppose that the seller chose M2, regardless of his type. This could not be an equilibrium, because a low-quality seller would want to deviate. Suppose he deviated and offered a contract almost like M1, except that $P_L = 29.99$ instead of 30 and $P_H = 49.99$ instead of 50. This new contract would yield positive expected payoff to the buyer whether the buyer believes the seller is low-quality or high-quality, and so it would be accepted. It would yield higher payoff to the low-quality seller than M2, and so the deviation would have been profitable. Similarly, if the seller chose M3 regardless of his type, a high-quality seller could profitably deviate in the same way.

The *Myerson Trading Game* is a good introduction to the flavor of the algebra in mechanism design problems. For more on this game, in a very different style of presentation, see Sections 6.4 and Chapter 10 of Myerson (1991). We will next go on to particular economic applications of mechanism design.⁵

10.4: Price Discrimination

When a firm has market power – most simply when it is a monopolist – it would like to charge different prices to different consumers. To the consumer who would pay up to

⁵xxx Think about risk sharing too, with risk aversion. Then it makes a big difference when the seller learns his type.

\$45,000 for a car, the firm would like to charge \$45,000; to the consumer who would pay up to \$36,000, the profit-maximizing price is \$36, 000. But how does the car dealer know how much each consumer is willing to pay?

He does not, and that is what makes this a problem of mechanism design under adverse selection. The consumer who would be willing to pay \$45,000 can hide under the guise of being a less intense consumer, and despite facing a monopolist he can end up retaining consumer surplus – an **informational rent**, a return to the consumer’s private information about his own type.⁶

Pigou was a contemporary of Keynes at Cambridge who usefully divided price discrimination into three types in 1920 but named them so obscurely that I relegate his names to the endnotes and use better ones here:

1 Interbuyer price discrimination. This is when the seller can charge different prices to different buyers. Smith’s price for a hamburger is \$4 per burger, but Jones’s is \$6.

2 Interquantity price discrimination or Nonlinear pricing. This is when the seller can charge different unit prices for different quantities. A consumer can buy a first sausage for \$9, a second sausage for \$4, and a third sausage for \$3. Rather than paying the “linear” total price of \$9 for one sausage, \$18 for two, and \$27 for three, he thus pays the nonlinear price of \$9 for one sausage, \$13 for two, and \$16 for three, the concave price path shown in Figure 3.

3 Perfect price discrimination. This combines interbuyer and interquantity price discrimination. When the seller does have perfect information and can charge each buyer that buyer’s reservation price for each unit bought, Smith might end up paying \$50 for his first hot dog and \$20 for his second, while next to him Jones pays \$4 for his first and \$3 for his second.

⁶In real life, a standard opening ploy of car salesman is simply to ask. “So, how much are you able to spend on a car today?” My recommendation: don’t tell him. This may sound obvious, but remember it the next time your department chairman asks you how high a salary it would take to keep you from leaving for another university.




Figure 3: Linear and Nonlinear Pricing

To illustrate price discrimination as mechanism design we will use a modified version of an example in Chapter 14 of Hal Varian's third edition (Varian, 1992).

Varian's Nonlinear Pricing Game

Players

One seller and two buyers, Smith and Jones.

The Order of Play

0 Nature assigns one of the two buyers to be Unenthusiastic with utility function u and the other to be Valuing with utility function v , Smith and Jones having equal probabilities of filling each role. The seller does not observe Nature's move.

1 The seller offers a price mechanism $r(x)$ under which a buyer can buy amount x for total price $r(x)$.

2 The buyers simultaneously choose to buy quantities x_u and x_v .

Payoffs

The seller has a constant marginal cost of c , so his payoff is $r(x_u) + r(x_v) - c \cdot (x_u + x_v)$. The buyers' payoffs are $u(x_u) - r(x_u)$ and $v(x_v) - r(x_v)$ if x is positive, and 0 if $x = 0$, with $u', v' > 0$ and $u'', v'' < 0$. The total and marginal willingnesses to pay are greater for the Valuing buyer. For all x ,

$$\begin{aligned} (a) \quad & u(x) < v(x) \text{ and} \\ (b) \quad & u'(x) < v'(x) \end{aligned} \tag{18}$$

Condition (18b) is known as the **single-crossing property**, since it implies that the indifference curves of the two agents cross at most one time (see also Section 11.1).

Combined with Condition (18a), it means they *never* cross – the Valuing buyer always has stronger demand. Figure 4 illustrates the single-crossing property in two different ways. Figure 4a on the next page directly illustrates assumptions (18a) and (18b) – the utility function of the Valuing player starts higher and rises more steeply. Since utility units can be rescaled, though, this assumption should make you uncomfortable – do we really want to assume that the Valuing player is a happier person at zero consumption than the Unenthusiastic player? I personally am not bothered, but many economists prefer to restrict themselves to models in which utility is only ordinal, not cardinal, and is in units that cannot be compared between people. We can do that here. Figure 4b illustrates the single-crossing property in goods-space, where it says that if we pick one indifference curve for each player, the two curves only cross once. In words, this says that the Valuing player always requires more extra money as compensation for reducing his consumption of the commodity we are studying than the Unenthusiastic player would. This approach avoids the issue of which player is happier, but the cost is that Figure 4b is harder to understand than Figure 4a.




Figure 4: The Single- Crossing Property

To ease into the difficult problem of solving for the equilibrium mechanism, let us solve for the equilibrium of two simpler versions of the game that limit it to (a) perfect price discrimination and (b) interbuyer discrimination.

Perfect Price Discrimination

The game would allow perfect price discrimination if the seller did know which buyer had which utility function. He can then just maximize profit subject to the participation

constraints for the two buyers:

$$\begin{aligned} & \text{Maximize} \\ & r(x_u), r(x_v), x_u, x_v \quad r(x_u) + r(x_v) - c \cdot (x_u + x_v). \end{aligned} \quad (19)$$

subject to

$$\begin{aligned} (a) \quad & u(x_u) - r(x_u) \geq 0 \quad \text{and} \\ (b) \quad & v(x_v) - r(x_v) \geq 0. \end{aligned} \quad (20)$$

The constraints will be satisfied as equalities, since the seller will charge all that the buyers will pay. Substituting for $r(x_u)$ and $r(x_v)$ into the maximand, the first order conditions become

$$\begin{aligned} (a) \quad & u'(x_u^*) - c = 0 \quad \text{and} \\ (b) \quad & v'(x_v^*) - c = 0. \end{aligned} \quad (21)$$

Thus, the seller will choose quantities so that each buyer's marginal utility equals the marginal cost of production, and will choose prices so that the entire consumer surpluses are eaten up: $r^*(x_u^*) = u(x_u^*)$ and $r^*(x_v^*) = v(x_v^*)$. Figure 5 shows this for the unenthusiastic buyer.




Figure 5: Perfect Price Discrimination

Interbuyer Price Discrimination

The interbuyer price discrimination problem arises when the seller knows which utility functions Smith and Jones have and can sell to them separately but he must charge each buyer a single price per unit and let the buyer choose the quantity. The seller's problem is

$$\begin{aligned} & \text{Maximize} \\ & x_u, x_v, p_u, p_v \quad p_u x_u + p_v x_v - c \cdot (x_u + x_v), \end{aligned} \quad (22)$$

subject to

$$\begin{aligned} (a) \quad & u(x_u) - p_u x_u \geq 0 \quad \text{and} \\ (b) \quad & v(x_v) - p_v x_v \geq 0 \end{aligned} \tag{23}$$

and

$$\begin{aligned} (a) \quad & x_u = \operatorname{argmax}[u(x_u) - p_u x_u] \quad \text{and} \\ (b) \quad & x_v = \operatorname{argmax}[v(x_v) - p_v x_v]. \end{aligned} \tag{24}$$

This should remind you of moral hazard. It is very like the problem of a principal designing two incentive contracts for two agents to induce appropriate effort levels given their different disutilities of effort.

The agents will solve their quantity choice problems in (24), yielding

$$\begin{aligned} (a) \quad & u'(x_u) - p_u = 0 \quad \text{and} \\ (b) \quad & v'(x_v) - p_v = 0. \end{aligned} \tag{25}$$

Thus, we can simplify the original problem in (22) to

$$\underset{x_u, x_v}{\operatorname{Maximize}} \quad u'(x_u)x_u + v'(x_v)x_v - c \cdot (x_u + x_v), \tag{26}$$

subject to

$$\begin{aligned} (a) \quad & u(x_u) - p_u x_u \geq 0 \quad \text{and} \\ (b) \quad & v(x_v) - p_v x_v \geq 0. \end{aligned} \tag{27}$$

The first order conditions are

$$\begin{aligned} (a) \quad & u''(x_u)x_u + u' = c \quad \text{and} \\ (b) \quad & v''(x_v)x_v + v' = c. \end{aligned} \tag{28}$$

This is just the ‘marginal revenue equals marginal cost’ condition that any monopolist uses, but one for each buyer instead of one for the entire market.

Back to Nonlinear Pricing

Neither the perfect price discrimination nor the interbuyer problems are mechanism design problems, since the seller is perfectly informed about the types of the buyers and has no need to worry about designing incentives to separate them. In the original game, however, separation is the seller’s main concern. He must satisfy not just the participation constraints, but self-selection constraints. The seller’s problem is

$$\underset{x_u, x_v, r(x_u), r(x_v)}{\operatorname{Maximize}} \quad r(x_u) + r(x_v) - c \cdot (x_u + x_v), \tag{29}$$

subject to the participation constraints,

$$\begin{aligned} (a) \quad & u(x_u) - r(x_u) \geq 0 \quad \text{and} \\ (b) \quad & v(x_v) - r(x_v) \geq 0, \end{aligned} \tag{30}$$

and the self-selection constraints,

$$\begin{aligned} (a) \quad & u(x_u) - r(x_u) \geq u(x_v) - r(x_v) \\ (b) \quad & v(x_v) - r(x_v) \geq v(x_u) - r(x_u). \end{aligned} \tag{31}$$

Not all of these constraints will be binding. If neither type had a binding participation constraint, the principal would be losing a chance to increase his profits, unless there were moral hazard in the model too and some kind of efficiency wage was at work. In a mechanism design problem like this, what always happens is that the contracts are designed so that one type of agent is pushed down to his reservation utility.

Suppose the optimal contract is in fact separating, and also that both types of agent accept a contract. I have shown that at least one type will have a binding participation constraint. The second type could accept that same contract and receive more than his reservation utility, so to separate the two types the principal must offer the second type a contract which also yields more than his reservation utility. The principal will not want to be overly generous to the second type, however, so he makes sure the second type gets no more utility from his assigned contract than from accepting the first type's contract. Thus, one type of agent will have a binding participation constraint, and the other will have a binding self-selection constraint, and the other two constraints will be nonbinding. The question is: which type of buyer has which constraint binding in Varian's Nonlinear Pricing Game?

Let us start with the premise that a given constraint is binding and see if we can use our data to find a contradiction. Assume that the Valuing participation constraint, (30b), is binding. Then $v(x_v) = r(x_v)$. Substituting for $v(x_v)$ in the self-selection constraint (31b) then yields

$$r(x_v) - r(x_v) \geq v(x_u) - r(x_u), \quad (32)$$

so $r(x_u) \geq v(x_u)$. It follows from assumption (18a), which says that $u(x) < v(x)$, that $r(x_u) \geq u(x_u)$. But the Unenthusiastic participation constraint, (30a), says that $r(x_u) \leq u(x_u)$, and since these are compatible only when $r(x_u) = u(x_u)$ and we have assumed that (30b) is the binding participation constraint, we have arrived at a contradiction. Our starting point must be false, and it is in fact (30a), not (30b), that is the binding participation constraint.

We could next start with the premise that self-selection constraint (31a) is binding and derive a contradiction using assumption (18b). But the reasoning above showed that if the participation constraint is binding for one type of agent then the self-selection constraint will be binding for the other, so we can jump to the conclusion that it is in fact self-selection constraint (31b) that is binding.

Rearranging our two binding constraints and setting them out as equalities yields:

$$\begin{aligned} (30a') \quad & r(x_u) = u(x_u) \\ \text{and} \\ (31b') \quad & r(x_v) = r(x_u) - v(x_u) + v(x_v) \end{aligned}$$

This allows us to reformulate the seller's problem from (29) as

$$\underset{x_u, x_v}{\text{Maximize}} \quad u(x_u) + u(x_u) - v(x_u) - v(x_v) - c \cdot (x_u + x_v), \quad (33)$$

which has the first-order conditions

$$\begin{aligned} (a) \quad & u'(x_u) - c + [u'(x_u) - v'(x_u)] = 0 \\ (b) \quad & v'(x_v) - c = 0 \end{aligned} \quad (34)$$

These first-order conditions could be solved for exact values of x_u and x_v if we chose particular functional forms, but they are illuminating even if we do not. Equation (34b) tells us that the Valuing type of buyer buys a quantity such that his last unit's marginal utility exactly equals the marginal cost of production; his consumption is at the efficient level. The Unenthusiastic type, however, buys less than his first-best amount, something we can deduce using the single-crossing property, assumption (18 b), that $u'(x) < v'(x)$, which implies from (34a) that $u'(x_u) - c > 0$ and the Unenthusiastic type has not bought enough to drive his marginal utility down to marginal cost. The intuition is that the seller must sell less than first-best optimal to the Unenthusiastic type so as not to make that contract too attractive to the Valuing type. On the other hand, making the Valuing type's contract more valuable to him actually helps separation, so x_v is chosen to maximize social surplus.

The single-crossing property has another important implication. Substituting from first-order condition (34b) into first-order condition (34a) yields

$$[u'(x_u) - v'(x_v)] + [u'(x_u) - v'(x_u)] = 0 \quad (35)$$

The second term in square brackets is negative by the single-crossing property. Thus, the first term must be positive. But since the single-crossing property tells us that $[u'(x_u) - v'(x_u)] < 0$, it must be true, since $v'' < 0$, that if $x_u \geq x_v$ then $[u'(x_u) - v'(x_v)] < 0$ – that is, that the first term is negative. We cannot have that without contradiction, so it must be that $x_u < x_v$. The Unenthusiastic buyer buys strictly less than the Valuing buyer. This accords with our intuition, and also lets us know that the equilibrium is separating, not pooling (though we still have not proven that the equilibrium involves both players buying a positive amount, something hard to prove elegantly since one player buying zero would be a corner solution to our maximization problem).

A Graphical Approach to the Same Problem

Under perfect price discrimination, the seller would charge $r_u = A + B$ and $r_v = A + B + J + K + L$ to the two buyers for quantities x_u^* and x_v^* , as shown in Figure 10.6a. An attempt to charge $r(x_u^*) = A + B$ and $r(x_v^*) = A + B + J + K + L$, however, would simply lead to both buyers choosing to buy x_u^* , which would yield the Valuing buyer a payoff of $J + K$ rather than the 0 he would get as a payoff from buying x_v^* .




Figure 6: The Varian Nonlinear Pricing Game

The seller could separate the two buyers by charging $r(x_u^*) = A + B$ and $r(x_v^*) = A + B$, since the Unenthusiastic buyer would have no reason to switch to the greater quantity, but that would not increase his profits any over pooling. Figure 6b shows that the seller would do better to slightly reduce the quantity sold to the Unenthusiastic buyer and reduce the price by the amount of the dark shading, while selling x_u^* to the Valuing buyer and raising the price to him by the light shaded area. The Valuing buyer will still not be tempted to buy the smaller quantity at the lower price.

The profit-maximizing mechanism found earlier is shown in Figure 10.6a by $r(x_u') = A$ and $r(x_v^*) = A + B + K + L$. The Unenthusiastic buyer is left with a binding participation constraint, because $r(x_u') = A = u(x_u')$. The Valuing buyer has a nonbinding participation constraint, because $r(x_v^*) = A + B + K + L < v(x_v^*) = A + B + J + K + L$. But the Valuing buyer does have a binding self selection constraint, because he is exactly indifferent between buying x_u' and x_v^* — $v(x_u') = r(x_u') = A$ and $v(x_v^*) - r(x_v^*) = (A + B + J + K + L) - (A + B + K + L) = J$. Thus, the diagram replicates the algebraic conclusions.

*10.5 Rate-of-Return Regulation and Government Procurement

The central idea in both government procurement and regulation of natural monopolies is that the government is trying to induce a private firm to efficiently provide a good to the public while covering the cost of production. If information is symmetric, this is

an easy problem; the government simply pays the firm the cost of producing the good efficiently, whether the good be a missile or electricity. Usually, however, the firm has better information about costs and demand than the government does.

The variety of ways the firm might have better information and the government might extract it has given rise to a large literature in which moral hazard with hidden actions, moral hazard with hidden knowledge, adverse selection, and signalling all put in appearances. Suppose the government wants a firm to provide cable television service to a city. The firm knows more about its costs before agreeing to accept the franchise (adverse selection), discovers more after accepting it and beginning operations (moral hazard with hidden knowledge), and exerts greater or smaller effort to keep costs low (moral hazard with hidden actions). The government's problem is to acquire cable service at the lowest cost. It wants to be generous enough to induce the firm to accept the franchise in the first place but no more generous than necessary. It cannot simply agree to cover the firm's costs, because the firm would always claim high costs and exert low effort. Instead, the government might auction off the right to provide the service, might allow the firm a maximum price (a **price cap**), or might agree to compensate the firm to varying degrees for different levels of cost (**rate-of- return regulation**).

The problems of regulatory franchises and government procurement are the same in many ways. If the government wants to purchase a cruise missile, it also has the problem of how much to offer the firm. Roughly speaking, the equivalent of a price cap is a flat price, and the equivalent of rate-of-return regulation is a cost-plus contract, although the details differ in interesting ways. (A price cap allows downwards flexibility in prices, and rate-of-return regulation allows an expected but not guaranteed profit, for example.)

Many of these situations are problems of moral hazard with hidden information, because one player is trying to design a contract that the other will accept that will then induce him to use his private information properly.

Although the literature on mechanism design can be traced back to Mirrlees (1971), its true blossoming has occurred since Baron & Myerson's 1982 article, "Regulating a Monopolist with Unknown Costs." McAfee & McMillan (1988), Spulber (1989) and Laffont & Tirole (1993) provide 168-page, 690-page, and 702-page treatments of the confusing array of possible models and policies in their books on government regulation. Here, we will look at a version of the model Laffont and Tirole use to introduce their book on pages 55 to 62. This is a two-type model in which a special cost characteristic and the effort of a firm is its private information but its realized cost is public and nonstochastic. The model combines moral hazard and adverse selection, but it will behave more like an adverse selection model. The government will reimburse the firm's costs, but also fixes a price (which if negative becomes a tax) that depend on the level of the firm's costs. The questions the model hopes to answer are (a) whether effort will be too high or too low and (b) whether the price is positive and rises with costs.

Procurement I: Perfect Information ⁷

⁷I have changed the notation from the 3rd edition of this book. The special problem variable x replaces the ability variable a ; p replaces s ; the type L firm becomes a special-cost firm.

Players

The government and the firm.

The Order of Play

0 Nature determines whether the firm has special problems that add costs of x , which has probability θ , or no special problems, which has probability $(1 - \theta)$. We will call these “special” and “normal” firms, with the understanding that “special” problems may be the norm in engineering projects. The government and the firm both observe this move.

1 The government offers a contract agreeing to cover the firm’s cost c of producing a cruise missile and specifying an additional price $p(c)$ for each cost level that the firm might report.

2 The firm accepts or rejects the contract.

3 If the firm accepts, it chooses effort level e , unobserved by the government.

4 The firm finishes the cruise missile at a cost of $c = c_0 + x - e$ or $c = c_0 - e$ which is observed by the government, plus an additional cost $f(e - c_0)$ that the government does not observe. The government reimburses c and pays $p(c)$.

Payoffs

Both firm and government are risk neutral and both receive payoffs of zero if the firm rejects the contract. If the firm accepts, its payoff is

$$\pi_{firm} = p - f(e - c_0), \quad (36)$$

where $f(e - c_0)$, the cost of effort, is increasing and convex, so $f' > 0$ and $f'' > 0$. Assume, too, for technical convenience, that f is increasingly convex, so $f''' > 0$.⁸ The government’s payoff is

$$\pi_{government} = B - (1 + \lambda)c - \lambda p - f, \quad (37)$$

where B is the benefit of the cruise missile and λ is the deadweight loss from the taxation needed for government spending.⁹

The model differs from other principal-agent models in this book because the principal cares about the welfare of the agent. If the government cared only about the value of the cruise missile and the cost to taxpayers, its payoff would be $[B - (1 + \lambda)c - (1 + \lambda)p]$. Instead, the payoff function maximizes social welfare, the sum of the welfares of the taxpayers and the firm. The welfare of the firm is $(p - f)$, and summing the two welfares yields equation (37). Either kind of government payoff function may be realistic, depending on the political balance in the country being modelled, and the model will have similar properties whichever one is used.

Assume for the moment that B is large enough that the government definitely wishes to build the missile (how large will become apparent later). Cost, not output, is the focus of this model. The optimal output is one cruise missile regardless of agency problems, but the government wants to minimize the cost of producing the missile.

⁸The argument of f is normalized to be $(c_0 - e)$ rather than just e to avoid clutter in the algebra later. The assumption that $f''' > 0$ allows the use of first-order conditions by making concave the maximand in (48), which is a difference of two concave functions. It will also make deterministic contracts superior to stochastic ones. See p. 58 of Laffont & Tirole (1993).

⁹Hausman & Poterba (1987) estimate this loss to be around \$0.30 for each \$1 of tax revenue raised at the margin for the United States.

In this first variant of the game, whether the firm has special problems is observed by the government, which can therefore specify a contract conditioned on the type of the firm. The government pays prices of p_N to a normal firm with the cost $\underline{c}$, p_S to a special firm with the cost $\bar{c}$, and a price of $p = 0$ to a firm that does not achieve its appropriate cost level.

The special firm exerts effort $e = c_0 + x - \bar{c}$, achieves $c = \bar{c}$, generating unobserved effort disutility $f(e - c_0) = f(x - \bar{c})$, so its participation constraint is:

$$\begin{aligned}\pi_S(S) &\geq 0 \\ p_S - f(x - \bar{c}) &\geq 0.\end{aligned}\tag{38}$$

Similarly, in equilibrium the normal firm exerts effort $e = c_0 - \underline{c}$, so its participation constraint is

$$\begin{aligned}\pi_N(N) &\geq 0 \\ p_N - f(-\underline{c}) &\geq 0\end{aligned}\tag{39}$$

To make a firm's payoff zero and reduce the deadweight loss from taxation, the government will provide prices that exactly cover the firm's disutility of effort. Since there is no uncertainty we can invert the cost equation and write it as $e = c_0 + x - c$ or $e = c_0 - c$. The prices will be $p_S = f(e - c_0) = f(x - \bar{c})$ and $p_N = f(e - c_0) = f(-\underline{c})$.

Suppose the government knows the firm has special problems. Substituting the price p_S into the government's payoff function, equation (37), yields

$$\pi_{government} = B - (1 + \lambda)\bar{c} - \lambda f(c_0 + x - \bar{c}) - f((x - \bar{c}) - c_0).\tag{40}$$

Since $f'' > 0$, the government's payoff function is concave, and standard optimization techniques can be used. The first-order condition for $\bar{c}$ is

$$\frac{\partial \pi_{government}}{\partial \bar{c}} = -(1 + \lambda) + (1 + \lambda)f'(x - \bar{c}) = 0,\tag{41}$$

so

$$f'(x - \bar{c}) = 1.\tag{42}$$

Since $f'(x - \bar{c}) = f'([c_0 + x - \bar{c}] - c_0)$ and $c_0 + x - \bar{c} = e$, equation (42) says that $\bar{c}$ should be chosen so that $f'(e - c_0) = 1$; at the optimal effort level, the marginal disutility of effort equals the marginal reduction in cost because of effort. This is the first-best efficient effort level, which we will denote by $e^* \equiv e : \{f'(e - c_0) = 1\}$.

Exactly the same is true for the normal firm, so $f'(x - \bar{c}) = f'(-\underline{c}) = 1$ and $\underline{c} = \bar{c} - x$. The cost targets assigned to each firm are $\bar{c} = c_0 + x - e^*$ and $\underline{c} = c_0 - e^*$. Since both types must exert the same effort, e^* , to achieve their different targets, $p_S = f(e^* - c_0) = p_N$. The two firms exert the same efficient effort level and are paid the same price to compensate for the disutility of effort. Let us call this price level p^* .

The assumption that B is sufficiently large can now be made more specific: it is that $B - (1 + \lambda)\bar{c} - \lambda f(e^* - c_0) - f(e^* - c_0) \geq 0$, which requires that $B - (1 + \lambda)(c_0 + x - e^*) - (1 + \lambda)p^* \geq 0$.

Procurement II: Incomplete Information

In the second variant of the game, the existence of special problems is not observed by the government, which must therefore provide incentives for the firm to volunteer its type if the normal firm is to produce at lower cost than the firm with special problems.

The government could use a pooling contract, simply providing a price of p^* for a cost of $c = c_0 + x - e^*$, enough to compensate the firm with special problems for its effort, with $p = 0$ for any other cost. Both types would accept this, but the normal firm could exert effort less than e^* and still get costs down enough to receive the price. (Notice that this is the cheapest possible pooling contract; any cheaper contract would be rejected by the firm with special problems.) Thus, if the government would build the cruise missile under full information knowing that the firm has special problems, it would also build it under incomplete information, when the firm might have special problems.

The pooling contract, however, is not optimal. Instead, the government could offer a choice between the contract $(p^*, c = c_0 + x - e^*)$ and a new contract that offers a higher price but requires reimbursable costs to be lower. By definition of e^* , $f'(c_0 + x - e^* - c_0) = 1$, so $f'(c_0 - e^* - c_0) < 1$, which is to say that the normal firm's marginal disutility of effort when it exerts just enough effort to get costs down to $c = c_0 + x - e^*$ is less than 1. This means that if the government can offer a new contract with slightly lower c but slightly higher p that will be acceptable to the normal firm but will have a lower combined expense of $(p + c)$. This tells us that a separating contract exists that is superior to the pooling contract.

Let us therefore find the optimal contract with values $(\underline{c}, p_N)$ and $(\bar{c}, p_S)$ and $p = 0$ for other cost levels. It will turn out that the $(\bar{c}, p_S)$ part of the optimal separating contract will not be the same as the pooling contract in the previous paragraph, because to find the optimal separating contract it is not enough to find the optimal "new contract;" we need to find the optimal *pair* of contracts, and by finding a new contract for the special-problems firm too, we will be able to reduce the government's expense from the normal firm's contract.

A separating contract must satisfy participation constraints and incentive compatibility constraints for each type of firm. The participation constraints are the same as in Procurement I, inequalities (38) and (39).

The incentive compatibility constraint for the special firm is

$$\pi_S(S) \geq \pi_S(N) \quad (43)$$

$$p_S - f(x - \bar{c}) \geq p_N - f(x - \underline{c}),$$

and for the normal firm it is

$$\pi_N(N) \geq \pi_N(S) \quad (44)$$

$$p_N - f(-\underline{c}) \geq p_S - f(-\bar{c}).$$

Since the normal firm can achieve the same cost level as the special firm with less effort, inequality (44) tells us that if we are to have $\underline{c} < \bar{c}$, as is necessary for us to have

a separating equilibrium, we need $P_N > P_S$. The second half of inequality (44) must be positive. If the special firm's participation constraint, inequality (38), is satisfied, then $p_S - f(-\bar{c}) > 0$. This, in turn implies that (39) is a strong inequality; the normal firm's participation constraint is nonbinding.

The special firm's participation constraint, (38), will be binding (and therefore satisfied as an equality), because the government will reduce the price as much as possible in order to avoid the deadweight loss of taxation. The normal firm's incentive compatibility constraint must also be binding, because if the pair $(\underline{c}, p_N)$ were strictly more attractive for the normal firm, the government could reduce the price p_N . Constraint (44) is therefore satisfied as an equality.¹⁰ Knowing that constraints (38) and (44) are binding, we can write from constraint (38),

$$p_S = f(x - \bar{c}) \quad (45)$$

and, making use of both (38) and (44),

$$p_N = f(-\underline{c}) + f(x - \bar{c}) - f(-\bar{c}). \quad (46)$$

From (37), the government's maximization problem under incomplete information is

$$\underset{\underline{c}, \bar{c}, p_N, p_S}{\text{Maximize}} \quad \theta [B - (1 + \lambda)\bar{c} - \lambda p_S - f(x - \bar{c})] + [1 - \theta] [B - (1 + \lambda)\underline{c} - \lambda p_N - f(-\underline{c})]. \quad (47)$$

Substituting for p_S and p_N from (45) and (46) reduces the problem to

$$\underset{\underline{c}, \bar{c}}{\text{Maximize}} \quad \theta [B - (1 + \lambda)\bar{c} - \lambda(f(x - \bar{c}) - f(x - \bar{c}))] + [1 - \theta] [B - (1 + \lambda)\underline{c} - \lambda f(-\underline{c}) - \lambda f(x - \bar{c}) + \lambda f(-\bar{c}) - f(-\underline{c})]. \quad (48)$$

(1) The first-order condition with respect to $\underline{c}$ is

$$(1 - \theta)[-(1 + \lambda) + \lambda f'(-\underline{c}) + f'(-\underline{c})] = 0, \quad (49)$$

which simplifies to

$$f'(-\underline{c}) = 1. \quad (50)$$

Thus, as earlier, $f'_N = 1$. The normal firm chooses the efficient effort level e^* in equilibrium, and $\underline{c}$ takes the same value as it did in Procurement I. Equation (46) can be rewritten as

$$p_N = p^* + f(x - \bar{c}) - f(-\bar{c}). \quad (51)$$

Because $f(x - \bar{c}) > f(-\bar{c})$, equation (51) shows that $p_N > p^*$. Incomplete information increases the price for the normal firm, which earns more than its reservation utility in the game with incomplete information. Since the firm with special problems will earn exactly zero, this means that the government is on average providing its supplier with an above-market rate of return, not because of corruption or political influence, but because that

¹⁰The same argument does not hold for the special firm, because if p_S were reduced, the participation constraint would be violated.

is the way to induce normal suppliers to reveal that they do not have special costs. This should be kept in mind as an alternative to the product quality model of Chapter 5 and the efficiency wage model of Section 8.1 for why above-average rates of return persist.

(2) The first-order condition with respect to $\bar{c}$ is

$$\theta [-(1 + \lambda) + \lambda f'(x - \bar{c}) + f'(x - \bar{c})] + [1 - \theta] [\lambda f'(x - \bar{c}) + f'(-\bar{c})] = 0. \quad (52)$$

This can be rewritten as

$$f'(x - \bar{c}) = 1 - \left(\frac{1 - \theta}{\theta(1 + \lambda)} \right) [\lambda f'(x - \bar{c}) + f'(-\bar{c})]. \quad (53)$$

Since the right-hand-side of equation (53) is less than one, the special firm has a lower level of f' than the normal firm, and must be exerting effort less than e^* since $f'' > 0$. Perhaps this explains the expression “good enough for government work”. Also since the special firm’s participation constraint, (38), is satisfied as an equality, it must also be true that $p_S < p^*$. The special firm’s price is lower than under full information, although since its effort is also lower, its payoff stays the same.

We must also see that the incentive compatibility constraint for the firm with special problems is satisfied as a weak inequality; the firm with special problems is not near being tempted to pick the normal firm’s contract. This is a bit subtle. Setting the left-hand-side of the incentive compatibility constraint (43) equal to zero because the participation constraint is binding for the firm with special problems, substituting in for p_N from equation (46) and rearranging yields

$$f(x - \underline{c}) - f(-\underline{c}) \geq f(x - \bar{c}) - f(-\bar{c}). \quad (54)$$

This is true, and true as a strict inequality, because $f'' > 0$ and the arguments of f on the left-hand-side of equation (54) take larger values than on the right-hand side, as shown in Figure 10.7.

Figure 7: The Disutility of Effort

To summarize, the government’s optimal contract will induce the normal firm to exert the first-best efficient effort level and achieve the first-best cost level, but will yield that firm

a positive profit. The contract will induce the firm with special costs to exert something less than the first-best effort level and result in a cost level higher than the first-best, but its profit will be zero.

There is a tradeoff between the government's two objectives of inducing the correct amount of effort and minimizing the subsidy to the firm. Even under complete information, the government cannot provide a subsidy of zero, or the firms will refuse to build the cruise missile. Under incomplete information, not only must the subsidies be positive but the normal firm earns **informational rents**; the government offers a contract that pays the normal firm with more than under complete information to prevent it from mimicking a firm with special problems by choosing an inefficiently low effort. The firm with special problems, however, does choose an inefficiently low effort, because if it were assigned greater effort it would have to be paid a greater subsidy, which would tempt the normal firm to imitate it. In equilibrium, the government has compromised by having some probability of an inefficiently high subsidy ex post, and some probability of inefficiently low effort.

In the last version of the game, the firm's type is not known to either player until after the contract is agreed upon. The firm, however, learns its type before it must choose its effort level.

Procurement III: Moral Hazard with Hidden Information

The Order of Play

- 1 The government offers a contract agreeing to cover the firm's cost c of producing a cruise missile and specifying an additional price $p(c)$ for each cost level that the firm might report.
- 2 The firm accepts or rejects the contract.
- 3 Nature determines whether the firm has special problems that add costs of x , which has probability θ , or no special problems, which has probability $(1 - \theta)$. We will call these "special" and "normal" firms, with the understanding that "special" problems may be the norm in engineering projects. The government and the firm both observe this move.
- 4 If the firm accepts, it chooses effort level e , unobserved by the government.
- 5 The firm finishes the cruise missile at a cost of $c = c_0 + x - e$ or $c = c_0 - e$ which is observed by the government, plus an additional cost $f(e - c_0)$ that the government does not observe. The government reimburses c and pays $p(c)$.

The contract must satisfy one overall participation constraint and two incentive compatibility constraints, one for each type of firm. The participation constraint is

$$\theta[p_S - f(x - \bar{c})] + [1 - \theta][p_N - f(-\underline{c})] \geq 0. \quad (55)$$

The incentive compatibility constraints are the same as before: for the special firm,

$$p_S - f(x - \bar{c}) \geq p_N - f(-x - \underline{c}), \quad (56)$$

and for the normal firm,

$$p_N - f(-\underline{c}) \geq p_S - f(-\bar{c}). \quad (57)$$

As before, constraint (55) will be binding (and therefore satisfied as an equality), because the government will reduce the price as much as possible in order to avoid the deadweight loss of taxation. The normal firm's incentive compatibility constraint must also be binding, because if the pair $(\bar{c}, p_N)$ were strictly more attractive for the normal firm, the government could reduce the price p_N . Constraint (57) is therefore satisfied as an equality.¹¹ Knowing that constraints (55) and (57) are binding, we can write from constraint (55),

$$p_S = f(x - \bar{c}) - \frac{[1 - \theta][p_N - f(-\underline{c})]}{\theta}. \quad (58)$$

Substituting from (58) for p_S into (57), we get

$$p_N - f(-\underline{c}) = f(x - \bar{c}) - \frac{[1 - \theta][p_N - f(-\underline{c})]}{\theta} - f(-\bar{c}). \quad (59)$$

This can be solved for p_N to yield

$$p_N = \theta[f(x - \bar{c}) - f(-\bar{c})] + f(-\underline{c}), \quad (60)$$

which when substituted into (58) yields

$$p_S = [1 - \theta][f(x - \bar{c}) - f(-\bar{c})]. \quad (61)$$

From (37), the government's maximization problem under incomplete information is

$$\underset{\underline{c}, \bar{c}, p_N, p_S}{\text{Maximize}} \quad \theta [B - (1 + \lambda)\bar{c} - \lambda p_S - f(x - \bar{c})] + [1 - \theta] [B - (1 + \lambda)\underline{c} - \lambda p_N - f(-\underline{c})]. \quad (62)$$

Substituting for p_N and p_S from (60) and (61) reduces the problem to

$$\underset{\underline{c}, \bar{c}, p_N, p_S}{\text{Maximize}} \quad \theta \{B - (1 + \lambda)\bar{c} - \lambda[1 - \theta][f(x - \bar{c}) - f(-\bar{c})] - f(x - \bar{c})\} \\ + [1 - \theta] [B - (1 + \lambda)\underline{c} - \lambda\{\theta[f(x - \bar{c}) - f(-\bar{c})] + f(-\underline{c})\} - f(-\underline{c})]. \quad (63)$$

(1) The first-order condition with respect to $\underline{c}$ is

$$(1 - \theta)[-(1 + \lambda) + \lambda f'(-\underline{c}) + f'(-\underline{c})] = 0, \quad (64)$$

just as under adverse selection, which simplifies to

$$f'(-\underline{c}) = 1. \quad (65)$$

Thus, as earlier, $f'_N = 1$. The normal firm chooses the efficient effort level e^* in equilibrium, and $\underline{c}$ takes the same value as it did in Procurement I. Equation (59) can be rewritten as

$$p_N = p^* + f(x - \bar{c}) - f(-\bar{c}). \quad (66)$$

¹¹The same argument does not hold for the firm with special costs, because if p_S were reduced, the participation constraint would be violated.xxx check this

Because $f(x - \bar{c}) > f(-\bar{c})$, equation (66) shows that $p_N > p^*$. The normal firm earns more than its reservation utility, even under complete information. The special firm must therefore earn less than its reservation utility, so that the overall participation constraint will be satisfied as an equality.

(2) The first-order condition with respect to $\bar{c}$ is

$$\theta \{-(1 + \lambda) - \lambda(1 - \theta)[-f'(x - \bar{c}) + f'(-\bar{c})] + f'(x - \bar{c})\} + \lambda[1 - \theta][f'(x - \bar{c}) - f'(-\bar{c})] = 0. \quad (67)$$

This can be rewritten as

$$xcxcvxcvcxf'(x - \bar{c}) = 1 - sdfsfdfdsf \quad (68)$$

Since the right-hand-side of equation (68) is less than one, the special firm has a lower level of f' than the normal firm, and must be exerting effort less than e^* , since $f'' > 0$. Also since the participation constraint, (55), is satisfied as an equality, it must also be true that $p_S < s^*$. The special firm's subsidy is lower than under full information, although since its effort is also lower, its payoff stays the same. sdfadfasdfsdf

We must also see that the incentive compatibility constraint for the special firm is satisfied as a weak inequality; it is not near being tempted to pick the normal firm's contract. Setting the left-hand-side of the incentive compatibility constraint (56) equal to zero because the participation constraint is binding the special firm, substituting in for p_N from equation (59) and rearranging yields zcddgafd

$$asdfsafddsfsf(-\underline{c}) - f(-x - \underline{c}) \geq f(-\bar{c}) - f(x - \bar{c}). \quad (69)$$

This is true, and true as a strict inequality, because $f'' > 0$ and the arguments of f on the left-hand-side of equation (69) take larger values than on the right-hand side.sdfasfd

xx Compare with PG II.

xxxxxx HERE, RESUME CHAPTER.

A little reflection will provide a host of additional ways to alter the Procurement Game. What if the firm discovers its costs only after accepting the contract? What if two firms bid against each other for the contract? What if the firm can bribe the government? What if the firm and the government bargain over the gains from the project instead of the government being able to make a take-it-or-leave-it contract offer? What if the game is repeated, so the government can use the information it acquires in the second period? If it is repeated, can the government commit to long-term contracts? Can it commit not to renegotiate? See Spulber (1989) and Laffont & Tirole (1993) if these questions interest you.

*10.6 The Groves Mechanism

Hidden knowledge is particularly important in public economics, the study of government spending and taxation. We have just seen in Section 10.5 how important the information of private citizens is to the government in trying to decide what price to pay businesses for public services. Much the same issues arise when the government is trying to decide how much to tax or how much to spend. In the optimal taxation literature in which Mirrlees (1971) is the classic article, citizens differ in their income-producing ability, and the government wishes to demand higher taxes from the more able citizens, a clear problem of hidden knowledge. An even purer hidden knowledge problem is choosing the level of public goods based on private preferences. The government must decide whether it is worthwhile to buy a public good based on the combined preferences of all the citizens, but it needs to discover those preferences. Unlike in the previous games in this chapter, a group of agents will now be involved, not just one agent. Moreover, unlike in most games but similarly to the regulating principal of Section 10.4, the government is an altruistic principal who cares directly about the utility of the agents, rather than a car buyer or an insurance seller who cares about the agents' utility only in order to satisfy self-selection and participation constraints.

The example below is adapted from p. 426 of Varian (1992). The mayor of a town is considering installing a streetlight costing \$100. Each of the five houses near the light would be taxed exactly \$20, but the mayor will only install it if he decides that the sum of the residents' valuations for it is greater than the cost.

The problem is to discover the valuations. If the mayor simply asks them, householder Smith might say that his valuation is \$5,000, and householder Brown might say that he likes the darkness and would pay \$5,000 to *not* have a streetlight, but all the mayor could conclude would be that Smith's valuation exceeded \$20 and Brown's did not. Talk is cheap, and the dominant strategy is to overreport or underreport.

The flawed mechanism just described can be denoted by

$$M_1 : \left(20, \sum_{i=1}^5 m_i \geq 100 \right), \quad (70)$$

which means that each resident pays \$20, and the light is installed if the sum of the valuations exceeds 100.

An alternative is to make resident i pay the amount of his message, or pay zero if it is negative. This mechanism is

$$M_2 : \left(\text{Max}\{m_i, 0\}, \sum_{j=1}^5 m_j \geq 100 \right). \quad (71)$$

Mechanism M_2 has no dominant strategy. Player i would announce $m_i = 0$ if he thought the project would go through without his support, but he would announce up to his valuation if necessary. There is a continuum of Nash equilibria that attain the efficient result. Most of these are asymmetric, and there is a problem of how the equilibrium to be played out becomes common knowledge. This is a simple mechanism, however, and it already teaches a lesson: that people are more likely to report their true political preferences if they must bear part of the costs themselves.

Instead of just ensuring that the correct decision is made in a Nash equilibrium, it may be possible to design a mechanism which makes truthfulness a **dominant-strategy mechanism**. Consider the mechanism

$$M_3 : \left(100 - \sum_{j \neq i} m_j, \sum_{j=1}^5 m_j \geq 100 \right). \quad (72)$$

Under mechanism M_3 , player i 's message does not affect his tax bill except by its effect on whether or not the streetlight is installed. If player i 's valuation is v_i , his full payoff is $v_i - 100 + \sum_{j \neq i} m_j$ if $m_i + \sum_{j \neq i} m_j \geq 100$, and zero otherwise. It is not hard to see that he will be truthful in a Nash equilibrium in which the other players are truthful, but we can go further: truthfulness is weakly dominant. Moreover, the players will tell the truth whenever lying would alter the mayor's decision.

Consider a numerical example. Suppose that Smith's valuation is 40 and the sum of the valuations is 110, so the project is indeed efficient. If the other players report their truthful sum of 70, Smith's payoff from truthful reporting is his valuation of 40 minus his tax of 30. Reporting more would not change his payoff, while reporting less than 30 would reduce it to 0.

If we are wondering whether Smith's strategy is dominant, we must also consider his best response when the other players lie. If they underreported, announcing 50 instead of the truthful 70, then Smith could make up the difference by overreporting 60, but his payoff would be $-10 (= 40 + 50 - 100)$ so he would do better to report the truthful 40, killing the project and leaving him with a payoff of 0. If the other players overreported, announcing 80 instead of the truthful 70, then Smith benefits if the project goes through, and he should report at least 20 to obtain his payoff of 40 minus 20. He is willing to report exactly 40, so there is an equilibrium with truth-telling.

The problem with a dominant-strategy mechanism like the one facing Smith is that it is not budget balancing. The government raises less in taxes than it spends on the project (in fact, the taxes would be negative). Lack of budget balancing is a crucial feature of dominant-strategy mechanisms. While the government deficit can be made either positive or negative, it cannot be made zero, unlike in the case of Nash mechanisms.

Notes

N10.1 The revelation principle and moral hazard with hidden knowledge

- The books by Fudenberg & Tirole (1991a), Laffont & Tirole (1993), Palfrey & Srivastava (1993), Spulber (1989), and Baron's chapter in the *Handbook of Industrial Organization* edited by Schmalensee and Willig (1989) are good places to look for more on mechanism design.
- Levmore (1982) discusses hidden knowledge problems in tort damages, corporate freezeouts, and property taxes in a law review article.
- The revelation principle was named by Myerson (1979) and can be traced back to Gibbard (1973). A further reference is Dasgupta, Hammond & Maskin (1979). Myerson's game theory book is, as one might expect, a good place to look for further details (Myerson [1991, pp. 258-63, 294-99]).
- Moral hazard with hidden knowledge is common in public policy. Should the doctors who prescribe drugs also be allowed to sell them? The question trades off the likelihood of over-prescription against the potentially lower cost and greater convenience of doctor-dispensed drugs. See "Doctors as Druggists: Good Rx for Consumers?" *Wall Street Journal*, June 25, 1987, p. 24.
- For a careful discussion of the unravelling argument for information revelation, see Milgrom (1981b).
- A hidden knowledge game requires that the state of the world matter to one of the players' payoffs, but not necessarily in the same way as in Production Game VII. The Salesman Game of Section 10.2 effectively uses the utility function $U(e, w, \theta)$ for the agent and $V(q-w)$ for the principal. The state of the world matters because the agent's disutility of effort varies across states. In other problems, his utility of money might vary across states.

N10.2 An example of moral hazard with hidden knowledge: *The Salesman Game*

- Sometimes students know more about their class rankings than the professor does. One professor of labor economics used a mechanism of the following kind for grading class discussion. Each student i reports a number evaluating other students in the class. Student i 's grade is an increasing function of the evaluations given i by other students and of the correlation between i 's evaluations and the other students'. There are many Nash equilibria, but telling the truth is a focal point.
- In dynamic games of moral hazard with hidden knowledge the **ratchet effect** is important: the agent takes into account that his information-revealing choice of contract this period will affect the principal's offerings next period. A principal might allow high prices to a public utility in the first period to discover that its costs are lower than expected, but in the next period the prices would be reduced. The contract is ratcheted irreversibly to be more severe. Hence, the company might not choose a contract which reveals its costs in the first period. This is modelled in Freixas, Guesnerie & Tirole (1985).

Baron (1989) notes that the principal might purposely design the equilibrium to be pooling in the first period so self selection does not occur. Having learned nothing, he can offer a more effective separating contract in the second period.

N10.4 Price Discrimination

- The names for price discrimination in Part 2, Chapter 17, Section 5 of Pigou (1920) are: (1) first-degree (perfect price discrimination), (2) second-degree (interquantity price discrimination), and (3) third-degree (interbuyer price discrimination). These arbitrary names have plagued generations of students of industrial organization, in parallel with the appalling Type I and Type II errors of statistics (better named as False Negatives or Rejections, and False Positives or Acceptances). I invented the terms **interbuyer price discrimination** and **interquantity price discrimination** for this edition, with the excuse that I think their meaning will be clear to anyone who already knows the concepts under their Pigouvian names.
- A narrower category of nonlinear pricing is the **quantity discount**, in which the price per unit declines with the quantity bought. Sellers are often constrained to this, since if the price per unit rises with the quantity bought, some means must be used to prevent a canny consumer from buying two batches of small quantities instead of one batch of a large quantity.
- Philips's 1983 book, *The Economics of Price Discrimination*, is a good reference on the subject.
- In Varian's Nonlinear Pricing Game the probabilities of types for each player are not independent, unlike in most games. This does not make the game more complicated, though. If the assumption were "Nature assigns each buyer a utility function u or v with independent probabilities of 0.5 for each type," then there would be not just two possible states of the world in this game— uv and vu for Smith and Jones's types—but four— uv, vu, uu , and vv . How would the equilibrium change?
- The careful reader will think, "How can we say that Buyer V always gets higher utility than Buyer U for given x ? Utility cannot be compared across individuals, and we could rescale Buyer V's utility function to make him always have lower utility without altering the essentials of the utility function." My reply is that more generally we could set up the utility functions as $v(x) + y$ and $u(x) + y$, with y denoting spending on all other goods (as Varian does in his book). Then to say that V always gets higher utility for a given x means that he always has a higher relative value than U does for good x relative to money. Rescaling to give V the utility function $.001v(x) + .001y$ would not alter that.
- The notation I used in Varian's Nonlinear Pricing Game is optimized for reading. If you wish to write this on the board or do the derivations for practice, use abbreviations like u_1 for $u_1(x_1)$, a for $v(x_1)$, and b for $v(x_2)$ to save writing. The tradeoff between brevity and transparency in notation is common, and must be made in light of whether you are writing on a blackboard or on a computer, for just yourself or for the generations.

N10.5 Rate-of-return regulation and government procurement

- I changed the notation from Laffont and Tirole and from my own previous edition. Rather than assign each type of firm a cost parameter β for a cost of $c = \beta - e$, I now assign each type of firm an ability parameter a , for a cost of $c = c_0 - a - e$. This will allow the desirable type of firm to be the one with the *High* value of the type parameter, as in most models.

N10.6 The Groves Mechanism

- Vickrey (1961) first suggested the nonbudget-balancing mechanism for revelation of preferences, but it was rediscovered later and became known as the Groves Mechanism (from Groves [1973]).

Problems

10.1. Unravelling

An elderly prospector owns a gold mine worth an amount θ drawn from the uniform distribution $U[0, 100]$ which nobody knows, including himself. He will certainly sell the mine, since he is too old to work it and it has no value to him if he does not sell it. The several prospective buyers are all risk neutral. The prospector can, if he desires, dig deeper into the hill and collect a sample of gold ore that will reveal the value of θ . If he shows the ore to the buyers, however, he must show genuine ore, since an unwritten Law of the West says that fraud is punished by hanging offenders from joshua trees as food for buzzards.

- (a) For how much can he sell the mine if he is clearly too feeble to have dug into the hill and examined the ore? What is the price in this situation if, in fact, the true value is $\theta = 70$?
- (b) For how much can he sell the mine if he can dig the test tunnel at zero cost? Will he show the ore? What is the price in this situation if, in fact, the true value is $\theta = 70$?
- (c) For how much can he sell the mine if, after digging the tunnel at zero cost and discovering θ , it costs him an additional 10 to verify the results for the buyers? What is his expected payoff?
- (d) Suppose that with probability 0.5 digging the test tunnel costs 5 for the prospector, but with probability 0.5 it costs him 120. Keep in mind that the 0-100 value of the mine is net of the buyer's digging cost. Denote the equilibrium price that buyers will pay for the mine after the prospector approaches them without showing ore by P . What is the buyer's posterior belief about the probability it costs 120 to dig the tunnel, as a function of P ? Denote this belief by $B(P)$ (Assume, as usual, that all these parameters are common knowledge, although only the prospector learns whether the cost is actually 0 or 120.)
- (e) What is the prospector's expected payoff in the conditions of part (d) if (i) the tunnel costs him 120, or (ii) the tunnel costs him 5?

10.2. Task Assignment

Table 1 shows the payoffs in the following game. Sally has been hired by Rayco to do either Job 1, to do Job 2, or to be a Manager. Rayco believes that Tasks 1 and 2 have equal probabilities of being the efficient ones for Sally to perform. Sally knows which task is efficient, but what she would like best is a job as Manager that gives her the freedom to choose rather than have the job designed for the task. The CEO of Rayco asks Sally which task is efficient. She can either reply "task 1," "task 2," or be silent. Her statement, if she makes one, is an example of "cheap talk," because it has no direct effect on anybody's payoff. See Farrell & Rabin (1996).

Table 1: The Right To Silence Game payoffs

		Sally's Job		
		Job 1	Job 2	Manager
Sally knows	Task 1 is efficient (0.5)	2, 5	1, -2	3, 3
	Task 2 is efficient (0.5)	1, -2	2, 5	3, 3
Payoffs to: (Sally, Rayco)				

- If Sally did not have the option of speaking, what would happen?
- There exist perfect Bayesian equilibria in which it does not matter how Sally replies. Find one of these in which Sally speaks at least some of the time, and explain why it is an equilibrium. You may assume that Sally is not morally or otherwise bound to speak the truth.
- There exists a perverse variety of equilibrium in which Sally always tells the truth and never is silent. Find an example of this equilibrium, and explain why neither player would have incentive to deviate to out-of-equilibrium behavior.

10.3. Agency Law

Mr. Smith is thinking of buying a custom-designed machine from either Mr. Jones or Mr. Brown. This machine costs \$5,000 to build, and it is useless to anyone but Smith. It is common knowledge that with 90 percent probability the machine will be worth \$10,000 to Smith at the time of delivery, one year from today, and with 10 percent probability it will only be worth \$2,000. Smith owns assets of \$1,000. At the time of contracting, Jones and Brown believe there is a 20 percent chance that Smith is actually acting as an “undisclosed agent” for Anderson, who has assets of \$50,000.

Find the price he would be under the following two legal regimes: (a) An undisclosed principal is not responsible for the debts of his agent; and (b) even an undisclosed principal is responsible for the debts of his agent. Also, explain (as part [c]) which rule a moral hazard model like this would tend to support.

10.4. Incentive Compatibility and Price Discrimination

Two consumers have utility functions $u_1(x_1, y_1) = a_1 \log(x_1) + y_1$ and $u_2(x_2, y_2) = a_2 \log(x_2) + y_2$, where $1 > a_2 > a_1$. The price of the y-good is 1 and each consumer has an initial wealth of 15. A monopolist supplies the x-good. He has a constant marginal cost of 1.2 up to his capacity constraint of 10. He will offer at most two price-quantity packages, (r_1, x_1) and (r_2, x_2) , where r_i is the total cost of purchasing x_i units. He cannot identify which consumer is which, but he can prevent resale.

- Write down the monopolist's profit maximization problem. You should have four constraints plus the capacity constraint.
- Which constraints will be binding at the optimal solution?

- (c) Substitute the binding constraints into the objective function. What is the resulting expression? What are the first-order conditions for profit maximization? What are the profit-maximizing values of x_1 and x_2 ?

10.5. The Groves Mechanism

A new computer costing 10 million dollars would benefit existing Divisions 1, 2, and 3 of a company with 100 divisions. Each divisional manager knows the benefit to his division (variables $v_i, i = 1, \dots, 3$), but nobody else does, including the company CEO. Managers maximize the welfare of their own divisions. What dominant strategy mechanism might the CEO use to induce the managers to tell the truth when they report their valuations? Explain why this mechanism will induce truthful reporting, and denote the reports by $x_i, i = 1, \dots, 3$. (You may assume that any budget transfers to and from the divisions in this mechanism are permanent— that the divisions will not get anything back later if the CEO collects more payments than he gives, for example.)

10.6. The Two-Part Tariff (Varian 14.10, modified)

One way to price discriminate is to charge a lump sum fee L to have the right to purchase a good, and then charge a per-unit charge p for consumption of the good after that. The standard example is an amusement park where the firm charges an entry fee and a charge for the rides inside the park. Such a pricing policy is known as a **two-part tariff**. Suppose that all consumers have identical utility functions given by $u(x)$ and that the cost of production is cx . If the monopolist sets a two-part tariff, will it produce the socially efficient level of output, too little, or too much?

10.7. Selling Cars

A car dealer must pay \$10,000 to the manufacturer for each car he adds to his inventory. He faces three buyers. From the point of view of the dealer, Smith's valuation is uniformly distributed between \$11,000 and \$21,000, Jones's is between \$9,000 and \$11,000, and Brown's is between \$4,000 and \$12,000. The dealer's policy is to make a separate take-it-or-leave-it offer to each customer, and he is smart enough to avoid making different offers to customers who could resell to each other. Use the notation that the maximum valuation is $\bar{V}$ and the range of valuations is R .

- What will the offers be?
- Who is most likely to buy a car? How does this compare with the outcome with perfect price discrimination under full information? How does it compare with the outcome when the dealer charges \$10,000 to each customer?
- What happens to the equilibrium prices if with probability 0.25 each buyer has a valuation of \$0, but the probability distribution remains otherwise the same? What happens to the equilibrium expected profit?
- What happens to the equilibrium price the seller offers to seller Jones if with probability 0.25 Jones has a valuation of \$30,000, but with probability 0.75 his valuation is uniformly distributed between \$9,000 and \$11,000 as before? Show the relation between price and profit on a rough graph.

11 Signalling

October 3, 1999. January 18, 2000. November 30, 2003. Eric Rasmusen, Erasmuse@indiana.edu.
[Http: //mypage.iu.edu/~erasmuse](http://mypage.iu.edu/~erasmuse). Footnotes starting with xxx are the author's notes to himself. Comments welcomed.

11.1: The Informed Player Moves First: Signalling

Signalling is a way for an agent to communicate his type under adverse selection. The signalling contract specifies a wage that depends on an observable characteristic — the signal — which the agent chooses for himself after Nature chooses his type. Figures 1d and 1e showed the extensive forms of two kinds of models with signals. If the agent chooses his signal before the contract is offered, he is signalling to the principal. If he chooses the signal afterwards, the principal is screening him. Not only will it become apparent that this difference in the order of moves is important, it will also be seen that signalling costs must differ between agent types for signalling to be useful, and the outcome is often inefficient.

We begin with signalling models in which workers choose education levels to signal their abilities. Section 11.1 lays out the fundamental properties of a signalling model, and Section 11.2 shows how the details of the model affect the equilibrium. Section 11.3 steps back from the technical detail to more practical considerations in applying the model to education. Section 11.4 turns the game into a screening model. Section 11.5 switches to diagrams and applies signalling to new stock issues to show how two signals need to be used when the agent has two unobservable characteristics. Section 11.6 addresses the rather different idea of signal jamming: strategic behavior a player uses to cover up information rather than to disclose it.

Spence (1973) introduced the idea of signalling in the context of education. We will construct a series of models which formalize the notion that education has no direct effect on a person's ability to be productive in the real world but useful for demonstrating his ability to employers. Let half of the workers have the type "high ability" and half "low ability," where ability is a number denoting the dollar value of his output. Output is assumed to be a noncontractible variable and there is no uncertainty. If output is contractible, it should be in the contract, as we have seen in Chapter 7. Lack of uncertainty is a simplifying assumption, imposed so that the contracts are functions only of the signals rather than a combination of the signal and the output.

Employers do not observe the worker's ability, but they do know the distribution of abilities, and they observe the worker's education. To simplify, we will specify that the players are one worker and two employers. The employers compete profits down to zero and the worker receives the gains from trade. The worker's strategy is his education level and his choice of employer. The employers' strategies are the contracts they offer giving wages as functions of education level. The key to the model is that the signal, education, is less costly for workers with higher ability.

In the first four variants of the game, workers choose their education levels before employers decide how pay should vary with education.

Education I

Players

A worker and two employers.

The Order of Play

- 0 Nature chooses the worker's ability $a \in \{2, 5.5\}$, the *Low* and *High* ability each having probability 0.5. The variable a is observed by the worker, but not by the employers.
- 1 The worker chooses education level $s \in \{0, 1\}$.
- 2 The employers each offer a wage contract $w(s)$.
- 3 The worker accepts a contract, or rejects both of them.
- 4 Output equals a .

Payoffs

The worker's payoff is his wage minus his cost of education, and the employer's is his profit.

$$\pi_{worker} = \begin{cases} w - 8s/a & \text{if the worker accepts contract } w. \\ 0 & \text{if the worker rejects both contracts.} \end{cases}$$

$$\pi_{employer} = \begin{cases} a - w & \text{for the employer whose contract is accepted.} \\ 0 & \text{for the other employer.} \end{cases}$$

The payoffs assume that education is more costly for a worker if his ability takes a lower value, which is what permits separation to occur.¹ As in any hidden knowledge game, we must think about both pooling and separating equilibria. Education I has both. In the pooling equilibrium, which we will call Pooling Equilibrium 1.1, both types of workers pick zero education and the employers pay the zero-profit wage of 3.75 regardless of the education level ($3.75 = [2+5.5]/2$).

$$\text{Pooling Equilibrium 1.1} \quad \begin{cases} s(Low) = s(High) = 0 \\ w(0) = w(1) = 3.75 \\ Prob(a = Low|s = 1) = 0.5 \end{cases}$$

Pooling Equilibrium 1.1 needs to be specified as a perfect Bayesian equilibrium rather than simply a Nash equilibrium because of the importance of the interpretation that the uninformed player puts on out-of-equilibrium behavior. The equilibrium needs to specify the employer's beliefs when he observes $s = 1$, since that is never observed in equilibrium. In Pooling Equilibrium 1.1, the beliefs are passive conjectures (see Section 6.2): employers believe that a worker who chooses $s = 1$ is *Low* with the prior probability, which is 0.5. Given this belief, both types of workers realize that education is useless, and the model reaches the unsurprising outcome that workers do not bother to acquire unproductive education.

Under other beliefs, the pooling equilibrium breaks down. Under the belief $Prob(a = Low|s = 1) = 0$, for example, employers believe that any worker who acquired education

¹xxx For each variant, formally ask whether the single-crossing property is satisfied.

is a *High*, so pooling is not Nash because the *High* workers are tempted to deviate and acquire education. This leads to the separating equilibrium for which signalling is best known, in which the high-ability worker acquires education to prove to employers that he really has high ability.

$$\textbf{Separating Equilibrium 1.2} \quad \begin{cases} s(\textit{Low}) = 0, s(\textit{High}) = 1 \\ w(0) = 2, w(1) = 5.5 \end{cases}$$

Following the method used in Chapters 7 to 10, we will show that Separating Equilibrium 1.2 is a perfect Bayesian equilibrium by using the standard constraints which an equilibrium must satisfy. A pair of separating contracts must maximize the utility of the *Highs* and the *Lows* subject to two constraints: (a) the participation constraints that the firms can offer the contracts without making losses; and (b) the self-selection constraints that the *Lows* are not attracted to the *High* contract, and the *Highs* are not attracted by the *Low* contract. The participation constraints for the employers require that

$$w(0) \leq a_L = 2 \quad \text{and} \quad w(1) \leq a_H = 5.5. \quad (1)$$

Competition between the employers makes the expressions in (1) hold as equalities. The self-selection constraint of the *Lows* is

$$U_L(s = 0) \geq U_L(s = 1), \quad (2)$$

which in Education I is

$$w(0) - 0 \geq w(1) - \frac{8(1)}{2}. \quad (3)$$

Since in Separating Equilibrium 1.2 the separating wage of the *Lows* is 2 and the separating wage of the *Highs* is 5.5 from (1), the self-selection constraint (3) is satisfied.

The self-selection constraint of the *Highs* is

$$U_H(s = 1) \geq U_H(s = 0), \quad (4)$$

which in Education I is

$$w(1) - \frac{8(1)}{5.5} \geq w(0) - 0. \quad (5)$$

Constraint (5) is satisfied by Separating Equilibrium 1.2.

There is another conceivable pooling equilibrium for Education I, in which $s(\textit{Low}) = s(\textit{High}) = 1$, but this turns out not to be an equilibrium, because the *Lows* would deviate to zero education. Even if such a deviation caused the employer to believe they were low - ability with probability 1 and reduce their wage to 2, the low - ability workers would still prefer to deviate, because

$$U_L(s = 0) = 2 \geq U_L(s = 1) = 3.75 - \frac{8(1)}{2}. \quad (6)$$

Thus, a pooling equilibrium with $s = 1$ would violate incentive compatibility for the *Low* workers.

Notice that we do not need to worry about a nonpooling constraint for this game, unlike in the case of the games of Chapter 9. One might think that because employers compete for workers, competition between them might result in their offering a pooling contract that the high-ability workers would prefer to the separating contract. The reason this does not matter is that the employers do not compete by offering contracts, but by reacting to workers who have acquired education. That is why this is signalling and not screening: the employers cannot offer contracts in advance that change the workers' incentives to acquire education.

We can test the equilibrium by looking at the best responses. Given the worker's strategy and the other employer's strategy, an employer must pay the worker his full output or lose him to the other employer. Given the employers' contracts, the *Low* has a choice between the payoff 2 for ignorance ($= 2 - 0$) and 1.5 for education ($= 5.5 - 8/2$), so he picks ignorance. The *High* has a choice between the payoff 2 for ignorance ($= 2 - 0$) and 4.05 for education ($= 5.5 - 8/5.5$, rounded), so he picks education.

Unlike the pooling equilibrium, the separating equilibrium does not need to specify beliefs. Either of the two education levels might be observed in equilibrium, so Bayes's Rule always tells the employers how to interpret what they see. If they see that an agent has acquired education, they deduce that his ability is *High* and if they see that he has not, they deduce that it is *Low*. A worker is free to deviate from the education level appropriate to his type, but the employers' beliefs will continue to be based on equilibrium behavior. If a *High* worker deviates by choosing $s = 0$ and tells the employers he is a *High* who would rather pool than separate, the employers disbelieve him and offer him the *Low* wage of 2 that is appropriate to $s = 0$, not the pooling wage of 3.75 or the *High* wage of 5.5.

Separation is possible because education is more costly for workers if their ability is lower. If education were to cost the same for both types of worker, education would not work as a signal, because the low-ability workers would imitate the high-ability workers. This requirement of different signalling costs is known as the **single-crossing property**, since when the costs are depicted graphically, as they will be in Figure 1, the indifference curves of the two types intersect a single time (see also Section 10.3).

A strong case can be made that the beliefs required for the pooling equilibria are not sensible. Harking back to the equilibrium refinements of Section 6.2, recall that one suggestion (from Cho & Kreps [1987]) is to inquire into whether one type of player could not possibly benefit from deviating, no matter how the uninformed player changed his beliefs as a result. Here, the *Low* worker could never benefit from deviating from Pooling Equilibrium 1.1. Under the passive conjectures specified, the *Low* has a payoff of 3.75 in equilibrium versus -0.25 ($= 3.75 - 8/2$) if he deviates and becomes educated. Under the belief that most encourages deviation – that a worker who deviates is *High* with probability one – the *Low* would get a wage of 5.5 if he deviated, but his payoff from deviating would only be 1.5 ($= 5.5 - 8/2$), which is less than 2. The more reasonable belief seems to be that a worker who acquires education is a *High*, which does not support the pooling equilibrium.

The nature of the separating equilibrium lends support to the claim that education *per se* is useless or even pernicious, because it imposes social costs but does not increase total output. While we may be reassured by the fact that Professor Spence himself thought it

worthwhile to become Dean of Harvard College, the implications are disturbing and suggest that we should think seriously about how well the model applies to the real world. We will do that in Section 11.3. For now, note that in the model, unlike most real-world situations, information about the agent's talent has no social value, because all agents would be hired and employed at the same task even under full information. Also, if side payments are not possible, Separating Equilibrium 1.2 is second-best efficient in the sense that a social planner could not make both types of workers better off. Separation helps the high-ability workers even though it hurts the low-ability workers.

11.2: Variants on the Signalling Model of Education

Although *Education I* is a curious and important model, it does not exhaust the implications of signalling. This section will start with *Education II*, which will show an alternative to the arbitrary assumption of beliefs in the perfect Bayesian equilibrium concept. *Education III* will be the same as *Education I* except for its different parameter values, and will have two pooling equilibrium rather than one separating and one pooling equilibrium. *Education IV* will allow a continuum of education levels, and will unify *Education I* and *Education III* by showing how all of their equilibria and more can be obtained in a model with a less restricted strategy space.

Education II: Modelling Trembles so Nothing is Out of Equilibrium

The pooling equilibrium of *Education I* required the modeller to specify the employers' out-of-equilibrium beliefs. An equivalent model constructs the game tree to support the beliefs instead of introducing them via the equilibrium concept. This approach was briefly mentioned in connection with the game of *PhD Admissions* in Section 6.2. The advantage is that the assumptions on beliefs are put in the rules of the game along with the other assumptions. So let us replace Nature's move in *Education I* and modify the payoffs as follows.

Education II

The Order of Play

0 Nature chooses worker ability $a \in \{2, 5.5\}$, each ability having probability 0.5. (a is observed by the worker, but not by the employer.) With probability 0.001, Nature endows a worker with free education.

...

Payoffs

$$\pi_{worker} = \begin{cases} w - 8s/a & \text{if the worker accepts contract } w \text{ (ordinarily)} \\ w & \text{if the worker accepts contract } w \text{ (with free education)} \\ 0 & \text{if the worker does not accept a contract} \end{cases}$$

With probability 0.001 the worker receives free education regardless of his ability. If the employer sees a worker with education, he knows that the worker might be one of this rare type, in which case the probability that the worker is *Low* is 0.5. Both $s = 0$

and $s = 1$ can be observed in any equilibrium and *Education II* has almost the same two equilibria as *Education I*, without the need to specify beliefs.² The separating equilibrium did not depend on beliefs, and remains an equilibrium. What was Pooling Equilibrium 1.1 becomes “almost” a pooling equilibrium — almost all workers behave the same, but the small number with free education behave differently. The two types of greatest interest, the *High* and the *Low*, are not separated, but the ordinary workers are separated from the workers whose education is free. Even that small amount of separation allows the employers to use Bayes’s Rule and eliminates the need for exogenous beliefs.

Education III: No Separating Equilibrium, Two Pooling Equilibria

Let us next modify *Education I* by changing the possible worker abilities from $\{2, 5.5\}$ to $\{2, 12\}$. The separating equilibrium vanishes, but a new pooling equilibrium emerges. In Pooling Equilibria 3.1 and 3.2, both pooling contracts pay the same zero-profit wage of 7 ($= [2 + 12]/2$), and both types of agents acquire the same amount of education, but the amount depends on the equilibrium.

$$\begin{array}{ll}
 \text{Pooling Equilibrium 3.1} & \left\{ \begin{array}{l} s(\text{Low}) = s(\text{High}) = 0 \\ w(0) = w(1) = 7 \\ \text{Prob}(a = \text{Low} | s = 1) = 0.5 \text{ (passive conjectures)} \end{array} \right. \\
 \\
 \text{Pooling Equilibrium 3.2} & \left\{ \begin{array}{l} s(\text{Low}) = s(\text{High}) = 1 \\ w(0) = 2, w(1) = 7 \\ \text{Prob}(a = \text{Low} | s = 0) = 1 \end{array} \right.
 \end{array}$$

Pooling Equilibrium 3.1 is similar to the pooling equilibrium in *Education I* and *II*, but Pooling Equilibrium 3.2 is inefficient. Both types of workers receive the same wage, but they incur the education costs anyway. Each type is frightened to do without education because the employer would pay him not as if his ability were average, but as if he were known to be *Low*.

Examination of Pooling Equilibrium 3.2 shows why a separating equilibrium no longer exists. Any separating equilibrium would require $w(0) = 2$ and $w(1) = 7$, but this is the contract that leads to Pooling Equilibrium 3.2. The self-selection and zero-profit constraints cannot be satisfied simultaneously, because the *Low* type is willing to acquire $s = 1$ to obtain the high wage.

It is not surprising that information problems create inefficiencies in the sense that first-best efficiency is lost. Indeed, the surprise is that in some games with asymmetric information, such as Broadway Game I in Section 7.4, the first-best can still be achieved by tricks such as boiling-in-oil contracts. More often, we discover that the outcome is second-best efficient: given the informational constraints, a social planner could not alter the equilibrium without hurting some type of player. Pooling Equilibrium 3.2 is not even second-best efficient, because Pooling Equilibrium 3.1 and Pooling Equilibrium 3.2 result in the exact same wages and allocation of workers to tasks. The inefficiency is purely

²xxx Use Bayes Rule to show exactly what hte beliefs are here.

a problem of unfortunate expectations, like the inefficiency from choosing the dominated equilibrium in Ranked Coordination.

Pooling Equilibrium 3.2 also illustrates a fine point of the definition of pooling, because although the two types of workers adopt the same strategies, the equilibrium contract offers different wages for different education. The implied threat to pay a low wage to an uneducated worker never needs to be carried out, so the equilibrium is still called a pooling equilibrium. Notice that perfectness does not rule out threats based on beliefs. The model imposes these beliefs on the employer, and he would carry out his threats, because he believes they are best responses. The employer receives a higher payoff under some beliefs than under others, but he is not free to choose his beliefs.

Following the approach of *Education II*, we could eliminate Pooling Equilibrium 3.2 by adding an exogenous probability 0.001 that either type is completely unable to buy education. Then no behavior is never observed in equilibrium and we end up with Pooling Equilibrium 3.1 because the only rational belief is that if $s = 0$ is observed, the worker has equal probability of being *High* or being *Low*. To eliminate Pooling Equilibrium 3.1 requires less reasonable beliefs; for example, a probability of 0.001 that a *Low* gets free education together with a probability of 0 that a *High* does.

These first three games illustrate the basics of signalling: (a) separating and pooling equilibria both may exist, (b) out-of-equilibrium beliefs matter, and (c) sometimes one perfect Bayesian equilibrium can Pareto dominate others. These results are robust, but Education IV will illustrate some dangers of using simplified games with binary strategy spaces instead of continuous and unbounded strategies. So far education has been limited to $s = 0$ or $s = 1$; Education IV allows it to take greater or intermediate values.

Education IV: Continuous Signals and Continua of Equilibria

Let us now return to *Education I*, with one change: that education s can take any level on the continuum between 0 and infinity.³

The game now has continua of pooling and separating equilibria which differ according to the value of education chosen. In the pooling equilibria, the equilibrium education level is s^* , where each s^* in the interval $[0, \bar{s}]$ supports a different equilibrium. The out-of-equilibrium belief most likely to support a pooling equilibrium is $Prob(a = Low|s \neq s^*) = 1$, so let us use this to find the value of $\bar{s}$, the greatest amount of education that can be generated by a pooling equilibrium. The equilibrium is Pooling Equilibrium 4.1, where $s^* \in [0, \bar{s}]$.

$$\text{Pooling Equilibrium 4.1} \quad \left\{ \begin{array}{l} s(Low) = s(High) = s^* \\ w(s^*) = 3.75 \\ w(s \neq s^*) = 2 \\ Prob(a = Low|s \neq s^*) = 1 \end{array} \right.$$

The critical value $\bar{s}$ can be discovered from the incentive compatibility constraint of

³xxx Givre the entire game description again.

the *Low* type, which is binding if $s^* = \bar{s}$. The most tempting deviation is to zero education, so that is the deviation that appears in the constraint.

$$U_L(s = 0) = 2 \leq U_L(s = \bar{s}) = 3.75 - \frac{8\bar{s}}{2}. \quad (7)$$

Equation (7) yields $\bar{s} = \frac{7}{16}$. Any value of s^* less than $\frac{7}{16}$ will also support a pooling equilibrium. Note that the incentive-compatibility constraint of the *High* type is not binding. If a *High* deviates to $s = 0$, he, too, will be thought to be a *Low*, so

$$U_H(s = 0) = 2 \leq U_H(s = \frac{7}{16}) = 3.75 - \frac{8\bar{s}}{5.5} \approx 3.1. \quad (8)$$

In the separating equilibria, the education levels chosen in equilibrium are 0 for the *Low*'s and s^* for the *High*'s, where each s^* in the interval $[\bar{s}, \bar{\bar{s}}]$ supports a different equilibrium. A difference from the case of separating equilibria in games with binary strategy spaces is that now there are possible out-of-equilibrium actions even in a separating equilibrium. The two types of workers will separate to two education levels, but that leaves an infinite number of out-of-equilibrium education levels. As before, let us use the most extreme belief for the employers' beliefs after observing an out-of-equilibrium education level: that $Prob(a = Low | s \neq s^*) = 1$. The equilibrium is Separating Equilibrium 4.2, where $s^* \in [\bar{s}, s]$.

$$\textbf{Separating Equilibrium 4.2} \quad \left\{ \begin{array}{l} s(Low) = 0, \quad s(High) = s^* \\ w(s^*) = 5.5 \\ w(s \neq s^*) = 2 \\ Prob(a = Low | s \notin \{0, s^*\}) = 1 \end{array} \right.$$

The critical value $\bar{s}$ can be discovered from the incentive-compatibility constraint of the *Low*, which is binding if $s^* = \bar{s}$.

$$U_L(s = 0) = 2 \geq U_L(s = \bar{s}) = 5.5 - \frac{8\bar{s}}{2}. \quad (9)$$

Equation (9) yields $\bar{s} = \frac{7}{8}$. Any value of s^* greater than $\frac{7}{8}$ will also deter the *Low* workers from acquiring education. If the education needed for the wage of 5.5 is too great, the *High* workers will give up on education too. Their incentive compatibility constraint requires that

$$U_H(s = 0) = 2 \leq U_H(s = \bar{\bar{s}}) = 5.5 - \frac{8\bar{\bar{s}}}{5.5}. \quad (10)$$

Equation (10) yields $\bar{\bar{s}} = \frac{77}{32}$. s^* can take any lower value than $\frac{77}{32}$ and the *High*'s will be willing to acquire education.

The big difference from *Education I* is that *Education IV* has Pareto-ranked equilibria. Pooling can occur not just at zero education but at positive levels, as in *Education III*, and the pooling equilibria with positive education levels are all Pareto inferior. Also, the separating equilibria can be Pareto ranked, since separation with $s^* = \bar{s}$ dominates

separation with $s^* = \bar{s}$. Using a binary strategy space instead of a continuum conceals this problem.

Education IV also shows how restricting the strategy space can alter the kinds of equilibria that are possible. *Education III* had no separating equilibrium because at the maximum possible signal, $s = 1$, the *Low*'s were still willing to imitate the *High*'s. *Education IV* would not have any separating equilibria either if the strategy space were restricted to allow only education levels less than $\frac{7}{8}$. Using a bounded strategy space eliminates possibly realistic equilibria.

This is not to say that models with binary strategy sets are always misleading. Education I is a fine model for showing how signalling can be used to separate agents of different types; it becomes misleading only when used to reach a conclusion such as "If a separating equilibrium exists, it is unique". As with any assumption, one must be careful not to narrow the model so much as to render vacuous the question it is designed to answer.

11.3 General Comments on Signalling in Education

Signalling and Similar Phenomena

The distinguishing feature of signalling is that the agent's action, although not directly related to output, is useful because it is related to ability. For the signal to work, it must be less costly for an agent with higher ability. Separation can occur in Education I because when the principal pays a greater wage to educated workers, only the *Highs*, whose utility costs of education are lower, are willing to acquire it. That is why a signal works where a simple message would not: actions speak louder than words.

Signalling is outwardly similar to other solutions to adverse selection. The high-ability agent finds it cheaper than the low-ability one to build a reputation, but the reputation-building actions are based directly on his high ability. In a typical reputation model he shows ability by producing high output period after period. Also, the nature of reputation is to require several periods of play, which signalling does not.

Another form of communication is possible when some observable variable not under the control of the worker is correlated with ability. Age, for example, is correlated with reliability, so an employer pays older workers more, but the correlation does not arise because it is easier for reliable workers to acquire the attribute of age. Because age is not an action chosen by the worker, we would not need game theory to model it.

Problems in Applying Signalling to Education

On the empirical level, the first question to ask of a signalling model of education is, "What is education?". For operational purposes this means, "In what units is education measured?". Two possible answers are "years of education" and "grade point average." If the sacrifice of a year of earnings is greater for a low-ability worker, years of education can serve as a signal. If less intelligent students must work harder to get straight As, then grade-point-average can also be a signal.

Layard & Psacharopoulos (1974) give three rationales for rejecting signalling as an

important motive for education. First, dropouts get as high a rate of return on education as those who complete degrees, so the signal is not the diploma, although it might be the years of education. Second, wage differentials between different education levels rise with age, although one would expect the signal to be less important after the employer has acquired more observations on the worker's output. Third, testing is not widely used for hiring, despite its low cost relative to education. Tests are available, but unused: students commonly take tests like the American SAT whose results they could credibly communicate to employers, and their scores correlate highly with subsequent grade point average. One would also expect an employer to prefer to pay an 18-year-old low wages for four years to determine his ability, rather than waiting to see what grades he gets as a history major.

Productive Signalling

Even if education is largely signalling, we might not want to close the schools. Signalling might be wasteful in a pooling equilibrium like Pooling Equilibrium 3.2, but in a separating equilibrium it can be second-best efficient for at least three reasons. First, it allows the employer to match workers with jobs suited to their talents. If the only jobs available were “professor” and “typist,” then in a pooling equilibrium, both *High* and *Low* workers would be employed, but they would be randomly allocated to the two jobs. Given the principle of comparative advantage, typing might improve, but I think, pridefully, that research would suffer.

Second, signalling keeps talented workers from moving to jobs where their productivity is lower but their talent is known. Without signalling, a talented worker might leave a corporation and start his own company, where he would be less productive but better paid. The naive observer would see that corporations hire only one type of worker (*Low*), and imagine there was no welfare loss.

Third, if ability is endogenous — moral hazard rather than adverse selection — signalling encourages workers to acquire ability. One of my teachers said that you always understand your next-to-last econometrics class. Suppose that solidly learning econometrics increases the student's ability, but a grade of A is not enough to show that he solidly learned the material. To signal his newly acquired ability, the student must also take “Time Series,” which he cannot pass without a solid understanding of econometrics. “Time Series” might be useless in itself, but if it did not exist, the students would not be able to show he had learned basic econometrics.

11.4: The Informed Player Moves Second: Screening

In screening games, the informed player moves second, which means that he moves in response to contracts offered by the uninformed player. Having the uninformed player make the offers is important because his offer conveys no information about himself, unlike in a signalling model.

Education V: Screening with a Discrete Signal

Players

A worker and two employers.

The Order of Play

0 Nature chooses worker ability $a \in \{2, 5.5\}$, each ability having probability 0.5. Employers do not observe ability, but the worker does.

1 Each employer offers a wage contract $w(s)$.

2 The worker chooses education level $s \in \{0, 1\}$.

3 The worker accepts a contract, or rejects both of them.

4 Output equals a .

Payoffs

$$\pi_{worker} = \begin{cases} w - \frac{8s}{a} & \text{if the worker accepts contract } w. \\ 0 & \text{if the worker rejects both contracts.} \end{cases}$$

$$\pi_{employer} = \begin{cases} a - w & \text{for the employer whose contract is accepted.} \\ 0 & \text{for the other employer.} \end{cases}$$

Education V has no pooling equilibrium, because if one employer tried to offer the zero profit pooling contract, $w(0) = 3.75$, the other employer would offer $w(1) = 5.5$ and draw away all the *Highs*. The unique equilibrium is

$$\textbf{Separating Equilibrium 5.1} \quad \begin{cases} s(Low) = 0, s(High) = 1 \\ w(0) = 2, w(1) = 5.5 \end{cases}$$

Beliefs do not need to be specified in a screening model. The uninformed player moves first, so his beliefs after seeing the move of the informed player are irrelevant. The informed player is fully informed, so his beliefs are not affected by what he observes. This is much like simple adverse selection, in which the uninformed player moves first, offering a set of contracts, after which the informed player chooses one of them. The modeller does not need to refine perfectness in a screening model. The similarity between adverse selection and screening is strong enough that Education V would not have been out of place in Chapter 9, but it is presented here because the context is so similar to the signalling models of education.

Education VI allows a continuum of education levels, in a game otherwise the same as Education V.

Education VI: Screening with a Continuous Signal

Players

A worker and two employers.

The Order of Play

0 Nature chooses worker ability $a \in \{2, 5.5\}$, each ability having probability 0.5. Employers do not observe ability, but the worker does.

1 Each employer offers a wage contract $w(s)$.

2 The worker choose education level $s \in [0, 1]$.

3 The worker chooses a contract, or rejects both of them.

4 Output equals a .

Payoffs.

$$\pi_{worker} = \begin{cases} w - 8s/a & \text{if the worker accepts contract } w. \\ 0 & \text{if the worker rejects both contracts.} \end{cases}$$

$$\pi_{employer} = \begin{cases} a - w & \text{for the employer whose contract is accepted.} \\ 0 & \text{for the other employer.} \end{cases}$$

Pooling equilibria generally do not exist in screening games with continuous signals, and sometimes separating equilibria in pure strategies do not exist either — recall Insurance Game III from Section 9.4. Education VI, however, does have a separating Nash equilibrium, with a unique equilibrium path.

$$\text{Separating Equilibrium 6.1} \quad \begin{cases} s(Low) = 0, s(High) = 0.875 \\ w = \begin{cases} 2 & \text{if } s < 0.875 \\ 5.5 & \text{if } s \geq 0.875 \end{cases} \end{cases}$$

In any separating contract, the *Lows* must be paid a wage of 2 for an education of 0, because this is the most attractive contract that breaks even. The separating contract for the *Highs* must maximize their utility subject to the constraints discussed in Education I. When the signal is continuous, the constraints are especially useful to the modeller for calculating the equilibrium. The participation constraints for the employers require that

$$w(0) \leq a_L = 2 \quad \text{and} \quad w(s^*) \leq a_H = 5.5, \quad (11)$$

where s^* is the separating value of education that we are trying to find. Competition turns the inequalities in (11) into equalities. The self selection constraint for the low-ability workers is

$$U_L(s = 0) \geq U_L(s = s^*), \quad (12)$$

which in Education VI is

$$w(0) - 0 \geq w(s^*) - \frac{8s^*}{2}. \quad (13)$$

Since the separating wage is 2 for the *Lows* and 5.5 for the *Highs*, constraint (13) is satisfied as an equality if $s^* = 0.875$, which is the crucial education level in Separating Equilibrium 6.1.

$$U_H(s = 0) = w(0) \leq U_H(s = s^*) = w(s^*) - \frac{8s^*}{5.5}. \quad (14)$$

If $s^* = 0.875$, inequality (14) is true, and it would also be true for higher values of s^* . Unlike the case of the continuous-strategy signalling game, Education IV, however, the equilibrium contract in Education VI is unique, because the employers compete to offer the most attractive contract that satisfies the participation and incentive compatibility constraints. The most attractive is the separating contract that Pareto dominates the other separating contracts by requiring the relatively low separating signal of $s^* = 0.875$.

Similarly, competition in offering attractive contracts rules out pooling contracts. The nonpooling constraint, required by competition between employers, is

$$U_H(s = s^*) \geq U_H(\text{pooling}), \quad (15)$$

which, for Education VI, is, using the most attractive possible pooling contract,

$$w(s^*) - \frac{8s^*}{5.5} \geq 3.75. \quad (16)$$

Since the payoff of *Highs* in the separating contract is 4.23 ($= 5.5 - 8 \cdot 0.875 / 5.5$, rounded), the nonpooling constraint is satisfied.

No Pooling Equilibrium in Education VI

Education VI lacks a pooling equilibrium, which would require the outcome $\{s = 0, w(0) = 3.75\}$, shown as C_1 in Figure 1. If one employer offered a pooling contract requiring more than zero education (such as the inefficient Pooling Equilibrium 3.2), the other employer could make the more attractive offer of the same wage for zero education. The wage is 3.75 to ensure zero profits. The rest of the wage function — the wages for positive education levels — can take a variety of shapes, so long as the wage does not rise so fast with education that the *Highs* are tempted to become educated.

But no equilibrium has these characteristics. In a Nash equilibrium, no employer can offer a pooling contract, because the other employer could always profit by offering a separating contract paying more to the educated. One such separating contract is C_2 in Figure 1, which pays 5 to workers with an education of $s = 0.5$ and yields a payoff of 4.89 to the *Highs* ($= 5 - [8 \cdot 0.5] / 5.5$, rounded) and 3 to the *Lows* ($= 5 - 8 \cdot 0.5 / 2$). Only *Highs* prefer C_2 to the pooling contract C_1 , which yields payoffs of 3.75 to both *High* and *Low*, and if only *Highs* accept C_2 , it yields positive profits to the employer.




Figure 1: Education VI: no pooling Nash equilibrium

Nonexistence of a pooling equilibrium in screening models without continuous strategy spaces is a general result. The linearity of the curves in Education VI is special, but in any screening model the *Lows* would have greater costs of education, which is equivalent to steeper indifference curves. This is the **single-crossing property** alluded to in Education I. Any pooling equilibrium must, like C_1 , lie on the vertical axis where education is zero and the wage equals the average ability. A separating contract like C_2 can always be found to the northeast of the pooling contract, between the indifference curves of the two types, and it will yield positive profits by attracting only the *Highs*.

Education VII: No Nash Equilibrium in Pure Strategies

In Education VI we showed that screening models have no pooling equilibria. In Education VII the parameters are changed a little to eliminate even the separating equilibrium in pure strategies. Let the proportion of *Highs* be 0.9 instead of 0.5, so the zero-profit pooling wage is 5.15 ($= 0.9[5.5] + 0.1[2]$) instead of 3.75. Consider the separating contracts C_3 and C_4 , shown in Figure 2, calculated in the same way as Separating Equilibrium 5.1. The pair (C_3, C_4) is the most attractive pair of contracts that separates *Highs* from *Lows*. *Low* workers accept contract C_3 , obtain $s = 0$, and receive a wage of 2, their ability. *Highs* accept contract C_4 , obtain $s = 0.875$, and receive a wage of 5.5, their ability. Education is not attractive to *Lows* because the *Low* payoff from pretending to be *High* is 2 ($= 5.5 - 8 \cdot 0.875/2$), no better than the *Low* payoff of 2 from C_3 ($= 2 - 8 \cdot 0/2$).




Figure 2: Education VII: No Nash Equilibrium

The wage of the pooling contract C_5 is 5.15, so that even the *Highs* strictly prefer C_5 to (C_3, C_4) . But our reasoning that no pooling equilibrium exists is still valid; some contract C_6 would attract all the *Highs* from C_5 . No Nash equilibrium in pure strategies exists, either separating or pooling.

A Summary of the Education Models

Because of signalling's complexity, most of this chapter has been devoted to elaboration of the education model. We began with Education I, which showed how with two types and two signal levels the perfect Bayesian equilibrium could be either separating or pooling. Education II took the same model and replaced the specification of out-of-equilibrium beliefs with an additional move by Nature, while Education III changed the parameters in Education I to increase the difference between types and to show how signalling could continue with pooling. Education IV changed Education I by allowing a continuum of education levels, which resulted in a continuum of inefficient equilibria, each with a different signal level. After a purely verbal discussion of how to apply signalling models, we looked at screening, in which the employer moves first. Education V was a screening reprise of Education I, while Education VI broadened the model to allow a continuous signal, which eliminates pooling equilibria. Education VII modified the parameters of Education VI to show that sometimes no pure-strategy Nash equilibrium exists at all.

Throughout it was implicitly assumed that all the players were risk neutral. Risk neutrality is unimportant, because there is no uncertainty in the model and the agents bear no risk. If the workers were risk averse and they differed in their degrees of risk aversion, the contracts could try to use the difference to support a separating equilibrium because willingness to accept risk might act as a signal. If the principal were risk averse he might offer a wage less than the average productivity in the pooling equilibrium, but he is under no risk at all in the separating equilibrium, because it is fully revealing. The models are also games of certainty, and this too is unimportant. If output were uncertain, agents would just make use of the expected payoffs rather than the raw payoffs and very little would change.

We could extend the education models further — allowing more than two levels of ability would be a high priority — but instead, let us turn to the financial markets and look graphically at a model with two continuous characteristics of type and two continuous signals.

***11.5. Two Signals: Game of Underpricing New Stock Issues**

One signal might not be enough when there is not one but two characteristics of an agent that he wishes to communicate to the principal. This has been generally analyzed in Engers (1987), and multiple signal models have been especially popular in financial economics, for example, the multiple signal model used to explain the role of investment bankers in new stock issues by Hughes (1986). We will use a model of initial public offerings of stock as the example in this section.

Empirically, it has been found that companies consistently issue stock at a price so low that it rises sharply in the days after the issue, an abnormal return estimated to average 11.4 percent (Copeland & Weston [1988], p. 377). The game of Underpricing New Stock Issues tries to explain this using the percentage of the stock retained by the original owner and the amount of underpricing as two signals. The two characteristics being signalled are the mean of the value of the new stock, which is of obvious concern to the potential buyers, and the variance, the importance of which will be explained later.

Underpricing New Stock Issues (Grinblatt & Hwang [1989])

Players

The entrepreneur and many investors.

The Order of Play

(See Figure 3a of Chapter 2 for a time line.)

0 Nature chooses the expected value (μ) and variance (σ^2) of a share of the firm using some distribution F .

1 The entrepreneur retains fraction α of the stock and offers to sell the rest at a price per share of P_0 .

2 The investors decide whether to accept or reject the offer.

3 The market price becomes P_1 , the investors' estimate of μ .

4 Nature chooses the value V of a share using some distribution G such that μ is the mean of V and σ^2 is the variance. With probability θ , V is revealed to the investors and becomes the market price.

5 The entrepreneur sells his remaining shares at the market price.

Payoffs

$$\pi_{\text{entrepreneur}} = U([1 - \alpha]P_0 + \alpha[\theta V + (1 - \theta)P_1]), \text{ where } U' > 0 \text{ and } U'' < 0.$$


$$\pi_{\text{investors}} = (1 - \alpha)(V - P_0) + \alpha(1 - \theta)(V - P_1).$$

The entrepreneur's payoff is the utility of the value of the shares he issues at P_0 plus the value of those he sells later at the price P_1 or V . The investors' payoff is the true value of the shares they buy minus the price they pay.

Underpricing New Stock Issues subsumes the simpler model of Leland & Pyle (1977), in which σ^2 is common knowledge and if the entrepreneur chooses to retain a large fraction of the shares, the investors deduce that the stock value is high. The one signal in that model is fully revealing because holding a larger fraction exposes the undiversified entrepreneur to a larger amount of risk, which he is unwilling to accept unless the stock value is greater than investors would guess without the signal.

If the variance of the project is high, that also increases the risk to the undiversified entrepreneur, which is important even though the investors are risk neutral and do not care directly about the value of σ^2 . Since the risk is greater when variance is high, the signal α is more effective and retaining a smaller amount allows the entrepreneur to sell the remainder at the same price as a larger amount for a lower-variance firm. Even though the investors are diversified and do not care directly about firm-specific risk, they are interested in the variance because it tells them something about the effectiveness of entrepreneur-retained shares as a signal of share value. Figure 3 shows the signalling schedules for two variance levels.


Figure 3: How the signal changes with the variance



In the game of Underpricing New Stock Issues, σ^2 is not known to the investors, so the signal is no longer fully revealing. An α equal to 0.1 could mean either that the firm has a low value with low variance, or a high value with high variance. But the entrepreneur can use a second signal, the price at which the stock is issued, and by observing α and P_0 , the investors can deduce μ and σ^2 .

I will use specific numbers for concreteness. The entrepreneur could signal that the stock has the high mean value, $\mu = 120$, in two ways: (a) retaining a high percentage, $\alpha = 0.4$, and making the initial offering at a high price of $P_0 = 90$, or (b) retaining a low percentage, $\alpha = 0.1$, and making the initial offering at a low price, $P_0 = 80$. Figure 4 shows the different combinations of initial price and fraction retained that might be used. If the stock has a high variance, he will want to choose behavior (b), which reduces his risk. Investors deduce that the stock of anyone who retains a low percentage and offers a low price actually has $\mu = 120$ and a high variance, so stock offered at the price of 80 rises in price. If, on the other hand, the entrepreneur retained $\alpha = .1$ and offered the high price $P_0 = 90$, investors would conclude that μ was lower than 120 but that variance was low also, so the stock would not rise in price. The low price conveys the information that this stock has a high mean and high variance rather than a low mean and low variance.

Figure 4: Different ways to signal a given μ .



This model explains why new stock is issued at a low price. The entrepreneur knows that the price will rise, but only if he issues it at a low initial price to show that the variance is high. The price discount shows that signalling by holding a large fraction of stock is unusually costly, but he is none the less willing to signal. The discount is costly because he is selling stock at less than its true value, and retaining stock is costly because he bears extra risk, but both are necessary to signal that the stock is valuable.

*11.6 Signal Jamming and Limit Pricing

This chapter has examined a number of models in which an informed player tries to convey information to an uninformed player by some means or other — by entering into an incentive contract, or by signalling. Sometimes, however, the informed party has the opposite problem: his natural behavior would convey his private information but he wants to keep it secret. This happens, for example, if one firm is informed about its poor ability to compete successfully, and it wants to conceal this information from a rival. The informed player may then engage in costly actions, just as in signalling, but now the costly action will be **signal jamming** (a term coined in Fudenberg & Tirole [1986c]): preventing information from appearing rather than generating information.

The model I will use to illustrate signal jamming is the limit pricing model of Rasmusen (1997). Limit pricing refers to the practice of keeping prices low to deter entry. Limit pricing can be explained in a variety of ways; notably, as a way for the incumbent to signal that he has low enough costs that rivals would regret entering, as in Problem 6.2 and Milgrom & Roberts [1982a]. Here, the explanation for limit pricing will be signal jamming: by keeping profits low, the incumbent keeps it unclear to the rival whether the market is big enough to accommodate two firms profitably. In the model, the incumbent can control S , a public signal of the size of the market. In the model below, this signal is the price that the incumbent charges, but it could equally well represent the incumbent's choice of

advertising or capacity. The reason the signal is important is that the entrant must decide whether to enter based on his belief as to the probability that the market is large enough to support two firms profitably.

Limit Pricing as Signal Jamming

Players

The incumbent and the rival.

The Order of Play

0 Nature chooses the market size M to be M_{Small} with probability θ and M_{Large} with probability $(1 - \theta)$, observed only by the incumbent.

1 The incumbent chooses the signal S to equal s_0 or s_1 for the first period if the market is small, s_1 or s_2 if it is large. This results in monopoly profit $\mu f(S) - C$, where $\mu > 2$. Both players observe the value of S .

2 The rival decides whether to be *In* or *Out* of the market.

3 If the rival chooses *In*, each player incurs cost C in the second period and they each earn the duopoly profit $M - C$. Otherwise, the incumbent earns $\mu f(S) - C$.

Payoffs

If the rival does not enter, the payoffs are $\pi_{incumbent} = (\mu f(S) - C) + (\mu f(S) - C)$ and $\pi_{rival} = 0$.

If the rival does enter, the payoffs are $\pi_{incumbent} = (\mu f(S) - C) + (M - C)$ and $\pi_{rival} = M - C$. Assume that $f(s_0) < f(s_1) = M_{Small} < f(s_2) = M_{Large}$, $M_{Large} - C > 0$, and $M_{Small} - C < 0$.

Thus, if the incumbent chooses s_1 , his profit will equal the maximum profit from a small market, even if the market is really large, but if he chooses s_2 , his profit will be the maximum value for a large market – but that choice will have revealed that the market is large. The duopoly profit in a large market is large enough to sustain two firms, but the duopoly profit in a small market will result in losses for both firms.




Figure 5: Signal jamming

There are four equilibria, each appropriate to a different parameter region in Figure 5. A small enough value of the parameter μ , which represents the value to being a monopoly, leads to a nonstrategic equilibrium exists, in which the incumbent simply maximizes profits in each period separately. This equilibrium is: (E1: Nonstrategic. $s_2|Large$, $s_1|Small$, $Out|s_0$, $Out|s_1$, $In|s_2$). The incumbent's equilibrium payoff in a large market is $\pi_I(s_2|Large) = (\mu M_{Large} - C) + (M_{Large} - C)$, compared with the deviation payoff of $\pi_I(s_1|Large) = (\mu M_{Small} - C) + (\mu M_{Large} - C)$. The incumbent has no incentive to deviate if $\pi_I(s_2|Large) - \pi_I(s_1|Large) = (1 + \mu)M_{Large} - \mu(M_{Small} + M_{Large}) \geq 0$, which is equivalent to

$$\mu \leq \frac{M_{Large}}{M_{Small}}, \quad (17)$$

as shown in Figure 5. The rival will not deviate, because the incumbent's choice fully reveals the size of the market.

Signal jamming occurs if monopoly profits are somewhat higher, and if the rival would refrain from entering the market unless he decides it is more profitable than his prior beliefs would indicate. The equilibrium is (E2: Pure Signal-Jamming. $s_1|Large$, $s_1|Small$, $Out|s_0$, $Out|s_1$, $In|s_2$). The rival's strategy is the same as in E1, so the incumbent's optimal behavior remains the same, and he chooses s_1 if the opposite of condition (17) is true. As for the rival, if he stays out, his second-period payoff is 0, and if he enters, its expected value is $\theta(M_{Small} - C) + (1 - \theta)(M_{Large} - C)$. Hence, as shown in Figure 5, he will follow the equilibrium behavior of staying out if

$$\theta \geq \frac{M_{Large} - C}{M_{Large} - M_{Small}}. \quad (18)$$

The intuition behind the signal jamming equilibrium is straightforward. The incumbent knows he will attract entry if he fully exploits the market when it is large, so he purposely dulls his efforts to conceal whether the market is large or small. If potential entrants place enough prior probability on the market being small, and are thus unwilling to enter without positive information that the market is large, the incumbent can thus deter entry.

Signal jamming shows up in other contexts. A wealthy man may refrain from buying a big house, landscaping his front yard, or wearing expensive clothing in order to avoid being a target for thieves or for political leaders in search of wealthy victims to tax or loot. A cabinet with shaky support may purposely take risky assertive positions because greater caution might hint to his rivals that his position was insecure and induce them to campaign actively against him. A general may advance his troops even when he is outnumbered, because to go on the defensive would provoke the enemy to attack. Note, however, that in each of these examples it is key that the uninformed player decide not to act aggressively if he fails to acquire any information.

A mixed form of signal jamming occurs if the probability of a small market is low, so if the signal of first-period revenues was jammed completely, the rival would enter anyway.

This equilibrium is (E3: Mixed Signal Jamming. $(s_1|Small, s_1|Large$ with probability α , $s_2|Large$ with probability $(1-\alpha)$, $Out|s_0, In|s_1$ with probability β , $Out|s_1$ with probability $(1-\beta)$, $In|s_2$). If the incumbent played $s_2|Large$ and $s_1|Small$, the rival would interpret s_1 as indicating a small market — an interpretation which would give the incumbent incentive to play $s_1|Large$. But if the incumbent always plays s_1 , the rival would enter even after observing s_1 , knowing there was a high probability that the market was really large. Hence, the equilibrium must be in mixed strategies, which is equilibrium E3, or the incumbent must convince the rival to stay out by playing s_0 , which is equilibrium E4.

For the rival to mix, he must be indifferent between the second-period payoffs of $\pi_E(In|s_1) = \frac{\theta}{\theta+(1-\theta)\alpha}(M_{Small} - C) + \frac{(1-\theta)\alpha}{\theta+(1-\theta)\alpha}(M_{Large} - C)$ and $\pi_E(Out|s_1) = 0$. Equating these two payoffs and solving for α yields $\alpha = \left(\frac{\theta}{1-\theta}\right) \left(\frac{C-M_{Small}}{M_{Large}-C}\right)$, which is always nonnegative, but avoids equalling one only if condition (18) is false.

For the incumbent to mix when the market is large, he must be indifferent between $\pi_I(s_2|Large) = (\mu M_{Large} - C) + (M_{Large} - C)$ and $\pi_I(s_1|Large) = (\mu M_{Small} - C) + \beta(M_{Large} - C) + (1-\beta)(\mu M_{Large} - C)$. Equating these two payoffs and solving for β yields $\beta = \frac{\mu M_{Small} - M_{Large}}{(\mu-1)M_{Large}}$, which is strictly less than one, and which is nonnegative if condition (18) is false.

If the market is small, the incumbent's alternatives are the equilibrium payoff, $\pi_I(s_1|Small) = (\mu M_{Small} - C) + \beta(M_{Small} - C) + (1-\beta)(\mu M_{Small} - C)$, and the deviation payoff, $\pi_I(f(s_0)|Small) = (\mu f(s_0) - C) + (\mu M_{Small} - C)$. The difference is

$$\pi_I(s_1|Small) - \pi_I(f(s_0)|Small) = [\mu M_{Small} + \beta M_{Small} + (1-\beta)\mu M_{Small}] - [\mu f(s_0) + \mu M_{Small}] \quad (19)$$

Expression (19) is nonnegative under either of two conditions, both of which are found by substituting the equilibrium value of β into expression (19). The first is if $f(s_0)$ is small enough, a sufficient condition for which is

$$f(s_0) \leq M_{Small} \left(1 - \frac{M_{Small}}{M_{Large}}\right). \quad (20)$$

The second is if μ is no greater than some amount Z^{-1} defined so that

$$\mu \leq \left(\frac{M_{Small}}{M_{Large}} - 1 + \frac{f(s_0)}{M_{Small}}\right)^{-1} = Z^{-1}. \quad (21)$$

If condition (20) is false, then $Z^{-1} > \frac{M_{Large}}{M_{Small}}$, because $Z < \frac{M_{Small}}{M_{Large}}$ and $Z > 0$. Thus, we can draw region E3 as it is shown in Figure 5.

It follows that if condition (21) is replaced by its converse, the unique equilibrium is for the incumbent to choose $s_0|Small$, and the equilibrium is (E4: Signalling. $s_0|Small, s_2|Large, Out|s_0, In|s_1, In|s_2$). Passive conjectures will support this pooling signalling equilibrium, as will the out-of-equilibrium belief that if the rival observes s_1 , he believes the market is large with probability $\frac{(1-\theta)\alpha}{\theta+(1-\theta)\alpha}$, as in equilibrium E3.

The signalling equilibrium is also an equilibrium for other parameter regions outside of E4, though less reasonable beliefs are required. Let the out-of-equilibrium belief be

$Prob(Large|s_1) = 1$. The equilibrium payoff is $\pi_I(s_0|Small) = (\mu f(s_0) - C) + (\mu M_{Small} - C)$ and the deviation payoff is $\pi_I(s_1|Small) = (\mu M_{Small} - C) + (M_{Small} - C)$. The signalling equilibrium remains an equilibrium so long as $\mu \geq \frac{M_{Small}}{f(s_0)}$.

The signalling equilibrium is an interesting one, because it turns the asymmetric information problem full circle. The informed player wants to conceal his private information by costly signal jamming if the information is *Large*, so when the information is *Small*, the player must signal at some cost that he is not signal jamming. If E4 is the equilibrium, the incumbent is hurt by the possibility of signal jamming; he would much prefer a simpler world in which it was illegal or nobody considered the possibility. This is often the case: strategic behavior can help a player in some circumstances, but given that the other players know he might be behaving strategically, everyone would prefer a world in which everyone is honest and nonstrategic.

N11.1 The informed player moves first: signalling

- The term “signalling” was introduced by Spence (1973). The games in this book take advantage of hindsight to build simpler and more rational models of education than in his original article, which used a rather strange equilibrium concept: a strategy combination from which no worker has incentive to deviate and under which the employer’s profits are zero. Under that concept, the firm’s incentives to deviate are irrelevant.

The distinction between signalling and screening has been attributed to Stiglitz & Weiss (1989). The literature has shown wide variation in the use of both terms, and “signal” is such a useful word that it is often used in models that have no signalling of the kind discussed in this chapter.

- One convention sometimes used in signalling models is to call the signalling player (the agent), the **sender** and the player signalled to (the principal), the **receiver**.
- The applications of signalling are too many to properly list. A few examples are the use of prices in Wilson (1980) and Stiglitz (1987), the payment of dividends in Ross (1977), bargaining (Section 12.5), and greenmail (Section 15.2). Banks (1990) has written a short book surveying signalling models in political science. Empirical papers include Layard & Psacharopoulos (1974) on education and Staten & Umbeck (1986) on occupational diseases.
- Legal bargaining is one area of application for signalling. See Grossman & Katz (1983). Reinganum (1988) has a nice example of the value of pre-commitment in legal signalling. In her model, a prosecutor who wishes to punish the guilty and release the innocent wishes, if parameters are such that most defendants are guilty, to commit to a pooling strategy in which his plea bargaining offer is the same whatever the probability that a particular defendant would be found guilty.
- The peacock’s tail may be a signal. Zahavi (1975) suggests that a large tail may benefit the peacock because, by hampering him, it demonstrates to potential mates that he is fit enough to survive even with a handicap.
- **Advertising.** Advertising is a natural application for signalling. The literature includes Nelson (1974), written before signalling was well known, Kihlstrom & Riordan (1984) and Milgrom & Roberts (1986). I will briefly describe a model based on Nelson’s. Firms are one of two types, low quality or high quality. Consumers do not know that a firm exists until they receive an advertisement from it, and they do not know its quality until they buy its product. They are unwilling to pay more than zero for low quality, but any product is costly to produce. This is not a reputation model, because it is finite in length and quality is exogenous.

If the cost of an advertisement is greater than the profit from one sale, but less than the profit from repeat sales, then high rates of advertising are associated with high product quality. A firm with low quality would not advertise, but a firm with high quality would.

The model can work even if consumers do not understand the market and do not make rational deductions from the firm’s incentives, so it does not have to be a signalling model. If consumers react passively and sample the product of any firm from whom they receive an advertisement, it is still true that the high quality firm advertises more, because the customers it attracts become repeat customers. If consumers do understand the firms’

incentives, signalling reinforces the result. Consumers know that firms which advertise must have high quality, so they are willing to try them. This understanding is important, because if consumers knew that 90 percent of firms were low quality but did not understand that only high quality firms advertise, they would not respond to the advertisements which they received. This should bring to mind Section 6.2's game of PhD Admissions.

- If there are just two workers in the population, the model is different depending on whether:
 - 1 Each is *High* ability with objective probability 0.5, so possibly both are *High* ability; or
 - 2 One of them is *High* and the other is *Low*, so only the subjective probability is 0.5.

The outcomes are different because in case (2) if one worker credibly signals he is *High* ability, the employer knows the other one must be *Low* ability.

Problems

11.1. Is Lower Ability Better?

Change Education I so that the two possible worker abilities are $a \in \{1, 4\}$.

- (a) What are the equilibria of this game? What are the payoffs of the workers (and the payoffs averaged across workers) in each equilibrium?
- (b) Apply the Intuitive Criterion (see N6.2). Are the equilibria the same?
- (c) What happens to the equilibrium worker payoffs if the high ability is 5 instead of 4?
- (d) Apply the Intuitive Criterion to the new game. Are the equilibria the same?
- (e) Could it be that a rise in the maximum ability reduces the average worker's payoff? Can it hurt all the workers?

11.2. Productive Education and Nonexistence of Equilibrium

Change Education I so that the two equally likely abilities are $a_L = 2$ and $a_H = 5$ and education is productive: the payoff of the employer whose contract is accepted is $\pi_{\text{employer}} = a + 2s - w$. The worker's utility function remains $U = w - \frac{8s}{a}$.

- (a) Under full information, what are the wages for educated and uneducated workers of each type, and who acquires education?
- (b) Show that with incomplete information the equilibrium is unique (except for beliefs and wages out of equilibrium) but unreasonable.

11.3. Price and Quality

Consumers have prior beliefs that Apex produces low-quality goods with probability 0.4 and high-quality with probability 0.6. A unit of output costs 1 to produce in either case, and it is worth 11 to the consumer if it is high quality and 0 if low quality. The consumer, who is risk neutral, decides whether to buy in each of two periods, but he does not know the quality until he buys. There is no discounting.

- (a) What is Apex's price and profit if it must choose one price, p^* , for both periods?
- (b) What is Apex's price and profit if it can choose two prices, p_1 and p_2 , for the two periods, but it cannot commit ahead to p_2 ?
- (c) What is the answer to part (b) if the discount rate is $r = 0.1$?
- (d) Returning to $r = 0$, what if Apex can commit to p_2 ?
- (e) How do the answers to (a) and (b) change if the probability of low quality is 0.95 instead of 0.4? (There is a twist to this question.)

11.4. Signalling with a Continuous Signal

Suppose that with equal probability a worker's ability is $a_L = 1$ or $a_H = 5$, and the worker chooses any amount of education $y \in [0, \infty)$. Let $U_{worker} = w - \frac{8y}{a}$ and $\pi_{employer} = a - w$.

- (a) There is a continuum of pooling equilibria, with different levels of y^* , the amount of education necessary to obtain the high wage. What education levels, y^* , and wages, $w(y)$, are paid in the pooling equilibria, and what is a set of out-of-equilibrium beliefs that supports them? What are the incentive compatibility constraints?
- (b) There is a continuum of separating equilibria, with different levels of y^* . What are the education levels and wages in the separating equilibria? Why are out-of-equilibrium beliefs needed, and what beliefs support the suggested equilibria? What are the self selection constraints for these equilibria?
- (c) If you were forced to predict one equilibrium which will be the one played out, which would it be?

11.5: Advertising.

Brydrex introduces a new shampoo which is actually very good, but is believed by consumers to be good with only a probability of 0.5. A consumer would pay 11 for high quality and 0 for low quality, and the shampoo costs 6 per unit to produce. The firm may spend as much as it likes on stupid TV commercials showing happy people washing their hair, but the potential market consists of 110 cold-blooded economists who are not taken in by psychological tricks. The market can be divided into two periods.

- (a) If advertising is banned, will Brydrex go out of business?
- (b) If there are two periods of consumer purchase, and consumers discover the quality of the shampoo if they purchase in the first period, show that Brydrex might spend substantial amounts on stupid commercials.
- (c) What is the minimum and maximum that Brydrex might spend on advertising, if it spends a positive amount?

11.6. Game Theory Books

In the Preface I explain why I listed competing game theory books by saying, "only an author quite confident that his book compares well with possible substitutes would do such a thing, and you will be even more certain that your decision to buy this book was a good one."

- (a) What is the effect of on the value of the signal if there is a possibility that I am an egotist who overvalues his own book?
- (b) Is there a possible non strategic reason why I would list competing game theory books?
- (c) If all readers were convinced by the signal of providing the list and so did not bother to even look at the substitute books, then the list would not be costly even to the author of a bad book, and the signal would fail. How is this paradox to be resolved? Give a verbal explanation.

- (d) Provide a formal model for part (c).

11.7. The Single-Crossing Property

If education is to be a good signal of ability,

- (a) Education must be inexpensive for all players.
- (b) Education must be more expensive for the high ability player.
- (c) Education must be more expensive for the low ability player.
- (d) Education must be equally expensive for all types of players.
- (e) Education must be costless for some small fraction of players.

11.8. A Continuum of Pooling Equilibria

Suppose that with equal probability a worker's ability is $a_L = 1$ or $a_H = 5$, and that the worker chooses any amount of education $y \in [0, \infty)$. Let $U_{worker} = w - \frac{8y}{a}$ and $\pi_{employer} = a - w$.

There is a continuum of pooling equilibria, with different levels of y^* , the amount of education necessary to obtain the high wage. What education levels, y^* , and wages, $w(y)$, are paid in the pooling equilibria, and what is a set of out-of-equilibrium beliefs that supports them? What are the self selection constraints?

11.9. Signal Jamming in Politics

A congressional committee has already made up its mind that tobacco should be outlawed, but it holds televised hearings anyway in which experts on both sides present testimony. Explain why these hearings might be a form of signalling, where the audience to be persuaded is congress as a whole, which has not yet made up its mind. You can disregard any effect the hearings might have on public opinion.

11.10. Salesman Clothing

Suppose a salesman's ability might be either $x = 1$ (with probability θ) or $x = 4$, and that if he dresses well, his output is greater, so that his total output is $x + 2s$ where s equals 1 if he dresses well and 0 if he dresses badly. The utility of the salesman is $U = w - \frac{8s}{x}$, where w is his wage. Employers compete for salesmen.

- (a) Under full information, what will the wage be for a salesman with low ability?
- (b) Show the self selection constraints that must be satisfied in a separating equilibrium under incomplete information.
- (c) Find all the equilibria for this game if information is incomplete.

11.11. Crazy Predators (adapted from Gintis [2000], Problem 12.10.)

Apex has a monopoly in the market for widgets, earning profits of m per period, but Brydoux has just entered the market. There are two periods and no discounting. Apex can either *Prey* on Brydoux with a low price or accept *Duopoly* with a high price, resulting in profits to Apex of $-p_a$ or d_a and to Brydoux of $-p_b$ or d_b . Brydoux must then decide whether to stay in the market for the second period, when Brydoux will make the same choices. If, however, Professor Apex, who owns 60 percent of the company's stock, is crazy, he thinks he will earn an amount $p^* > d_a$ from preying on Brydoux (and he does not learn from experience). Brydoux initially assesses the probability that Apex is crazy at θ .

- (a) Show that under the following condition, the equilibrium will be separating, i.e., Apex will behave differently in the first period depending on whether the Professor is crazy or not:

$$-p_a + m < 2d_a \quad (22)$$

- (b) Show that under the following condition, the equilibrium can be pooling, i.e., Apex will behave the same in the first period whether the Professor is crazy or not:

$$\theta \geq \frac{d_b}{p_b + d_b} \quad (23)$$

- (c) If neither two condition (22) nor (23) apply, the equilibrium is hybrid, i.e., Apex will use a mixed strategy and Brydox may or may not be able to tell whether the Professor is crazy at the end of the first period. Let α be the probability that a sane Apex preys on Brydox in the first period, and let β be the probability that Brydox stays in the market in the second period after observing that Apex chose *Prey* in the first period. Show that the equilibrium values of α and β are:

$$\alpha = \frac{\theta p_b}{(1 - \theta)d_b} \quad (24)$$

$$\beta = \frac{-p_a + m - 2d_a}{m - d_a} \quad (25)$$

- (d) Is this behavior related to any of the following phenomenon?– Signalling, Signal Jamming, Reputation, Efficiency Wages.

11.12. Actions and Strategies

Explain the difference between an “action” and a “strategy,” using a signal jamming game as an example.

11.13. Monopoly Quality

A consumer faces a monopoly. He initially believes that the probability that the monopoly has a high-quality product is H , and that a high-quality monopoly would be able to send him an advertisement at zero cost. With probability $(1-H)$, though, the monopoly has low quality, and it would cost the firm A to send an ad. The firm does send an ad, offering the product at price P . The consumer’s utility from a high-quality product is $X > P$, but from a low quality product it is 0. The production cost is C for the monopolist regardless of quality, where $C < P - A$. If the consumer does not buy the product, the seller does not incur the production cost.

You may assume that the high-quality firm always sends an ad, that the consumer will not buy unless he receives an ad, and that P is exogenous.

- Draw the extensive form for this game.
- What is the equilibrium if H is sufficiently high?
- If H is low enough, the equilibrium is in mixed strategies. The high-quality firm always advertises, the low quality firm advertises with probability M , and the consumer buys with probability N . Show using Bayes Rule how the consumer’s posterior belief R that the firm is high-quality changes once he receives an ad.
- Explain why the equilibrium is not in pure strategies if H is too low (but H is still positive).
- Find the equilibrium probability of M . (You don’t have to figure out N .)

PART III Applications

October 3, 1999. January 17, 2000. November 30, 2003. Eric Rasmusen, Erasmuse@indiana.edu.
[Http: //mypage.iu.edu/~erasmuse](http://mypage.iu.edu/~erasmuse). Footnotes starting with xxx are the author's notes to himself. Comments welcomed.

12 Bargaining

12.1 The Basic Bargaining Problem: Splitting a Pie

Part III of this book is designed to stretch your muscles by providing more applications of the techniques from Parts I and II. The next four chapters may be read in any order. They concern three ways that prices might be determined. Chapter 12 is about bargaining—where both sides exercise market power. Chapter 13 is about auctions—where the seller has market power, but sells a limited amount of a good and wants buyers to compete against each other. Chapter 14 is about fixed-price models with a variety of different features such as differentiated or durable goods. One thing all these chapters have in common is that they use new theory to answer old questions.

Bargaining theory attacks a kind of price determination ill described by standard economic theory. In markets with many participants on one side or the other, standard theory does a good job of explaining prices. In competitive markets we find the intersection of the supply and demand curves, while in markets monopolized on one side we find the monopoly or monopsony output. The problem is when there are few players on each side. Early in one's study of economics, one learns that under bilateral monopoly (one buyer and one seller), standard economic theory is inapplicable because the traders must bargain. In the chapters on asymmetric information we would have come across this repeatedly except for our assumption that either the principal or the agent faced competition.

Sections 12.1 and 12.2 introduce the archetypal bargaining problem, Splitting a Pie, ever more complicated versions of which make up the rest of the chapter. Section 12.2, where we take the original rules of the game and apply the Nash bargaining solution, is our one dip into cooperative game theory in this book. Section 12.3 looks at bargaining as a finitely repeated process of offers and counteroffers, and Section 12.4 views it as an infinitely repeated process. Section 12.5 returns to a finite number of repetitions (two, in fact), but with incomplete information. Finally, Section 12.6 approaches bargaining from a different level: how could people construct a mechanism for bargaining, a pre-arranged set of rules that would maximize their expected surplus.

Splitting a Pie

Players

Smith and Jones.

The Order of Play

The players choose shares θ_s and θ_j of the pie simultaneously.

Payoffs

If $\theta_s + \theta_j \leq 1$, each player gets the fraction he chose:
$$\begin{cases} \pi_s = & \theta_s. \\ \pi_j = & \theta_j. \end{cases}$$

If $\theta_s + \theta_j > 1$, then $\pi_s = \pi_j = 0$.

Splitting a Pie resembles the game of Chicken except that it has a continuum of Nash equilibria: any strategy combination (θ_s, θ_j) such that $\theta_s + \theta_j = 1$ is Nash. The Nash

concept is at its worst here, because the assumption that the equilibrium being played is common knowledge is very strong when there is a continuum of equilibria. The idea of the focal point (section 1.5) might help to choose a single Nash equilibrium. The strategy space of Chicken is discrete and it has no symmetric pure-strategy equilibrium, but the strategy space of Splitting a Pie is continuous, which permits a symmetric pure-strategy equilibrium to exist. That equilibrium is the even split, $(0.5, 0.5)$, which is a focal point.

If the players moved in sequence, the game would have a tremendous first-mover advantage. If Jones moved first, the unique Nash outcome would be $(0,1)$, although only weakly, because Smith would be indifferent as to his action. (This is the same open-set problem that was discussed in Section 4.3.) Smith accepts the offer if he chooses θ_s to make $\theta_s + \theta_j = 1$, but if we added only an epsilon-worth of ill will to the model he would pick $\theta_s > 0$ and reject the offer.

In many applications this version of Splitting a Pie is unacceptably simple, because if the two players find their fractions add to more than 1 they have a chance to change their minds. In labor negotiations, for example, if manager Jones makes an offer which union Smith rejects, they do not immediately forfeit the gains from combining capital and labor. They lose a week's production and make new offers. The recent trend in research has been to model such a sequence of offers, but before we do that let us see how cooperative game theory deals with the original game.¹


12.2 The Nash Bargaining Solution²

When game theory was young a favorite approach was to decide upon some characteristics an equilibrium should have based on notions of fairness or efficiency, mathematicize the characteristics, and maybe add a few other axioms to make the equilibrium turn out neatly. Nash (1950a) did this for the bargaining problem in what is perhaps the best-known application of cooperative game theory. Nash's objective was to pick axioms that would characterize the agreement the two players would anticipate making with each other. He used a game only a little more complicated than Splitting a Pie. In the Nash model, the two players can have different utilities if they do not come to an agreement, and the utility functions can be nonlinear in terms of shares of the pie. Figures 1a and 1b compare the two games.

Figure 1: (a) Nash Bargaining Game; (b) Splitting a Pie

¹xxx Relate this to the take-it-or-leave-it offer idea of earlier chapters, and note chaptr 7 on agency especailly.

²xxx This needs numerical examples.



In Figure 1, the shaded region denoted by X is the set of feasible payoffs, which we will assume to be convex. The disagreement point is $\bar{U} = (\bar{U}_s, \bar{U}_j)$. The Nash bargaining solution, $U^* = (U_s^*, U_j^*)$, is a function of $\bar{U}$ and X . The axioms that generate the concept are as follow

1 *Invariance*. For any strictly increasing linear function F ,

$$U^*[F(\bar{U}), F(X)] = F[U^*(\bar{U}, X)]. \quad (1)$$

This says that the solution is independent of the units in which utility is measured.

2 *Efficiency*. The solution is Pareto optimal, so that not both players can be made better off. In mathematical terms,

$$(U_s, U_j) > U^* \Rightarrow (U_s, U_j) \notin X. \quad (2)$$

3 *Independence of Irrelevant Alternatives*. If we drop some possible utility combinations from X , leaving the smaller set Y , then if U^* was not one of the dropped points, U^* does not change.

$$U^*(\bar{U}, X) \in Y \subseteq X \Rightarrow U^*(\bar{U}, Y) = U^*(\bar{U}, X). \quad (3)$$

4 *Anonymity (or Symmetry)*. Switching the labels on players Smith and Jones does not affect the solution.

The axiom of Independence of Irrelevant Alternatives is the most debated of the four, but if I were to complain, it would be about the axiomatic approach, which depends heavily on the intuition behind the axioms. Everyday intuition says that the outcome should be efficient and symmetric, so that other outcomes can be ruled out a priori. But most of the games in the earlier chapters of this book turn out to have reasonable but inefficient outcomes, and games like Chicken have reasonable asymmetric outcomes.

Whatever their drawbacks, these axioms fully characterize the Nash solution. It can be proven that if U^* satisfies the four axioms above, then it is the unique strategy combination

such that

$$U^* = \underset{U \in X, U \geq \bar{U}}{\operatorname{argmax}} (U_s - \bar{U}_s)(U_j - \bar{U}_j). \quad (4)$$

Splitting a Pie is a simple enough game that not all the axioms are needed to generate a solution. If we put the game in this context, however, problem (12.4) becomes

$$\begin{aligned} & \text{Maximize} && (\theta_s - 0)(\theta_j - 0), \\ & \theta_s, \theta_j \mid \theta_s + \theta_j \leq 1 \end{aligned} \quad (5)$$

which generates the first order conditions

$$\theta_s - \lambda = 0, \quad \text{and} \quad \theta_j - \lambda = 0, \quad (6)$$

where λ is the Lagrange multiplier on the constraint. From (12.6) and the constraint, we obtain $\theta_s = \theta_j = 1/2$, the even split that we found as a focal point of the noncooperative game.

Although Nash's objective was simply to characterize the anticipations of the players, I perceive a heavier note of morality in cooperative game theory than in noncooperative game theory. Cooperative outcomes are neat, fair, beautiful, and efficient. In the next few sections we will look at noncooperative bargaining models that, while plausible, lack every one of those features. Cooperative game theory may be useful for ethical decisions, but its attractive features are often inappropriate for economic situations, and the spirit of the axiomatic approach is very different from the utility maximization of economic theory.

It should be kept in mind, however, that the ethical component of cooperative game theory can also be realistic, because people are often ethical, or pretend to be. People very often follow the rules they believe represent virtuous behavior, even at some monetary cost. In bargaining experiments in which one player is given the ability to make a take-it-or-leave-it offer, it is very commonly found that he offers a 50-50 split. Presumably this is because either he wishes to be fair or he fears a spiteful response from the other player to a smaller offer. If the subjects are made to feel that they had "earned" the right to be the offering party, they behave much more like the players in noncooperative game theory (Hoffman & Spitzer [1985]). Frank (1988) and Thaler (1992) describe numerous occasions where simple games fail to describe real-world or experimental results. People's payoffs include more than their monetary rewards, and sometimes knowing the cultural disutility of actions is more important than knowing the dollar rewards. This is one reason why it is helpful to a modeller to keep his games simple: when he actually applies them to the real world, the model must not be so unwieldy that he cannot combine it with his knowledge of the particular setting.

12.3 Alternating Offers over Finite Time

In the games of the next two sections, the actions are the same as in Splitting a Pie, but with many periods of offers and counteroffers. This means that strategies are no longer just actions, but rather rules for choosing actions based on the actions chosen in earlier periods.

Alternating Offers

Players

Smith and Jones.

The Order of Play

1 Smith makes an offer θ_1 .

1* Jones accepts or rejects.

2 Jones makes an offer θ_2 .

2* Smith accepts or rejects.

...

T Smith offers θ_T .

T* Jones accepts or rejects.

Payoffs

The discount factor is $\delta \leq 1$.

If Smith's offer is accepted by Jones in round m ,

$$\begin{aligned}\pi_s &= \delta^m \theta_m, \\ \pi_j &= \delta^m (1 - \theta_m).\end{aligned}$$

If Jones's offer is accepted, reverse the subscripts.

If no offer is ever accepted, both payoffs equal zero.

When a game has many rounds we need to decide whether discounting is appropriate. If the discount rate is r then the discount factor is $\delta = 1/(1+r)$, so, without discounting, $r = 0$ and $\delta = 1$. Whether discounting is appropriate to the situation being modelled depends on whether delay should matter to the payoffs because the bargaining occurs over real time or the game might suddenly end (section 5.2). The game Alternating Offers can be interpreted in either of two ways, depending on whether it occurs over real time or not. If the players made all the offers and counteroffers between dawn and dusk of a single day, discounting would be inconsequential because, essentially, no time has passed. If each offer consumed a week of time, on the other hand, the delay before the pie was finally consumed would be important to the players and their payoffs should be discounted.

Consider first the game without discounting. There is a unique subgame perfect outcome — Smith gets the entire pie — which is supported by a number of different equilibria. In each equilibrium, Smith offers $\theta_s = 1$ in each period, but each equilibrium is different in terms of when Jones accepts the offer. All of them are weak equilibria because Jones is indifferent between accepting and rejecting, and they differ only in the timing of Jones's final acceptance.

Smith owes his success to his ability to make the last offer. When Smith claims the entire pie in the last period, Jones gains nothing by refusing to accept. What we have here is not really a first-mover advantage, but a last-mover advantage in offering, a difference not apparent in the one-period model.

In the game with discounting, the total value of the pie is 1 in the first period, δ in the second, and so forth. In period T , if it is reached, Smith would offer 0 to Jones, keeping 1 for himself, and Jones would accept under our assumption on indifferent players. In period

$T - 1$, Jones could offer Smith δ , keeping $(1 - \delta)$ for himself, and Smith would accept, although he could receive a greater share by refusing, because that greater share would arrive later and be discounted.

By the same token, in period $T - 2$, Smith would offer Jones $\delta(1 - \delta)$, keeping $1 - \delta(1 - \delta)$ for himself, and Jones would accept, since with a positive share Jones also prefers the game to end soon. In period $T - 3$, Jones would offer Smith $\delta[1 - \delta(1 - \delta)]$, keeping $1 - \delta[1 - \delta(1 - \delta)]$ for himself, and Smith would accept, again to prevent delay. Table 1 shows the progression of Smith's shares when $\delta = 0.9$.

Table 1: Alternating offers over finite time

Round	Smith's share	Jones's share	Total value	Who offers?
$T - 3$	0.819	0.181	0.9^{T-4}	Jones
$T - 2$	0.91	0.09	0.9^{T-3}	Smith
$T - 1$	0.9	0.1	0.9^{T-2}	Jones
T	1	0	0.9^{T-1}	Smith

As we work back from the end, Smith always does a little better when he makes the offer than when Jones does, but if we consider just the class of periods in which Smith makes the offer, Smith's share falls. If we were to continue to work back for a large number of periods, Smith's offer in a period in which he makes the offer would approach $\frac{1}{1+\delta}$, which equals about 0.53 if $\delta = 0.9$. The reasoning behind that precise expression is given in the next section. In equilibrium, the very first offer would be accepted, since it is chosen precisely so that the other player can do no better by waiting.

12.4 Alternating Offers over Infinite Time

The Folk theorem of Section 5.2 says that when discounting is low and a game is repeated an infinite number of times, there are many equilibrium outcomes. That does not apply to the bargaining game, however, because it is not a repeated game. It ends when one player accepts an offer, and only the accepted offer is relevant to the payoffs, not the earlier proposals. In particular, there are no out-of-equilibrium punishments such as enforce the Folk Theorem's outcomes.

Let players Smith and Jones have discount factors of δ_s and δ_j which are not necessarily equal but are strictly positive and no greater than one. In the unique subgame perfect outcome for the infinite-period bargaining game, Smith's share is

$$\theta_s = \frac{1 - \delta_j}{1 - \delta_s \delta_j}, \quad (7)$$

which, if $\delta_s = \delta_j = \delta$, is equivalent to

$$\theta_s = \frac{1}{1 + \delta}. \quad (8)$$

If the discount rate is high, Smith gets most of the pie: a 1,000 percent discount rate ($r = 10$) makes $\delta = 0.091$ and $\theta_s = 0.92$ (rounded), which makes sense, since under such extreme discounting the second period hardly matters and we are almost back to the simple game of Section 12.1. At the other extreme, if r is small, the pie is split almost evenly: if $r = 0.01$, then $\delta \approx 0.99$ and $\theta_s \approx 0.503$.

It is crucial that the discount rate be strictly greater than 0, even if only by a little. Otherwise, the game has the same continuum of perfect equilibria as in Section 12.1. Since nothing changes over time, there is no incentive to come to an early agreement. When discount rates are equal, the intuition behind the result is that since a player's cost of delay is proportional to his share of the pie, if Smith were to offer a grossly unequal split, such as (0.7, 0.3), Jones, with less to lose by delay, would reject the offer. Only if the split is close to even would Jones accept, as we will now prove.

Proposition 1 (Rubinstein [1982]) *In the discounted infinite game, the unique perfect equilibrium outcome is $\theta_s = \frac{1-\delta_j}{(1-\delta_s\delta_j)}$, where Smith is the first mover.*

Proof

We found that in the T -period game Smith gets a larger share in a period in which he makes the offer. Denote by M the maximum nondiscounted share, taken over all the perfect equilibria that might exist, that Smith can obtain in a period in which he makes the offer. Consider the game starting at t . Smith is sure to get no more than M , as noted in Table 2. (Jones would thus get $1 - M$, but that is not relevant to the proof.)

Table 2: Alternating offers over infinite time

Round	Smith's share	Jones's share	Who offers?
$T - 2$	$1 - \delta_j(1 - \delta_s M)$		Smith
$T - 1$		$1 - \delta_s M$	Jones
T	M		Smith

The trick is to find a way besides M to represent the maximum Smith can obtain. Consider the offer made by Jones at $t - 1$. Smith will accept any offer which gives him more than the discounted value of M received one period later, so Jones can make an offer of $\delta_s M$ to Smith, retaining $1 - \delta_s M$ for himself. At $t - 2$, Smith knows that Jones will turn down any offer less than the discounted value of the minimum Jones can look forward to receiving at $t - 1$. Smith, therefore, cannot offer any less than $\delta_j(1 - \delta_s M)$ at $t - 2$.

Now we have two expressions for “the maximum which Smith can receive,” which we can set equal to each other:

$$M = 1 - \delta_j (1 - \delta_s M). \quad (9)$$

Solving equation (9) for M , we obtain

$$M = \frac{1 - \delta_j}{1 - \delta_s \delta_j}. \quad (10)$$

We can repeat the argument using m , the minimum of Smith’s share. If Smith can expect at least m at t , Jones cannot receive more than $1 - \delta_s m$ at $t - 1$. At $t - 2$ Smith knows that if he offers Jones the discounted value of that amount, Jones will accept, so Smith can guarantee himself $1 - \delta_j (1 - \delta_s m)$, which is the same as the expression we found for M . The smallest perfect equilibrium share that Smith can receive is the same as the largest, so the equilibrium outcome must be unique.

No Discounting, but a Fixed Bargaining Cost

There are two ways to model bargaining costs per period: as proportional to the remaining value of the pie (the way used above), or as fixed costs each period, which is analyzed next (again following Rubinstein [1982]). To understand the difference, think of labor negotiations during a construction project. If a strike slows down completion, there are two kinds of losses. One is the loss from delay in renting or selling the new building, a loss proportional to its value. The other is the loss from late-completion penalties in the contract, which often take the form of a fixed penalty each week. The two kinds of costs have very different effects on the bargaining process.

To represent this second kind of cost, assume that there is no discounting, but whenever Smith or Jones makes an offer, he incurs the cost c_s or c_j . In every subgame perfect equilibrium, Smith makes an offer and Jones accepts, but there are three possible cases.

Delay costs are equal

$$c_s = c_j = c.$$

The Nash indeterminacy of Section 12.1 remains almost as bad; any fraction such that each player gets at least c is supported by some perfect equilibrium.

Delay hurts Jones more

$$c_s < c_j.$$

Smith gets the entire pie. Jones has more to lose than Smith by delaying, and delay does not change the situation except by diminishing the wealth of the players. The game is stationary, because it looks the same to both players no matter how many periods have already elapsed. If in any period t Jones offered Smith x , in period $(t - 1)$ Smith could offer Jones $(1 - x - c_j)$, keeping $(x + c_j)$ for himself. In period $(t - 2)$, Jones would offer Smith

$(x + c_j - c_s)$, keeping $(1 - x - c_j + c_s)$ for himself, and in periods $(t - 4)$ and $(t - 6)$ Jones would offer $(1 - x - 2c_j + 2c_s)$ and $(1 - x - 3c_j + 3c_s)$. As we work backwards, Smith's advantage rises to $\gamma(c_j - c_s)$ for an arbitrarily large integer γ . Looking ahead from the start of the game, Jones is willing to give up and accept zero.

Delay hurts Smith more

$$c_s > c_j.$$

Smith gets a share worth c_j and Jones gets $(1 - c_j)$. The cost c_j is a lower bound on the share of Smith, the first mover, because if Smith knows Jones will offer $(0,1)$ in the second period, Smith can offer $(c_j, 1 - c_j)$ in the first period and Jones will accept.

12.5 Incomplete Information

Instant agreement has characterized even the multiperiod games of complete information discussed so far. Under incomplete information, knowledge can change over the course of the game and bargaining can last more than one period in equilibrium, a result that might be called inefficient but is certainly realistic. Models with complete information have difficulty explaining such things as strikes or wars, but if over time an uninformed player can learn the type of the informed player by observing what offers are made or rejected, such unfortunate outcomes can arise. The literature on bargaining under incomplete information is vast. For this section, I have chosen to use a model based on the first part of Fudenberg & Tirole (1983), but it is only a particular example of how one could construct such a model, and not a good indicator of what results are to be expected from bargaining.

Let us start with a one-period game. We will denote the price by p_1 because we will carry the notation over to a two-period version.

One-Period Bargaining with Incomplete Information

Players

A seller, and a buyer called Buyer₁₀₀ or Buyer₁₅₀ depending on his type.

The Order of Play

0 Nature picks the buyer's type, his valuation of the object being sold, which is $b = 100$ with probability γ and $b = 150$ with probability $(1 - \gamma)$.

1 The seller offers price p_1 .

2 The buyer accepts or rejects p_1 .

Payoffs

The seller's payoff is p_1 if the buyer accepts the offer, and otherwise 0.

The buyer's payoff is $b - p_1$ if he accepts the offer, and otherwise 0.

Equilibrium:

Buyer₁₀₀: accept if $p_1 \leq 100$.

Buyer₁₅₀: accept if $p_1 \leq 150$.

Seller: offer $p_1 = 100$ if $\gamma \geq 1/3$ and $p_1 = 150$ otherwise.

Both types of buyers have a dominant strategy for the last move: accept any offer $p_1 < b$. Accepting any offer $p_1 \leq b$ is a weakly best response to the seller's equilibrium strategy. No equilibrium exists in which a buyer rejects an offer of $p_1 = b$, because we would fall into the open-set problem: there would be no greatest offer in $[0, b)$ that the buyer would accept, and so we could not find a best response for the seller.

The only two strategies that make sense for the seller are $p_1 = 100$ and $p_1 = 150$, since prices lower than 100 would lead to a sale with the same probability as $p_1 = 100$, prices in $(100, 150]$ would have the same probability as $p_1 = 150$, and prices greater than 150 would yield zero profits. The seller will choose $p_1 = 150$ if it yields a higher payoff than $p_1 = 100$; that is, if

$$\pi(p_1 = 100) = \gamma(100) + (1 - \gamma)(100) < \pi(p_1 = 150) = \gamma(0) + (1 - \gamma)(150), \quad (11)$$

which requires that

$$\gamma < 1/3. \quad (12)$$

Thus, if less than a third of buyers have a valuation of 100, the seller will charge 150, gambling that he is not facing such a buyer.

This means, of course, that if $\gamma < 1/3$, sometimes no sale will be made. This is the most interesting feature of the model. By introducing incomplete information into a bargaining model, we have explained why bargaining sometimes breaks down and efficient trades fail to be carried out. This suggests that when wars occur because nations cannot agree, or strikes occur because unions and employers cannot agree, we should look to information asymmetry for an explanation.

Note also that this has some similarity to a mechanism design problem. It is crucial that only one offer can be made. Once the offer $p_1 = 150$ is made and rejected, the seller realizes that $b = 100$. At that point, he would like to make a second offer, of $p_1 = 100$. But of course if he could do that, then rejection of the first offer would not convey the information that $b = 100$.

Now let us move to a two-period version of the same game. This will get quite a bit more complex, so let us restrict ourselves to the case of $\gamma = 1/6$. Also, we will need to make an assumption on discounting—the loss that results from a delay in agreement. Let us assume that each player loses a fixed amount $D = 4$ if there is no agreement in the first period. (Notice that this means a player can end up with a negative payoff by playing this game, something experienced bargainers will find realistic.)

Two-Period Bargaining with Incomplete Information

Players

A seller, and a buyer called Buyer₁₀₀ or Buyer₁₅₀ depending on his type.

The Order of Play

0 Nature picks the buyer's type, his valuation of the object being sold, which is $b = 100$ with probability $1/6$ and $b = 150$ with probability $5/6$.

1 The seller offers price p_1 .

2 The buyer accepts or rejects p_1 .

3 The seller offers price p_2 .

4 The buyer accepts or rejects p_2 .

Payoffs

The seller's payoff is p_1 if the buyer accepts the first offer, $(p_2 - 4)$ if he accepts the second offer, and -4 if he accepts no offer.

The buyer's payoff is $(b - p_1)$ if he accepts the first offer, $(b - p_2 - 4)$ if he accepts the second offer, and -4 if he accepts no offer.

Equilibrium Behavior (Separating, in mixed strategies)

Buyer₁₀₀: Accept if $p_1 \leq 104$. Accept if $p_2 \leq 100$.

Buyer₁₅₀: Accept with probability 0.6 if $p_1 = 150$. Accept if $p_2 \leq 150$.

Seller: Offer $p_1 = 150$. If $p_1 = 150$ is rejected, offer $p_2 = 100$ with probability $\phi = 0.08$ and $p_2 = 150$ with probability 0.92.

If Buyer₁₀₀ deviates and rejects an offer of p_1 less than 104, his payoff will be -4 , which is worse than $100 - p_1$. Rejecting p_2 does not result in any extra transactions cost, so he rejects any $p_2 < 100$.

Buyer₁₅₀'s equilibrium payoff is either

$$\pi_{Buyer\ 150}(Accept\ p_1 = 150) = b - 150 = 0, \quad (13)$$

or

$$\pi_{Buyer\ 150}(Reject\ p_1 = 150) = -4 + \phi(b - 100) + (1 - \phi)(b - 150) = -4 + \phi(b - 100), \quad (14)$$

which also equals zero if $\phi = 0.08$. Thus, Buyer₁₅₀ is indifferent and is willing to mix in the first period. In the second period, the game is just like the one-period game, so he will accept any offer of $p_2 \leq 150$.

To check on whether the seller has any incentive to deviate, let us work back from the end. If the game has reached the second period, he knows that the fraction of Buyer₁₀₀'s has increased, since there was some probability that a Buyer₁₅₀ would have accept $p_1 = 150$. The prior probability was $Prob(Buyer_{100}) = 1/6$, but the posterior is

$$\begin{aligned} Prob(Buyer_{100}|Rejected\ p_1 = 150) &= \frac{Prob(Rejected\ p_1=150|Buyer_{100})Prob(Buyer_{100})}{Prob(Rejected\ p_1=150)} \\ &= \frac{Prob(Rejected\ p_1=150|Buyer_{100})Prob(Buyer_{100})}{Prob(Rej|100)Prob(100)+Prob(Rej|150)Prob(150)} \\ &= \frac{(1)(1/6)}{(1)(1/6)+(0.4)(5/6)} \\ &= \frac{1}{3}. \end{aligned} \quad (15)$$

From the equilibrium of the one-period game, we know that if the probability of $b = 100$ is $1/3$, the seller is indifferent between $p_2 = 100$ and $p_2 = 150$. Thus, he is willing to mix, and in particular to choose $p_2 = 100$ with probability .08.

How about in the first period? Well, I cheated a bit there in describing the equilibrium. I did not say what Buyer₁₅₀ would do if the seller deviated to $p_1 \in (100, 150)$. What has to happen there is that if the seller deviates in that way, then some but not all of the Buyer₁₅₀'s accept the offer, and in the second period there is some probability of $p_2 = 100$ and some probability of $p_2 = 150$. The mixing probabilities, however, are not quite the same as the equilibrium mixing probabilities. Buyer₁₅₀ is now comparing

$$\pi_{Buyer\ 150}(Accept\ p_1) = 150 - p_1 \quad (16)$$

with

$$\pi_{Buyer\ 150}(Reject\ p_1) = -4 + \phi(p_1)(150 - 100) + (1 - \phi(p_1))(150 - 150) = -4 + 50\phi(p_1), \quad (17)$$

which are equal if

$$\phi(p_1) = \frac{154 - p_1}{50}. \quad (18)$$

Thus, if p_1 is close to 100, $\phi(p_1)$ is close to 1, and the seller must almost certainly charge $p_2 = 100$.

The seller is willing to mix in the second period so long as 0.6 of the Buyer₁₅₀'s accept p_1 , so that part of the strategy doesn't have to change. But if it doesn't, that means a deviation to $p_1 < 150$ won't change the seller's second-period payoff, or the probability that p_1 is accepted, so it will simply reduce his first-period revenues. Thus, neither player has incentive to deviate from the proposed equilibrium.^{3 4}

The most important lesson of this model is that bargaining can lead to inefficiency. Some of the Buyer₁₅₀s delay their transactions until the second period, which is inefficient since the payoffs are discounted. Moreover, there is at least some probability that the Buyer₁₀₀s never buy at all, as in the one-period game, and the potential gains from trade are lost.

Note, too, that this is a model in which prices fall over time as bargaining proceeds. The first period price is definitely $p_1 = 150$, but the second period price might fall to $p_1 = 100$. This can happen because high-valuation buyers know that though the price might fall if they wait, on the other hand it might not fall, and they will just incur an extra delay cost. This result has close parallels to the durable monopoly pricing problem that will be discussed in Chapter 14.

The price the buyer pays depends heavily on the seller's equilibrium beliefs. If the seller thinks that the buyer has a high valuation with probability 0.5, the price is 100, but

³xxx COMMITMENT EQUILIBRIUM . Suppose seller could commit to a second-period price. Would it affect anything? Yes. Mechanism design.

⁴xxx I think there are other equilibria with first period prices of more than 150. They will not be very interesting equilibria though.

if he thinks the probability is 0.05, the price rises to 150. This implies that a buyer is unfortunate if he is part of a group which is believed to have high valuations more often; even if his own valuation is low, what we might call his bargaining power is low when he is part of a high-valuing group. Ayres (1991) found that when he hired testers to pose as customers at car dealerships, their success depended on their race and gender even though they were given identical predetermined bargaining strategies to follow. Since the testers did as badly even when faced with salesmen of their own race and gender, it seems likely that they were hurt by being members of groups that usually can be induced to pay higher prices.

***12.6 Setting up a Way to Bargain: The Myerson-Satterthwaite Mechanism**

Let us now think about a different way to approach bargaining under incomplete information. This will not be a different methodology, for we will stay with noncooperative game theory, but now let us ask what would happen under different sets of formalized rules – different mechanisms.

Mechanisms were the topic of chapter 10, and bargaining is a good setting for considering them. We have seen in section 12.5 that under incomplete information it may easily happen that inefficiency arises in bargaining. This inefficiency surely varies depending on the rules of the game. Thus, if feasible, the players would like to bind themselves in advance to follow whichever rules are best at avoiding inefficiency.

Suppose a group of players in a game are interacting in some way. They would like to set up some rules for their interaction in advance, and this set of rules is what we call a mechanism. Usually, models analyze different mechanisms without asking how the players would agree upon them, taking that as exogenous to the model. This is reasonable – the mechanism may be assigned by history as an institution of the market. If it is not, then there is bargaining over which mechanism to use, a difficult extra layer of complexity.

Let us now consider the situation of two people trying to exchange a good under various mechanisms. The mechanism must do two things:

- 1 Tell under what circumstances the good should be transferred from seller to buyer; and
- 2 Tell the price at which the good should be transferred, if it is transferred at all.

Usually these two things are made to depend on **reports** of the two players – that is, on statements they make.

The first mechanisms we will look at are simple.

Bilateral Trading I: Complete Information

Players

A buyer and a seller.

The Order of Play

0 Nature independently chooses the seller to value the good at v_s and the buyer at v_b using the uniform distribution between 0 and 1. Both players observe these values.

1 The seller reports p_s .

2 The buyer chooses “Accept” or “Reject.”

3 The good is allocated to the buyer if the buyer accepts and to the seller otherwise. The price at which the trade takes place, if it does, is $p = p_s$.

Payoffs

If there is no trade, both players have payoffs of 0. If there is trade, the seller’s payoff is $p - v_s$ and the buyer’s is $v_b - p$.

I have normalized the payoffs so that each player’s payoff is zero if no trade occurs. I could instead have normalized to $\pi_s = v_s$ and to $\pi_b = 0$ if no trade occurred, a common alternative.

The unique subgame perfect Nash equilibrium of this game is for the seller to report $p_s = v_b$ and for the buyer to accept if $p_s \leq v_b$. Note that although it is Nash, it is not subgame perfect for the seller to charge $p_s = v_b$ and for the buyer to accept only if $p_s < v_b$, which would result in trade never occurring. The seller would deviate to offering a slightly higher price.

This is an efficient allocation mechanism, in the sense that the good ends up with the player who values it most highly. The problem is that the buyer would be unlikely to agree to this mechanism in the first place, before the game starts, because although it is efficient it always gives all of the social surplus to the seller.

Note, too, that this is not a truth-telling mechanism, in the sense that the seller does not reveal his own value in his report of p_s . An example of a truth-telling mechanism for this game would replace move (3) with

(3′) The good is allocated to the seller if the buyer accepts and to the buyer otherwise. The price at which the trade takes place, if it does, is $p = v_s + \frac{v_b - v_s}{2}$.

This new mechanism makes the game a bit silly. In it, the seller’s action is irrelevant, since p_s does not affect the transaction price. Instead, the buyer decides whether or not trade is efficient and accepts if it is, at a price which splits the surplus evenly. In one equilibrium, the buyer accepts if $v_b \geq v_s$ and rejects otherwise, and the seller reports $p_s = v_s$. (In other equilibria, the buyer might choose to accept only if $v_b > v_s$, and the seller chooses other reports for p_s , but those other equilibria are the same as the first one in all important respects.) The mechanism would be acceptable to both players in advance, however, it gives no incentive to anyone to lie, and it makes sense if the game really has complete information.

Let us next look at a game of incomplete information and a mechanism which does depend on the players' actions.

Bilateral Trading II: Incomplete Information

Players

A buyer and a seller.

The Order of Play

0 Nature independently chooses the seller to value the good at v_s and the buyer at v_b using the uniform distribution between 0 and 1. Each player's value is his own private information.

1 The seller reports p_s and the buyer reports p_b .

2 The buyer accepts or rejects the seller's offer. The price at which the trade takes place, if it does, is p_s .

Payoffs

If there is no trade, the seller's payoff is 0 and the buyer's is 0.

If there is trade, the seller's payoff is $p_s - v_s$ and the buyer's is $v_b - p_s$.

This mechanism does not use the buyer's report at all, and so perhaps it is not surprising that the result is inefficient. It is easy to see, working back from the end of the game, that the buyer's equilibrium strategy is to accept the offer if $v_b \geq p_s$ and to reject it otherwise. If the buyer does that, the seller's expected payoff is

$$[p_s - v_s] [Prob\{v_b \geq p_s\}] + 0 [Prob\{v_b \leq p_s\}] = [p_s - v_s] [1 - p_s]. \quad (19)$$

Differentiating this with respect to p_s and setting equal to zero yields the seller's equilibrium strategy of

$$p_s = \frac{1 + v_s}{2}. \quad (20)$$

This is not efficient because if v_b is just a little bigger than v_s , trade will not occur even though gains from trade do exist. In fact, trade will fail to occur whenever $v_b < \frac{1+v_s}{2}$.

Let us try another simple mechanism, which at least uses the reports of both players, replacing move (2) with (2')

(2') The good is allocated to the seller if $p_s > p_b$ and to the buyer otherwise. The price at which the trade takes place, if it does, is p_s .

Suppose the buyer truthfully reports $p_b = v_b$. What will the seller's best response be? The seller's expected payoff for the p_s he chooses is now

$$[p_s - v_s] [Prob\{p_b(v_b) \geq p_s\}] + 0 [Prob\{p_b(v_b) \leq p_s\}] = [p_s - v_s] [1 - p_s]. \quad (21)$$

where the expectation has to be taken over all the possible values of v_b , since p_b will vary with v_b .


Maximizing this, the seller's strategy will solve the first-order condition $1 - 2p_s + v_s = 0$, and so will again be

$$p_s(v_s) = \frac{1 + v_s}{2} = \frac{1}{2} + \frac{v_s}{2}. \quad (22)$$

Will the buyer's best response to this strategy be $p_b = v_b$? Yes, because whenever $v_b \geq \frac{1}{2} + \frac{v_s}{2}$ the buyer is willing for trade to occur, and the size of p_b does not affect the transactions price, only the occurrence or nonoccurrence of trade. The buyer needs to worry about causing trade to occur when $v_b < \frac{1}{2} + \frac{v_s}{2}$, but this can be avoided by using the truthtelling strategy. The buyer also needs to worry about preventing trade from occurring when $v_b > \frac{1}{2} + \frac{v_s}{2}$, but choosing $p_b = v_b$ prevents this from happening either.

Thus, it seems that either mechanism (2) or (2') will fail to be efficient. Often, the seller will value the good less than the buyer, but trade will fail to occur and the seller will end up with the good anyway – whenever $v_b > \frac{1+v_s}{2}$. Figure 2 shows when trades will be completed based on the parameter values.

Figure 2: Trades in Bilateral Trading II



As you might imagine, one reason this is an inefficient mechanism is that it fails to make effective use of the buyer's information. The next mechanism will do better. Its trading rule is called the **double auction mechanism**. The problem is like that of the Groves Mechanism because we are trying to come up with an action rule (allocate the object to the buyer or to the seller) based on the agents' reports (the prices they suggest), under the condition that each player has private information (his value).

Bilateral Trading III: The Double Auction Mechanism

Players

A buyer and a seller.

The Order of Play

0 Nature independently chooses the seller to value the good at v_s and the buyer at v_b using the uniform distribution between 0 and 1. Each player's value is his own private information.

1 The buyer and the seller simultaneously decide whether to try to trade or not.

2 If both agree to try, the seller reports p_s and the buyer reports p_b simultaneously.

3 The good is allocated to the seller if $p_s \geq p_b$ and to the buyer otherwise. The price at which the trade takes place, if it does, is $p = \frac{(p_b + p_s)}{2}$.

Payoffs

If there is no trade, the seller's payoff is 0 and the buyer's is zero. If there is trade, then the seller's payoff is $p - v_s$ and the buyer's is $v_b - p$.

The buyer's expected payoff for the p_b he chooses is

$$\left[v_b - \frac{p_b + E[p_s | p_b \geq p_s]}{2} \right] [Prob\{p_b \geq p_s\}], \quad (23)$$

where the expectation has to be taken over all the possible values of v_s , since p_s will vary with v_s .

The seller's expected payoff for the p_s he chooses is

$$\left[\frac{p_s + E(p_b | p_b \geq p_s)}{2} - v_s \right] [Prob\{p_b \geq p_s\}], \quad (24)$$

where the expectation has to be taken over all the possible values of v_b , since p_b will vary with v_b .

The game has lots of Nash equilibria. Let's focus on two of them, a **one-price equilibrium** and the unique **linear equilibrium**.

In the **one-price equilibrium**, the buyer's strategy is to offer $p_b = x$ if $v_b \geq x$ and $p_b = 0$ otherwise, for some value $x \in [0, 1]$. The seller's strategy is to ask $p_s = x$ if $v_s \leq x$ and $p_s = 1$ otherwise. Figure 3 illustrates the one-price equilibrium for a particular value of x . Suppose $x = 0.7$. If the seller were to deviate and ask prices lower than 0.7, he would just reduce the price he receives. If the seller were to deviate and ask prices higher than 0.7, then $p_s > p_b$ and no trade occurs. So the seller will not deviate. Similar reasoning applies to the buyer, and to any value of x , including 0 and 1, where trade never occurs.




Figure 3: Trade in the one-price equilibrium

The **linear equilibrium** can be derived very neatly. Suppose the seller uses a linear strategy, so $p_s(v_s) = \alpha_s + c_s v_s$. Then from the buyer's point of view, p_s will be uniformly distributed from α_s to $\alpha_s + c_s$ with density $1/c_s$, as v_s ranges from 0 to 1. Since $E_b[p_s | p_b \geq p_s] = E_b(p_s | p_s \in [a_s, p_b]) = \frac{a_s + p_b}{2}$, the buyer's expected payoff (23) becomes

$$\left[v_b - \frac{p_b + \frac{\alpha_s + p_b}{2}}{2} \right] \left[\frac{p_b - \alpha_s}{c_s} \right]. \quad (25)$$

Maximizing with respect to p_b yields

$$p_b = \frac{2}{3}v_b + \frac{1}{3}\alpha_s. \quad (26)$$

Thus, if the seller uses a linear strategy, the buyer's best response is a linear strategy too! We are well on our way to a Nash equilibrium.

If the buyer uses a linear strategy $p_b(v_b) = \alpha_b + c_b v_b$, then from the seller's point of view p_b is uniformly distributed from α_b to $\alpha_b + c_b$ with density $1/c_b$ and the seller's payoff function, expression (24), becomes, since $E_s(p_b | p_b \geq p_s) = E_s(p_b | p_b \in [p_s, \alpha_b + c_b]) = \frac{p_s + \alpha_b + c_b}{2}$,

$$\left[\frac{p_s + \frac{p_s + \alpha_b + c_b}{2}}{2} - v_s \right] \left[\frac{\alpha_b + c_b - p_s}{c_b} \right]. \quad (27)$$

Maximizing with respect to p_s yields

$$p_s = \frac{2}{3}v_s + \frac{1}{3}(\alpha_b + c_b). \quad (28)$$

Solving equations (26) and (28) together yields

$$p_b = \frac{2}{3}v_b + \frac{1}{12} \quad (29)$$

and

$$p_s = \frac{2}{3}v_s + \frac{1}{4}. \quad (30)$$

So we have derived a linear equilibrium. Manipulation of the equilibrium strategies shows that trade occurs if and only if $v_b \geq v_s + (1/4)$, which is to say, trade occurs if the valuations differ enough. The linear equilibrium does not make all efficient trades, because sometimes $v_b > v_s$ and no trade occurs, but it does make all trades with joint surpluses of $1/4$ or more. Figure 4 illustrates this.




Figure 4: Trade in the linear equilibrium

One detail about equation (29) should bother you. The equation seems to say that if $v_b = 0$, the buyer chooses $p_b = 1/12$. If that happens, though, the buyer is bidding more than his value! The reason this can be part of the equilibrium is that it is only a weak Nash equilibrium. Since the seller never chooses lower than $p_s = 1/4$, the buyer is safe in choosing $p_b = 1/12$; trade never occurs anyway when he makes that choice. He could just as well bid 0 instead of $1/12$, but then he wouldn't have a linear strategy.

The linear equilibrium is not a truth-telling equilibrium. The seller does not report his true value v_s , but rather reports $p_s = (2/3)v_s + 1/4$. But we could replicate the outcome in a truth-telling equilibrium. We could have the buyer and seller agree that they would make reports r_s and r_b to a neutral mediator, who would then choose the trading price p . He would agree in advance to choose the trading price p by (a) mapping r_s onto p_s just as in the equilibrium above, (b) mapping r_b onto p_b just as in the equilibrium above, and (c) using p_b and p_s to set the price just as in the double auction mechanism. Under this mechanism, both players would tell the truth to the mediator. Let us compare the original linear mechanism with a truth-telling mechanism.

The Chatterjee-Samuelson mechanism. *The good is allocated to the seller if $p_s \geq p_b$ and to the buyer otherwise. The price at which the trade takes place, if it does, is*

$$p = \frac{(p_b + p_s)}{2}$$

A direct incentive-compatible mechanism. *The good is allocated to the seller if $\frac{2}{3}p_s + \frac{1}{4} \geq \frac{2}{3}p_b + \frac{1}{12}$, which is to say, if $p_s \geq p_b - 1/4$, and to the buyer otherwise. The price at which the trade takes place, if it does, is*

$$p = \frac{(\frac{2}{3}p_b + \frac{1}{12}) + (\frac{2}{3}p_s + \frac{1}{4})}{2} = \frac{p_b + p_s}{3} + \frac{1}{6} \quad (31)$$

What I have done is substituted the equilibrium strategies of the two players into the mechanism itself, so now they will have no incentive to set their reports different from the truth. The mechanism itself looks odd, because it says that trade cannot occur unless v_b is more than $1/4$ greater than v_s , but we cannot use the rule of trading if $v_b > v_s$ because then the players would start misreporting again. The truth-telling mechanism only works because it does not penalize players for telling the truth, and in order not to penalize them, it cannot make full use of the information to achieve efficiency.

In this game we have imposed a trading rule on the buyer and seller, rather than letting them decide for themselves what is the best trading rule. Myerson & Satterthwaite (1983) prove that of all the equilibria and all the mechanisms that are budget balancing, the linear equilibrium of the double auction mechanism yields the highest expected payoff to the players, the expectation being taken ex ante, before Nature has chosen the types. The mechanism is not optimal when viewed after the players have been assigned their types, and a player might not be happy with the mechanism once he knew his type. He will, however, at least be willing to participate.

What mechanism would players choose, ex ante, if they knew they would be in this game? If they had to choose after they were informed of their type, then their proposals for mechanisms could reveal information about their types, and we would have a model of bargaining under incomplete information that would resemble signalling models. But what if they chose a mechanism before they were informed of their type, and did not have the option to refuse to trade if after learning their type they did not want to use the mechanism?

In general, mechanisms have the following parts.

- 1 Each agent i simultaneously makes a report p_i .
- 2 A rule $x(p)$ determines the action (such as who gets the good, whether a bridge is built, etc.) based on the p .
- 3 Each agent i receives an incentive transfer a_i that in some way depends on his own report.
- 4 Each agent receives a budget-balancing transfer b_i that does not depend on his own report.

We will denote the agent's total transfer by t_i , so $t_i = a_i + b_i$.

In Bilateral Trading III, the mechanism had the following parts.

- 1 Each agent i simultaneously made a report p_i .
- 2 If $p_s \geq p_b$, the good was allocated to the seller, but otherwise to the buyer.
- 3 If there was no trade, then $a_s = a_b = 0$. If there was trade, then $a_s = \frac{(p_b + p_s)}{2}$ and $a_b = -(\frac{(p_b + p_s)}{2})$.
- 4 No further transfer b_i was needed, because the incentive transfers balanced the budget by themselves.

It turns out that if the players in Bilateral Trading can settle their mechanism and agree to try to trade in advance of learning their types, an efficient budget-balancing mechanism exists that can be implemented as a Nash equilibrium. The catch will be that after discovering his type, a player will sometimes regret having entered into this mechanism.

This would actually be part of a subgame perfect Nash equilibrium of the game as a whole. The mechanism design literature tends not to look at the entire game, and asks “Is there a mechanism which is efficient when played out as the rules of a game?” rather than “Would the players choose a mechanism that is efficient?”

Bilateral Trading IV: The Expected Externality Mechanism

Players

A buyer and a seller.

The Order of Play

-1 Buyer and seller agree on a mechanism $(x(p), t(p))$ that makes decisions x based on reports p and pays t to the agents, where p and t are 2-vectors and x allocated the good either to the buyer or the seller.

0 Nature independently chooses the seller to value the good at v_s and the buyer at v_b using the uniform distribution between 0 and 1. Each player’s value is his own private information.

1 The seller reports p_s and the buyer reports p_b simultaneously.

2 The mechanism uses $x(p)$ to decide who gets the good, and $t(p)$ to make payments.

Payoffs

Player i ’s payoff is $v_i + t_i$ if he is allocated the good, t_i otherwise.

I was vague on how the two parties agree on a mechanism. The mechanism design literature is also very vague, and focuses on efficiency rather than payoff-maximization. To be more rigorous, we should have one player propose the mechanism and the other accept or reject. The proposing player would add an extra transfer to the mechanism to reduce the other player’s expected payoff to his reservation utility.

Let me use the term **action surplus** to denote the utility an agent gets from the choice of action.

The **expected externality mechanism** has the following objectives for each of the parts of the mechanism.

1 Induce the agents to make truthful reports.

2 Choose the efficient action.

3 Choose the incentive transfers to make the agents choose truthful reports in equilibrium.

4 Choose the budget-balancing transfers so that the incentive transfers add up to zero.

First I will show you a mechanism that does this. Then I will show you how I came up with that mechanism. Consider the following three- part mechanism:

1 The seller announces p_s . The buyer announces p_b . The good is allocated to the seller if $p_s \geq p_b$, and to the buyer otherwise.

2 The seller gets transfer $t_s = \frac{(1-p_s^2)}{2} - \frac{(1-p_b^2)}{2}$.

3 The buyer gets transfer $t_b = \frac{(1-p_b^2)}{2} - \frac{(1-p_s^2)}{2}$.

Note that this is budget-balancing:

$$\frac{(1-p_s^2)}{2} - \frac{(1-p_b^2)}{2} + \frac{(1-p_b^2)}{2} - \frac{(1-p_s^2)}{2} = 0. \quad (32)$$

The seller's expected payoff as a function of his report p_s is the sum of his expected action surplus and his expected transfer. We have already computed his transfer, which is not conditional on the action taken.

The seller's action surplus is 0 if the good is allocated to the buyer, which happens if $v_b > p_s$, where we use v_b instead of p_b because in equilibrium $p_b = v_b$. This has probability $1 - p_s$. The seller's action surplus is v_s if the good is allocated to the seller, which has probability p_s . Thus, the expected action surplus is $p_s v_s$.

The seller's expected payoff is therefore

$$p_s v_s + \frac{(1-p_s^2)}{2} - \frac{(1-p_b^2)}{2}. \quad (33)$$

Maximizing with respect to his report, p_s , the first order condition is

$$v_s - p_s = 0, \quad (34)$$

so the mechanism is incentive compatible – the seller tells the truth.

The buyer's expected action surplus is v_b if his report is higher, e.g. if $p_b > v_s$, and zero otherwise, so his expected payoff is

$$p_b v_b + \frac{(1-p_b^2)}{2} - \frac{(1-p_s^2)}{2} \quad (35)$$

Maximizing with respect to his report, p_b , the first order condition is

$$v_b - p_b = 0, \quad (36)$$

so the mechanism is incentive compatible – the buyer tells the truth.

Now let's see how to come up with the transfers. The expected externality mechanism relies on two ideas.

The first idea is that to get the incentives right, each agent's incentive transfer is made equal to the sum of the expected action surpluses of the other agents, where the expectation is calculated conditionally on (a) the other agents reporting truthfully, and (b) our agent's report. This makes the agent internalize the effect of his externalities on

the other agents. His expected payoff comes to equal the expected social surplus. Here, this means, for example, that the seller's incentive transfer will equal the buyer's expected action surplus. Thus, denoting the uniform distribution by F ,

$$\begin{aligned} a_s &= \int_0^{p_s} 0 dF(v_b) + \int_{p_s}^1 v_b dF(v_b) \\ &= 0 + \left. \frac{v_b^2}{2} \right|_{p_s}^1 \\ &= \frac{1}{2} - \frac{p_s^2}{2}. \end{aligned} \tag{37}$$

The first integral is the expected buyer action surplus if no transfer is made because the buyer's value v_b is less than the seller's report p_s , so the seller keeps the good and the buyer's action surplus is zero. The second integral is the surplus if the buyer gets the good, which occurs whenever the buyer's value, v_b (and hence his report p_b), is greater than the seller's report, p_s .

We can do the same thing for the buyer's incentive, finding the seller's expected surplus.

$$\begin{aligned} a_b &= \int_0^{p_b} 0 dF(v_s) + \int_{p_b}^1 v_s dF(v_s) \\ &= 0 + \left. \frac{v_s^2}{2} \right|_{p_b}^1 \\ &= \frac{1}{2} - \frac{p_b^2}{2}. \end{aligned} \tag{38}$$

If the seller's value v_s is low, then it is likely that the buyer's report of p_b is higher than v_s , and the seller's action surplus is zero because the trade will take place. If the seller's value v_s is high, then the seller will probably have a positive action surplus.

The second idea is that to get budget balancing, each agent's budget-balancing transfer is chosen to help pay for the other agents' incentive transfers. Here, we just have two agents, so the seller's budget-balancing transfer has to pay for the buyer's incentive transfer. That is very simple: just set the seller's budget-balancing transfer b_s equal to the buyer's incentive transfer a_b (and likewise set b_b equal to a_s).

The intuition and mechanism can be extended to N agents. There are now N reports $p_1, \dots, p_N$. Let the action chosen be $x(p)$, where p is the N -vector of reports, and the action surplus of agent i is $W_i(x(p), v_i)$. To make each agent's incentive transfer equal to the sum of the expected action surpluses of the other agents, choose it so

$$a_i = E(\sum_{j \neq i} W_j(x(p), v_j)). \tag{39}$$

The budget balancing transfers can be chosen so that each agent's incentive transfer is paid for by dividing the cost equally among the other $N - 1$ agents:

$$b_i = \frac{1}{N-1} (\sum_{j \neq i} E(\sum_{k \neq j} W_k(x(p), v_k))). \tag{40}$$

There are other ways to divide the costs that will still allow the mechanism to be incentive compatible, but equal division is easiest to think about.

The expected externality mechanism does have one problem: the participation constraint. If the seller knows that $v_s = 1$, he will not want to enter into this mechanism. His expected transfer would be $t_s = 0 - (1 - 0.5)^2/2 = -0.125$. Thus, his payoff from the

mechanism is $1 - 0.125 = 0.875$, whereas he could get a payoff of 1 if he refused to participate. We say that this mechanism fails to be **interim incentive compatible**, because at the point when the agents discover their own types, but not those of the other agents, the agents might not want to participate in the mechanism or choose the actions we desire.

N12.1 The Nash bargaining solution

- See Binmore, Rubinstein, & Wolinsky (1986) for a comparison of the cooperative and noncooperative approaches to bargaining. For overviews of cooperative game theory see Luce & Raiffa (1957) and Shubik (1982).
- While the Nash bargaining solution can be generalized to n players (see Harsanyi [1977], p. 196), the possibility of interaction between coalitions of players introduces new complexities. Solutions such as the Shapley value (Shapley [1953b]) try to account for these complexities.

The **Shapley value** satisfies the properties of invariance, anonymity, efficiency, and linearity in the variables from which it is calculated. Let S_i denote a **coalition** containing player i ; that is, a group of players including i that makes a sharing agreement. Let $v(S_i)$ denote the sum of the utilities of the players in coalition S_i , and $v(S_i - \{i\})$ denote the sum of the utilities in the coalition created by removing i from S_i . Finally, let $c(s)$ be the number of coalitions of size s containing player i . The Shapley value for player i is then

$$\phi_i = \frac{1}{n} \sum_{s=1}^n \frac{1}{c(s)} \sum_{S_i} [v(S_i) - v(S_i - \{i\})]. \quad (41)$$

where the S_i are of size s . The motivation for the Shapley value is that player i receives the average of his marginal contributions to different coalitions that might form. Gul (1989) has provided a noncooperative interpretation.

N12.2 Alternating offers over infinite time

- The proof of Proposition 12.1 is not from the original Rubinstein (1982), but is adapted from Shaked & Sutton (1984). The maximum rather than the supremum can be used because of the assumption that indifferent players always accept offers.
- In extending alternating offers to three players, there is no obviously best way of specifying how players make and accept offers. Haller (1986) shows that for at least one specification, the outcome is not similar to the Rubinstein (1982) outcome, but rather is a return to the indeterminacy of the game without discounting.

N12.3 Incomplete information.

- Bargaining under asymmetric information has inspired a large literature. In early articles, Fudenberg & Tirole (1983) uses a two-period model with two types of buyers and two types of sellers. Sobel & Takahashi (1983) builds a model with either T or infinite periods, a continuum of types of buyers, and one type of seller. Crampton (1984) uses an infinite number of periods, a continuum of types of buyers, and a continuum of types of sellers. Rubinstein (1985a) uses an infinite number of periods, two types of buyers, and one type of seller, but the types of buyers differ not in their valuations, but in their discount rates. Rubinstein (1985b) puts emphasis on the choice of out-of-equilibrium conjectures. Samuelson (1984) looks at the case where one bargainer knows the size of the pie better than the other bargainer. Perry (1986) uses a model with fixed bargaining costs and asymmetric

information in which each bargainer makes an offer in turn, rather than one offering and the other accepting or rejecting. For overviews, see excellent surveys of Sutton (1986) and Kennan & Wilson (1993).

- The asymmetric information model in Section 12.5 has **one-sided** asymmetry in the information: only the buyer's type is private information. Fudenberg & Tirole (1983) and others have also built models with **two-sided** asymmetry, in which buyers' and sellers' types are both private information. In such models a multiplicity of perfect Bayesian equilibria can be supported for a given set of parameter values. Out-of-equilibrium beliefs become quite important, and provided much of the motivation for the exotic refinements mentioned in section 6.2.
- There is no separating equilibrium if, instead of discounting, the asymmetric information model has fixed-size per-period bargaining costs, unless the bargaining cost is higher for the high-valuation buyer than for the low-valuation. If, for example, there is no discounting, but a cost of c is incurred each period that bargaining continues, no separating equilibrium is possible. That is the typical signalling result. In a separating equilibrium the buyer tries to signal a low valuation by holding out, which fails unless it really is less costly for a low-valuation buyer to hold out. See Perry (1986) for a model with fixed bargaining costs which ends after one round of bargaining.

N12.4 Setting up a way to bargain: the Myerson- Satterthwaite mechanism

- The Bilateral Trading model originated in Chatterjee & Samuelson (1983, p. 842), who also analyze the more general mechanism with $p = \theta p_s + (1 - \theta)p_b$. I have adapted this description from Gibbons (1992, p.158).
- Discussions of the general case can be found in Fudenberg & Tirole (1991a, p. 273), and Mas-Colell, Whinston & Green (1994, p. 885). It is also possible to add extra costs that depend on the action chosen (for example, a transactions tax if the good is sold from buyer to seller). See Fudenberg and Tirole, p. 274. I have taken the term "expected externality mechanism" from MWG. Fudenberg and Tirole use "AGV mechanism" for the same thing, because the idea was first published in Arrow (1979) and D'Aspremont & Varet (1979). Myerson (1991) is also worth looking into.

Problems

12.1. A Fixed Cost of Bargaining and Incomplete Information

Smith and Jones are trying to split 100 dollars. In bargaining round 1, Smith makes an offer at cost c , proposing to keep S_1 for himself. Jones either accepts (ending the game) or rejects. In round 2, Jones makes an offer of S_2 for Smith, at cost 10, and Smith either accepts or rejects. In round 3, Smith makes an offer of S_3 at cost c , and Jones either accepts or rejects. If no offer is ever accepted, the 100 dollars goes to a third player, Parker.

- (a) If $c = 0$, what is the equilibrium outcome?
- (b) If $c = 80$, what is the equilibrium outcome?
- (c) If Jones' priors are that $c = 0$ and $c = 80$ are equally likely, but only Smith knows the true value, what are the players' equilibrium strategies in rounds 2 and 3? (that is: what are S_2 and S_3 , and what acceptance rules will each player use?)
- (d) If Jones' priors are that $c = 0$ and $c = 80$ are equally likely, but only Smith knows the true value, what are the equilibrium strategies for round 1? (Hint: the equilibrium uses mixed strategies.)

12.2: Selling Cars

A car dealer must pay \$10,000 to the manufacturer for each car he adds to his inventory. He faces three buyers. From the point of view of the dealer, Smith's valuation is uniformly distributed between \$12,000 and \$21,000, Jones's is between \$9,000 and \$12,000, and Brown's is between \$4,000 and \$12,000. The dealer's policy is to make a single take-it-or-leave-it offer to each customer, and he is smart enough to avoid making different offers to customers who could resell to each other. Use the notation that the maximum valuation is $\bar{V}$ and the range of valuations is R .

- (a)] What will the offers be?
- (b) Who is most likely to buy a car? How does this compare with the outcome with perfect price discrimination under full information? How does it compare with the outcome when the dealer charges \$10,000 to each customer?
- (c) What happens to the equilibrium prices if, with probability 0.25, each buyer has a valuation of \$0, but the probability distribution remains otherwise the same?

12.3. The Nash Bargaining Solution

Smith and Jones, shipwrecked on a desert island, are trying to split 100 pounds of cornmeal and 100 pints of molasses, their only supplies. Smith's utility function is $U_s = C + 0.5M$ and Jones's is $U_j = 3.5C + 3.5M$. If they cannot agree, they fight to the death, with $U = 0$ for the loser. Jones wins with probability 0.8.

- (a) What is the threat point?
- (b) With a 50-50 split of the supplies, what are the utilities if the two players do not recontract? Is this efficient?

- (c) Draw the threat point and the Pareto frontier in utility space (put U_s on the horizontal axis).
- (d) According to the Nash bargaining solution, what are the utilities? How are the goods split?
- (e) Suppose Smith discovers a cookbook full of recipes for a variety of molasses candies and corn muffins, and his utility function becomes $U_s = 10C + 5M$. Show that the split of goods in part (d) remains the same despite his improved utility function.

12.4.: Price Discrimination and Bargaining

A seller with marginal cost constant at c faces a continuum of consumers represented by the linear demand curve $Q^d = a - bP$, where $a > c$. Demand is at a rate of one or zero units per consumer, so if all consumers between points 1 and 2.5 on the consumer continuum make purchases at a price of 13, we say that a total of 1.5 units are sold at a price of 13 each.

- (a) What is the seller's profit if he chooses one take-it-or-leave-it price?

Answer. This is the simple monopoly pricing problem. Profit is

$$\pi = Q(P - C) = Q(a/b - Q/b - c).$$

Differentiating with respect to Q yields

$$\frac{d\pi}{dQ} = a/b - 2Q/b - c = 0,$$

which can be solved to give us

$$Q_m = \frac{a - bc}{2}.$$

The price is then, using the demand curve,

$$P_m = \frac{a/b + c}{2},$$

which is to say that the price will be halfway between marginal cost and price which drives demand to zero. Profit is

$$\pi_m = \left(\frac{a/b - c}{2} \right) \left(\frac{a - cb}{2} \right).$$

- (b) What is the seller's profit if he chooses a continuum of take-it-or-leave-it prices at which to sell, one price for each consumer? (You should think here of a pricing function, since each consumer is infinitesimal).

Answer. Under perfect price discrimination, the seller captures the entire area under the demand curve and over the marginal cost curve, because he charges each consumer exactly the reservation price. Since the price at which quantity demanded falls to zero is a/b and the quantity when price equals marginal cost is $a - bc$, the area of this profit triangle is

$$\pi_{ppd} = (1/2)(a/b - c)(a - bc)$$

Note that this is exactly twice the monopoly profit found earlier.

- (c) What is the seller's profit if he bargains separately with each consumer, resulting in a continuum of prices? You may assume that bargaining costs are zero and that buyer and seller have equal bargaining power.

Answer. In this case, which I call "isoperfect price discrimination," profits are exactly half of what they are under perfect price discrimination, since the price charged to a consumer will exactly split the surplus he would have if the price equalled marginal cost. Thus, the profit is the same as using the simple monopoly price.

12.5. A Fixed Cost of Bargaining and Incomplete Information

Smith and Jones are trying to split 100 dollars. In bargaining round 1, Smith makes an offer at cost c , proposing to keep S_1 for himself. Jones either accepts (ending the game) or rejects. In round 2, Jones makes an offer of S_2 for Smith, at cost 10, and Smith either accepts or rejects. In round 3, Smith makes an offer of S_3 at cost c , and Jones either accepts or rejects. If no offer is ever accepted, the 100 dollars goes to a third player, Parker.

- (a) If $c = 0$, what is the equilibrium outcome?
- (b) If $c = 80$, what is the equilibrium outcome?
- (c) If Jones' priors are that $c = 0$ and $c = 80$ are equally likely, but only Smith knows the true value, what is the equilibrium outcome? (Hint: the equilibrium uses mixed strategies.)

12.6. A Fixed Bargaining Cost, Again

Apex and Brydax are entering into a joint venture that will yield 500 million dollars, but they must negotiate the split first. In bargaining round 1, Apex makes an offer at cost 0, proposing to keep A_1 for itself. Brydax either accepts (ending the game) or rejects. In Round 2, Brydax makes an offer at cost 10 million of A_2 for Apex, and Apex either accepts or rejects. In Round 3, Apex makes an offer of A_3 at cost c , and Brydax either accepts or rejects. If no offer is ever accepted, the joint venture is cancelled.

- (a) If $c = 0$, what is the equilibrium? What is the equilibrium outcome?
- (b) If $c = 10$, what is the equilibrium? What is the equilibrium outcome?
- (c) If $c = 300$, what is the equilibrium? What is the equilibrium outcome?

12.7. Myerson-Satterthwaite

The owner of a tract of land values his land at v_s and a potential buyer values it at v_b . The buyer and seller do not know each other's valuations, but guess that they are uniformly distributed between 0 and 1. The seller and buyer suggest p_s and p_b simultaneously, and they have agreed that the land will be sold to the buyer at price $p = \frac{(p_b + p_s)}{2}$ if $p_s \leq p_b$.

The actual valuations are $v_s = 0.2$ and $v_b = 0.8$. What is one equilibrium outcome given these valuations and this bargaining procedure? Explain why this can happen.

12.8. Negotiation (Rasmusen [2002])

Two parties, the Offeror and the Acceptor, are trying to agree to the clauses in a contract. They have already agreed to a basic contract, splitting a surplus 50- 50, for a surplus of Z for each player. The offeror can at cost C offer an additional clause which the acceptor can accept outright, inspect carefully (at cost M), or reject outright. The additional clause is either “genuine,” yielding the Offeror X_g and the Acceptor Y_g if accepted, or “misleading,” yielding the Offeror X_m (where $X_m > X_g > 0$) and the Acceptor $-Y_m < 0$.

What will happen in equilibrium?

13 Auctions

13.1 Auction Classification

Because auctions are stylized markets with well-defined rules, modelling them with game theory is particularly appropriate. Moreover, several of the motivations behind auctions are similar to the motivations behind the asymmetric information contracts of Part II of this book. Besides the mundane reasons such as speed of sale that make auctions important, auctions are useful for a variety of informational purposes. Often the buyers know more than the seller about the value of what is being sold, and the seller, not wanting to suggest a price first, uses an auction as a way to extract information. Art auctions are a good example, because the value of a painting depends on the buyer's tastes, which are known only to himself.

Auctions are also useful for agency reasons, because they hinder dishonest dealing between the seller's agent and the buyer. If the mayor were free to offer a price for building the new city hall and accept the first contractor who showed up, the lucky contractor would probably be the one who made the biggest political contribution. If the contract is put up for auction, cheating the public is more costly, and the difficulty of rigging the bids may outweigh the political gain.

We will spend most of this chapter on the effectiveness of different kinds of auction rules in extracting surplus from buyers, which requires considering the strategies with which they respond to the rules. Section 13.1 classifies auctions based on the relationships between different buyers' valuations of what is being auctioned, and explains the possible auction rules and the bidding strategies optimal for each rule. Section 13.2 compares the outcomes under the various rules. Section 13.3 looks at what happens when bidders do not know the value of the object to themselves, but that value is private, uncorrelated with the value to anyone else. Section 13.4 discusses optimal strategies under common-value information, which can lead bidders into the "winner's curse" if they are not careful. Section 13.5 is about information asymmetry in common-value auctions.

Private-Value and Common-Value Auctions

Auctions differ enough for an intricate classification to be useful. One way to classify auctions is based on differences in the values buyers put on what is being auctioned. We will call the dollar value of the utility that player i receives from an object its value to him, V_i , and we will call his *estimate* of its value his **valuation**, $\hat{V}_i$.

In a **private-value** auction, each player's valuation is independent of those of the other players. An example is the sale of antique chairs to people who will not resell them. Usually a player's value equals his valuation in a private-value auction.

¹xxx Footnotes starting with xxx are the author's notes to himself. Comments are welcomed.

If an auction is to be private-value, it cannot be followed by costless resale of the object. If there were resale, a bidder's valuation would depend on the price at which he could resell, which would depend on the other players' valuations.

What is special about a private-value auction is that a player cannot extract any information about his own value from the valuations of the other players. Knowing all the other bids in advance would not change his valuation, although it might well change his bidding strategy. The outcomes would be similar even if he had to estimate his own value, so long as the behavior of other players did not help him to estimate it, so this kind of auction could just as well be called "private-valuation auction."

In a **common-value** auction, the players have identical values, but each player forms his own valuation by estimating on the basis of his private information. An example is bidding for US Treasury bills. A player's valuation would change if he could sneak a look at the other players' valuations, because they are all trying to estimate the same true value.

The values in most real-world auctions are a combination of private-value and common-value, because the valuations of the different players are correlated but not identical. This is sometimes called the **affiliated values** case. As always in modelling, we trade off descriptive accuracy against simplicity. It is common for economists to speak of mixed auctions as "common-value" auctions, since their properties are closer to those of common-value auctions.

Auction Rules and Private-Value Strategies

Auctions have as many different sets of rules as poker games do. We will begin with four different sets of auction rules, and in the private-value setting, since it is simplest. In teaching this material, I ask each student to pick a valuation between 80 and 100, after which we conduct the various kinds of auctions. I advise the reader to try this. Pick two valuations and try out sample strategy combinations for the different auctions as they are described. Even though the values are private, it will immediately become clear that the best-response bids still depend on the strategies the bidder thinks the other players have adopted.

The types of auctions to be described are:

- 1 English (first price open cry);
- 2 First price sealed bid;
- 3 Second price sealed bid (Vickrey);
- 4 Dutch (descending).

English (first price open cry)

Rules

Each bidder is free to revise his bid upwards. When no bidder wishes to revise his bid further, the highest bidder wins the object and pays his bid.

Strategies

A player's strategy is his series of bids as a function of (1) his value, (2) his prior estimate of other players' valuations, and (3) the past bids of all the players. His bid can therefore be updated as his information set changes.

Payoffs

The winner's payoff is his value minus his highest bid. The losers' payoffs are zero.

A player's dominant strategy in a private-value English auction is to keep bidding some small amount ϵ more than the previous high bid until he reaches his valuation, and then to stop. This is optimal because he always wants to buy the object if the price is less than its value to him, but he wants to pay the lowest price possible. All bidding ends when the price reaches the valuation of the player with the second-highest valuation. The optimal strategy is independent of risk neutrality if players know their own values with certainty rather than having to estimate them, although risk-averse players who must estimate their values should be more conservative in bidding.

In common-value open-cry auctions, the bidding procedure is important. Possible procedures include (1) the auctioneer to raise prices at a constant rate, (2) the bidders to raise prices as specified in the rules above, and (3) the **open-exit** auction, in which the price rises continuously and players must publicly announce that they are dropping out (and cannot re-enter) when the price becomes unacceptably high. In an open-exit auction the players have more evidence available about each others' valuations than when they can drop out secretly.

First-Price Sealed-Bid

Rules

Each bidder submits one bid, in ignorance of the other bids. The highest bidder pays his bid and wins the object.

Strategies

A player's strategy is his bid as a function of his value and his prior beliefs about other players' valuations.

Payoffs

The winner's payoff is his value minus his bid. The losers' payoffs are zero.

Suppose Smith's value is 100. If he bid 100 and won when the second bid was 80, he would wish that he had bid only less. If it is common knowledge that the second-highest value is 80, Smith's bid should be $80 + \epsilon$. If he is not sure about the second-highest value, the problem is difficult and no general solution has been discovered. The tradeoff is between bidding high—thus winning more often—and bidding low—thus benefiting more if the bid wins. The optimal strategy, whatever it may be, depends on risk neutrality and beliefs about the other bidders, so the equilibrium is less robust than the equilibria of English and second-price auctions.

Nash equilibria can be found for more specific first-price auctions. Suppose there are N risk-neutral bidders independently assigned values by Nature using a uniform density from 0 to some amount $\bar{v}$. Denote player i 's value by v_i , and let us consider the strategy for player 1. If some other player has a higher value, then in a symmetric equilibrium, player 1 is going to lose the auction anyway, so he can ignore that possibility in finding his optimal bid. Player 1's equilibrium strategy is to bid ϵ above his expectation of the second-highest value, conditional on his bid being the highest (i.e., assuming that no other bidder has a value over v_1).

If we assume that v_1 is the highest value, the probability that Player 2's value, which is uniformly distributed between 0 and v_1 , equals v is $1/v_1$, and the probability that v_2 is less than or equal to v is v/v_1 . The probability that v_2 equals v and is the second-highest value is

$$Prob(v_2 = v) \cdot Prob(v_3 \leq v) \cdot Prob(v_4 \leq v) \cdots Prob(v_N \leq v), \quad (1)$$

which equals

$$\left(\frac{1}{v_1}\right) \left(\frac{v}{v_1}\right)^{N-2}. \quad (2)$$

Since there are $N - 1$ players besides player 1, the probability that one of them has the value v , and v is the second-highest is $N - 1$ times expression (2). Let us define the value of the second-highest value to be $v_{(2)}$ (as distinct from " v_2 ," the value of the second bidder). The expectation of $v_{(2)}$ is the integral of v over the range 0 to v_1 ,

$$\begin{aligned} Ev_{(2)} &= \int_0^{v_1} v(N-1)(1/v_1)[v/v_1]^{N-2} dv \\ &= (N-1) \frac{1}{v_1^{N-1}} \int_0^{v_1} v^{N-1} dv \\ &= \frac{(N-1)v_1}{N}. \end{aligned} \quad (3)$$

Thus we find that player 1 ought to bid a fraction $\frac{N-1}{N}$ of his own value, plus ϵ . If there are 2 bidders, he should bid $\frac{1}{2}v_1$, but if there are 10 he should bid $\frac{9}{10}v_1$.

A Mixed-Strategy Equilibrium in a First-Price Auction. The previous example is an elegant result but not a general rule. Often the equilibrium in a first-price auction is not even in pure strategies. Consider an auction in which each bidder's private value v is either 10 or 16 with equal probability and is known only to himself.

If the auction is first-price sealed bid, a bidder's optimal strategy is to set bid $b(v = 10) = 10$ and if $v = 16$ to use a mixed strategy, mixing over the support $[x, y]$, where it will turn out that $x = 10$ and $y = 13$, and his expected payoff will be

$$\begin{aligned} E\pi(v = 10) &= 0 \\ E\pi(v = 16) &= 3. \end{aligned} \tag{4}$$

This will serve as an illustration of how to find an equilibrium mixed strategy when players mix over a continuum of pure strategies rather than just between two of them. The first step is to see why the equilibrium cannot be in pure strategies (though it can include pure strategies following some moves by Nature). Suppose both players are using the strategy $b(v = 10) = 10$, and $b_1(v_1 = 16) = z_1$ and $b_2(v_2 = 16) = z_2$, where z_1 and z_2 are in $(10, 16]$ because a bid of $10 + \epsilon$ will always defeat the $b(v = 10) = 10$ where a bid of 10 or less would not, and a bid of over 16 would yield a negative payoff. Either $z_1 = z_2$, or $z_1 \neq z_2$. If $z_1 = z_2$, then each player has incentive to deviate to $z_1 - \epsilon$ and win always instead of tying. If $z_1 < z_2$, then Player 2 will deviate to bid just $z_1 + \epsilon$. If he does that, however, Player 1 would have incentive to deviate to bid $z_1 + 2\epsilon$, so he could win always at trivially higher cost. The same holds true if $z_2 < z_1$. Thus, there is no equilibrium in pure strategies.

The second step is to figure out what pure strategies will be mixed between by a player with $v = 16$. The bid, b , will not be less than 10 (which would always lose) or greater than 16 (which would always win and yield a negative payoff). In fact, since the pure strategy of $b = 10 + \epsilon | v = 16$ will win with probability of at least 0.50 (because the other player happens to have $v = 10$, thus yielding a payoff of at least $0.50(16 - 10) = 3$, the upper bound y must be no greater than 13.

Consider the following two possible features of an equilibrium mixing distribution. Suppose that Bidder 2's density has one or both of these features.

(a) The mixing density has an atom at point a in $[10, 13]$ – some particular point which has positive probability that we will denote by $T(a)$.

(b) The mixing density has a gap $[g, h]$ somewhere in $[10, 13]$ – that is, there is zero probability of a bid in $[g, h]$.

(refutation of a) If Player 2 puts positive probability on point $a > 10$, then Player 1 should respond by putting positive probability on point $a - \epsilon$.xxx unfinished.

(refutation of b) xxx unfinished Suppose $x > 10$ instead of $x = 10$. The bidder's expected payoff is, if the other player has no atom at x in his strategy,

$$0.5(16 - x) + 0.5 \text{Prob.}(\text{win with } x | v_2 = 16)(16 - x) = 0.5(16 - x). \quad (5)$$

If the bidder deviates to bidding 10 instead of x , his probability of winning is virtually unchanged, but his payoff increases to:

$$0.5(16 - 10) + 0.5 \text{Prob.}(\text{win with } x | v_2 = 16)(16 - x) = 0.5(16 - 10). \quad (6)$$

Thus, we can conclude that the mixing density $m(b)$ is positive over the entire interval $[10, 13]$, with no atoms. What will it look like? Let us confine ourselves to looking for a symmetric equilibrium, in which both players use the same function $m(b)$. We know the expected payoff from any bid b in the support must equal the payoff from $b = 10$ or $b = 13$, which is 3. Therefore,

$$0.5(16 - b) + 0.5M(b)(16 - b) = 3. \quad (7)$$

This implies that $(16 - b) + M(b)(16 - b) = 6$, so

$$M(b) = \frac{6}{16 - b} - 1 = \frac{6 - 16 + b}{16 - b} = \frac{b - 10}{16 - b}, \quad (8)$$

so the mixing distribution is uniform, with density $m(b) = \frac{1}{16 - b}$ on the support $[10, 13]$.

If the auction were second-price sealed-bid or English, a bidder's optimal strategy is to set bid or bid ceiling $b(v = 10) = 10$ and $b(v = 16) = 16$. His expected payoff is then

$$\begin{aligned} E\pi(v = 10) &= 0 \\ E\pi(v = 16) &= 0.5(16 - 10) + 0.5(16 - 10) = 3. \end{aligned} \quad (9)$$

The expected price is 13, the same as in the first-price auction, so the seller is indifferent about the rules. The buyer's expected payoff is also the same: 3 if $v = 16$ and 0 if $v = 10$. The only difference is that now the buyer's payoff ranges over the continuum $[0, 3]$ rather than being either 0 or 3.

Second-Price Sealed-Bid (Vickrey)

Rules

Each bidder submits one bid, in ignorance of the other bids. The bids are opened, and the highest bidder pays the amount of the second-highest bid and wins the object.

Strategies

A player's strategy is his bid as a function of his value and his prior belief about other players' valuations.

Payoffs

The winner's payoff is his value minus the second-highest bid that was made. The losers' payoffs are zero.

Second-price auctions are similar to English auctions. They are rarely used in reality, but are useful for modelling. Bidding one's valuation is the dominant strategy: a player who bids less is more likely to lose the auction, but pays the same price if he does win. The structure of the payoffs is reminiscent of the Groves mechanism of section 10.6, because in both games a player's strategy affects some major event (who wins the auction or whether the project is undertaken), but his strategy affects his own payoff only via that event. In the auction's equilibrium, each player bids his value and the winner ends up paying the second-highest value. If players know their own values, the outcome does not depend on risk neutrality.²

Dutch (descending)

Rules

The seller announces a bid, which he continuously lowers until some buyer stops him and takes the object at that price.

Strategies

A player's strategy is when to stop the bidding as a function of his valuation and his prior beliefs as to other players' valuations.

Payoffs

The winner's payoff is his value minus his bid. The losers' payoffs are zero.

The Dutch auction is **strategically equivalent** to the first-price sealed-bid auction, which means that there is a one-to-one mapping between the strategy sets and the equilibria of the two games. The reason for the strategic equivalence is that no relevant information is disclosed in the course of the auction, only at the end, when it is too late to change anybody's behavior. In the first-price auction a player's bid is irrelevant unless it is the highest, and in the Dutch auction a player's stopping price is also irrelevant unless it is the highest. The equilibrium price is calculated the same way for both auctions.

Dutch auctions are actually used. One example is the Ontario tobacco auction, which uses a clock four feet in diameter marked with quarter-cent gradations. Each of six or so buyers has a stop button. The clock hand drops a quarter cent at a time, and the stop buttons are registered so that ties cannot occur (tobacco buyers need reflexes like race-car drivers). The farmers who are selling their tobacco watch from an adjoining room and can later reject the bids if they feel they are too low (a form of reserve price); 2,500,000 lb. a day can be sold using the clock (Cassady [1967] p. 200).

²xxx I should make note of the odd 2nd price asymmetric equilibrium with zero price, in private value auctions.

Dutch auctions are common in less obvious forms. Filene's is one of the biggest stores in Boston, and Filene's Basement is its most famous department. In the basement are a variety of marked-down items formerly in the regular store, each with a price and date attached. The price customers pay at the register is the price on the tag minus a discount which depends on how long ago the item was dated. As time passes and the item remains unsold, the discount rises from 10 to 50 to 70 percent. The idea of predictable time discounting has also recently been used by bookstores ("Waldenbooks to Cut Some Book Prices in Stages in Test of New Selling Tactic," *Wall Street Journal*, March 29, 1988, p. 34).

13.2 Comparing Auction Rules

When one mentions auction theory to an economic theorist, the first thing that springs to his mind is the idea, in some sense, that different kinds of auctions are really the same. Milgrom & Weber (1982) give a good summary of how and why this is true. Regardless of the information structure, the Dutch and first-price sealed-bid auctions are the same in the sense that the strategies and the payoffs associated with the strategies are the same. That equivalence does not depend on risk neutrality, but let us assume that all players are risk neutral for the next few paragraphs.

In private, independent-value auctions, the second-price sealed-bid and English auctions are the same in the sense that the bidder who values the object most highly wins and pays the valuation of the bidder who values it the second highest, but the strategies are different in the two auctions. In all four kinds of private independent-value auctions discussed, the seller's expected price is the same. This fact is the biggest result in auction theory: the **revenue equivalence theorem** (Vickrey [1961]). We will show it using the following game, a mechanism design approach.

The Auctions Mechanism Game

Players: A seller and N buyers.

Order of Play:

0. Nature chooses buyer i 's value for the object, v_i , using the strictly positive, atomless, density $f(v)$ on the interval $[v, \bar{v}]$.
1. The seller chooses a mechanism $[x(a), t(a)]$ that takes payment t and gives the object with probability x to a player (himself, or one of the N buyers) who announces that his value is a .³ He also chooses the procedure in which players select a (sequentially, simultaneously, etc.).
2. Each buyer simultaneously chooses to participate in the auction or to stay out.
3. The buyers and seller choose a according to the mechanism procedure.
4. The object is allocated and transfers are paid according to the mechanism, if it was accepted by all players.

Payoffs:

The seller's payoff is

$$\pi_s = \sum_{i=1}^N t(a_i, a_{-i}) \quad (10)$$

Buyer i 's payoff is

$$\pi_i = x(a_i, a_{-i})v_i - t(a_i, a_{-i}) \quad (11)$$

Many auction procedures fit the mechanism paradigm. The x function could allocate the good with 70% probability to the high bidder and with 30% probability to the lowest bidder, for example; each bidder could be made to pay the amount he bids, even if he loses; t could include an entry fee; there could be a "reserve price" which is the minimum bid for which the seller will surrender the good. The seller must choose a mechanism that for each type v_i satisfies a participation constraint (bidder v_i will join the auction, so, for example, the entry fee is not too large) and an incentive compatibility constraint (the bidder will truthfully reveal his type). The game has multiple equilibria, because there is more than one mechanism that maximizes the seller's payoff.

³In equilibrium, the probability that more than one player will announce the same a is zero, so we will not bother to specify a tie-breaking rule. More properly, the seller should do so, but in this game, any rule would do equally well.

If the other $(N - 1)$ buyers each choose $a_j = v_j$, as they will in a direct mechanism, which induces truthtelling, let us denote the expected maximized payoff of a buyer with value v_i as $\pi(v_i)$, so that

$$\pi(v_i) \equiv \max_{a_i} \{E_{v_{-i}}[x(a_i, v_{-i})v_i - t(a_i, v_{-i})]\}. \quad (12)$$

Another way to write $\pi(v_i)$ is as a base level, $\pi(\underline{v})$, plus the integral of its derivatives from $\underline{v}$ to v_i :

$$\pi(v_i) = \pi(\underline{v}) + \int_{v=\underline{v}}^{v_i} \frac{d\pi(v)}{dv} dv. \quad (13)$$

The seller will not give the lowest-valuing buyer type, $\underline{v}$, a positive expected payoff because then t_i could be increased by a constant amount for all types— an entry fee— without violating the participation or incentive compatibility constraints. Thus, $\pi(\underline{v}) = 0$ and we can rewrite the payoff as

$$\pi(v_i) = \int_{v=\underline{v}}^{v_i} \frac{d\pi(v)}{dv} dv. \quad (14)$$

To connect equations (12) and (14), our two expressions for $\pi(v_i)$, we will use a trick discovered by Mirrlees (1971) for mechanism design problems generally and use the Envelope Theorem to eliminate the t transfers. Differentiating $\pi(v_i)$ with respect to v_i , the Envelope Theorem says that we can ignore the indirect effect of v_i on π via its effect on a_i , since $\frac{da_i}{dv_i} = 0$ in the maximized payoff. therefore,

$$\begin{aligned} \frac{d\pi(v_i)}{dv_i} &= E_{v_{-i}} \left[\frac{\partial \pi(v_i)}{\partial v_i} + \frac{\partial \pi(v_i)}{\partial a_i} \frac{da_i}{dv_i} \right] \\ &= E_{v_{-i}} \frac{\partial \pi(v_i)}{\partial v_i} \\ &= E_{v_{-i}} x(a_i, v_{-i}), \end{aligned} \quad (15)$$

where the last line takes the partial derivative of equation (12).

Substituting back into equation (14) and using the fact that in a truthful direct mechanism $a_i = v_i$, we arrive at

$$\pi(v_i) = E_{v_{-i}} \int_{v=\underline{v}}^{v_i} x(v, v_{-i}) dv. \quad (16)$$

Now we can rearrange (12) and use our new $\pi(v_i)$ expression in (16) to solve for the expected transfer from a buyer of type v_i to the seller:

$$\begin{aligned} E_{v_{-i}} t(v_i, v_{-i}) &= E_{v_{-i}} x(v_i, v_{-i}) v_i - \pi(v_i) \\ &= E_{v_{-i}} \left[x(v_i, v_{-i}) v_i - \int_{v=\underline{v}}^{v_i} x(v, v_{-i}) dv \right]. \end{aligned} \quad (17)$$

Let $\pi_s(i)$ denote the seller's expected revenue from buyer i , not yet knowing any of the buyers' types, so

$$\begin{aligned} \pi_s(i) &= E_{v_i} E_{v_{-i}} t(v_i, v_{-i}) \\ &= E_{v_i} E_{v_{-i}} \left[x(v_i, v_{-i}) v_i - \int_{m=\underline{v}}^{v_i} x(m, v_{-i}) dm \right] \\ &= E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} \left[x(v, v_{-i}) v - \int_{m=\underline{v}}^v x(m, v_{-i}) dm \right] f(v) dv \\ &= E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} \left[x(v, v_{-i}) f(v) v - \int_{m=\underline{v}}^v x(m, v_{-i}) f(v) dm \right] dv \end{aligned} \quad (18)$$

At this point, we need to integrate by parts to deal with $\int x(m, v_{-i}) f(v) dm$. The formula for integration by parts is $\int_{z=a}^b g(z) h'(z) dz = g(z) h(z) \Big|_{z=a}^b - \int_{z=a}^b h(z) g'(z) dz$. Let $z = v$, $g(z) = \int_{m=\underline{v}}^v x(m, v_{-i}) dm$ and $h'(z) = f(v)$. It follows that $g'(z) = x(v, v_{-i})$ and $h(z) = F(v)$ and we can write $\pi_s(i)$ as

$$\begin{aligned} &E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) f(v) v dv - E_{v_{-i}} \left(F(v) \int_{m=\underline{v}}^v x(m, v_{-i}) dm \Big|_{v=\underline{v}}^{\bar{v}} - \int_{v=\underline{v}}^{\bar{v}} F(v) x(v, v_{-i}) dv \right) \\ &= E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) f(v) v dv - \\ &E_{v_{-i}} \left(F(\bar{v}) \int_{m=\underline{v}}^{\bar{v}} x(m, v_{-i}) dm - F(\underline{v}) \int_{m=\underline{v}}^{\underline{v}} x(m, v_{-i}) dm - \int_{v=\underline{v}}^{\bar{v}} F(v) x(v, v_{-i}) dv \right). \end{aligned} \quad (19)$$

Since $F(\bar{v}) = 1$ and $F(\underline{v}) = 0$ by definition, and since we can divide by f because of the assumption that the density is strictly positive, we can further rewrite $\pi_s(i)$ as

$$\begin{aligned} &= E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) f(v) v dv - E_{v_{-i}} \left(\int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) dv - \int_{v=\underline{v}}^{\bar{v}} F(v) x(v, v_{-i}) dv \right) \\ &= E_{v_{-i}} \int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) f(v) \left(v - \frac{1 - F(v)}{f(v)} \right) dv \end{aligned} \quad (20)$$

The seller's expected payoff from all N bidders sums $\pi_s(i)$ up over i :

$$E_v \pi_s = \sum_{i=1}^N \int_{v=\underline{v}}^{\bar{v}} x(v, v_{-i}) f(v) \left(v - \frac{1 - F(v)}{f(v)} \right) dv \quad (21)$$

The seller wants to choose $x()$ so as to maximize (21). The way to do this is to set $x = 1$ for the v which has the biggest $\left(v - \frac{1 - F(v)}{f(v)} \right)$, which is what we will next examine.

MR approach

Think of a firm facing a demand curve made up of a continuum of bidders drawn from $f(v)$. The marginal revenue from this demand curve is⁴

$$MR(v) = v - \frac{1 - F(v)}{f(v)}. \quad (22)$$

If the value of the object to the seller is v_{seller} , then he should act like a monopolist with constant marginal cost of $v_{seller} = 0$ facing the demand curve based on $f(v)$, which means he should set a reserve price at the quantity where marginal revenue equals marginal cost. Figure 1 illustrates this. The analog to price is, of course, the value to the buyer, and the analog to quantity is the probability of sale at that price, which is $1 - F(v)$. This gives a good explanation for why the reserve price should be above the seller's use value for the object. In particular, if the seller gets no benefit from retaining the object ($v_{seller} = 0$), he should still set a reserve price where marginal revenue equals marginal cost. And the optimal reserve price is independent of the number of bidders—it just depends on the possible bidder values. He must publicize this, of course, because the point of the reserve price is then just to induce the buyers to bid higher.

⁴See Bulow and Roberts [1989] on risk-neutral private values; Bulow & Klemperer [1996] on common values and risk aversion.




Figure 1: Auctions and Marginal Revenue

This expression has economic meaning, as Bulow & Klemperer (1996) show.⁵ Suppose the seller had to make a single take-it-or-leave-it bargaining offer, instead of holding an auction. We can think of v as being like a price on a demand curve and $[1 - F(v)]$ as being like a quantity, since at price $p = v$ the single unit has probability $[1 - F(v)]$ of being sold. Revenue is then $R = pq = v[1 - F(v)]$ and marginal revenue is $\frac{dR}{dq} = p + q \frac{dp}{dq} = p + \frac{q}{dq/dp} = v + \frac{[1-F(v)]}{-f(v)}$. Thus, $\left(v_i - \frac{1-F(v_i)}{f(v_i)}\right)$ is like marginal revenue, and it makes sense to choose a price or set up an auction rule that elicits a price such that marginal revenue is maximized. Figure 1 illustrates this.⁶

Notice that t is not present in equation (21). It is there implicitly, however, because we have assumed we have found a truthful mechanism, and to satisfy the participation and incentive compatibility constraints, t has to be chosen appropriately. We can use equation (21) to deduce some properties of the optimal mechanism, and then we will find a number of mechanisms that satisfy those properties and all achieve the same expected payoff for the seller.

Suppose $\left(v_i - \frac{1-F(v_i)}{f(v_i)}\right)$ is increasing in v_i . This is a reasonable assumption, satisfied if the **monotone hazard rate** condition that $\frac{1-F(v_i)}{f(v_i)}$ decreases in v_i is true. In that case, the seller's best auction sells with probability one to the buyer with the biggest value of v_i . Thus, we have proved:

⁵xxx Also: conditional expectation of biggest minus second- biggest bid.

⁶In price theory, the monopolist chooses q or p so that marginal revenue equals not zero, but marginal cost. Here, we have assumed the seller gets zero value from not selling the good. If he did, there would be an opportunity cost to selling the good which would play the role of the marginal cost.

THE REVENUE EQUIVALENCE THEOREM. Suppose all bidders's valuations are drawn from the same strictly positive and atomless density $f(v)$ over $[\underline{v}, \bar{v}]$ and that f satisfies the monotone hazard rate condition. Any auction in which type $\underline{v}$ has zero expected surplus and the winner is the bidder with the highest value will have the same expected profit for the seller.

In particular, the following forms are optimal and yield identical expected profits:

Ascending auction. Everyone pays an entry fee (to soak up the surplus of the $\underline{v}$ type). The winner is the highest-value bidder, and he is refunded his entry fee but pays the value of the second-highest valuer.

Second-price sealed bid. The same.xxx

First-price sealed bid. The winner is the person who bids the highest. Is he the highest valuer? A bidder's expected payoff is $\pi(v_i) = P(b_i)(v_i - b_i) - T$, where T is the entry fee and $P(b_i)$ is the probability of winning with bid b_i . The first order condition is $\frac{d\pi(v_i)}{db_i} = P'(v_i - b_i) - P = 0$, with second order condition $\frac{d^2\pi(v_i)}{db_i^2} \leq 0$. Using the implicit function theorem and the fact that $\frac{d^2\pi(v_i)}{db_i dv_i} = P' > 0$, we can conclude that $\frac{db_i}{dv_i} \geq 0$. But it cannot be that $\frac{db_i}{dv_i} = 0$, because then there would be values v_1 and v_2 such that $b_1 = b_2 = b$ and $\frac{d\pi(v_1)}{db_1} = P'(b)(v_1 - b) - P(b) = 0 = \frac{d\pi(v_2)}{db_2} = P'(b)(v_2 - b) - P(b)$, which cannot be true. So bidders with higher values bid higher, and the highest valuer will win the auction.

Dutch (descending) auction. Same as the first-price sealed-bid.

All-pay sealed-bid. Each player pays his bid. Whoever bids highest is awarded the object.xxx

Although the different auctions have the same expected payoff for the seller, they do not have the same realized payoff. In the first-price-sealed-bid auction, for example, the winning buyer's payment depends entirely on his own valuation. In the second-price-sealed-bid auction the winning buyer's payment depends entirely on the second-highest valuation, which is sometimes close to his own valuation and sometimes much less.

In the Dutch and first-price sealed-bid auctions, the winning bidder has estimated the value of the second-highest bidder, and that estimate, while correct on average, is above or below the true value in particular realizations. The variance of the price is higher in those auctions because of the additional estimation, which means that a risk-averse seller should use the English or second-price auction.

Hazard Rates

Suppose we have a density $f(v)$ for a bidder's value for an object being sold, with cumulative distribution $F(v)$ on support $[\underline{v}, \bar{v}]$. The hazard rate $h(v)$ is defined as

$$h(v) = \frac{f(v)}{1 - F(v)} \quad (23)$$

What this means is that $h(v)$ is the probability density of v for the distribution which is like $F(v)$ except cutting off the lower values, so its support is limited to $[v, \bar{v}]$. In economic terms, $h(v)$ is the probability density for the valuing equalling v given that we know that value equals at least v .

For most distributions we use, the hazard rate is increasing, including the uniform, normal, logistic, and exponential distributions, and any distribution with increasing density over its support (see Bagnoli & Bergstrom [1989]). Figure 2 shows three of them




Figure 2: Three Densities to Illustrate Hazard Rates

Myerson (1981) and Bulow & Roberts (1989) look at optimal auctions when bidders are asymmetric or higher values do not imply that $(v_i - \frac{1-F(v_i)}{f(v_i)})$ is higher. Then the revenue-maximizing auction might not allocate the good to the bidder with the highest value. This is because that bidder might not be the bidder with the highest *expected* value. If there were two bidders, it could happen that $(\frac{1-F_1(v)}{f_1(v)}) > (\frac{1-F_2(v)}{f_2(v)})$, in which case the auction should be biased in favor of Bidder 2 who should sometimes win even if $v_2 < v_1$.

Whether the auction is private-value or not, the Dutch and first-price sealed-bid auctions are strategically equivalent. If the auction is correlated-value and there are three or more bidders, the open-exit English auction leads to greater revenue than the second-price sealed-bid auction, and both yield greater revenue than the first-price sealed-bid auction (Milgrom & Weber [1982]). If there are just two bidders, however, the open-exit English auction is no better than the second-price sealed-bid auction, because the open-exit feature – knowing when nonbidding players drop out – is irrelevant.

A question of less practical interest is whether an auction form is Pareto optimal; that is, does the auctioned object end up in the hands of whoever values it most? In a common-value auction this is not an interesting question, because all bidders value the object equally. In a private-value auction, all the five auction forms – first-price, second-price, Dutch, English, and all-pay – are Pareto optimal. They are also optimal in a correlated-value auction if all players draw their information from the same distribution and if the equilibrium is in symmetric strategies.

Hindering Buyer Collusion

As I mentioned at the start of this chapter, one motivation for auctions is to discourage collusion between players. Some auctions are more vulnerable to this than others. Robinson (1985) has pointed out that whether the auction is private-value or common-value, the first-price sealed-bid auction is superior to the second-price sealed-bid or English auctions for deterring collusion among bidders.

Consider a buyer's cartel in which buyer Smith has a private value of 20, the other buyers' values are each 18, and they agree that everybody will bid 5 except Smith, who will bid 6. (We will not consider the rationality of this choice of bids, which might be based on avoiding legal penalties.) In an English auction this is self-enforcing, because if somebody cheats and bids 7, Smith is willing to go all the way up to 20 and the cheater will end up with no gain from his deviation. Enforcement is also easy in a second-price sealed-bid auction, because the cartel agreement can be that Smith bids 20 and everyone else bids 6.

In a first-price sealed-bid auction, however, it is hard to prevent buyers from cheating on their agreement in a one-shot game. Smith does not want to bid 20, because he would have to pay 20, but if he bids anything less than the other players' value of 18 he risks them overbidding him. The buyer will end up paying a price of 18, rather than the 6 he would receive in an English auction with collusion. The seller therefore will use the first-price sealed-bid auction if he fears collusion.⁷

13.3 Risk and Uncertainty over Values

⁷Even then, his fears may be realized. Bidding rings are not uncommon even though they are illegal and even though first-price auctions are used. See Sultan (1974) for the Electric Conspiracy, one famous example.

In a private-value auction, does it matter what the seller does, given the revenue equivalence theorem? Yes, because of risk aversion, which invalidates the theorem. Risk aversion makes it important which auction rule is chosen, because the seller should aim to reduce uncertainty, even in a private value auction. (In a common value auction, reducing uncertainty has the added benefit of ameliorating the winner's curse.)

Consider the following question:

If the seller can reduce bidder uncertainty over the value of the object being auctioned, should he do so?

We must assume that this is a precommitment on the part of the seller, since otherwise he would like to reveal favorable information and conceal unfavorable information. But it is often plausible that the seller can set up an auction system which reduces uncertainty - say, by a regular policy of allowing bidders to examine the goods before the auction begins. Let us build a model to show the effect of such a policy.

Suppose there are N bidders, each with a private value, in an ascending open-cry auction. Each measures his private value v with an independent error ϵ . This error is with equal probability $-x$, $+x$ or 0 . The bidders have diffuse priors, so they take all values of v to be equally likely, ex ante. Let us denote the measured value by $\hat{v} = v + \epsilon$, which is an unbiased estimate of v . What should bidder i bid up to?

If bidder i is risk neutral, he should bid up to $p = \hat{v}$. If he pays $\hat{v}$, his expected utility is,

$$\pi(\text{risk neutral}, p = \hat{v}) = \frac{([\hat{v} - x] - \hat{v})}{3} + \frac{(\hat{v} - \hat{v})}{3} + \frac{(\hat{v} + x - \hat{v})}{3} = 0. \quad (24)$$

If bidder i is risk averse, however, and wins with bid p , his expected utility is, if his utility function is $\pi = U(v - p)$ for concave U ,

$$\pi(\text{risk averse}, p) = \frac{U([\hat{v} - x] - p)}{3} + \frac{U(\hat{v} - p)}{3} + \frac{U([\hat{v} + x] - p)}{3} \quad (25)$$

Since the utility function U is concave,

$$\frac{U([\hat{v} - x] - p)}{3} + \frac{U([\hat{v} + x] - p)}{3} < \frac{2}{3}U(\hat{v} - p). \quad (26)$$

The implication is that a fair gamble of x has less utility than no gamble. This means that the middle term in equation (25) must be positive if it is to be true that $\pi = U(0)$, which means that $\hat{v} - p > 0$. In other words, bidder i will have a negative expected payoff unless his maximum bid is strictly less than his valuation.

Other auctions with risk-averse bidders are more difficult to analyze. The problem is that in a first-price sealed-bid auction or a Dutch auction, there is risk not only from uncertainty over the value but over how much the other players will bid. One finding is that in a private-value auction the first-price sealed-bid auction yields a greater expected revenue than the English or second-price auctions. That is because by increasing his bid from the level optimal for a risk-neutral bidder, the risk-averse bidder insures himself. If he wins, his surplus is slightly less because of the higher price, but he is more likely to win and avoid a surplus of zero. Thus, the buyers' risk aversion helps the seller.

Notice that the seller does not have control over all the elements of the model. The seller can often choose the auction rules unilaterally. This includes not just how bids are made, but such things as whether the bidders get to know how many potential bidders are in the auction, whether the seller himself is allowed to bid, and so forth. Also, the seller can decide how much information to release about the goods. The seller cannot, however, decide whether the bidders are risk averse or not, or whether they have common values or private values – no more than he can choose what their values are for the good he is selling. All of those things concern the utility functions of the bidders. At best, the seller can do things such as choose to produce goods to sell at the auction which have common values instead of private values.⁸

An error I have often seen is to think that the presence of uncertainty in valuations always causes the winner's curse. It does not, unless the auction is a common-value one. Uncertainty in one's valuation is a necessary but not sufficient condition for the winner's curse. It is true that risk-averse bidders should not bid as high as their valuations if they are uncertain about their valuations, even if the auction is a private-value one. That sounds a lot like a winner's curse, but the reason for the discounted bids is completely different, depending as it does on risk aversion. If bidders are uncertain about valuations but they are risk-neutral, their dominant strategy is still to bid up to their valuations. If the winner's curse is present, even if a bidder is risk-neutral he discounts his bid because if he wins on average his valuation will be greater than the value.

Risk Aversion in Private-Value Auctions

When buyers are risk averse, the ranking of seller revenues is:

1st price sealed-bid and Dutch: best and identical to each other
English: next
2nd-price sealed bid: worst

To understand this, consider the following example. Each bidder's private value is either 10 or 16, with equal probability, and is known only to himself.

⁸xxx Here add reference to my work on buying info on one's value.

If the auction is second-price sealed-bid or English, a bidder's optimal strategy is to set bid or bid ceiling $b(v = 10) = 10$ and $b(v = 16) = 16$. His expected payoff is then

$$\begin{aligned} E\pi(v = 10) &= 0 \\ E\pi(v = 16) &= 0.5(16 - 10) + 0.5(16 - 10) = 3. \end{aligned} \tag{27}$$

If the auction is first-price sealed bid or Dutch, a bidder's optimal strategy is to set bid $b(v = 10) = 10$ and if $v = 16$ to use a mixed strategy, mixing over the support $[x, y]$, where it will turn out that $x = 10$ and $y = 13$. His expected payoff will be

$$\begin{aligned} E\pi(v = 10) &= 0 \\ E\pi(v = 16) &= 3. \end{aligned} \tag{28}$$

Here is how one finds the optimal strategy in a symmetric equilibrium. Suppose $x > 10$ instead of $x = 10$. The bidder's expected payoff is, if the other player has no atom at x in his strategy,

$$0.5(16 - x) + 0.5\text{Prob.}(\text{win with } x|v_2 = 16)(16 - x) = 0.5(16 - x). \tag{29}$$

If the bidder deviates to bidding 10 instead of x , his probability of winning is virtually unchanged, but his payoff increases to:

$$0.5(16 - 10) + 0.5\text{Prob.}(\text{win with } x|v_2 = 16)(16 - x) = 0.5(16 - 10). \tag{30}$$

What if Bidder 2 did have an atom at x , and bids x with probability P ? Then Bidder 1 should have an atom at $x + \epsilon$ instead, for some small ϵ . Thus, it must be that $x = 10$.

Suppose $y = 16$. Then the bidder's expected payoff is zero. This, however, is a lower expected payoff than bidding 12, which has an expected payoff of at least $0.5(16 - 12) > 0$ because its payoff is positive if the other player's value is just 10. To find the exact value of y , note that in a mixed strategy equilibrium, all pure strategies in the mixing support must have the same expected payoff. If 10 is in the mixing support, then its expected payoff (since it will almost surely win) is $\pi(b = 10) = 0.5(16 - 10) = 0.5(16 - 10)$. The expected payoff from the pure strategy of bidding y , which always wins (being at the top of the support) is $\pi(b = 16) = 0.5(16 - y) = 0.5(16 - y)$. Equating these yields $y = 13$.

The expected price is 13, the same as in the second-price auction, so the seller is indifferent about the rules, as the Revenue Equivalence Theorem says. The buyer's expected payoff is also the same: 3 if $v = 16$ and 0 if $v = 10$. The only difference is that now the buyer's payoff ranges over the continuum $[0, 3]$ rather than being either 0 or 3.

xxxI need to state the equilibrium and then derive it. What is the mixing distribution?

Now let us make the players risk averse. The optimal second-price strategy does not change.

But now the utility of 6 has shrunk relative to the utility of 3, so the optimal 1st-price strategy does change. That strategy made the expected payoff— not the expected dollar profit— equal for each pure strategy on the mixing support. As before, the optimal strategy if $v_1 = 16$ is to mix over $[x, y]$. with $x = 10$ and $y < 16$, but now y will increase. The payoff from bidding 10 is, if we normalize by setting $U(0) \equiv 0$,

$$\pi(b = 10) = 0.5U(16 - 10) + 0.5U(0) = 0.5U(6). \quad (31)$$

The payoff from bidding y is

$$\pi(b = 10) = 0.5U(16 - y) + 0.5U(16 - y). \quad (32)$$

This requires that y be set so that

$$0.5U(6) = U(16 - y) \quad (33)$$

The utility of winning with $b = 10$ has to equal twice the utility of winning with $b = y$. This means, given concave $U()$, that the dollar profit from winning with $b = 10$ has to equal *more* than twice the dollar profit of winning with $b = y$. Thus, we need $(16 - y) < 6$, and y must increase. Intuitively, if the buyer is risk averse, he becomes less willing to take a chance of losing the auction by bidding low, and more willing to bid high and get a lower payoff but with greater probability.

13.4: Common-Value Auctions and the Winner's Curse

In Section 13.2 we distinguished private-value auctions from common-value auctions, in which the values of the players are identical but their valuations may differ. All four sets of rules discussed there can be used for common-value auctions, but the optimal strategies are different. In common-value auctions, each player can extract useful information about the object's value to himself from the bids of the other players. Surprisingly enough, a buyer can use the information from other buyers' bids even in a sealed-bid auction, as will be explained below.

When I teach this material I bring a jar of pennies to class and ask the students to bid for it in an English auction. All but two of the students get to look at the jar before the bidding starts, and everybody is told that the jar contains more than 5 and less than 100 pennies. Before the bidding starts, I ask each student to write down his best guess of the number of pennies. The two students who do not get to see the jar are like “technical analysts,” those peculiar people who try to forecast stock prices using charts showing the past movements of the stock while remaining ignorant of the stock's “fundamentals.”

A common-value auction in which all the bidders knew the value would not be very interesting, but more commonly, as in the penny jar example, the bidders must estimate the common value. The obvious strategy, especially following our discussion of private-value auctions, is for a player to bid up to his unbiased estimate of the number of pennies in the jar. But this strategy makes the winner's payoff negative, because the winner is the bidder who has made the largest positive error in his valuation. The bidders who underestimated the number of pennies, on the other hand, lose the auction, but their payoff is limited to a downside value of zero, which they would receive even if the true value were common knowledge. Only the winner suffers from overbidding: he has stumbled into the **winner's curse**. When other players are better informed, it is even worse for an uninformed player to win. Anyone, for example, who wins an auction against 50 experts should worry about why they all bid less.

To avoid the winner's curse, players should scale down their estimates in forming their bids. The mental process is a little like deciding how much to bid in a private-value, first-price sealed-bid auction, in which bidder Smith estimates the second-highest value conditional upon himself having the highest value and winning. In the common-value auction, Smith estimates his own value, not the second-highest, conditional upon himself winning the auction. He knows that if he wins using his unbiased estimate, he probably bid too high, so after winning with such a bid he would like to retract it. Ideally, he would submit a bid of $[X \text{ if } I \text{ lose, but } (X - Y) \text{ if } I \text{ win}]$, where X is his valuation conditional upon losing and $(X - Y)$ is his lower valuation conditional upon winning. If he still won with a bid of $(X - Y)$ he would be happy; if he lost, he would be relieved. But Smith can achieve the same effect by simply submitting the bid $(X - Y)$ in the first place, since the size of losing bids is irrelevant.

Another explanation of the winner's curse can be devised from the Milgrom definition of "bad news" (Milgrom [1981b], appendix B). Suppose that the government is auctioning off the mineral rights to a plot of land with common value V and that bidder i has valuation $\hat{V}_i$. Suppose also that the bidders are identical in everything but their valuations, which are based on the various information sets Nature has assigned them, and that the equilibrium is symmetric, so the equilibrium bid function $b(\hat{V}_i)$ is the same for each player. If Bidder 1 wins with a bid $b(\hat{V}_1)$ that is based on his prior valuation $\hat{V}_1$, his posterior valuation $\tilde{V}_1$ is

$$\tilde{V}_1 = E(V | \hat{V}_1, b(\hat{V}_2) < b(\hat{V}_1), \dots, b(\hat{V}_n) < b(\hat{V}_1)). \quad (34)$$

The news that $b(\hat{V}_2) < \infty$ would be neither good nor bad, since it conveys no information, but the information that $b(\hat{V}_2) < b(\hat{V}_1)$ is bad news, since it rules out values of b more likely to be produced by large values of $\hat{V}_2$. In fact, the lower the value of $b(\hat{V}_1)$, the worse is the news of having won. Hence,

$$\tilde{V}_1 < E(V | \hat{V}_1) = \hat{V}_1, \quad (35)$$

and if Bidder 1 had bid $b(\hat{V}_1) = \hat{V}_1$ he would immediately regret having won. If his winning bid were enough below $\hat{V}_1$, however, he would be pleased to win.

Deciding how much to scale down the bid is a hard problem because the amount depends on how much all the other players scale down. In a second-price auction a player calculates the value of $\tilde{V}_1$ using equation (34), but that equation hides considerable complexity under the disguise of the term $b(\hat{V}_2)$, which is itself calculated as a function of $b(\hat{V}_1)$ using an equation like (34).⁹

Oil Tracts and the Winner's Curse

The best known example of the winner's curse is from bidding for offshore oil tracts. Offshore drilling can be unprofitable even if oil is discovered, because something must be paid to the government for the mineral rights. Capen, Clapp & Campbell (1971) suggest that bidders' ignorance of the winner's curse caused overbidding in US government auctions of the 1960s. If the oil companies had bid close to what their engineers estimated the tracts were worth, rather than scaling down their bids, the winning companies would have lost on their investments. The hundredfold difference in the sizes of the bids in the sealed-bid auctions shown in Table 1 lends some plausibility to the view that this is what happened.

Table 1 Bids by serious competitors in oil auctions

⁹xxxx Here it would be nice to do a numerical examlpe of how much to scale down. Nobody seems to have one, thoguh— not Milgrom, Cramton, Dixit-Skeath, Myerson, Gintis. Some of them use the Wallet Game, which might be worth adding. Simply using the uniform distirubtion for f is tough because order statistics are tricky— this involves finding the expected value, an integral, given that I have valuation x_1 and the second-highest valuation is x_2 . And the first-rpice bidding function is just characetrizied by a differnetial equation, not found by anyone. Deriving the advantage of Enlgish over second-price over first-rpice is easier, and might be worth doing.

Offshore Louisiana 1967	Santa Barbara Channel 1968	Offshore Texas 1968	Alaska North Slope 1969
Tract SS 207	Tract 375	Tract 506	Tract 253
32.5	43.5	43.5	10.5
17.7	32.1	15.5	5.2
11.1	18.1	11.6	2.1
7.1	10.2	8.5	1.4
5.6	6.3	8.1	0.5
4.1		5.6	0.4
3.3		4.7	
		2.8	
		2.6	
		0.7	
		0.7	
		0.4	

Later studies such as Mead, Moseidjord & Sorason (1984) that actually looked at profitability conclude that the rates of return from offshore drilling were not abnormally low, so perhaps the oil companies did scale down their bids rationally. The spread in bids is surprisingly wide, but that does not mean that the bidders did not properly scale down their estimates. Although expected profits are zero under optimal bidding, realized profits could be either positive or negative. With some probability, one bidder makes a large overestimate which results in too high a bid even after rationally adjusting for the winner's curse. The knowledge of how to bid optimally does not eliminate bad luck; it only mitigates its effects.

Another consideration is the rationality of the other bidders. If bidder Apex has figured out the winner's curse, but bidders Brydox and Central have not, what should Apex do? Its rivals will overbid, which affects Apex's best response. Apex should scale down its bid even further than usual, because the winner's curse is intensified against overoptimistic rivals. If Apex wins against a rival who usually overbids, Apex has very likely overestimated the value.

Risk aversion affects bidding in a surprisingly similar way. If all the players are equally risk averse, the bids would be lower, because the asset is a gamble, whose value is lower for the risk averse. If Smith is more risk averse than Brown, then Smith should be more cautious for two reasons. The direct reason is that the gamble is worth less to Smith. The indirect reason is that when Smith wins against a rival like Brown who regularly bids more, Smith probably overestimated the value. Parallel reasoning holds if the players are risk neutral, but the private value of the object differs among them.

Asymmetric equilibria can even arise when the players are identical. Second-price, two-person, common-value auctions usually have many asymmetric equilibria besides the symmetric equilibrium we have been discussing (see Milgrom [1981c] and Bikhchandani [1988]). Suppose that Smith and Brown have identical payoff functions, but Smith thinks Brown is going to bid aggressively. The winner's curse is intensified for Smith, who would probably have overestimated if he won against an aggressive bidder like Brown, so Smith bids more cautiously. But if Smith bids cautiously, Brown is safe in bidding aggressively, and there is an asymmetric equilibrium. For this reason, acquiring a reputation for aggressiveness is valuable.

Oddly enough, if there are three or more players the sealed-bid, second-price, common-value auction has a unique equilibrium, which is also symmetric. The open-exit auction is different: it has asymmetric equilibria, because after one bidder drops out, the two remaining bidders know that they are alone together in a subgame which is a two-player auction. Regardless of the number of players, first-price sealed-bid auctions do not have this kind of asymmetric equilibrium. Threats in a first-price auction are costly because the high bidder pays his bid even if his rival decides to bid less in response. Thus, a bidder's aggressiveness is not made safer by intimidation of another bidder.

The winner's curse crops up in situations seemingly far removed from auctions. An employer must beware of hiring a worker passed over by other employers. Someone renting an apartment must hope that he is not the first visitor who arrived when the neighboring trumpeter was asleep. A firm considering a new project must worry that the project has been considered and rejected by competitors. The winner's curse can even be applied to political theory, where certain issues keep popping up. Opinions are like estimates, and one interpretation of different valuations is that everyone gets the same data, but they analyze it differently.

On a more mundane level, in 1987 there were four major candidates – Bush, Kemp, Dole, and Other – running for the Republican nomination for President of the United States. Consider an entrepreneur auctioning off four certificates, each paying one dollar if its particular candidate wins the nomination. If every bidder is rational, the entrepreneur should receive a maximum of one dollar in total revenue from these four auctions, and less if bidders are risk averse. But holding the auction in a bar full of partisans, how much do you think he would actually receive?

Wallet Game

Players

Smith and Jones.

Order of Play

(0) Nature chooses the amounts s_1 and s_2 of the money in Smith's wallet and Jones's, using density functions $f_1(s_1)$ and $f_2(s_2)$. Each player observes only his own wallet contents.

(1) Each player chooses a bid ceiling p_1 or p_2 . An auctioneer auctions off the two wallets by gradually raising the price until either p_1 or p_2 is reached and one player exits.

Payoffs:

The player that exits first gets zero. The winning player has a payoff of $(s_1 + s_2 - \text{Min}(p_1, p_2))$.

One equilibrium is for bidder i to choose bid ceiling $p_i = 2s_i$. This is an equilibrium because if he wins at that price, the value of the wallets is at least $2s_i$, since player j 's signal must be $s_j = s_i$. If he were to choose a ceiling any lower, then he might pass up a chance to get the wallet at a price less than its value; if he chooses a ceiling any higher, the other player might drop out and i would overpay.

There are other equilibria, though—asymmetric ones.

In general, asymmetric equilibria are common in common-value auctions. That is because the severity of the winner's curse facing player i ¹⁰ depends on the bidding behavior of the other players. If other players bid aggressively, then if i wins anyway, he must have a big overestimate of the value of the object. So the more aggressive are the other players, the more conservative ought i to be—which in turn will make the other players more aggressive.

Here, another equilibrium is $(p_1 = 10s_1, p_2 = \frac{10}{9}s_2)$. If the two players tie, having chosen $p = p_1 = p_2$, then $10s_1 = \frac{10}{9}s_2$, which implies that $s_1 = \frac{1}{9}s_2$, which implies that $s_1 + s_2 = 10s_1 = p$, and $v = p$. This has a worse equilibrium payoff for bidder 2, because he hardly ever wins, and when he does win it is because s_1 was very low—so there is hardly any money in Bidder 1's wallet. Being the aggressive bidder in an equilibrium is valuable. If there is a sequence of auctions, this means establishing a reputation for aggressiveness can be worthwhile (see Bikchandani [1988]).

Common Values

Milgrom and Weber found that when there is a common value element in an auction game (“affiliated values”), then the ranking of seller revenues is:

English: best

2nd-price sealed bid: next best

1st price sealed-bid and Dutch: worst and identical

It is actually hard to come up with classroom examples for common value auctions. Paul Klemperer has done so, however, at p. 69 of his 1999 survey.

¹⁰xxx A term invented by Robert Wilson in the 60's, I think. Find source.

Suppose that n signals are independently drawn from the uniform distribution on $[\underline{s}, \bar{s}]$. Note that the expectation of the k th highest value is

$$Es_{(k)} = \underline{s} + \left(\frac{n+1-k}{n+1} \right) (\bar{s} - \underline{s}) \quad (36)$$

In particular, this means the expectation of the second-highest value is

$$Es_{(2)} = \underline{s} + \left(\frac{n+1-2}{n+1} \right) (\bar{s} - \underline{s}) = \underline{s} + \left(\frac{n-1}{n+1} \right) (\bar{s} - \underline{s}) \quad (37)$$

and the expectation of the lowest value is

$$Es_{(n)} = \underline{s} + \left(\frac{n+1-n}{n+1} \right) (\bar{s} - \underline{s}) = \underline{s} + \left(\frac{1}{n+1} \right) (\bar{s} - \underline{s}). \quad (38)$$

Suppose n risk-neutral bidders, $i = 1, 2, \dots, n$ each receive a signal s_i independently drawn from the uniform distribution on $[v - m, v + m]$, where v is the true value of the object to each of them. Assume that they have “diffuse priors” on v , which means they think any value is equally likely.

Denote the j^{th} highest signal by $s_{(j)}$. Note that

$$Ev|(s_1, s_2, \dots, s_n) = \frac{s_{(n)} + s_{(1)}}{2}. \quad (39)$$

This is a remarkable property of the uniform distribution: if you observe signals 6, 7, 11, and 24, the expected value of the object is 15 ($= [6+24]/2$), well above the mean of 12 and the median of 9, because only the extremes of 6 and 24 are useful information. A density that had a peak, like the normal density, would yield a different result, but here all we can tell from the data is that all values of v between $(6 + m)$ and $(24 - m)$ are equally probable.

Figure 3 illustrates this. Someone who saw just signals $s_{(n)}$ and $s_{(1)}$ could deduce that v could not be less than $(s_{(1)} - m)$ or greater than $(s_{(n)} + m)$. Learning $s_{(2)}$, for example, would be unhelpful, because the only information it conveys is that $v \leq (s_{(2)} + m)$ and $v \geq (s_{(2)} - m)$, and our observer has already figured that out.




Figure 3: Extracting Information From Uniformly Distributed Signals

What are the strategies in symmetric equilibria for the different auction rules? (We will ignore possible asymmetric equilibria.)

English, ascending, open-cry, open-exit auction

Equilibrium: If nobody else has quit yet, drop out when the price rises to s_i . Otherwise, drop out when the price rises to $p_i = \frac{p_{(n)} + s_i}{2}$, where $p_{(n)}$ is the price at which the first dropout occurred.

If nobody else has quit yet, then bidder i is safe in agreeing to pay the price. Either (a) he has the lowest signal, and will lose the auction, or (b) everybody else has signal s_i too, and they will all drop out at the same time, or (c) he will never drop out, and he will win. In case (b), his estimate of the value is s_i , and that is where he should drop out.

Once one person has dropped out at $p_{(n)}$, the other bidders can guess that he had the lowest signal, so they know that signal $s_{(n)}$ must equal $p_{(n)}$. Suppose bidder i has signal $s_i > s_{(n)}$. Either (a) someone else has a higher signal and bidder i will lose the auction, or (b) everybody else who has not yet dropped out has signal s_i too, and they will all drop out at the same time, or (c) he will never drop out and he will win. In case (b), his estimate of the value is $p_{(i)} = \frac{p_{(n)} + s_i}{2}$, since $p_{(n)}$ and s_i are the extreme signal values, and that is where he should drop out.

The price paid by the winner will be the price at which the second- highest bidder drops out, which is $\frac{s_{(n)}+s_{(2)}}{2}$. These expected values are

$$Es_{(n)} = (v - m) + \left(\frac{n+1-n}{n+1} \right) ((v+m) - (v-m)) = v + \left(\frac{2-n}{n+1} \right) 2m. \quad (40)$$

and

$$Es_{(2)} = (v - m) + \left(\frac{n+1-2}{n+1} \right) ((v+m) - (v-m)) = v + \left(\frac{n-3}{n+1} \right) 2m. \quad (41)$$

Averaging them yields the expected winning price,

$$Ep_{(2)} = \frac{\left[v + \left(\frac{2-n}{n+1} \right) 2m \right] + \left[v + \left(\frac{n-3}{n+1} \right) 2m \right]}{2} = v - \frac{1}{2} \left(\frac{1}{n+1} \right) 2m. \quad (42)$$

Notice that the bigger m is, the lower the expected seller revenue. Also notice that the higher is n , the greater is the expected seller revenue. This will be true for all three auction rules we examine here.

2nd-Price Sealed-bid Auction

Equilibrium: Bid $p_i = s_i - \left(\frac{n-2}{n} \right) m$.

Bidder i thinks of himself as being tied for winner with one other bidder, and so having to pay his bid. So he thinks he is the highest of $(n-1)$ bidders drawn from $[v-m, v+m]$ and tied with one other, so on average if this happens, $s_i = (v-m) + \left(\frac{[(n-1)+1-1]}{[n-1]+1} \right) ([v+m] - [v-m]) = (v-m) + \left(\frac{n-1}{n} \right) (2m) = v + \frac{n-2}{n}(m)$. He will bid this value, which is, solving for v , $p_i = s_i - \left(\frac{n-2}{n} \right) (m)$.

On average, the second-highest bidder actually has the signal $Es_{(2)} = v + \left(\frac{n-3}{n+1} \right) m$, as found earlier. So the expected price, and hence the expected revenue from the auction is

$$Ep_{(2)} = \left[v + \left(\frac{n-3}{n+1} \right) m \right] - \left(\frac{n-2}{n} \right) (m) = v + \left(\frac{n(n-3) - (n+1)(n-2)}{(n+1)n} \right) m, \quad (43)$$

which equals

$$v - \left(\frac{n-1}{n} \right) \left(\frac{1}{n+1} \right) 2m. \quad (44)$$

Note that in this example, the expected revenue is lower. Why? It is because bidder 2 does not know the lower bound is so low when he makes his bid. He has to guess at the lower bound.

1st-price sealed-bid auction, Dutch descending auction

Equilibrium: Bid $(s_i - m)$.

Bidder i bids $(s_i - z)$ for some amount z that does not depend on his signal, because given the assumption of diffuse priors, he does not know whether his signal is a high one or a low one.¹¹ Define T so that $s_i \equiv v - m + T$. Bidder i has the highest signal and wins the auction if T is big enough, which has probability $\left(\frac{T}{2m}\right)^{n-1}$, because it is the probability that the $(n - 1)$ other signals are all less than $(v - m + T)$. He earns v minus his bid of $(s_i - z)$ if he wins, which equals $(z + m - T)$. If, instead, he deviates and bids a small amount ϵ higher, he would win $(z + m - (T - \epsilon))$ with additional probability. Using a Taylor expansion ($g(T + \epsilon) \approx g(T) + g'(T)\epsilon$) tells us that

$$\left(\frac{T + \epsilon}{2m}\right)^{n-1} - \left(\frac{T}{2m}\right)^{n-1} \approx (n - 1)T^{n-2} \left(\frac{1}{2m}\right)^{n-1} \epsilon. \quad (45)$$

The disadvantage of bidding higher is that Bidder i would pay an additional ϵ in the $\left(\frac{T}{2m}\right)^{n-1}$ cases in which he would have won anyway. In equilibrium, he is indifferent about this small deviation,¹² so

$$\int_{T=0}^{2m} \left[\left((n - 1)T^{n-2} \left(\frac{1}{2m}\right)^{n-1} \epsilon \right) (z + m - T) - \epsilon \left(\frac{T}{2m}\right)^{n-1} \right] dT = 0. \quad (46)$$

This implies that

$$\left(\frac{\epsilon}{2m}\right)^{n-1} \int_{T=0}^{2m} \left[\left((n - 1)T^{n-2} \right) (z + m) - (n - 1)T^{n-1} - T^{n-1} \right] dT = 0. \quad (47)$$

which in turn implies that

$$\left(\frac{\epsilon}{2m}\right)^{n-1} \left|_{T=0}^{2m} T^{n-1} (z + m) - T^n = 0, \quad (48)$$

so $(2m)^{n-1}(z + m) - (2m)^n - 0 + 0 = 0$ and $z = m$. Bidder i 's optimal strategy in the symmetric equilibrium is to bid $p_i = s_i - m$. The winning bid is set by the bidder with the highest signal, and that highest signal's expected value is

$$\begin{aligned} Es_{(1)} &= \underline{s} + \left(\frac{n+1-1}{n+1}\right) (\bar{s} - \underline{s}) \\ &= v - m + \left(\frac{n}{n+1}\right) ((v + m) - (v - m)) \\ &= v - m + \left(\frac{n}{n+1}\right) (2m) \end{aligned} \quad (49)$$

The expected revenue is therefore

$$Ep_{(1)} = v - (1) \left(\frac{1}{n+1}\right) 2m. \quad (50)$$

Here, the revenue is even lower than under the first two auction rules.

¹¹xxx This does not explain why he does not, for example, shrink his signal, bidding zs_i . Think about that.

¹²xxx why? Property of continuous density?

A Mechanism to Extract All the Surplus (see Myerson [1981])

Ask bidder i to declare s_i , allocate the good to the high bidder at the price $\frac{s_{(1)} + s_{(n)}}{2}$, which is an unbiased estimate of v , and ensure truthtelling by the boiling-in-oil punishment of additional transfers of $t = -\infty$ if the reports are such that $s_{(n)} + m < s_{(1)}$, which cannot possibly occur if all bidders tell the truth.

Matthews (1987) takes the buyer's viewpoint and show that buyers with increasing absolute risk aversion and private values prefer first-price auctions to second-price, even though the prices are higher, because they are also less risky.

Myerson (1981) shows that if the bidders' private information is correlated, the seller can construct a mechanism that extracts all the information and all the surplus.

Bidders' signals are **affiliated** if a high value of one bidder's signal makes high values of the other bidders' signals more likely, roughly.

13.5 Information in Common-Value Auctions

The Seller's Information

Milgrom & Weber (1982) have found that honesty is the best policy as far as the seller is concerned. If it is common knowledge that he has private information, he should release it before the auction. The reason is not that the bidders are risk averse (though perhaps this strengthens the result), but the "No news is bad news" result of section 8.1. If the seller refuses to disclose something, buyers know that the information must be unfavorable, and an unravelling argument tells us that the quality must be the very worst possible.

Quite apart from unravelling, another reason to disclose information is to mitigate the winner's curse, even if the information just reduces uncertainty over the value without changing its expectation. In trying to avoid the winner's curse, bidders lower their bids, so anything which makes it less of a danger raises their bids.

Asymmetric Information among the Buyers

Suppose that Smith and Brown are two of many bidders in a common-value auction. If Smith knows he has uniformly worse information than Brown (that is, if his information partition is coarser than Brown's), he should stay out of the auction: his expected payoff is negative if Brown expects zero profits.

If Smith's information is not uniformly worse, he can still benefit by entering the auction. Having independent information, in fact, is more valuable than having good information. Consider a common-value, first-price, sealed-bid auction with four bidders. Bidders Smith and Black have the same good information, Brown has that same information plus an extra signal, and Jones usually has only a poor estimate, but one different from any other bidder's. Smith and Black should drop out of the auction – they can never beat Brown without overpaying. But Jones will sometimes win, and his expected surplus is positive. If, for example, real estate tracts are being sold, and Jones is quite ignorant of land values, he can still do well if, on rare occasions, he has inside information concerning the location of a new freeway, even though ordinarily he should refrain from bidding. If Smith and Black both use the same appraisal formula, they will compete each other's profits away, and if Brown uses the formula plus extra private information, he drives their profits negative by taking some of the best deals from them and leaving the worst ones.

In general, a bidder should bid less if there are more bidders or his information is absolutely worse (that is, if his information partition is coarser). He should also bid less if parts of his information partition are coarser than those of his rivals, even if his information is not uniformly worse. These considerations are most important in sealed-bid auctions, because in an open-cry auction information is revealed by the bids while other bidders still have time to act on it.

N13.1 Auction classification and private-value strategies

- McAfee & McMillan (1987) and Milgrom (1987) are excellent surveys of the literature and theory of auctions. Both articles take some pains to relate the material to models of asymmetric information. More recent is Klemperer (1999). Milgrom & Weber (1982) is a classic article that covers many aspects of auctions. Paul Milgrom's consulting firm, Agora Market Design, has a website with many good working papers that can be found via www.market-design.com. Klemperer (2000) collects many of the most important published papers.
- Cassady (1967) is an excellent source of institutional detail on auctions. The appendix to his book includes advertisements and sets of auction rules, and he cites numerous newspaper articles. See also *The New York Times*, (July 26, 1985) and pp. A1, C20 (July 31, 1985); *The Wall Street Journal*, pp. 1, 16 (August 24, 1984); and "The Crackdown on Colluding Roadbuilders," *Fortune*, p. 79 (October 3, 1983).
- Bargaining and auctions are two extremes in ways to sell goods. In between are various mixtures such as bargaining with an outside option of going to someone else, auctions with reserve prices, and so forth. For one readable comparison of the two sale methods, see Bulow & Klemperer (1996).
- **The Dollar Auction.** Auctions look like tournaments in that the winner is the player who chooses the largest amount for some costly variable, but in auctions the losers generally do not incur costs proportional to their bids. Shubik (1971), however, has suggested an auction for a dollar bill in which both the first- and the second-highest bidders pay the second price. If both players begin with infinite wealth, the game illustrates why equilibrium might not exist if strategy sets are unbounded. Once one bidder has started bidding against another, both of them do best by continuing to bid so as to win the dollar as well as pay the bid. This auction may seem absurd, but it has some similarity to patent races (see section 15.1) and arms races.

N13.2 Comparing auction rules

- Many (all?) leading auction theorists were involved in the seven- billion dollar spectrum auction by the United States government in 1994, either helping the government sell spectrum or helping bidders decide how to buy it. Paul Milgrom's 1999 book, *Auction Theory for Privatization*, tells the story. See also McAfee & McMillan (1996). Interesting institutional details have come in the spectrum auctions and stimulated new theoretical research. Ayres & Cramton (1996), for example, explore the possibility that affirmative action provisions designed to help certain groups of bidders may have actually increased the revenue raised by the seller by increasing the amount of competition in the auction.
- One might think that an ascending second-price, open-cry auction would come to the same results as an ascending first-price, open-cry auction, because if the price advances by ϵ at each bid, the first and second bids are practically the same. But the second-price auction can be manipulated. If somebody initially bids \$10 for something worth \$80, another bidder could safely bid \$1,000. No one else would bid more, and he would pay only the second price: \$10.
- In one variant of the English auction, the auctioneer announces each new price and a bidder can hold up a card to indicate he is willing to bid that price. This set of rules is more practical to administer in large crowds and it also allows the seller to act strategically during the course of the auction. If, for example, the two highest valuations are 100 and 130, this kind of auction could yield a price of 110, while the usual rules would only allow a price of $100 + \epsilon$.
- Vickrey (1961) notes that a Dutch auction could be set up as a second-price auction. When the first bidder presses his button, he primes a buzzer that goes off when a second bidder presses a button.
- Auctions are especially suitable for empirical study because they are so stylized and generate masses of data. Hendricks & Porter (1988) is the best-known comparison of auction theory with data. Tenorio (1993) is another nice example of empirical work using data from real auctions, in his case, the Zambian foreign exchange market. See Laffont (1997) for a survey of empirical work.

- Second-price auctions have actually been used in a computer operating system. An operating system must assign a computer's resources to different tasks, and researchers at Xerox Corporation designed the Spawn system, under which users allocate "money" in a second-price sealed bid auction for computer resources. See "Improving a Computer Network's Efficiency," *The New York Times*, (March 29, 1989) p. 35.
- After the last bid of an open-cry art auction in France, the representative of the Louvre has the right to raise his hand and shout "pre-emption de l'etat," after which he takes the painting at the highest price bid (*The Economist*, May 23, 1987, p. 98). How does that affect the equilibrium strategies? What would happen if the Louvre could resell?
- **Share Auctions.** In a share auction each buyer submits a bid for both a quantity and a price. The bidder with the highest price receives the quantity for which he bid at that price. If any of the product being auctioned remains, the bidder with the second-highest price takes the quantity he bid for, and so forth. The rules of a share auction can allow each buyer to submit several bids, often called a **schedule** of bids. The details of share auctions vary, and they can be either first-price or second-price. Models of share auctions are very complicated; see Wilson (1979).
- **Reserve prices.** A reserve price is one below which the seller refuses to sell. Reserve prices can increase the seller's revenue, and their effect is to make the auction more like a regular fixed-price market. For discussion, see Milgrom & Weber (1982). They are also useful when buyers collude, a situation of bilateral monopoly. See "At Many Auctions, Illegal Bidding Thrives as a Longtime Practice Among Dealers," *Wall Street Journal*, February 19, 1988 p. 21. In some real-world English auctions, the auctioneer does not announce the reserve price in advance, and he starts the bidding below it. This can be explained as a way of allowing bidders to show each other that their valuations are greater than the starting price, even though it may turn out that they are all lower than the reserve price.
- Concerning auctions with risk-averse players, see Maskin & Riley (1984).

- Che & Gale (1998) point out that if bidders differ in their willingness to pay in a private value auction because of budget constraints rather than tastes then the revenue equivalence theorem can fail. The following example from page 2 of their paper shows this. Suppose two budget-constrained bidders are bidding for one object. In auction 1, each buyer has a budget of 2 and knows only his own value, which is drawn uniformly from $[0,1]$. The budget constraints are never binding, and it turns out that the expected price is $1/3$ under either a first-price or a second-price auction. In auction 2, however, each buyer knows only his own budget, which is drawn uniformly from $[0,1]$, and both have values for the object of 2. The budget constraint is always binding, and the equilibrium strategy is to bid one's entire budget under either set of auction rules. The expected price is still $1/3$ in the second-price auction, but now it is $2/3$ in the first-price auction. The seller therefore prefers to use a first-price auction.

N13.3 Common-value auctions and the winner's curse

- The winner's curse and the idea of common values versus private values have broad application. The winner's curse is related to the idea of regression to the mean discussed in section 2.4. Kaplow & Shavell (1996) use the idea to discuss property versus liability rules, one of the standard rule choices in law-and-economics. If someone violates a property rule, the aggrieved party can undo the violation, as when a thief is required to surrender stolen property. If someone violates a liability rule, the aggrieved party can only get monetary compensation, as when someone who breaches a contract is required to pay damages to the aggrieved party. Kaplow and Shavell argue that if a good has independent values, a liability rule is best because it gives efficient incentives for rule violation; but if it has common value and courts make errors in measuring the common value, a property rule may be better. See especially around page 761 of their article.

N13.4 Information in common-value auctions

- Even if valuations are correlated, the optimal bidding strategies can still be the same as in private-value auctions if the values are independent. If everyone overestimates their values by 10 percent, a player can still extract no information about his value by seeing other players' valuations.

- “Getting carried away” may be a rational feature of a common-value auction. If a bidder has a high private value and then learns from the course of the bidding that the common value is larger than he thought, he may well end up paying more than he had planned, although he would not regret it afterwards. Other explanations for why bidders seem to pay too much are the winner’s curse and the fact that in every auction all but one or two of the bidders think that the winning bid is greater than the value of the object.
- Milgrom & Weber (1982) use the concept of **affiliated** variables in classifying auctions. Roughly speaking, random variables X and Y are affiliated if a larger value of X means that a larger value of Y is more likely, or at least, no less likely. Independent random variables are affiliated.

Problems

13.1. Rent Seeking

Two risk-neutral neighbors in sixteenth century England, Smith and Jones, have gone to court and are considering bribing a judge. Each of them makes a gift, and the one whose gift is the largest is awarded property worth £2,000. If both bribe the same amount, the chances are 50 percent for each of them to win the lawsuit. Gifts must be either £0, £900, or £2,000.

- (a) What is the unique pure-strategy equilibrium for this game?
- (b) Suppose that it is also possible to give a £1500 gift. Why does there no longer exist a pure-strategy equilibrium?
- (c) What is the symmetric mixed-strategy equilibrium for the expanded game? What is the judge's expected payoff?
- (d) In the expanded game, if the losing litigant gets back his gift, what are the two equilibria? Would the judge prefer this rule?

13.2. The Founding of Hong Kong

The Tai-Pan and Mr. Brock are bidding in an English auction for a parcel of land on a knoll in Hong Kong. They must bid integer values, and the Tai-Pan bids first. Tying bids cannot be made, and bids cannot be withdrawn once they are made. The direct value of the land is 1 to Brock and 2 to the Tai-Pan, but the Tai-Pan has said publicly that he wants it, so if Brock gets it, he receives 5 in “face” and the Tai-Pan loses 10. Moreover, Brock hates the Tai-Pan and receives 1 in utility for each 1 that the Tai-Pan pays out to get the land.

- (a) First suppose there were no “face” or “hate” considerations, just the direct values. What are the equilibria if the Tai-pan bids first?
- (b) Continue supposing there were no “face” or “hate” considerations, just the direct values. What are the three possible equilibria if Mr. Brock bids first? (Hint: in one of them, Brock wins; in the other two, the Tai-pan wins.)
- (c) Now fill in the entries in Table 13.2.

Table 13.2: The Tai-Pan Game

Winning bid:	1	2	3	4	5	6	7	8	9	10	11	12
If Brock wins:												
π_{Brock}												
$\pi_{Tai-Pan}$												
If Brock loses:												
π_{Brock}												
$\pi_{Tai-Pan}$												

- (d) In equilibrium, who wins, and at what bid?
- (e) What happens if the Tai-Pan can precommit to a strategy?
- (f) What happens if the Tai-Pan cannot precommit, but he also hates Brock, and gets 1 in utility for each 1 that Brock pays out to get the land?

13.3. Government and Monopoly

Incumbent Apex and potential entrant Brydoux are bidding for government favors in the widget market. Apex wants to defeat a bill that would require it to share its widget patent rights with Brydoux. Brydoux wants the bill to pass. Whoever offers the chairman of the House Telecommunications Committee more campaign contributions wins, and the loser pays nothing. The market demand curve is $P = 25 - Q$, and marginal cost is constant at 1.

- (a) Who will bid higher if duopolists follow Bertrand behavior? How much will the winner bid?
- (b) Who will bid higher if duopolists follow Cournot behavior? How much will the winner bid?
- (c) What happens under Cournot behavior if Apex can commit to giving away its patent freely to everyone in the world if the entry bill passes? How much will Apex bid?

13.4. An Auction with Stupid Bidders

Smith's value for an object has a private component equal to 1 and component common with Jones and Brown. Jones's and Brown's private components both equal zero. Each player estimates the common component Z independently, and player i 's estimate is either x_i above the true value or x_i below, with equal probability. Jones and Brown are naive and always bid their valuations. The auction is English. Smith knows X_i , but not whether his estimate is too high or too low.

- (a) If $x_{Smith} = 0$, what is Smith's dominant strategy if his estimate of Z equals 20?
- (b) If $x_i = 8$ for all players and Smith estimates $Z = 20$, what are the probabilities that he puts on different values of Z ?
- (c) If $x_i = 8$ but Smith knows that $Z = 13$ with certainty, what are the probabilities he puts on the different combinations of bids by Jones and Brown?
- (d) Why is 9 a better upper limit on bids for Smith than 21, if his estimate of Z is 20, and $x_i = 8$ for all three players?
- (e) Suppose Smith could pay amount 0.001 to explain optimal bidding strategy to his rival bidders, Jones and Brown. Would he do so?

13.5. A Teapot Auction with Incomplete Information

Smith believes that Brown's value v_b for a teapot being sold at auction is 0 or 100 with equal probability. Smith's value of $v_s = 400$ is known by both players.

- (a) What are the players' equilibrium strategies in an open cry auction? You may assume that in case of ties, Smith wins the auction.
- (b) What are the players' equilibrium strategies in a first-price sealed-bid auction? You may

14 Pricing

January 17, 2000. March 6, 2003. December 19, 2003. Eric Rasmusen, Erasmuse@indiana.edu. [Http://mypage.iu.edu/~erasmuse](http://mypage.iu.edu/~erasmuse). Footnotes starting with xxx are the author's notes to himself. Comments welcomed.

14.1 Quantities as Strategies: Cournot Equilibrium Revisited

Chapter 14 is about how firms with market power set prices. Section 14.1 generalizes the Cournot Game of Section 3.5 in which two firms choose the quantities they sell, while Section 14.2 sets out the Bertrand model of firms choosing prices.¹ Section 14.3 goes back to the origins of product differentiation, and develops two Hotelling location models. Section 14.4 shows how to do comparative statics in games, using the differentiated Bertrand model as an example and supermodularity and the implicit function theorem as tools. Section 14.5 shows that even if a firm is a monopolist, if it sells a durable good it suffers competition from its future self.

Cournot Behavior with General Cost and Demand Functions

In the next few sections, sellers compete against each other while moving simultaneously. We will start by generalizing the Cournot Game of Section 3.5 from linear demand and zero costs to a wider class of functions. The two players are firms Apex and Brydax, and their strategies are their choices of the quantities q_a and q_b . The payoffs are based on the total cost functions, $c(q_a)$ and $c(q_b)$, and the demand function, $p(q)$, where $q = q_a + q_b$. This specification says that only the sum of the outputs affects the price. The implication is that the firms produce an identical product, because whether it is Apex or Brydax that produces an extra unit, the effect on the price is the same.

Let us take the point of view of Apex. In the Cournot-Nash analysis, Apex chooses its output of q_a for a given level of q_b as if its choice did not affect q_b . From its point of view, q_a is a function of q_b , but q_b is exogenous. Apex sees the effect of its output on price as

$$\frac{\partial p}{\partial q_a} = \frac{dp}{dq} \frac{\partial q}{\partial q_a} = \frac{dp}{dq}. \quad (1)$$

Apex's payoff function is

$$\pi_a = p(q)q_a - c(q_a). \quad (2)$$

To find Apex's reaction function, we differentiate with respect to its strategy to obtain

$$\frac{d\pi_a}{dq_a} = p + \frac{dp}{dq}q_a - \frac{dc}{dq_a} = 0, \quad (3)$$

which implies

$$q_a = \frac{\frac{dc}{dq_a} - p}{\frac{dp}{dq}}, \quad (4)$$

¹xxxx This intro needs reworking because of moved sections.

or, simplifying the notation,

$$q_a = \frac{c' - p}{p'}. \quad (5)$$

If particular functional forms for $p(q)$ and $c(q_a)$ are available, equation (5) can be solved to find q_a as a function of q_b . More generally, to find the change in Apex's best response for an exogenous change in Brydax's output, differentiate (5) with respect to q_b , remembering that q_b exerts not only a direct effect on $p(q_a + q_b)$, but possibly an indirect effect via q_a .

$$\frac{dq_a}{dq_b} = \frac{(p - c')(p'' + p'' \frac{dq_a}{dq_b})}{p'^2} + \frac{c'' \frac{dq_a}{dq_b} - p' - p' \frac{dq_a}{dq_b}}{p'}. \quad (6)$$

Equation (6) can be solved for $\frac{dq_a}{dq_b}$ to obtain the slope of the reaction function,

$$\frac{dq_a}{dq_b} = \frac{(p - c')p'' - p'^2}{2p'^2 - c''p' - (p - c')p''} \quad (7)$$

If both costs and demand are linear, as in section 3.5, then $c'' = 0$ and $p'' = 0$, so equation (7) becomes

$$\frac{dq_a}{dq_b} = -\frac{p'^2}{2p'^2} = -\frac{1}{2}. \quad (8)$$

The general model faces two problems that did not arise in the linear model: nonuniqueness and nonexistence. If demand is concave and costs are convex, which implies that $p'' < 0$ and $c'' > 0$, then all is well as far as existence goes. Since price is greater than marginal cost ($p > c'$), equation (7) tells us that the reaction functions are downward sloping, because $2p'^2 - c''p' - (p - c')p''$ is positive and both $(p - c')p''$ and $-p'^2$ are negative. If the reaction curves are downward sloping, they cross and an equilibrium exists, as was shown in Chapter 3's Figure 1 for the linear case represented by equation (8). We usually do assume that costs are at least weakly convex, since that is the result of diminishing or constant returns, but there is no reason to believe that demand is either concave or convex. If the demand curves are not linear, the contorted reaction functions of equation (7) might give rise to multiple Cournot equilibria as in the present chapter's Figure 1.




Figure 1: Multiple Cournot-Nash Equilibria

If demand is convex or costs are concave, so $p'' > 0$ or $c'' < 0$, the reaction functions can be upward sloping, in which case they might never cross and no equilibrium would exist. The problem can also be seen from Apex's payoff function, equation (2). If $p(q)$ is convex, the payoff function might not be concave, in which case standard maximization techniques break down. The problems of the general Cournot model teach a lesson to modellers: sometimes simple assumptions such as linearity generate atypical results.

Many Oligopolists³

Let us return to the simpler game in which production costs are zero and demand is linear. For concreteness, we will use the particular inverse demand function

$$p(q) = 120 - q. \quad (9)$$

Using (9), the payoff function, (2), becomes

$$\pi_a = 120q_a - q_a^2 - q_bq_a. \quad (10)$$

In section 3.5, firms picked outputs of 40 apiece given demand function (9). This generated a price of 40. With n firms instead of two, the demand function is

$$p\left(\sum_{i=1}^n q_i\right) = 120 - \sum_{i=1}^n q_i, \quad (11)$$

³xxx Use positive marginal costs throughout.

and firm j 's payoff function is

$$\pi_j = 120q_j - q_j^2 - q_j \sum_{i \neq j} q_i. \quad (12)$$

Differentiating j 's payoff function with respect to q_j yields

$$\frac{d\pi_j}{dq_j} = 120 - 2q_j - \sum_{i \neq j} q_i = 0. \quad (13)$$

The first step in finding the equilibrium is to guess that it is symmetric, so that $q_j = q_i$, ($i = 1, \dots, n$). This is an educated guess, since every player faces a first-order condition like (13). By symmetry, equation (13) becomes $120 - (n+1)q_j = 0$, so that

$$q_j = \frac{120}{n+1}. \quad (14)$$

Consider several different values for n . If $n = 1$, then $q_j = 60$, the monopoly optimum; and if $n = 2$ then $q_j = 40$, the Cournot output found in section 3.5. If $n = 5$, $q_j = 20$; and as n rises, individual output shrinks to zero. Moreover, the total output of $nq_j = \frac{120n}{n+1}$ gradually approaches 120, the competitive output, and the market price falls to zero, the marginal cost of production. As the number of firms increases, profits fall.

14.2 Prices as Strategies

Here we will explore the Bertrand model more.

Capacity Constraints: the Edgeworth Paradox

Let us start by altering the Bertrand model by constraining each firm to sell no more than $K = 70$ units. The industry capacity of 140 exceeds the competitive output, but do profits continue to be zero?

When capacities are limited we require additional assumptions because of the new possibility that a firm with a lower price might attract more consumers than it can supply. We need to specify a **rationing rule** telling which consumers are served at the low price and which must buy from the high-price firm. The rationing rule is unimportant to the payoff of the low-price firm, but crucial to the high-price firm. One possible rule is

Intensity rationing. *The consumers able to buy from the firm with the lower price are those who value the product most.*

The inverse demand function from equation (9) is $p = 120 - q$, and under intensity rationing the K consumers with the strongest demand buy from the low-price firm. Suppose

that Brydcox is the low-price firm, charging a price of 30, so that 90 consumers wish to buy from it, though only K can do so. The residual demand facing Apex is then

$$q_a = 120 - p_a - K. \quad (15)$$

This is the demand curve in Figure 2a.




Figure 2: Rationing Rules when $p_b = 30$, $p_a > 30$, and $K = 70$

Under **intensity rationing**, if $K = 70$ the payoff functions are

$$\pi_a = \begin{cases} p_a \cdot \text{Min}\{120 - p_a, 70\} & \text{if } p_a < p_b & (a) \\ \frac{p_a(120 - p_a)}{2} & \text{if } p_a = p_b & (b) \\ 0 & \text{if } p_a > p_b, p_b \geq 50 & (c) \\ p_a(120 - p_a - 70) & \text{if } p_a > p_b, p_b < 50 & (d) \end{cases} \quad (16)$$

Here is why equations (16c) and (16d) look the way they do. If Brydcox has the lower price, all consumers will want to buy from Brydcox if they buy at all, but only 70 will be able to. If Brydcox's price is more than 50, then less than 70 will want to buy at all, and so 0 consumers will be left for Apex – which is equation (16c). If Brydcox's price is less than 50, then Brydcox will sell 70 units, and the residual demand curve facing Apex is as in equation (15), yielding equation (16d).

The appropriate rationing rule depends on what is being modelled. Intensity rationing is appropriate if buyers with more intense demand make greater efforts to obtain low prices. If the intense buyers are wealthy people who are unwilling to wait in line, the least intense buyers might end up at the low-price firm which is the case of **inverse-intensity**

rationing. An intermediate rule is proportional rationing, under which every type of consumer is equally likely to be able to buy at the low price.

Proportional rationing. *Each consumer has the same probability of being able to buy from the low-price firm.*

Under proportional rationing, if $K = 70$ and 90 consumers wanted to buy from Brydax, $2/9 (= \frac{q(p_b) - K}{q(p_b)})$ of each type of consumer will be forced to buy from Apex (for example, $2/9$ of the type willing to pay 120). The residual demand curve facing Apex, shown in Figure 14.2b and equation (17), intercepts the price axis at 120, but slopes down at a rate three times as fast as market demand because there are only $2/9$ as many remaining consumers of each type.

$$q_a = (120 - p_a) \left(\frac{120 - p_b - K}{120 - p_b} \right) \quad (17)$$

The capacity constraint has a very important effect: $(0,0)$ is no longer a Nash equilibrium in prices. Consider Apex's best response when Brydax charges a price of zero. If Apex raises his price above zero, he retains most of his consumers (because Brydax is already producing at capacity), but his profits rise from zero to some positive number, regardless of the rationing rule. In any equilibrium, both players must charge prices within some small amount ϵ of each other, or the one with the lower price would deviate by raising his price. But if the prices are equal, then both players have unused capacity, and each has an incentive to undercut the other. No pure-strategy equilibrium exists under either rationing rule. This is known as the **Edgeworth paradox** after Edgeworth (1897, 1922).

Suppose that demand is linear, with the highest reservation price being $p = 100$ and the maximum market quantity $Q = 100$ at $p = 0$. Suppose also that there are two firms, Apex and Brydax, each having a constant marginal cost of 0 up to capacity of $Q = 80$ and infinity thereafter. We will assume intensity rationing of buyers.

Note that industry capacity of 160 exceeds market demand of 100 if price equals marginal cost. Note also that the monopoly price is 50, which with quantity of 50 yields industry profit of 2,500. But what will be the equilibrium?

Prices of $(p_a = 0, p_b = 0)$ are not an equilibrium. Apex's profit would be zero in that strategy combination. If Apex increased its price to 5, what would happen? Brydax would immediately sell $Q = 80$, and to the most intense 80 percent of buyers. Apex would be left with all the buyers between $p = 20$ and $p = 5$ on the demand curve, for $Q_a = 15$ and profit of $\pi_a = (5)(15) = 75$. So deviation by Apex is profitable. (Of course, $p = 5$ is not necessarily the most profitable deviation – but we do not need to check that; I looked for an *easy* deviation.)

Equal prices of (p_a, p_b) with $p_a = p_b > 0$ are not an equilibrium. Even if the price is close to 0, Apex would sell at most 50 units as its half of the market, which is less than its capacity of 80. Apex could deviate to just below p_b and have a discontinuous jump in sales for an increase in profit, just as in the basic Bertrand game.

Unequal prices of (p_a, p_b) are not an equilibrium. Without loss of generality, suppose $p_a > p_b$. So long as p_b is less than the monopoly price of 50, Brydax would deviate to a

new price even close to but not exceeding p_a . And this is not *just* the open-set problem. Once Brydoux is close enough to Apex, Apex would deviate by jumping to a price just below Brydoux.

If capacities are large enough, the Edgeworth paradox disappears. Consider capacities of 150 per firm, for example. The argument made above for why equal prices of 0 is not an equilibrium fails, because if Apex were to deviate to a positive price, Brydoux would be fully capable of serving the entire market, leaving Apex with no consumers.

If capacities are small enough, the Edgeworth paradox also disappears, but so does the Bertrand paradox. Suppose each firm has a capacity of 20. They each will choose to sell at a price of 60, in which case they will each sell 20 units, their entire capacities. Apex will have a payoff of 1,200. If Apex deviates to a lower price, it will not sell any more, so that would be unprofitable. If Apex deviates to a higher price, it will sell fewer, and since the monopoly price is 50, its profit will be lower; note that a price of 61 and a quantity of 19 yields profits of 1,159, for example.

We could have expanded the model to explain why the firms have small capacities by adding a prior move in which they choose capacity subject to a cost per unit of capacity, foreseeing what will happen later in the game.

A mixed-strategy equilibrium does exist, calculated using intensity rationing by Levitan & Shubik (1972) and analyzed in Dasgupta & Maskin (1986b). Expected profits are positive, because the firms charge positive prices. Under proportional rationing, as under intensity rationing, profits are positive in equilibrium, but the high-price firm does better with proportional rationing. The high-price firm would do best with **inverse-intensity rationing**, under which the consumers with the least intense demand are served at the low-price firm, leaving the ones willing to pay more at the mercy of the high-price firm.

Even if capacity were made endogenous, the outcome would be inefficient, either because firms would charge prices higher than marginal cost (if their capacity were low), or they would invest in excess capacity (even though they price at marginal cost).

14.3 Location Models

In Section 14.2 we analyzed the Bertrand model with differentiated products using demand functions whose arguments were the prices of both firms. Such a model is suspect because it is not based on primitive assumptions. In particular, the demand functions might not be generated by maximizing any possible utility function. A demand curve with a constant elasticity less than one, for example, is impossible because as the price goes to zero, the amount spent on the commodity goes to infinity. Also, demand curves (??) and (??) were restricted to prices below a certain level, and it would be good to be able to justify that restriction.

Location models construct demand functions like (??) and (??) from primitive assumptions. In location models, a differentiated product's characteristics are points in a space. If cars differ only in their mileage, the space is a one-dimensional line. If acceleration is also important, the space is a two-dimensional plane. An easy way to think about this approach is to consider the location where a product is sold. The product “gasoline sold at the corner of Wilshire and Westwood,” is different from “gasoline sold at the corner of Wilshire and Fourth.” Depending on where consumers live, they have different preferences over the two, but, if prices diverge enough, they will be willing to switch from one gas station to the other.

Location models form a literature in themselves. We will look at the first two models analyzed in the classic article of Hotelling (1929), a model of price choice and a model of location choice. Figure 4 shows what is common to both. Two firms are located at points x_a and x_b along a line running from zero to one, with a constant density of consumers throughout. In the Hotelling Pricing Game, firms choose prices for given locations. In the Hotelling Location Game, prices are fixed and the firms choose the locations.




Figure 4: Location Models

The Hotelling Pricing Game (Hotelling [1929])

Players

Sellers Apex and Brydox, located at x_a and x_b , where $x_a < x_b$, and a continuum of buyers indexed by location $x \in [0, 1]$.

The Order of Play

- 1 The sellers simultaneously choose prices p_a and p_b .
- 2 Each buyer chooses a seller.

Payoffs

Demand is uniformly distributed on the interval $[0,1]$ with a density equal to one (think of each consumer as buying one unit). Production costs are zero. Each consumer always buys, so his problem is to minimize the sum of the price plus the linear transport cost, which is θ per unit distance travelled.

$$\pi_{\text{buyer at } x} = V - \text{Min}\{\theta|x_a - x| + p_a, \theta|x_b - x| + p_b\}. \quad (18)$$

$$\pi_a = \begin{cases} p_a(0) = 0 & \text{if } p_a - p_b > \theta(x_b - x_a) & \text{(a)} \\ & \text{(Brydax captures entire market)} \\ p_a(1) = p_a & \text{if } p_b - p_a > \theta(x_b - x_a) & \text{(b)} \\ & \text{(Apex captures entire market)} \\ p_a(\frac{1}{2\theta} [(p_b - p_a) + \theta(x_a + x_b)]) & \text{otherwise (the market is divided)} & \text{(c)} \end{cases} \quad (19)$$

Brydax has analogous payoffs.

The payoffs result from buyer behavior. A buyer's utility depends on the price he pays and the distance he travels. Price aside, Apex is most attractive of the two sellers to the consumer at $x = 0$ ("consumer 0") and least attractive to the consumer at $x = 1$ ("consumer 1"). Consumer 0 will buy from Apex so long as

$$V - (\theta x_a + p) > V - (\theta x_b + p_b), \quad (20)$$

which implies that

$$p_a - p_b < \theta(x_b - x_a), \quad (21)$$

which yields payoff (19a) for Apex. Consumer 1 will buy from Brydax if

$$V - [\theta(1 - x_a) + p_a] < V - [\theta(1 - x_b) + p_b], \quad (22)$$

which implies that

$$p_b - p_a < \theta(x_b - x_a), \quad (23)$$

which yields payoff (19b) for Apex.

Very likely, inequalities (21) and (23) are both satisfied, in which case Consumer 0 goes to Apex and Consumer 1 goes to Brydax. This is the case represented by payoff (19c), and the next task is to find the location of consumer x^* , defined as the consumer who is at the boundary between the two markets, indifferent between Apex and Brydax. First, notice that if Apex attracts Consumer x_b , he also attracts all $x > x_b$, because beyond x_b the consumers' distances from both sellers increase at the same rate. So we know that if there is an indifferent consumer he is between x_a and x_b . Knowing this, (18) tells us that

$$V - [\theta(x^* - x_a) + p_a] = V - -[\theta(x_b - x^*) + p_b], \quad (24)$$

so that

$$p_b - p_a = \theta(2x^* - x_a - x_b), \quad (25)$$

and

$$x^* = \frac{1}{2\theta} [(p_b - p_a) + \theta(x_a + x_b)]. \quad (26)$$

Do remember that equation (26) is valid only if there really does exist a consumer who is indifferent – if such a consumer does not exist, equation (26) will generate a number for x^* , but that number is meaningless.

Since Apex keeps all the consumers between 0 and x^* , equation (26) is the demand function facing Apex so long as he does not set his price so far above Brydoux's that he loses even consumer 0. The demand facing Brydoux equals $(1 - x^*)$. Note that if $p_b = p_a$, then from (26), $x^* = \frac{x_a + x_b}{2}$, independent of θ , which is just what we would expect. Demand is linear in the prices of both firms, and looks similar to demand curves (??) and (??), which were used in Section 3.xxx for the Bertrand game with differentiated products.⁴

Now that we have found the demand functions, the Nash equilibrium can be calculated in the same way as in Section 14.2, by setting up the profit functions for each firm, differentiating with respect to the price of each, and solving the two first-order conditions for the two prices. If there exists an equilibrium in which the firms are willing to pick prices to satisfy inequalities (21) and (23), then it is

$$p_a = \frac{(2 + x_a + x_b)\theta}{3}, \quad p_b = \frac{(4 - x_a - x_b)\theta}{3}. \quad (27)$$

From (27) one can see that Apex charges a higher price if a large x_a gives it more safe consumers or a large x_b makes the number of contestable consumers greater. The simplest case is when $x_a = 0$ and $x_b = 1$, when (27) tells us that both firms charge a price equal to θ . Profits are positive and increasing in the transportation cost.

We cannot rest satisfied with the neat equilibrium of equation (27), because the assumption that there exists an equilibrium in which the firms choose prices so as to split the market on each side of some boundary consumer x^* is often violated. Hotelling did not notice this, and fell into a common mathematical trap. Economists are used to models in which the calculus approach gives an answer that is both the local optimum and the global optimum. In games like this one, however, the local optimum is not global, because of the discontinuity in the objective function. Vickrey (1964) and D'Aspremont, Gabszewicz & Thisse (1979) have shown that if x_a and x_b are close together, no pure-strategy equilibrium exists, for reasons similar to why none exists in the Bertrand model with capacity constraints. If both firms charge nonrandom prices, neither would deviate to a slightly different price, but one might deviate to a much lower price that would capture every single consumer. But if both firms charged that low price, each would deviate by raising his price slightly. It turns out that if, for example, Apex and Brydoux are located symmetrically around the center of the interval, $x_a \geq 0.25$, and $x_b \leq 0.75$, no pure-strategy equilibrium exists (although a mixed-strategy equilibrium does, as Dasgupta & Maskin [1986b] show).

Hotelling should have done some numerical examples. And he should have thought about the comparative statics carefully. Equation (27) implies that Apex should choose a higher price if both x_a and x_b increase, but it is odd that if the firms are locating closer

⁴xxx Make the link to Section 14.2 clearer.

together, say at 0.90 and 0.91, that Apex should be able to charge a higher price, rather than suffering from more intense competition. This kind of odd result is a typical clue that the result has a logical flaw somewhere. Until the modeller can figure out an intuitive reason for his odd result, he should suspect an error. For practice, let us try a few numerical examples, illustrated in Figure 5.




Figure 5: Numerical examples for Hotelling pricing

Example 1. Everything works out simply

Try $x_a = 0$, $x_b = 0.7$ and $\theta = 0.5$. Then equation (27) says $p_a = (2+0+0.7)0.5/3 = 0.45$ and $p_b = (4-0-0.7)0.5/3 = 0.55$. Equation (26) says that $x^* = \frac{1}{2 \cdot 0.5} [(0.55 - 0.45) + 0.5(0.0 + 0.7)] = 0.45$.

In Example 1, there is a pure strategy equilibrium and the equations generated sensible numbers given the parameters we chose. But it is not enough to calculate just one numerical example.

Example 2. Same location – but different prices?

Try $x_a = 0.9$, $x_b = 0.9$ and $\theta = 0.5$. Then equation (27) says $p_a = (2.0 + 0.9 + 0.9)0.5/3 \approx 0.63$ and $p_b = (4.0 - 0.9 - 0.9)0.5/3 \approx 0.37$.

Example 2 shows something odd happening. The equations generate numbers that seem innocuous until one realizes that if both firms are located at 0.9, but $p_a = 0.63$ and $p_b = 0.37$, then Brydcox will capture the entire market! The result is nonsense, because equation (27)'s derivation relied on the assumption that $x_a < x_b$, which is false in this example.

Example 3. Locations too close to each other.

$x^* < x_a < x_b$. Try $x_a = 0.7, x_b = 0.9$ and $\theta = 0.5$. Then equation (27) says that $p_a = (2.0 + 0.7 + 0.9)0.5/3 = 0.6$ and $p_b = (4 - 0.7 - 0.9)0.5/3 = 0.4$. As for the split of the market, equation (26) says that $x^* = \frac{1}{2 \cdot 0.5} [(0.4 - 0.6) + 0.5(0.7 + 0.9)] = 0.6$.

Example 3 shows a serious problem. If the market splits at $x^* = 0.6$ but $x_a = 0.7$ and $x_b = 0.9$, the result violates our implicit assumption that the players split the market. Equation (26) is based on the premise that there does exist some indifferent consumer, and when that is a false premise, as under the parameters of Example 3, equation (26) will still spit out a value of x^* , but the value will not mean anything. And in fact the consumer at $x = 0.6$ is not really indifferent between Apex and Brydcox. He could buy from Apex at a total cost of $0.6 + 0.1(0.5) = 0.65$ or from Brydcox, at a total cost of $0.4 + 0.3(0.5) = 0.55$. In fact, there exists no consumer who strictly prefers Apex. Even Apex's 'home' consumer at $x = 0.7$ would have a total cost of buying from Brydcox of $0.4 + 0.5(0.9 - 0.7) = 0.5$ and would prefer Brydcox. Similarly, the consumer at $x = 0$ would have a total cost of buying from Brydcox of $0.4 + 0.5(0.9 - 0.0) = 0.85$, compared to a cost from Apex of $0.6 + 0.5(0.7 - 0.0) = 0.95$, and he, too, would prefer Brydcox.

The problem in examples 2 and 3 is that the firm with the higher price would do better to deviate with a discontinuous price cut to just below the other firm's price. Equation (27) was derived by calculus, with the implicit assumption that a local profit maximum was also a global profit maximum, or, put differently, that if no small change could raise a firm's payoff, then it had found the optimal strategy. Sometimes a big change will increase a player's payoff even though a small change would not. Perhaps this is what they mean in business by the importance of "nonlinear thinking" or "thinking out of the envelope." The everyday manager or scientist as described by Schumpeter (1934) and Kuhn (1970) concentrates on analyzing incremental changes and only the entrepreneur or genius breaks through with a discontinuously new idea, the profit source or paradigm shift.

Let us now turn to the choice of location. We will simplify the model by pushing consumers into the background and imposing a single exogenous price on all firms.

The Hotelling Location Game (Hotelling [1929])

Players

n Sellers.

The Order of Play

The sellers simultaneously choose locations $x_i \in [0, 1]$.

Payoffs

Consumers are distributed along the interval $[0, 1]$ with a uniform density equal to one. The price equals one, and production costs are zero. The sellers are ordered by their location so $x_1 \leq x_2 \leq \dots \leq x_n$, $x_0 \equiv 0$ and $x_{n+1} \equiv 1$. Seller i attracts half the consumers from the gaps on each side of him, as shown in figure 14.6, so that his payoff is

$$\pi_1 = x_1 + \frac{x_2 - x_1}{2}, \quad (28)$$

$$\pi_n = \frac{x_n - x_{n-1}}{2} + 1 - x_n, \quad (29)$$

or, for $i = 2, \dots, n-1$,

$$\pi_i = \frac{x_i - x_{i-1}}{2} + \frac{x_{i+1} - x_i}{2}. \quad (30)$$




Figure 6: Payoffs in the Hotelling Location Game

With **one seller**, the location does not matter in this model, since the consumers are captive. If price were a choice variable and demand were elastic, we would expect the monopolist to locate at $x = 0.5$.

With **two sellers**, both firms locate at $x = 0.5$, regardless of whether or not demand is elastic. This is a stable Nash equilibrium, as can be seen by inspecting Figure 4 and imagining best responses to each other's location. The best response is always to locate ε closer to the center of the interval than one's rival. When both firms do this, they end up splitting the market since both of them end up exactly at the center.




Figure 7: Nonexistence of pure strategies with three players

With **three sellers** the model does not have a Nash equilibrium in pure strategies. Consider any strategy combination in which each player locates at a separate point. Such a strategy combination is not an equilibrium, because the two players nearest the ends would edge in to squeeze the middle player's market share. But if a strategy combination has any two players at the same point a , as in Figure 7, the third player would be able to acquire a share of at least $(0.5 - \epsilon)$ by moving next to them at b ; and if the third player's share is that large, one of the doubled-up players would deviate by jumping to his other side and capturing his entire market share. The only equilibrium is in mixed strategies.




Figure 8: The Equilibrium Mixed-Strategy Density in the Three-Player Location Game

Suppose all three players use the same mixing density, with $m(x)$ the probability density for location x , and positive density on the support $[g, h]$, as depicted in Figure 8. We will need the density for the distribution of the minimum of the locations of Players 2 and 3. Player 2 has location x with density $m(x)$, and Player 3's location is greater than that with probability $1 - M(x)$, letting M denote the cumulative distribution, so the density for Player 2 having location x and it being smaller is $m(x)[1 - M(x)]$. The density for either Player 2 or Player 3 choosing x and it being smaller than the other firm's location is then $2m(x)[1 - M(x)]$.

If Player 1 chooses $x = g$ then his expected payoff is

$$\pi_1(x_1 = g) = g + \int_g^h 2m(x)[1 - M(x)] \left(\frac{x - g}{2} \right) dx, \quad (31)$$

where g is the safe set of consumers to his left, $2m(x)[1 - M(x)]$ is the density for x being the next biggest location of a firm, and $\frac{x-g}{2}$ is Player 1's share of the consumers between his own location of g and the next biggest location.

If Player 1 chooses $x = h$ then his expected payoff is, similarly,

$$\pi_1(x_1 = h) = (1 - h) + \int_g^h 2m(x)M(x) \left(\frac{h - x}{2} \right) dx, \quad (32)$$

where $(1 - h)$ is the set of safe consumers to his right

In a mixed strategy equilibrium, Player 1's payoffs from these two pure strategies must be equal, and they are also equal to his payoff from a location of 0.5, which we can plausibly

guess is in the support of his mixing distribution. Going on from this point, the algebra and calculus start to become fierce. Shaked (1982) has computed the symmetric mixing probability density $m(x)$ to be as shown in Figure 9,

$$m(x) = \begin{cases} 2 & \text{if } \frac{1}{4} \leq x \leq \frac{3}{4} \\ 0 & \text{otherwise} \end{cases} \quad (33)$$

I do not know how Shaked came to his answer, but I would tackle the problem by guessing that $M(x)$ was a uniform distribution and seeing if it worked, which was perhaps his method too. (You can check that using this mixing density, the payoffs in equation (31) and (32) do equal each other.) Note also that this method has only shown what the symmetric equilibrium is like; it turns out that asymmetric equilibria also exist (Osborne & Pitchik [1986]).




Figure 9: The Equilibrium Mixing Density for Location

Strangely enough, three is a special number. With **more than three sellers**, an equilibrium in pure strategies does exist if the consumers are uniformly distributed, but this is a delicate result (Eaton & Lipsey [1975]). Dasgupta & Maskin (1986b), as amended by Simon (1987), have also shown that an equilibrium, possibly in mixed strategies, exists for any number of players n in a space of any dimension m .

Since prices are inflexible, the competitive market does not achieve efficiency. A benevolent social planner or a monopolist who could charge higher prices if he located his outlets closer to more consumers would choose different locations than competing firms. In particular, when two competing firms both locate in the center of the line, consumers are no better off than if there were just one firm. The average distance of a consumer from a seller would be minimized by setting $x_1 = 0.25$ and $x_2 = 0.75$, the locations that would be chosen either by the social planner or the monopolist.⁵

The Hotelling Location Model, however, is very well suited to politics. Often there is just one dimension of importance in political races, and voters will vote for the candidate

⁵xxxInsert fig14.efficiency.jpg here.

closest to their own position, so there is no analog to price. The Hotelling Location Model predicts that the two candidates will both choose the same position, right on top of the median voter. This seems descriptively realistic; it accords with the common complaint that all politicians are pretty much the same.

14.4 Comparative Statics and Supermodular Games

Comparative statics is the analysis of what happens to endogenous variables in a model when the exogenous variable change. This is a central part of economics. When wages rises, for example, we wish to know how the price of steel will change in response. Game theory presents special problems for comparative statics, because when a parameter changes, not only does Smith's equilibrium strategy change in response, but Jones's strategy changes as a result of Smith's change as well. A small change in the parameter might produce a large change in the equilibrium because of feedback between the different players' strategies.

Let us use a differentiated Bertrand game as an example. Suppose there are N firms, and for firm j the demand curve is

$$Q_j = \text{Max}\{\alpha - \beta_j p_j + \sum_{i \neq j} \gamma_i p_i, 0\}, \quad (34)$$

with $\alpha \in (0, \infty)$, $\beta_i \in (0, \infty)$, and $\gamma_i \in (0, \infty)$ for $i = 1, \dots, N$. Assume that the effect of p_j on firm j 's sales is larger than the effect of the other firms' prices, so that

$$\beta_j > \sum_{i \neq j} \gamma_i. \quad (35)$$

Let firm i have constant marginal cost κc_i , where $\kappa \in \{1, 2\}$ and $c_i \in (0, \infty)$, and let us assume that each firm's costs are low enough that it does operate in equilibrium. (The shift variable κ could represent the effect of the political regime on costs.)

The payoff function for firm j is

$$\pi_j = (p_j - \kappa c_j)(\alpha - \beta_j p_j + \sum_{i \neq j} \gamma_i p_i). \quad (36)$$

Firms choose prices simultaneously.

Does this game have an equilibrium? Does it have several equilibria? What happens to the equilibrium price if a parameter such as c_j or κ changes? These are difficult questions because if c_j increases, the immediate effect is to change firm j 's price, but the other firms will react to the price change, which in turn will affect j 's price. Moreover, this is not a symmetric game – the costs and demand curves differ from firm to firm, which could make algebraic solution of the Nash equilibrium quite messy. It is not even clear whether the equilibrium is unique.

Two approaches to comparative statics can be used here: the implicit function theorem, and supermodularity. We will look at each in turn.

The Implicit Function Theorem

The implicit-function theorem says that if $f(y, z) = 0$, where y is endogenous and z is exogenous, then

$$\frac{dy}{dz} = - \left(\frac{\frac{\partial f}{\partial z}}{\frac{\partial f}{\partial y}} \right). \quad (37)$$

It is worth knowing how to derive this. We start with $f(y, z) = 0$, which can be rewritten as $f(y(z), z) = 0$, since y is endogenous. Using the calculus chain rule,

$$\frac{df}{dz} = \frac{\partial f}{\partial z} + \frac{\partial f}{\partial y} \frac{dy}{dz} = 0. \quad (38)$$

where the expression equals zero because after a small change in z , f will still equal zero after y adjusts. Solving for $\frac{dy}{dz}$ yields equation (37).

The implicit function theorem is especially useful if y is a choice variable and z a parameter, because then we can use the first-order condition to set $f(y, z) \equiv \frac{\partial \pi}{\partial y} = 0$ and the second-order condition tells us that $\frac{\partial f}{\partial y} = \frac{\partial^2 \pi}{\partial y^2} \leq 0$. One only has to make certain that the solution is an interior solution, so the first- and second- order conditions are valid.

We do have a complication if the model is strategic: there will be more than one endogenous variable, because more than one player is choosing variable values. Suppose that instead of simply $f(y, z) = 0$, our implicit equation has two endogenous and two exogenous variables, so $f(y_1, y_2, z_1, z_2) = 0$. The extra z_2 is no problem; in comparative statics we are holding all but one exogenous variable constant. But the y_2 does add something to the mix. Now, using the calculus chain rule yields not equation (38) but

$$\frac{df}{dz_1} = \frac{\partial f}{\partial z_1} + \frac{\partial f}{\partial y_1} \frac{dy_1}{dz_1} + \frac{\partial f}{\partial y_2} \frac{dy_2}{dz_1} = 0. \quad (39)$$

Solving for $\frac{dy_1}{dz_1}$ yields

$$\frac{dy_1}{dz_1} = - \left(\frac{\frac{\partial f}{\partial z_1} + \frac{\partial f}{\partial y_2} \frac{dy_2}{dz_1}}{\frac{\partial f}{\partial y_1}} \right). \quad (40)$$

It is often unsatisfactory to solve out for $\frac{dy_1}{dz_1}$ as a function of both the exogenous variables z_1 and z_2 and the endogenous variable y_2 (though it is okay if all you want is to discover whether the change is positive or negative), but ordinarily the modeller will also have available an optimality condition for Player 2 also: $g(y_1, y_2, z_1, z_2) = 0$. This second condition yields an equation like (40), so that two equations can be solved for the two unknowns.

We can use the differentiated Bertrand game to see how this works out. Equilibrium prices will lie inside the interval $(c_j, \bar{p})$ for some large number $\bar{p}$, because a price of c_j would yield zero profits, rather than the positive profits of a slightly higher price, and $\bar{p}$ can be chosen to yield zero quantity demanded and hence zero profits. The equilibrium

or equilibria are, therefore, interior solutions, in which case they satisfy the first-order condition

$$\frac{\partial \pi_j}{\partial p_j} = \alpha - 2\beta_j p_j + \sum_{i \neq j} \gamma_i p_i + \kappa c_j \beta_j = 0, \quad (41)$$

and the second-order condition,

$$\frac{\partial^2 \pi_j}{\partial p_j^2} = -2\beta_j < 0. \quad (42)$$

Next, apply the implicit function theorem by using p_i and c_i , $i = 1, \dots, N$, instead of y_i and z_i , $i = 1, 2$, and by letting $\frac{\partial \pi_j}{\partial p_j} = 0$ from equation (41) be our $f(y_1, y_2, z_1, z_2) = 0$. The chain rule yields

$$\frac{df}{dc_j} = -2\beta_j \frac{dp_j}{dc_j} + \sum_{i \neq j} \gamma_i \frac{dp_i}{dc_j} + \kappa \beta_j = 0, \quad (43)$$

so

$$\frac{dp_j}{dc_j} = \frac{\sum_{i \neq j} \gamma_i \frac{dp_i}{dc_j} + \kappa \beta_j}{2\beta_j}. \quad (44)$$

Just what is $\frac{dp_i}{dc_j}$? For each i , we need to find the first-order condition for firm i and then use the chain rule again. The first-order condition for Player i is that the derivative of π_i with respect to p_i (*not* p_j) equals zero, so

$$g^i \equiv \frac{\partial \pi_i}{\partial p_i} = \alpha - 2\beta_i p_i + \sum_{k \neq i} \gamma_k p_k + \kappa c_i \beta_i = 0. \quad (45)$$

The chain rule yields (keeping in mind that it is a change in c_j that interests us, *not* a change in c_i),

$$\frac{dg^i}{dc_j} = -2\beta_i \frac{dp_i}{dc_j} + \sum_{k \neq i} \gamma_k \frac{dp_k}{dc_j} = 0. \quad (46)$$

With equation (44), the $(N - 1)$ equations (46) give us N equations for the N unknowns $\frac{dp_i}{dc_j}$, $i = 1, \dots, N$.

It is easier to see what is going on if there are just two firms, j and i . Equations (44) and (46) are then

$$\frac{dp_j}{dc_j} = \frac{\gamma_i \frac{dp_i}{dc_j} + \kappa \beta_j}{2\beta_j}. \quad (47)$$

and

$$-2\beta_i \frac{dp_i}{dc_j} + \gamma_j \frac{dp_j}{dc_j} = 0. \quad (48)$$

Solving these two equations for $\frac{dp_j}{dc_j}$ and $\frac{dp_i}{dc_j}$ yields

$$\frac{dp_j}{dc_j} = \frac{2\beta_i \beta_j \kappa}{4\beta_i \beta_j - \gamma_i \gamma_j} \quad (49)$$

and

$$\frac{dp_i}{dc_j} = \frac{\gamma_j \beta_j \kappa}{4\beta_i \beta_j - \gamma_i \gamma_j}. \quad (50)$$

Keep in mind that the implicit function theorem only tells about infinitesimal changes, not finite changes. If c_n increases enough, then the nature of the equilibrium changes drastically, because firm n goes out of business. Even if c_n increases a finite amount, the implicit function theorem is not applicable, because then the change in p_n will cause changes in the prices of other firms, which will in turn change p_n again.

We cannot go on to discover the effect of changing κ on p_n , because κ is a discrete variable, and the implicit function theorem only applies to continuous variables. The implicit function theorem is none the less very useful when it does apply. This is a simple example, but the approach can be used even when the functions involved are very complicated. In complicated cases, knowing that the second-order condition holds allows the modeller to avoid having to determine the sign of the denominator if all that interests him is the sign of the relationship between the two variables.

Supermodularity

The second approach uses the idea of the supermodular game, an idea related to that of strategic complements (Chapter 3.6). Suppose that there are N players in a game, subscripted by i and j , and that player i has a strategy consisting of $\bar{s}^i$ elements, subscripted by s and t , so his strategy is the vector $y^i = (y_1^i, \dots, y_{\bar{s}^i}^i)$. Let his strategy set be S^i and his payoff function be $\pi^i(y^i, y^{-i}; z)$, where z represents a fixed parameter. We say that the game is a **smooth supermodular game** if the following four conditions are satisfied for every player $i = 1, \dots, N$:

A1' The strategy set is an interval in $R^{\bar{s}^i}$:

$$S^i = [\underline{y}^i, \bar{y}^i]. \quad (51)$$

A2' π^i is twice continuously differentiable on S^i .

A3' (Supermodularity) Increasing one component of player i 's strategy does not decrease the net marginal benefit of any other component: for all i , and all s and t such that $1 \leq s < t \leq \bar{s}^i$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial y_t^i} \geq 0. \quad (52)$$

A4' (Increasing differences in one's own and other strategies) Increasing one component of i 's strategy does not decrease the net marginal benefit of increasing any component of player j 's strategy: for all $i \neq j$, and all s and t such that $1 \leq s \leq \bar{s}^i$ and $1 \leq t \leq \bar{s}^j$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial y_t^j} \geq 0. \quad (53)$$

In addition, we will be able to talk about the comparative statics of smooth supermodular games if a fifth condition is satisfied, the increasing differences condition, (A5').

A5': (Increasing differences in one's own strategies and parameters) Increasing parameter z does not decrease the net marginal benefit to player i of any component of his own strategy: for all i , and all s such that $1 \leq s \leq \bar{s}^i$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial z} \geq 0. \quad (54)$$

The heart of supermodularity is in assumptions A3' and A4'. Assumption A3' says that the components of player i 's strategies are all **complementary inputs**; when one component increases, it is worth increasing the other components too. This means that even if a strategy is a complicated one, one can still arrive at qualitative results about the strategy, because all the components of the optimal strategy will move in the same direction together. Assumption A4' says that the strategies of players i and j are **strategic complements**; when player i increases a component of his strategy, player j will want to do so also. When the strategies of the players reinforce each other in this way, the feedback between them is less tangled than if they undermined each other.

I have put primes on the assumptions because they are the special cases, for smooth games, of the general definition of supermodular games in the Mathematical Appendix. Smooth games use differentiable functions, but the supermodularity theorems apply more generally. One condition that is relevant here is condition A5:

A5: π^i has increasing differences in y^i and z for fixed y^{-i} ; for all $y^i \geq y^{i'}$, the difference $\pi^i(y^i, y^{-i}, z) - \pi^i(y^{i'}, y^{-i}, z)$ is nondecreasing with respect to z .

Is the differentiated Bertrand game supermodular? The strategy set can be restricted to $[c_i, \bar{p}]$ for player i , so A1' is satisfied. π_i is twice continuously differentiable on the interval $[c_i, \bar{p}]$, so A2' is satisfied. A player's strategy has just one component, p_i , so A3' is immediately satisfied. The following inequality is true,

$$\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} = \gamma_j > 0, \quad (55)$$

so A4' is satisfied. And it is also true that

$$\frac{\partial^2 \pi_i}{\partial p_i \partial c_i} = \kappa \beta_i > 0, \quad (56)$$

so A5' is satisfied for c_i .

From equation (41), $\frac{\partial \pi_i}{\partial p_i}$ is increasing in κ , so $\pi_i(p_i, p_{-i}, \kappa) - \pi_i(p'_i, p_{-i}, \kappa)$ is nondecreasing in κ for $p_i > p'_i$, and A5 is satisfied for κ .

Thus, all the assumptions are satisfied. This being the case, a number of theorems can be applied, including the following two.

Theorem 1. *If the game is supermodular, there exists a largest and a smallest Nash equilibrium in pure strategies.*

Theorem 2. *If the game is supermodular and assumption (A5) or (A5') is satisfied, then the largest and smallest equilibrium are nondecreasing functions of the parameter z .*

Applying Theorems 1 and 2 yields the following results for the differentiated Bertrand game:

1. There exists a largest and a smallest Nash equilibrium in pure strategies (Theorem 1).
2. The largest and smallest equilibrium prices for firm i are nondecreasing functions of the cost parameters c_i and κ (Theorem 2).

Note that supermodularity, unlike the implicit function theorem, has yielded comparative statics on κ , the discrete exogenous variable. It yields weaker comparative statics on c_i , however, because it just finds the effect of c_i on p_i^* to be nondecreasing, rather than telling us its value or whether it is actually increasing.

14.5 Vertical Differentiation⁶

Vertical Differentiation I: Monopoly Quality Choice

Players

A seller and a continuum of buyers.

The Order of Play

- 0 There is a continuum of buyers of length 1 parametrized by quality desire θ_i distributed by Nature uniformly on $[0, 1]$.
- 1 The seller picks quality s_1 from the interval $[0, \bar{s}]$.
- 2 The seller picks prices p_1 from the interval $[0, \infty)$.
- 3 Buyer i chooses one unit of a good, or refrains from buying. The seller produces it at constant marginal cost c , which does not vary with quality.

Payoffs

The seller maximizes

$$(p_1 - c)q_1 \tag{57}$$

Buyer i 's payoff is zero if he does not buy. If he does buy, it is

$$(\alpha + \theta_i)s - p, \tag{58}$$

where the parameter $\alpha \in (0, 1)$ is the same for all buyers.

The participation constraint for consumer i is

$$(\alpha + \theta_i)s_1 - p_1 \geq 0, \tag{59}$$

⁶xxx From Tirole, p. 296 and chapter 3 Price disc. Shaked-Sutton (1983) *Econometrica*. "Natural Oligopolies."

which will be binding for some buyer type θ^* for which

$$(\alpha + \theta^*)s_1 = p_1, \quad (60)$$

so $\theta^* = \frac{p_1}{s_1} - \alpha$, and

$$q_i = (1 - \theta^*) = 1 + \alpha - \frac{p_1}{s_1}. \quad (61)$$

The seller maximizes

$$\pi = (p_1 - c)q_1 = (p_1 - c)\left[1 + \alpha - \frac{p_1}{s_1}\right]. \quad (62)$$

This is maximized at the corner solution of $s_1 = \bar{s}$ since high quality has no extra cost. Then the first order condition for choosing p_1 is

$$\frac{d\pi}{dp_1} = 1 + \alpha - 2\frac{p_1}{\bar{s}} + \frac{c}{\bar{s}} = 0, \quad (63)$$

so

$$p_1 = \frac{(\alpha + 1)\bar{s} + c}{2}; \quad (64)$$

that is, the price is halfway between c and $(\alpha + 1)\bar{s}$, which is the valuation of the most quality-valuing buyer.

The seller's profit is then

$$\pi = (p_1 - c)q_1 = \left(\frac{(\alpha + 1)\bar{s} + c}{2} - c\right)\left[1 + \alpha - \frac{\frac{(\alpha + 1)\bar{s} + c}{2}}{s_1}\right]. \quad (65)$$

Next we will allow the seller to use two quality levels. A social planner would just use one, the maximal one of $s = \bar{s}$, since it is no cheaper to produce lower quality. The monopoly seller might use two, however, because it can help him to price discriminate between buyers who value quality differently.

Vertical Differentiation II: Price Discrimination Using Quality

Players

A seller and a continuum of buyers.

The Order of Play

0 There is a continuum of buyers of length 1 parametrized by quality desire θ_i distributed by Nature uniformly on $[0, 1]$.

1 The seller picks qualities s_1 and s_2 from the interval $[0, \bar{s}]$.

2 The seller picks prices p_1 and p_2 from the interval $[0, \infty)$.

3 Buyer i chooses one unit of a good, or refrains from buying. The seller produces it at constant marginal cost c , which does not vary with quality.

Payoffs

The seller maximizes

$$(p_1 - c)q_1 + (p_2 - c)q_2 \quad (66)$$

Buyer i 's payoff is zero if he does not buy. If he does buy, it is

$$(\alpha + \theta_i)s - p, \quad (67)$$

where the parameter $\alpha \in (0, 1)$ is the same for all buyers.

This is a problem of mechanism design and price discrimination by quality. The seller needs to pick s_1, s_2, p_1 , and p_2 to satisfy incentive compatibility and participation constraints if he wants to offer two qualities with positive sales of both, and he also needs to decide if that is more profitable than offering just one quality (e.g., choosing s_1 and p_1 to satisfy participation constraints and choosing s_2 and p_2 to violate them).

We already solved the one-quality problem, in Vertical Differentiation I.

Let us assume that the seller constructs his mechanism so that every type of buyer does make a purchase (xxx CHeck on this later.)

When two qualities are used, buyers with $\theta \in [0, q_1]$ buy the low quality and buyers with $\theta \in [q_1, 1]$ buy the high quality. Demand for the high-quality good will be $q_2 = 1 - q_1$. The high quality will be set to $s_2 = \bar{s}$, since if any lower s_2 were picked the seller could increase s_2 and p_2 could rise, increasing his payoff.

The buyer with $\theta \in [0, q_1]$ who has the greatest temptation to buy high quality instead of low is the one with $\theta = q_1$. He must satisfy the self-selection constraint

$$(\alpha + q_1)s_1 - p_1 \geq (\alpha + q_1)s_2 - p_2. \quad (68)$$

The buyer with $\theta \in [q_1, 1]$ who has the greatest temptation to buy low quality instead of high is the one with $\theta = q_1$. He must satisfy the self-selection constraint

$$(\alpha + q_1)s_1 - p_1 \leq (\alpha + q_1)s_2 - p_2. \quad (69)$$

Since this is the same buyer type, we see that the constraints have to be satisfied as equalities. Substituting $s_2 = \bar{s}$, we get

$$(\alpha + q_1)s_1 - p_1 = (\alpha + q_1)\bar{s} - p_2. \quad (70)$$

Which participation constraint will be satisfied as an equality? That for $\theta = 0$. If the lowest of types had a positive surplus from buying the low-quality good, p_1 being less than his valuation of $(\alpha + 0)s_1$, then so would all the other types. Thus, the seller could increase p_1 without losing any customers. Since the customer with $\theta = 0$ gets zero surplus, we know that

$$\alpha s_1 - p_1 = 0. \quad (71)$$

Putting together equations (70) and (71) gives us

$$(\alpha + q_1)\frac{p_1}{\alpha} - p_1 = (\alpha + q_1)\bar{s} - p_2, \quad (72)$$

which when solved for q_1 yields us a sort of demand curve. Here's the algebra.

$$p_1 - p_1 + p_2 = (\alpha + q_1)\bar{s} - \frac{q_1 p_1}{\alpha} \quad (73)$$

so

$$p_2 - \alpha\bar{s} = q_1 \left(\bar{s} - \frac{p_1}{\alpha} \right) \quad (74)$$

and

$$q_1 = \frac{(p_2 - \alpha)\alpha}{\bar{s}\alpha - p_1}. \quad (75)$$

We can rewrite profit from (66) as

$$\begin{aligned} \pi &= (p_1 - c)q_1 + (p_2 - c)q_2 \\ &= (p_1 - c)q_1 + (p_2 - c)(1 - q_1) \\ &= (p_1 - p_2)q_1 + p_2 - c \\ &= (p_1 - p_2) \left(\frac{(p_2 - \alpha)\alpha}{\bar{s}\alpha - p_1} \right) + p_2 - c \end{aligned} \quad (76)$$

The seller maximizes his payoff function by choice of p_1 and p_2 , giving us two first order conditions to solve for those two variables:

$$\frac{\partial \pi}{\partial p_1} = \left(\frac{(p_2 - \alpha)\alpha}{\bar{s}\alpha - p_1} \right) + (p_1 - p_2) \left(\frac{(p_2 - \alpha)\alpha}{(\bar{s}\alpha - p_1)^2} \right) = 0, \quad (77)$$

so

$$(\bar{s}\alpha - p_1) + (p_1 - p_2) = 0, \quad (78)$$

and

$$p_2 = \bar{s}\alpha. \quad (79)$$

$$\frac{\partial \pi}{\partial p_2} = - \left(\frac{2p_2 - \alpha}{\bar{s}\alpha - p_1} \right) + 1 + (p_1 - p_2) \left(\frac{\alpha}{(\bar{s}\alpha - p_1)^2} \right) = 0, \quad (80)$$

so

$$\frac{\partial \pi}{\partial p_2} = -(2p_2 - \alpha)\alpha + (\bar{s}\alpha - p_1) + p_1\alpha = 0, \quad (81)$$

so

$$2p_2 - \alpha = (\bar{s}\alpha - p_1) + p_1 \quad (82)$$

This problem is mathematically identical to price discriminating by quantity purchases.

ASSUMPTION 1:

$$\underline{\theta} + 1 \geq \underline{\theta} \quad (83)$$

or

$$\underline{\theta} \leq 1. \quad (84)$$

This says that there is enough consumer heterogeneity.

If Assumption 1 is violated, then one firm takes over the market.

ASSUMPTION 2:

$$c + \frac{\underline{\theta} + 1 - 2\underline{\theta}}{3}(s_2 - s_1) \leq \underline{\theta}s_1. \quad (85)$$

This will ensure that the market is covered—that every consumer does buy from one firm or the other.

Vertical Differentiation III: Duopoly Quality Choice

Players

Two sellers and a continuum of buyers.

The Order of Play

0 There is a continuum of buyers of length 1 parametrized by quality desire θ_i distributed by Nature uniformly on $[0, 1]$.

1 Sellers 1 and 2 simultaneously choose qualities s_1 and s_2 from the interval $[0, \bar{s}]$.

2 Sellers 1 and 2 simultaneously pick prices p_1 and p_2 from the interval $[0, \infty)$.

3 Buyer i chooses one unit of a good, or refrains from buying. The sellers produce at constant marginal cost c , which does not vary with quality.

Payoffs

Seller j 's payoff is

$$(p_j - c)q_j. \quad (86)$$

Buyer i 's payoff is zero if he does not buy. If he does buy, from seller j , it is

$$(\alpha + \theta_i)s_j - p_j, \quad (87)$$

where the parameter $\alpha \in (0, 1)$ is the same for all buyers.

Work back from the end of the game. Let us assume that the firms have chosen qualities so that each has some sales. Then there is an indifferent consumer type θ such that the consumer's payoff from each firm's good is equal. That consumer's type is q_1 , since all consumers in $[0, q_1]$ will buy Good 1. For $\theta = q_1$,

$$(\alpha + q_1)s_1 - p_1 = (\alpha + q_1)s_2 - p_2, \quad (88)$$

so

$$\alpha(s_1 - s_2) - p_1 + p_2 = q_1(s_2 - s_1) \quad (89)$$

The demands will thus be ⁷

$$q_1 = \frac{p_2 - p_1}{s_2 - s_1} - \alpha \quad (90)$$

⁷xxx This should be the same as in the two-good monopoly case.

and

$$q_2 = 1 - q_1 = 1 - \frac{p_2 - p_1}{s_2 - s_1} + \alpha \quad (91)$$

When firms maximize their payoffs by choice of price, the reaction curves turn out to be

$$p_1 = \frac{p_2 + c - \alpha(s_2 - s_1)}{2} \quad (92)$$

and

$$p_2 = \frac{p_1 + c + (\alpha + 1)(s_2 - s_1)}{2} \quad (93)$$

From the reaction curves we can see that prices are strategic complements.

When solved for equilibrium prices, it turns out that $p_2 > p_1$. Substituting from the last two equations,

$$\begin{aligned} p_1 &= \frac{\frac{p_1 + c + (\alpha + 1)(s_2 - s_1)}{2} + c - \alpha(s_2 - s_1)}{2} \\ &= \frac{p_1}{4} + \frac{c + (\alpha + 1)(s_2 - s_1)}{4} + \frac{c - \alpha(s_2 - s_1)}{2} \\ &= \left(\frac{4}{3}\right) \left(\frac{c + (\alpha + 1)(s_2 - s_1)}{4} + \frac{c - \alpha(s_2 - s_1)}{2}\right) \\ &= \frac{c + (\alpha + 1)(s_2 - s_1) + 2c - 2\alpha(s_2 - s_1)}{3} \\ &= c + \frac{(1 - \alpha)(s_2 - s_1)}{3} \end{aligned} \quad (94)$$

Substituting into the reaction curve for Firm 2, this yields

$$\begin{aligned} p_2 &= \frac{c + \frac{(1 - \alpha)(s_2 - s_1)}{3} + c + (\alpha + 1)(s_2 - s_1)}{2} \\ &= \frac{3c + (1 - \alpha)(s_2 - s_1) + 3c + 3(\alpha + 1)(s_2 - s_1)}{6} \\ &= c + \frac{(1 - \alpha + 3\alpha + 3)(s_2 - s_1)}{6} \\ &= c + \frac{(4 + 2\alpha)(s_2 - s_1)}{6} \\ &= c + \frac{(2 + \alpha)(s_2 - s_1)}{3} \end{aligned} \quad (95)$$

The price p_2 is larger than p_1 : $p_2 - p_1 = \frac{1 + 2\alpha(s_2 - s_1)}{3} > 0$.

The quantities are then

$$q_1 = \frac{p_2 - p_1}{s_2 - s_1} - \alpha = \frac{1 + 2\alpha}{3} - \alpha = \frac{1 - \alpha}{3}, \quad (96)$$

and

$$q_2 = 1 - q_1 = \frac{2 + \alpha}{3}. \quad (97)$$

With a greater price and quantity than Firm 1, and the same average cost, Firm 2 has higher profits. The profits are

$$\pi_1 = \frac{(1 - \alpha)^2(s_2 - s_1)}{9} \quad (98)$$

and

$$\pi_2 = \frac{(\alpha + 2)^2(s_2 - s_1)}{9} \quad (99)$$

Note that for both firms, profits are increasing in $(s_2 - s_1)$.

How about quality choice, in the first stage? Well, both firms benefit from having more different qualities. So in equilibrium, they will be separated as far as possible—

$$s_1 = 0 \quad s_2 = \bar{s}. \quad (100)$$

Qualities are neither strategic complements nor substitutes, if those concepts are defined in terms of the marginal benefit one's own strategy when the other player increases his strategy. That is because a player's quality is a corner solution:

$$\frac{\partial \pi_1}{\partial s_1} = \frac{-(1 - \alpha)^2}{9} < 0. \quad (101)$$

It is nonetheless true that

$$\frac{\partial \pi_1}{\partial s_2} = \frac{(1 - \alpha)^2}{9} > 0, \quad (102)$$

which is to say that Firm 1 is helped if Firm 2 chooses a higher quality level. This is a bit unusual: ordinarily this is the case only when the strategies are complements.

This game is a battle of the sexes. There are two Nash equilibria in pure strategies, in each of which a different one of the two firms gets to be "Firm 2" with its higher profits.⁸

14.5 Durable Monopoly⁹

Introductory economics courses are vague on the issue of the time period over which transactions take place. When a diagram shows the supply and demand for widgets, the x -axis is labelled "widgets," not "widgets per week" or "widgets per year." Also, the diagram splits off one time period from future time periods, using the implicit assumption that supply and demand in one period is unaffected by events of future periods. One problem with this on the demand side is that the purchase of a good which lasts for more than one use is an

⁸xxx What is the mixed-strategy equilibrium in which the firm mix over qualities— so they would ompete head to head sometiems? Probably htye mix over lots of qualities.

⁹xxx I need to relate this closely to the auctions Auction Equivalence Theorem adn to the Bargaining chapter. There is a deep linkage.

investment; although the price is paid now, the utility from the good continues into the future. If Smith buys a house, he is buying not just the right to live in the house tomorrow, but the right to live in it for many years to come, or even to live in it for a few years and then sell the remaining years to someone else. The continuing utility he receives from this durable good is called its **service flow**. Even though he may not intend to rent out the house, it is an investment decision for him because it trades off present expenditure for future utility. Since even a shirt produces a service flow over more than an instant of time, the durability of goods presents difficult definitional problems for national income accounts. Houses are counted as part of national investment (and an estimate of their service flow as part of services consumption), automobiles as durable goods consumption, and shirts as nondurable goods consumption, but all are to some extent durable investments.

In microeconomic theory, “durable monopoly” refers not to monopolies that last a long time, but to monopolies that sell durable goods. These present a curious problem. When a monopolist sells something like a refrigerator to a consumer, that consumer drops out of the market until the refrigerator wears out. The demand curve is, therefore, changing over time as a result of the monopolist’s choice of price, which means that the modeller should not make his decisions in one period and ignore future periods. Demand is not **time separable**, because a rise in price at time t_1 affects the quantity demanded at time t_2 .

The durable monopolist has a special problem because in a sense he does have a competitor – himself in the later periods. If he were to set a high price in the first period, thereby removing high-demand buyers from the market, he would be tempted to set a lower price in the next period to take advantage of the remaining consumers. But if it were known he would lower the price, the high-demand buyers would not buy at a high price in the first period. The threat of the future low price forces the monopolist to keep his current price low.

To formalize this situation, let the seller have a monopoly on a durable good which lasts two periods. He must set a price for each period, and the buyer must decide what quantity to buy in each period. Because this one buyer is meant to represent the entire market demand, the moves are ordered so that he has no market power, as in the principal-agent models in Section 7.3. Alternatively, the buyer can be viewed as representing a continuum of consumers (see Coase [1972] and Bulow [1982]). In this interpretation, instead of “the buyer” buying q_1 in the first period, q_1 of the buyers each buy one unit in the first period.

Durable Monopoly

Players

A buyer and a seller.

The Order of Play

- 1 The seller picks the first-period price, p_1 .
- 2 The buyer buys quantity q_1 and consumes service flow q_1 .
- 3 The seller picks the second-period price, p_2 .
- 4 The buyer buys additional quantity q_2 and consumes service flow $(q_1 + q_2)$.

Payoffs

Production cost is zero and there is no discounting. The seller's payoff is his revenue, and the buyer's payoff is the sum across periods of his benefits from consumption minus his expenditure. His benefits arise from his being willing to pay as much as

$$B(q_t) = 60 - \frac{q_t}{2} \quad (103)$$

for the marginal unit service flow consumed in period t , as shown in Figure 10. The payoffs are therefore

$$\pi_{seller} = q_1 p_1 + q_2 p_2 \quad (104)$$

and

$$\begin{aligned} \pi_{buyer} &= [consumer\ surplus_1] + [consumer\ surplus_2] \\ &= [total\ benefit_1 - expenditure_1] + [total\ benefit_2 - expenditure_2] \\ &= \left[\frac{(60 - B(q_1))q_1}{2} + B(q_1)q_1 - p_1 q_1 \right] \\ &\quad + \left[\frac{60 - B(q_1 + q_2)}{2} (q_1 + q_2) + B(q_1 + q_2)(q_1 + q_2) - p_2 q_2 \right] \end{aligned} \quad (105)$$

Thinking about durable monopoly is hard because we are used to one-period models in which the demand curve, which relates the price to the quantity demanded, is identical to the marginal-benefit curve, which relates the marginal benefit to the quantity consumed. Here, the two curves are different. The marginal benefit curve is the same each period, since it is part of the rules of the game, relating consumption to utility. The demand curve will change over time and depends on the equilibrium strategies, depending as it does on the number of periods left in which to consume the good's services, expected future prices, and the quantity already owned. Marginal benefit is a given for the buyer; quantity demanded is his strategy.

The buyer's total benefit in period 1 is the dollar value of his utility from his purchase of q_1 , which equals the amount he would have been willing to pay to rent q_1 . This is composed of the two areas shown in figure 14.10a, the upper triangle of area $\left(\frac{1}{2}\right)(q_1 + q_2)(60 - B(q_1 + q_2))$ and the lower rectangle of area $(q_1 + q_2)B(q_1 + q_2)$. From this must be subtracted his expenditure in period 1, $p_1 q_1$, to obtain what we might call his consumer surplus in the first period. Note that $p_1 q_1$ will not be the lower rectangle, unless by some strange accident, and the "consumer surplus" might easily be negative, since the expenditure in period 1 will also yield utility in period 2 because the good is durable.




Figure 10: The Buyer's Marginal Benefit per Period in the Game of Durable Monopoly

To find the equilibrium price path one cannot simply differentiate the seller's utility with respect to p_1 and p_2 , because that would violate the sequential rationality of the seller and the rational response of the buyer. Instead, one must look for a subgame perfect equilibrium, which means starting in the second period and discovering how much the buyer would purchase given his first-period purchase of q_1 , and what second-period price the seller would charge given the buyer's second-period demand function.

In the first period, the marginal unit consumed was the $q_1 - th$. In the second period, it will be the $(q_1 + q_2) - th$. The residual demand curve after the first period's purchases is shown in Figure 10b. It is a demand curve very much like the demand curve resulting from intensity rationing in the capacity-constrained Bertrand game of Section 14.2, as shown in Figure 2a. The most intense portion of the buyer's demand, up to q_1 units, has already been satisfied, and what is left begins with a marginal benefit of $B(q_1)$, and falls at the same slope as the original marginal benefit curve. The equation for the residual demand is therefore, using equation (103),

$$p_2 = B(q_1) - \frac{1}{2}q_2 = 60 - \frac{1}{2}q_1 - \frac{1}{2}q_2. \quad (106)$$

Solving for the monopoly quantity, q_2^* , the seller maximizes $q_2 p_2$, solving the problem

$$\underset{q_2}{\text{Maximize}} \quad q_2 \left(60 - \frac{q_1 + q_2}{2} \right), \quad (107)$$

which generates the first-order condition

$$60 - q_2 - \frac{1}{2}q_1 = 0, \quad (108)$$

so that

$$q_2^* = 60 - \frac{1}{2}q_1. \quad (109)$$

From equations (106) and (109), it can be seen that $p_2^* = 30 - q_1/4$.

We must now find q_1^* . In period one, the buyer looks ahead to the possibility of buying in period two at a lower price. Buying in the first period has two benefits: consumption of the service flow in the first period and consumption of the service flow in the second period. The price he would pay for a unit in period one cannot exceed the marginal benefit from the first-period service flow in period one plus the foreseen value of p_2 , which from (109) is $30 - q_1/4$. If the seller chooses to sell q_1 in the first period, therefore, he can do so at the price

$$\begin{aligned} p_1(q_1) &= B(q_1) + p_2 \\ &= (60 - \tfrac{1}{2}q_1) + (30 - \tfrac{1}{4}q_1), \\ &= 90 - \tfrac{3}{4}q_1. \end{aligned} \tag{110}$$

Knowing that in the second period he will choose q_2 according to (109), the seller combines (109) with (110) to give the maximand in the problem of choosing q_1 to maximize profit over the two periods, which is

$$\begin{aligned} (p_1q_1 + p_2q_2) &= (90 - \tfrac{3}{4}q_1)q_1 + (30 - \tfrac{1}{4}q_1)(60 - \tfrac{1}{2}q_1) \\ &= 1800 + 60q_1 - \tfrac{5}{8}q_1^2, \end{aligned} \tag{111}$$

which has the first-order condition

$$60 - \tfrac{5}{4}q_1 = 0, \tag{112}$$

so that

$$q_1^* = 48 \tag{113}$$

and, making use of (110), $p_1^* = 54$.

It follows from (109) that $q_2^* = 36$ and $p_2 = 18$. The seller's profits over the two periods are $\pi_s = 3,240$ ($= 54(48) + 18(36)$).

The purpose of these calculations is to compare the situation with three other market structures: a competitive market, a monopolist who rents instead of selling, and a monopolist who commits to selling only in the first period.

A *competitive market* bids down the price to the marginal cost of zero. Then, $p_1 = 0$ and $q_1 = 120$ from (103), and profits equal zero.

If the monopolist *rents* instead of selling, then equation (103) is like an ordinary demand equation, because the monopolist is effectively selling the good's services separately each period. He could rent a quantity of 60 each period at a rental fee of 30 and his profits would sum to $\pi_s = 3,600$. That is higher than 3,240, so profits are higher from renting than from selling outright. The problem with selling outright is that the first-period price cannot be very high or the buyer knows that the seller will be tempted to lower the price once the buyer has bought in the first period. Renting avoids this problem.

If the monopolist can *commit to not producing in the second period*, he will do just as well as the monopolist who rents, since he can sell a quantity of 60 at a price of 60, the

sum of the rents for the two periods. An example is the artist who breaks the plates for his engravings after a production run of announced size. We must also assume that the artist can convince the market that he has broken the plates. People joke that the best way an artist can increase the value of his work is by dying, and that, too, fits the model.

If the modeller ignored sequential rationality and simply looked for the Nash equilibrium that maximized the payoff of the seller by his choice of p_1 and p_2 , he would come to the commitment result. An example of such an equilibrium is ($p_1 = 60$, $p_2 = 200$, *Buyer purchases according to* $q_1 = 120 - p_1$, *and* $q_2 = 0$). This is Nash because neither player has incentive to deviate given the other's strategy, but it fails to be subgame perfect, because the seller should realize that if he deviates and chooses a lower price once the second period is reached, the buyer will respond by deviating from $q_2 = 0$ and will buy more units.

With more than two periods, the difficulties of the durable-goods monopolist become even more striking. In an infinite-period model without discounting, if the marginal cost of production is zero, the equilibrium price for outright sale instead of renting is constant – at zero! Think about this in the context of a model with many buyers. Early consumers foresee that the monopolist has an incentive to cut the price after they buy, in order to sell to the remaining consumers who value the product less. In fact, the monopolist would continue to cut the price and sell more and more units to consumers with weaker and weaker demand until the price fell to marginal cost. Without discounting, even the high-valuation consumers refuse to buy at a high price, because they know they could wait until the price falls to zero. And this is not a trick of infinity: a large number of periods generates a price close to zero.

We can also use the durable monopoly model to think about the durability of the product. If the seller can develop a product so flimsy that it only lasts one period, that is equivalent to renting. A consumer is willing to pay the same price to own a one-hoss shay that he knows will break down in one year as he would pay to rent it for a year. Low durability leads to the same output and profits as renting, which explains why a firm with market power might produce goods that wear out quickly. The explanation is not that the monopolist can use his market power to inflict lower quality on consumers— after all, the price he receives is lower too— but that the lower durability makes it credible to high-valuation buyers that the seller expects their business in the future and will not lower his price.

N14.1 Quantities as Strategies: the Cournot Equilibrium Revisited

- Articles on the existence of a pure-strategy equilibrium in the Cournot model include Novshek (1985) and Roberts & Sonnenschein (1976).
- **Merger in a Cournot model.** A problem with the Cournot model is that a firm's best policy is often to split up into separate firms. Apex gets half the industry profits in a duopoly game. If Apex split into firms $Apex_1$ and $Apex_2$, it would get two thirds of the profit in the Cournot triopoly game, even though industry profit falls.

This point was made by Salant, Switzer & Reynolds (1983) and is the subject of problem 14.2. It is interesting that nobody noted this earlier, given the intense interest in Cournot models. The insight comes from approaching the problem from asking whether a player could improve his lot if his strategy space were expanded in reasonable ways.

- An ingenious look at how the number of firms in a market affects the price is Bresnahan & Reiss (1991), which looks empirically at a number of very small markets with one, two, three or more competing firms. They find a big decline in the price from one to two firms, a smaller decline from two to three, and not much change thereafter.

Exemplifying theory, as discussed in the Introduction to this book, lends itself to explaining particular cases, but it is much less useful for making generalizations across industries. Empirical work associated with exemplifying theory tends to consist of historical anecdote rather than the linear regressions to which economics has become accustomed. Generalization and econometrics are still often useful in industrial organization, however, as Bresnahan & Reiss (1991) shows. The most ambitious attempt to connect general data with the modern theory of industrial organization is Sutton's 1991 book, *Sunk Costs and Market Structure*, which is an extraordinarily well-balanced mix of theory, history, and numerical data.

N14.2 Prices as strategies: the Bertrand equilibrium

- As Morrison (1998) points out, Cournot actually does (in Chapter 7) analyze the case of price competition with imperfect substitutes, as well as the quantity competition that bears his name. It is convenient to continue to contrast "Bertrand" and "Cournot" competition, however, though a case can be made for simplifying terminology to "price" and "quantity" competition instead. For the history of how the Bertrand name came to be attached to price competition, see Dimand & Dore (1999).
- Intensity rationing has also been called **efficient rationing**. Sometimes, however, this rationing rule is inefficient. Some low-intensity consumers left facing the high price decide not to buy the product even though their benefit is greater than its marginal cost. The reason intensity rationing has been thought to be efficient is that it is efficient if the rationed-out consumers are unable to buy at any price.
- OPEC has tried both price and quantity controls ("OPEC, Seeking Flexibility, May Choose Not to Set Oil Prices, but to Fix Output," *Wall Street Journal*, October 8, 1987, p. 2; "Saudi King Fahd is Urged by Aides To Link Oil Prices to Spot Markets," *Wall Street Journal*, October 7, 1987, p. 2). Weitzman (1974) is the classic reference on price versus quantity control by regulators, although he does not use the context of oligopoly. The decision

rests partly on enforceability, and OPEC has also hired accounting firms to monitor prices (“Dutch Accountants Take On a Formidable Task: Ferreting Out ‘Cheaters’ in the Ranks of OPEC,” *Wall Street Journal*, February 26, 1985, p. 39).

- Kreps & Scheinkman (1983) show how capacity choice and Bertrand pricing can lead to a Cournot outcome. Two firms face downward-sloping market demand. In the first stage of the game, they simultaneously choose capacities, and in the second stage they simultaneously choose prices (possibly by mixed strategies). If a firm cannot satisfy the demand facing it in the second stage (because of the capacity limit), it uses intensity rationing (the results depend on this). The unique subgame perfect equilibrium is for each firm to choose the Cournot capacity and price.
- Haltiwanger & Waldman (1991) have suggested a dichotomy applicable to many different games between players who are **responders**, choosing their actions flexibly, and those who are **nonresponders**, who are inflexible. A player might be a nonresponder because he is irrational, because he moves first, or simply because his strategy set is small. The categories are used in a second dichotomy, between games exhibiting **synergism**, in which responders choose to do whatever the majority do (upward sloping reaction curves), and games exhibiting **congestion**, in which responders want to join the minority (downward sloping reaction curves). Under synergism, the equilibrium is more like what it would be if all the players were nonresponders; under congestion, the responders have more influence. Haltiwanger and Waldman apply the dichotomies to network externalities, efficiency wages, and reputation.
- Section 14.3 shows how to generate demand curves (??) and (??) using a location model, but they can also be generated directly by a quadratic utility function. Dixit (1979) states with respect to three goods 0, 1, and 2, the utility function

$$U = q_0 + \alpha_1 q_1 + \alpha_2 q_2 - \frac{1}{2} (\beta_1 q_1^2 + 2\gamma q_1 q_2 + \beta_2 q_2^2) \quad (114)$$

(where the constants $\alpha_1, \alpha_2, \beta_1$, and β_2 are positive and $\gamma^2 \leq \beta_1 \beta_2$) generates the inverse demand functions

$$p_1 = \alpha_1 - \beta_1 q_1 - \gamma q_2 \quad (115)$$

and

$$p_2 = \alpha_2 - \beta_2 q_2 - \gamma q_1. \quad (116)$$

- There are many ways to specify product differentiation. This chapter looks at horizontal differentiation where all consumers agree that products A and B are more alike than A and C, but they disagree as to which is best. Another way horizontal differentiation might work is for each consumer to like a particular product best, but to consider all others as equivalent. See Dixit & Stiglitz (1977) for a model along those lines. Or, differentiation might be vertical: all consumers agree that A is better than B and B is better than C but they disagree as to how *much* better A is than B. Firms therefore offer different qualities at different prices. Shaked & Sutton (1983) have explored this kind of vertical differentiation.

N14.3 Location Models

- For a booklength treatment of location models, see Greenhut & Ohta (1975).

- Vickrey notes the possible absence of a pure-strategy equilibrium in Hotelling's model in pp.323-324 of his 1964 book *Microstatics*. D'Aspremont, Gabszewicz & Thirise (1979) work out the mixed-strategy equilibrium for the case of quadratic transportation costs, and Osborne & Pitchik (1987) do the same for Hotelling's original model.
- Location models and switching cost models are attempts to go beyond the notion of a market price. Antitrust cases are good sources for descriptions of the complexities of pricing in particular markets. See, for example, Sultan's 1974 book on electrical equipment in the 1950s, or antitrust opinions such as *US v. Addyston Pipe & Steel Co.*, 85 F. 271 (1898).
- It is important in location models whether the positions of the players on the line are moveable. See, for example, Lane (1980).
- The location games in this chapter model use a one-dimensional space with end points, i.e., a line segment. Another kind of one-dimensional space is a circle (not to be confused with a disk). The difference is that no point on a circle is distinctive, so no consumer preference can be called extreme. It is, if you like, Peoria versus Berkeley. The circle might be used for modelling convenience or because it fits a situation: e.g., airline flights spread over the 24 hours of the day. With two players, the Hotelling location game on a circle has a continuum of pure-strategy equilibria that are one of two types: both players locating at the same spot, versus players separated from each other by 180°. The three-player model also has a continuum of pure-strategy equilibria, each player separated from another by 120°, in contrast to the nonexistence of a pure-strategy equilibrium when the game is played on a line segment.
- Characteristics such as the color of cars could be modelled as location, but only on a player-by-player basis, because they have no natural ordering. While Smith's ranking of (red=1, yellow=2, blue=10) could be depicted on a line, if Brown's ranking is (red=1, blue=5, yellow=6) we cannot use the same line for him. In the text, the characteristic was something like physical location, about which people may have different preferences but agree on what positions are close to what other positions.

N14.6 Durable monopoly

- The proposition that price falls to marginal cost in a durable monopoly with no discounting and infinite time is called the "Coase Conjecture," after Coase (1972). It is really a proposition and not a conjecture, but alliteration was too strong to resist.
- Gaskins (1974) has written a well-known article on the problem of the durable monopolist who foresees that he will be creating his own future competition in the future because his product can be recycled, using the context of the aluminum market.
- Leasing by a durable monopoly was the main issue in the antitrust case *US v. United Shoe Machinery Corporation*, 110 F. Supp. 295 (1953), but not because it increased monopoly profits. The complaint was rather that long-term leasing impeded entry by new sellers of shoe machinery, a curious idea when the proposed alternative was outright sale. More likely, leasing was used as a form of financing for the machinery consumers; by leasing, they did not need to borrow as they would have to do if it was a matter of financing a purchase. See Wiley, Ramseyer, and Rasmusen (1990).

- Another way out of the durable monopolist's problem is to give best-price guarantees to consumers, promising to refund part of the purchase price if any future consumer gets a lower price. Perversely, this hurts consumers, because it stops the seller from being tempted to lower his price. The "most-favored-consumer" contract, which is the analogous contract in markets with several sellers, is analyzed by Holt & Scheffman (1987), for example, who demonstrate how it can maintain high prices, and Png & Hirshleifer (1987), who show how it can be used to price discriminate between different types of buyers.
- The durable monopoly model should remind you of bargaining under incomplete information. Both situations can be modelled using two periods, and in both situations the problem for the seller is that he is tempted to offer a low price in the second period after having offered a high price in the first period. In the durable monopoly model this would happen if the high-valuation buyers bought in the first period and thus were absent from consideration by the second period. In the bargaining model this would happen if the buyer rejected the first-period offer and the seller could conclude that he must have a low valuation and act accordingly in the second period. With a rational buyer, neither of these things can happen, and the models' complications arise from the attempt of the seller to get around the problem.

In the durable-monopoly model this would happen if the high-valuation buyers bought in the first period and thus were absent from consideration by the second period. In the bargaining model this would happen if the buyer rejected the first-period offer and the seller could conclude that he must have a low valuation and act accordingly in the second period. For further discussion, see the survey by Kennan & Wilson (1993).

Problems

14.1. Differentiated Bertrand with Advertising

Two firms that produce substitutes are competing with demand curves

$$q_1 = 10 - \alpha p_1 + \beta p_2 \quad (117)$$

and

$$q_2 = 10 - \alpha p_2 + \beta p_1. \quad (118)$$

Marginal cost is constant at $c = 3$. A player's strategy is his price. Assume that $\alpha > \beta/2$.

- (a) What is the reaction function for Firm 1? Draw the reaction curves for both firms.
- (b) What is the equilibrium? What is the equilibrium quantity for Firm 1?
- (c) Show how Firm 2's reaction function changes when β increases. What happens to the reaction curves in the diagram?
- (d) Suppose that an advertising campaign could increase the value of β by one, and that this would increase the profits of each firm by more than the cost of the campaign. What does this mean? If either firm could pay for this campaign, what game would result between them?

14.2. Cournot Mergers (See Salant, Switzer, & Reynolds [1983])

There are three identical firms in an industry with demand given by $P = 1 - Q$, where $Q = q_1 + q_2 + q_3$. The marginal cost is zero.

- (a) Compute the Cournot equilibrium price and quantities.
- (b) How do you know that there are no asymmetric Cournot equilibria, in which one firm produces a different amount than the others?
- (c) Show that if two of the firms merge, their shareholders are worse off.

14.3. Differentiated Bertrand

Two firms that produce substitutes have the demand curves

$$q_1 = 1 - \alpha p_1 + \beta(p_2 - p_1) \quad (119)$$

and

$$q_2 = 1 - \alpha p_2 + \beta(p_1 - p_2), \quad (120)$$

where $\alpha > \beta$. Marginal cost is constant at c , where $c < 1/\alpha$. A player's strategy is his price.

- (a) What are the equations for the reaction curves $p_1(p_2)$ and $p_2(p_1)$? Draw them.
- (b) What is the pure-strategy equilibrium for this game?

- (c) What happens to prices if α , β , or c increase?
- (d) What happens to each firm's price if α increases, but only firm 2 realizes it (and firm 2 knows that firm 1 is uninformed)? Would firm 2 reveal the change to firm 1?

Problem 14.4. Asymmetric Cournot Duopoly

Apex has variable costs of q_a^2 and a fixed cost of 1000, while Brydax has variable costs of $2q_b^2$ and no fixed cost. Demand is $p = 115 - q_a - q_b$.

- (a) What is the equation for Apex's Cournot reaction function?
- (b) What is the equation for Brydax' Cournot reaction function?
- (c) What are the outputs and profits in the Cournot equilibrium?

Problem 14.5. Omitted.

Problem 14.6. Price Discrimination

A seller faces a large number of buyers whose market demand is given by $P = \alpha - \beta Q$. Production marginal cost is constant at c .

- (a) What is the monopoly price and profit?
- (b) What are the prices under perfect price discrimination if the seller can make take-it-or-leave-it offers? What is the profit?
- (c) What are the prices under perfect price discrimination if the buyer and sellers bargain over the price and split the surplus evenly? What is the profit?

January 18, 2000. September 7, 2003.

Eric Rasmusen, Erasmuse@indiana.edu. Myfile.iu.edu/~erasmuse.¹

References and Name Index

Forthcoming and unpublished articles and books have been assigned the years “(forthcoming)” and “(unpublished)”. The page numbers where a reference is mentioned in the text are listed after the reference. The date of first publication, which may differ from the date of the printing cited, follows the author’s name. Some publications (e.g., *The Wall Street Journal*) cited in footnotes in the main text but are not the in the bibliography. In the case of newspapers, keep in mind that although page numbers are given, the exact page may differ among regional and time-of-day editions.

Abreu, Dilip, David Pearce & Ennio Stacchetti (1986) “Optimal Cartel Equilibria with Imperfect Monitoring,” *Journal of Economic Theory*, 39: 251-269 (June 1986).

Abreu, Dilip, David Pearce & Ennio Stacchetti (1990) “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58: 1041-1064 (September 1990).

Akerlof, George (1970) “The Market for Lemons: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84: 488-500 (August 1970). Reprinted in Rasmusen (2001). .

Akerlof, George (1976) “The Economics of Caste and of the Rat Race and Other Woeful Tales,” *Quarterly Journal of Economics*, 90: 599-617 (November 1976).

Akerlof, George (1980) “A Theory of Social Custom, of which Unemployment may be One Consequence,” *Quarterly Journal of Economics*, 94: 749-775 (June 1980).

Akerlof, George (1983) “Loyalty Filters,” *American Economic Review*, 73: 54- 63 (March 1983).

Akerlof, George & Janet Yellen, eds. (1986) *Efficiency Wage Models of the Labor Market*, Cambridge: Cambridge University Press, ISBN: 0521312841.

Alchian, Armen. See Klein et al. (1978).

Alchian, Armen & Harold Demsetz (1972) “Production, Information Costs and Economic Organization,” *American Economic Review*, 62: 777-795 (December 1972).

Aliprantis, Charalambos & Subir Chakrabarti (1999) *Games and Decisionmaking*, Oxford: Oxford University Press, ISBN: 0195126092.

Antle, Rick & Abbie Smith (1986) “An Empirical Investigation of the Relative Performance Evaluation of Corporate Executives,” *Journal of Accounting Research*, 24: 1-39 (Spring 1986).

¹xxx Footnotes starting with xxx are the author’s notes to himself. Comments are welcomed. This section is xx pages long.

- Arrow, Kenneth (1979) "The Property Rights Doctrine and Demand Revelation Under Incomplete Information," in *Economics and Human Welfare*, ed. by Michael Boskin, New York: Academic Press, ISBN: 0121188507.
- Arrow, Kenneth (1985) "The Economics of Agency," pp. 37-51 of *Principals and Agents: The Structure of Business*, edited by John Pratt & Richard Zeckhauser. Boston: Harvard Business School Press, ISBN: 0875841643.
- Ashenfelter, Orley & David Bloom (1984) "Models of Arbitrator Behavior: Theory and Evidence," *American Economic Review*, 74: 111-124 (March 1984).
- Aumann, Robert (1964a) "Markets with a Continuum of Traders," *Econometrica*, 32: 39-50 (January/April 1964).
- Aumann, Robert (1964b) "Mixed and Behavior Strategies in Infinite Extensive Games," pp. 627-650 of *Annals of Mathematics Studies, No. 52*, Princeton: Princeton University Press.
- Aumann, Robert (1974) "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1: 67-96 (March 1974).
- Aumann, Robert (1976) "Agreeing to Disagree," *Annals of Statistics*, 4: 1236-1239 (November 1976).
- Aumann, Robert (1981) "Survey of Repeated Games," in *Essays in Game Theory and Mathematical Economics in Honor of Oscar Morgenstern*, edited by Robert Aumann, Mannheim: Bibliographisches Institut, ISBN: 3411016094.
- Aumann, Robert (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55: 1-18 (January 1987).
- Aumann, Robert (1988) *Lectures on Game Theory*, (Underground Classics in Economics) Boulder: Westview Press, ASIN: 0813375789.
- Aumann, Robert (1997) "On the State of the Art in Game Theory," in *Understanding Strategic Interaction*, edited by Wulf Albers, Werner Guth, Peter Hammerstein, Benny Moldovanu, and Eric van Damme, Berlin: Springer Verlag, ISBN: 3540614907.
- Aumann, Robert & Sergiu Hart (1992) *Handbook of Game Theory with Economic Applications*, New York: North-Holland, ISBN: 0444880984.
- Avenhaus, Rudolf, Bernhard von Stengel & Shmuel Zamir (19xxx), "Inspection Games," Chapter xxx of Volume III of the *Handbook of Game Theory*, edited by Robert Aumann & Sergiu Hart, xxx.
- Axelrod, Robert (1984) *The Evolution of Cooperation*, New York: Basic Books, ISBN: 0465021220.
- Axelrod, Robert & William Hamilton (1981) "The Evolution of Cooperation," *Science*, 211: 1390-96 (March 1981). Reprinted in Rasmusen (2001).

- Ayres, Ian (1990) "Playing Games with the Law," *Stanford Law Review*, 42: 1291-1317 (May 1990).
- Ayres, Ian (1991) "Fair Driving: Gender and Race Discrimination in Retail Car Negotiations," *Harvard Law Review*, 104: 817-872 (February 1991).
- Ayres, Ian & Peter Cramton (1996) "Deficit Reduction Through Diversity: How Affirmative Action at the FCC Increased Auction Competition," *Stanford Law Review*, 48: 761-814 (April 1996).
- Ayres, Ian & Robert Gertner (1992) "Strategic Contractual Inefficiency and the Optimal Choice of Legal Rules," *Yale Law Journal*, 101: 729-773 (January 1992).
- Ayres, Ian. See Brown & Ayres (1994).
- Bagchi, Arunabha (1984) *Stackelberg Differential Games in Economic Models*, Berlin: Springer-Verlag, ASIN: 0387135871.
- Bagehot, Walter (1971) "The Only Game in Town" *Financial Analysts Journal*, 27: 12-22 (March/April 1971). Reprinted in Rasmusen (2001).
- Bagnoli, Mark & Theodore Bergstrom (1994) "Log-Concave Probability and its Applications," working paper, <http://ideas.repec.org/p/wpa/wuwpmi/9410002.html> (viewed August 28, 2003).
- Baiman, Stanley (1982) "Agency Research in Managerial Accounting: A Survey," *Journal of Accounting Literature*, 1: 154-213 (Spring 1982).
- Baird, Douglas, Gertner, Robert & Randal Picker (1994) *Strategic Behavior and the Law: The Role of Game Theory and Information Economics in Legal Analysis*, Cambridge, Mass: Harvard University Press, ISBN: xxx.
- Baker, George, Michael Jensen, & Kevin J. Murphy (1988) "Compensation and Incentives: Practice vs. Theory," *Journal of Finance*, 43: 593-616 (July 1988).
- Baldwin, B. & G. Meese (1979) "Social Behavior in Pigs Studied by Means of Operant Conditioning," *Animal Behavior*, 27: 947-957 (August 1979).
- Banks, Jeffrey (1991) *Signalling Games in Political Science*, Chur, Switzerland: Harwood Publishers, ISBN: 3718650878.
- Baron, David (1989) "Design of Regulatory Mechanisms and Institutions," in Schmalensee & Willig (1989).
- Baron, David & David Besanko (1984) "Regulation, Asymmetric Information, and Auditing," *Rand Journal of Economics*, 15: 447-470 (Winter 1984).
- Baron, David & Robert Myerson (1982) "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50: 911-930 (July 1982).
- Barzel, Yoram (1968) "Optimal Timing of Innovations," *Review of Economics and Statistics*, 50: 348-355 (August 1968).

- Basar, Tamar & Geert Olsder (1999) *Dynamic Noncooperative Game Theory*, 2nd edn., revised, Philadelphia: Society for Industrial and Applied Mathematics (1st edn. 1982, 2nd edn. 1995), ISBN: 089871429X.
- Basil Blackwell (1985) *Guide for Authors*, Oxford: Basil Blackwell, ISBN: 0631137076.
- Basu, Kaushik (1993) *Lectures in Industrial Organization Theory*, Oxford: Blackwell Publishers, ASIN: 1557863431.
- Baumol, William & Stephen Goldfeld (1968) *Precursors in Mathematical Economics: An Anthology*, London: London School of Economics and Political Science.
- Bazerman and Samuelson (1983) xxx I won the Auction but donw't wnat hte price," *Journal of Conflict Resolution*, 27: 618-34.
- Becker, Gary (1968) "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76: 169-217 (March/April 1968).
- Becker, Gary & George Stigler (1974) "Law Enforcement, Malfeasance and Compensation of Enforcers," *Journal of Legal Studies*, 3: 1-18 (January 1974).
- Benoit, Jean-Pierre & Vijay Krishna (1985) "Finitely Repeated Games," *Econometrica*, 17: 317-320 (July 1985).
- Bernanke, Benjamin (1983) "Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression," *American Economic Review*, 73: 257-276 (June 1983).
- Bernheim, B. Douglas (1984a) "Rationalizable Strategic Behavior," *Econometrica*, 52: 1007-28 (July 1984).
- Bernheim, B. Douglas (1984b) "Strategic Deterrence of Sequential Entry into an Industry," *Rand Journal of Economics*, 15: 1-11 (Spring 1984).
- Bernheim, B. Douglas, Bezalel Peleg, & Michael Whinston (1987) "Coalition-Proof Nash Equilibria I: Concepts," *Journal of Economic Theory*, 42: 1-12 (June 1987).
- Bernheim, B. Douglas & Michael Whinston (1987) "Coalition-Proof Nash Equilibria II: Applications," *Journal of Economic Theory*, 42: 13-29 (June 1987).
- Bertrand, Joseph (1883) "Rechercher sur la theorie mathematique de la richesse," *Journal des Savants*, 48: 499-508 (September 1883).
- Besanko, David, David Dranove & Mark Shanley (1996) *Economics of Strategy*, New York: John Wiley and Sons, ISBN: 0471598496.
- Besanko, David. See Baron & Besanko (1984).
- Bierman, H. Scott & Fernandez, Luis (1998) *Game Theory with Economic Applications*, 2nd edn., Reading,: Addison Wesley (1st edn. 1993), ISBN: 0201562987.
- Bikhchandani, Sushil (1988) "Reputations in Repeated Second Price Auctions," *Journal of Economic Theory*, 46: 97-119 (October 1988).

- Bikhchandani, Sushil, David Hirshleifer & Ivo Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100: 992-1026 (October 1992).
- Binmore, Ken (1990) *Essays on the Foundations of Game Theory*, Oxford: Basil Blackwell, ASIN: 0631168664.
- Binmore, Ken (1992) *Fun and Games: A Text on Game Theory*, Lexington: D. C. Heath, ISBN: 0669246034.
- Binmore, Ken & Partha Dasgupta, eds. (1986) *Economic Organizations as Games*, Oxford: Basil Blackwell, ISBN: 063114255X.
- Binmore, Ken, Ariel Rubinstein, & Asher Wolinsky (1986) "The Nash Bargaining Solution in Economic Modelling," *Rand Journal of Economics*, 17: 176-188 (Summer 1986).
- Blanchard, Olivier (1979) "Speculative Bubbles, Crashes, and Rational Expectations," *Economics Letters*, 3: 387-389 (1979).
- Bloom, David. See Ashenfelter & Bloom (1984).
- Bond, Eric (1982) "A Direct Test of the 'Lemons' Model: The Market for Used Pickup Trucks," *American Economic Review*, 72: 836-40 (September 1982).
- Border, Kim (1985) *Fixed Point Theorems with Applications to Economics and Game Theory*, Cambridge: Cambridge University Press, ASIN: 0521265649.
- Border, Kim & Joel Sobel (1987) "Samurai Accountant: A Theory of Auditing and Plunder," *Review of Economic Studies*, 54: 525-540 (October 1987).
- Bowersock, G. (1985) "The Art of the Footnote," *American Scholar*, 52: 54-62 (Winter 1983/84).
- Boyd, Robert & Jeffrey Lorberbaum (1987) "No Pure Strategy is Evolutionarily Stable in the Repeated Prisoner's Dilemma Game," *Nature*, 327: 58-59 (May 1987).
- Boyd, Robert & Peter Richerson (1985) *Culture and the Evolutionary Process*, Chicago: University of Chicago Press, ISBN: 0226069311.
- Brams, Steven (1980) *Biblical Games: A Strategic Analysis of Stories in the Old Testament*, Cambridge: MIT Press, ISBN: 0262021447.
- Brams, Steven (1983) *Superior Beings: If They Exist, How Would We Know?*, New York: Springer-Verlag, ASIN: 0387912231.
- Brams, Steven (1994) "Game Theory and Literature," *Games and Economic Behavior*, 6: 32-54 (January 1994).
- Brams, Steven & D. Marc Kilgour (1988) *Game Theory and National Security*, Oxford: Basil Blackwell, ASIN: 1557860041.

- Brandenburger, Adam (1992) "Knowledge and Equilibrium in Games," *Journal of Economic Perspectives*, 6: 83-102 (Fall 1992).
- Bresnahan, Timothy & Peter Reiss (1991) "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99: 977-1009 (October 1991).
- Brown, Jennifer & Ian Ayres (1994) "Economic Rationales for Mediation," *Virginia Law Review*, 80: 323-401 (March 1994).
- Bulow, Jeremy (1982) "Durable-Goods Monopolists," *Journal of Political Economy*, 90: 314-332 (April 1982).
- Bulow, Jeremy, John Geanakoplos & Paul Klemperer (1985) "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93: 488-511 (June 1985).
- Bulow, Jeremy & Paul Klemperer (1996) "Auctions versus Negotiations, " *American Economic Review*, 86, 1: 180-194 (March 1996).
- Bulow, Jeremy & John Roberts (1989) "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97, 5: 1060-1090 (October 1989).
- Calfee, John. See Craswell & Calfee (1986).
- Campbell, Richmond & Lanning Sowden (1985) *Paradoxes of Rationality and Cooperation: Prisoner's Dilemma and Newcomb's Problem*, Vancouver: University of British Columbia Press, ASIN: 0774802154
- Campbell, W. See Capen et al. (1971).
- Canzoneri, Matthew & Dale Henderson (1991) *Monetary Policy in Interdependent Economies*, Cambridge: MIT Press, ISBN: 0262031787.
- Capen, E., R. Clapp, & W. Campbell (1971) "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23: 641-653 (June 1971).
- Cass, David & Karl Shell (1983) "Do Sunspots Matter?" *Journal of Political Economy*, 91: 193-227 (April 1983).
- Cassady, Ralph (1967) *Auctions and Auctioneering*, Berkeley: California University Press.
- Chakrabarti, Subir. See Aliprantis & Chakrabarti (1999).
- Chammah, Albert. See Rapoport & Chammah (1965).
- Chatterjee, Kalyan and William Samuelson (1983) "Bargaining Under Incomplete Information," *Operations Research*, 31: 835-851 (September/October 1983).
- Che Yeon-Koo & Ian Gale (1998) "Standard Auctions with Financially Constrained Bidders," *Review of Economic Studies*, 65: 1-21 (January 1998).
- Che Yeon-Koo. See Polinsky & Che (1991).

- Chiang, Alpha (1984) *Fundamental Methods of Mathematical Economics*, 3rd edn., New York: McGraw-Hill (1984, 1st edn. 1967), ISBN: 0070108137.
- Chiappori, P. A., S. Levitt & T. Groseclose, (2002) "Testing Mixed Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer," *American Economic Review*, 92, 4: 1138-1151 (September 2002).
- Cho, In-Koo & David Kreps (1987) "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102: 179-221 (May 1987).
- Clapp, R. See Capen et al. (1971).
- Clavell, James (1966) *Tai-Pan*, New York: Athenum (1966). ISBN: 1580601073
- Coase, Ronald (1960) "The Problem of Social Cost," *Journal of Law & Economics* 3: 1-44 (October 1960).
- Coase, Ronald (1972) "Durability and Monopoly," *Journal of Law and Economics*, 15: 143-149 (April 1972).
- Cooper, Russell (1999) *Coordination Games: Complementarities and Macroeconomics*, Cambridge: Cambridge University Press, ISBN: 0521578965.
- Cooter, Robert & Peter Rappoport (1984) "Were the Ordinalists Wrong about Welfare Economics?" *Journal of Economic Literature*, 22: 507-30 (June 1984).
- Cooter, Robert & Daniel Rubinfeld (1989) "Economic Analysis of Legal Disputes and Their Resolution," *Journal of Economic Literature*, 27: 1067-1097 (September 1989).
- Copeland, Thomas & J. Fred Weston (1988) *Financial Theory and Corporate Policy*, 3rd edn. Reading, MA: Addison-Wesley (1st edn., 1983), ISBN: 0201106485.
- Cournot, Augustin (1838) *Recherches sur les Principes Mathematiques de la Theorie des Richesses*, Paris: M. Riviere & C. (1838). Translated in *Researches into the Mathematical Principles of Wealth*, New York: A.M. Kelly (1960).
- Cox, David & David Hinkley (1974) *Theoretical Statistics*, London: Chapman and Hall, ASIN: 0470181443.
- Cramton, Peter (1984) "Bargaining with Incomplete Information: An Infinite Horizon Model with Two-Sided Uncertainty," *Review of Economic Studies*, 51: 579-593 (October 1984).
- Cramton, Peter, Robert Gibbons & Paul Klemperer (1987) "Dissolving a Partnership Efficiently," *Econometrica*, 55: 615-632.
- Cramton, Peter. See Ayres & Cramton (1996).
- Crawford, Robert. See Klein et al. (1978).
- Crawford, Vincent (1982) "Compulsory Arbitration, Arbitral Risk and Negotiated Settlements: A Case Study in Bargaining under Imperfect Information," *Review of Economic Studies*, 49: 69-82 (January 1982).

- Crawford, Vincent & Hans Haller (1990) "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, 58: 571-597 (May 1990).
- Crawford, Vincent & Joel Sobel (1982) "Strategic Information Transmission," *Econometrica*, 50: 1431-1452 (November 1982).
- Crocker, Keith. See Masten & Crocker (1985).
- Dalkey, Norman (1953) "Equivalence of Information Patterns and Essentially Determinate Games," pp. 217-243 of Kuhn & Tucker (1953).
- Dasgupta, Partha. See Binmore & Dasgupta (1986).
- Dasgupta, Partha, Peter Hammond, & Eric Maskin (1979) "The Implementation of Social Choice Rules; Some General Rules on Incentive Compatibility," *Review of Economic Studies*, 46: 185-216 (April 1979).
- Dasgupta, Partha & Eric Maskin (1986a) "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," *Review of Economic Studies*, 53: 1-26 (January 1986).
- Dasgupta, Partha & Eric Maskin (1986b) "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies*, 53: 27-41 (January 1986).
- Dasgupta, Partha & Joseph Stiglitz (1980) "Uncertainty, Industrial Structure, and the Speed of R&D," *Bell Journal of Economics*, 11: 1-28 (Spring 1980).
- D'Aspremont, Claude, J. Gabszewicz & Jacques Thisse (1979) "On Hotelling's 'Stability of Competition'," *Econometrica*, 47: 1145-1150 (September 1979). Reprinted in Rasmusen (2000a).
- D'Aspremont, Claude & L. Gerard Varet (1979) "Incentives and Incomplete Information," *Journal of Public Economics*, 11: 25-45 (February 1979).
- David, Paul (1985) "CLIO and the Economics of QWERTY," *AEA Papers and Proceedings*, 75: 332-337 (May 1985).
- Davis, Morton (1970) *Game Theory: A Nontechnical Introduction*, New York: Basic Books, ISBN: 0465026265.
- Davis, Philip, Reuben Hersh & Elena Marchisotto (1981) *The Mathematical Experience*, Boston: Birkhauser, ISBN: 376433018X.
- Dawes, Robyn (1988) *Rational Choice in an Uncertain World*, Fort Worth, Texas: Harcourt Brace, ISBN: 0155752154.
- Dawkins, Richard (1989) *The Selfish Gene*, 2nd edn., Oxford: Oxford University Press (1st edn. 1976), ISBN: 0192177737.
- Debreu, Gerard (1952) "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences*, 38: 886-893 (1952).

- Debreu, Gerard (1959) *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, New Haven: Yale University Press, ISBN: 0300015593.
- Debreu, Gerard. See Arrow & Debreu (1954).
- Debreu, Gerard & Herbert Scarf (1963) "A Limit Theorem on the Core of an Economy," *International Economic Review*, 4: 235-246 (September 1963).
- DeBrock, Lawrence & J. Smith (1983) "Joint Bidding, Information Pooling, and the Performance of Petroleum Lease Auctions," *Bell Journal of Economics*, 14: 395-404 (Autumn 1983).
- Demsetz, Harold. See Alchian & Demsetz (1972).
- Deneckere, Raymond J. & R. Preston McAfee (1996) "Damaged Goods," *Journal of Economics and Management Strategy*, 5, 2: 149-174 (Summer 1996).
- Dewatripont, M. (1989) "Renegotiation and Information Revelation over Time in Optimal Labor Contracts," *Quarterly Journal of Economics*, 104: 589-620 (August 1989).
- Diamond, Douglas (1984) "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51: 393-414 (July 1984).
- Diamond, Douglas (1989) "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, 97: 828-862 (August 1989).
- Diamond, Douglas W. & P. Dybvig (1983), "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91: 401-419.
- Diamond, Peter (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90: 881-894.
- Diamond, Peter & Michael Rothschild, eds. (1978) *Uncertainty in Economics: Readings and Exercises*, New York: Academic Press, ISBN: 0122148509.
- Dimand, Mary Ann & Robert Dimand (1996) *A History of Game Theory*, London: Routledge, ISBN: 0415072573.
- Dimand, Mary Ann & Robert Dimand (1997) *The Foundations of Game Theory*, 3 vol., Cheltenham, England: Edward Elgar Publishing, ISBN: 1858982979.
- Dimand, Robert and Mohammed Dore (1999) "Cournot, Bertrand, and Game Theory: A Further Note," *Atlantic Economic Journal*, 27: 325-333 (September 1999).
- DiMona, Joseph. See Haldeman & DiMona (1978).
- Dixit, Avinash (1979) "A Model of Duopoly Suggesting a Theory of Entry Barriers," *Bell Journal of Economics*, 10: 20-32 (Spring 1979).
- Dixit, Avinash (1980) "The Role of Investment in Entry Deterrence," *Economic Journal*, 90: 95-106 (March 1980).

- Dixit, Avinash & Barry Nalebuff (1991) *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*, New York: Norton, ASIN: 0393029239.
- Dixit, Avinash & Susan Skeath (1998) *Games of Strategy*, New York: Norton, ISBN: xxx.
- Dixit, Avinash & Joseph Stiglitz (1977) "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67: 297-308 (June 1977).
- Dore, Mohammed. See Dimand & Dore (1999).
- Dranove, David. See Besanko, Dranove & Shanley (1996).
- Dresher, Melvin, Albert Tucker, & Philip Wolfe, eds. (1957) *Contributions to the Theory of Games, Volume III, Annals of Mathematics Studies, No. 39*, Princeton: Princeton University Press.
- Dubey, Pradeep, Ori Haimanko & Andriy Zapechelnyuk (2002) "Strategic Substitutes and Potential Games," working paper, SUNY Stonybrook, Dept. of Economics (11 April 2002).
- Dugatkin, Lee & Hudson Reeve, eds. (1998) *Game Theory & Animal Behavior*, Oxford: Oxford University Press, ASIN: 0195096924.
- Dutta, Prajit (1999) *Strategies and Games: Theory and Practice*, Cambridge: MIT Press, ISBN: 0262041693.
- Dybvig, P. see Diamond and Dybvig (1983).
- Dyer, Douglas, John H. Kagel, & Dan Levin (1989) "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *Economic Journal*, 99,394: 108-115 (March 1989).
- Eaton, C. & Richard Lipsey (1975) "The Principle of Minimum Differentiation Reconsidered: Some New Developments in the Theory of Spatial Competition," *Review of Economic Studies*, 42: 27-49 (January 1975).
- Eatwell, John, Murray Milgate & Peter Newman (1989) *The New Palgrave: Game Theory*, New York: W.W. Norton & Co., ASIN: 0393027333.
- Edgeworth, Francis (1897) "La Teoria Pura del Monopolio," *Giornale Degli Economisti*, 40: 13-31 (1925). Translated in pp. 111-142 of Edgeworth, Francis, *Papers Relating to Political Economy*, Vol. I. London: Macmillan (1925).
- Edgeworth, Francis (1922) "The Mathematical Economics of Professor Amoroso," *Economic Journal*, 30: 400-407 (September 1922). Reprinted in Rasmusen (2001).
- Ehrenberg, Ronald G. & Michael L. Bognanno (1990) "The Incentive Effects of Tournaments Revisited: Evidence From the European PGA Tour," *Industrial and Labor Relations Review*, 43: 74-88S.

- Eichberger, Jurgen (1993) *Game Theory for Economists*, San Diego: Academic Press, ISBN: 0122336208.
- Engers, Maxim (1987) "Signalling with Many Signals," *Econometrica*, 55: 663-674 (May 1987).
- Engers, Maxim & Luis Fernandez (1987) "Market Equilibrium with Hidden Knowledge and Self-Selection," *Econometrica*, 55: 425-439 (March 1987).
- Fama, Eugene (1980) "Banking in the Theory of Finance," *Journal of Monetary Economics*, 6: 39-57 (January 1980).
- Farrell, Joseph (forthcoming) "Monopoly Slack and Competitive Rigor: A Simple Model," MIT mimeo (February 1983). Published in Rasmusen (2001).
- Farrell, Joseph (1987) "Cheap Talk, Coordination, and Entry," *Rand Journal of Economics*, 18: 34-39 (Spring 1987). Reprinted in Rasmusen (2001).
- Farrell, Joseph & Matthew Rabin (1996) "Cheap Talk," *Journal of Economic Perspectives*, 10:103-118 (Summer 1996).
- Farrell, Joseph & Garth Saloner (1985) "Standardization, Compatibility, and Innovation," *Rand Journal of Economics*, 16: 70-83 (Spring 1985).
- Farrell, Joseph & Carl Shapiro (1988) "Dynamic Competition with Switching Costs," *Rand Journal of Economics*, 19: 123-137 (Spring 1988).
- Fernandez, Luis. See Engers & Fernandez (1987).
- Fisher, D. C. & J. Ryan (1992) "Optimal Strategies for a Generalized 'Scissors, Paper, and Stone' Game," *American Mathematical Monthly*, 99: 935-942.
- Fisher, Franklin (1989) "Games Economists Play: A Noncooperative View," *Rand Journal of Economics*, 20: 113-124 (Spring 1989).
- Flanagan, Thomas (1998) *Game Theory and Canadian Politics*, Toronto: University of Toronto Press, ISBN: 0802040942.
- Forgo, F. See Szep & Forgo (1978).
- Fowler, Henry (1965) *A Dictionary of Modern English Usage*, 2nd edn. New York: Oxford University Press, ASIN: 0195001540.
- Fowler, Henry & F. Fowler (1931) *The King's English*, 3rd edn. Oxford: Clarendon Press (1949).xxx
- Frank, Robert (1988) *Passions within Reason: The Strategic Role of the Emotions*, New York: Norton, ASIN: 0393026043.
- Franks, Julian, Robert Harris, & Colin Mayer (1988) "Means of Payment in Takeovers: Results for the U.K. and U.S.," in *Corporate Takeovers: Causes and Consequences*, Alan Auerbach, ed. Chicago: University of Chicago Press, ASIN: 0226032116.

- Freixas, Xavier, Roger Guesnerie, & Jean Tirole (1985) "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52: 173-191 (April 1985).
- Friedman, Daniel (1991) "Evolutionary Games in Economics," *Econometrica*, 59: 637-666 (May 1991).
- Friedman, David (2000) *Law's Order: An Economic Account*, Princeton: Princeton University Press, http://www.daviddfriedman.com/laws_order/index.shtml (viewed August 31, 2003).
- Friedman, James (1990) *Game Theory with Applications to Economics*, New York: Oxford University Press (1st edn. 1986).
- Friedman, Milton (1953) *Essays in Positive Economics*, Chicago: University of Chicago Press.
- Fudenberg, Drew & David Levine (1983) "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31: 251-268 (December 1983).
- Fudenberg, Drew & David Levine (1986) "Limit Games and Limit Equilibria," *Journal of Economic Theory*, 38: 261-279 (April 1986).
- Fudenberg, Drew & Eric Maskin (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54: 533-554 (May 1986).
- Fudenberg, Drew & Jean Tirole (1983) "Sequential Bargaining with Incomplete Information," *Review of Economic Studies*, 50: 221-247 (April 1983).
- Fudenberg, Drew & Jean Tirole (1986a) *Dynamic Models of Oligopoly*, Chur, Switzerland: Harwood Academic Publishers, ISBN: 3718602792.
- Fudenberg, Drew & Jean Tirole (1986b) "A Theory of Exit in Duopoly," *Econometrica*, 54: 943-960 (July 1986).
- Fudenberg, Drew & Jean Tirole (1986c) "A Signal-Jamming Theory of Predation," *Rand Journal of Economics*, 17: 366-376 (Autumn 1986).
- Fudenberg, Drew & Jean Tirole (1991a) *Game Theory*, Cambridge: MIT Press, ISBN: 0262061414.
- Fudenberg, Drew & Jean Tirole (1991b) "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53: 236-260 (April 1991).
- Gabszewicz, J. See d'Aspremont et al. (1979).
- Galbraith, John Kenneth (1954) *The Great Crash*, Boston: Houghton Mifflin (1954).
- Gale, Ian. See Che & Gale (1998).
- Gal-Or, Esther (1985) "First Mover and Second Mover Advantages," *International Economic Review*, 26: 649-653 (October 1985).

- Gardner, Roy, *Games for Business and Economics*, New York: John Wiley and Sons, ISBN: 0471311502.
- Gaskins, Darius (1974) "Alcoa Revisited: The Welfare Implications of a Second-Hand Market," *Journal of Economic Theory*, 7: 254-271 (March 1974).
- Gates, Scott & Brian Humes (1997) *Games, Information, and Politics: Applying Game Theoretic Models to Political Science*, Ann Arbor: University of Michigan Press, ISBN: 0472065645.
- Gaudet, Gerard & Stephen Salant (1991) "Uniqueness of Cournot Equilibrium: New Results from Old Methods," *Review of Economic Studies*, 58: 399-404 (April 1991).
- Gaver, Kenneth & Jerold Zimmerman (1977) "An Analysis of Competitive Bidding on BART Contracts," *Journal of Business*, 50: 279-295 (July 1977).
- Geanakoplos, John (1992) "Common Knowledge," *Journal of Economic Perspectives*, 6: 53-82 (Fall 1992).
- Geanakoplos, John. See Bulow et al. (1985).
- Geanakoplos, John & Heraklis Polemarchakis (1982) "We Can't Disagree Forever," *Journal of Economic Theory*, 28: 192- 200 (October 1982).
- Gelman, Judith & Steven Salop (1983) "Judo Economics: Capacity Limitation and Coupon Competition," *Bell Journal of Economics*, 14: 315-325 (Autumn 1983).
- Gertner, Robert. See Ayres & Gertner (1992), Baird, Gertner & Picker (1994).
- Ghemawat, Pankaj (1997) *Games Businesses Play: Cases and Models*, Cambridge: MIT Press, ISBN: 0262071827.
- Ghemawat, Pankaj & Barry Nalebuff (1985) "Exit," *Rand Journal of Economics*, 16: 184-194 (Summer 1985).
- Gibbard, Allan (1973) "Manipulation of Voting Schemes: A General Result," *Econometrica*, 41: 587-601 (July 1973).
- Gibbons, Robert (1992) *Game Theory for Applied Economists*, Princeton: Princeton University Press , ISBN: 0691003955.
- Gilbert, Richard & David Newbery (1982) "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 72: 514-526 (June 1982).
- Gillies, Donald (1953) "Locations of Solutions," pp 12-12 of *Report of an Informal Conference on the Theory of n-Person Games*, Princeton Mathematics mimeo (1953).
- Gintis, Herbert (2000) *Game Theory Evolving*, Princeton: Princeton University Press , ISBN: 0691009430.
- Gjesdal, Froystein (1982) "Information and Incentives: The Agency Information Problem," *Review of Economic Studies*, 49: 373-390 (July 1982).

- Glicksberg, Irving (1952) "A Further Generalization of the Kakutani Fixed Point Theorem with Application to Nash Equilibrium Points," *Proceedings of the American Mathematical Society*, 3: 170-174 (February 1952).
- Glosten, Lawrence & Paul Milgrom (1985) "Bid, Ask, and Transaction Prices in a Specialist Model with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14: 71-100 (March 1985).
- Goldfeld, Stephen. See Baumol & Goldfeld (1968).
- Gonik, Jacob (1978) "Tie Salesmen's Bonuses to their Forecasts," *Harvard Business Review*, 56: 116-123 (May/June 1978). Reprinted in Rasmusen (2001).
- Gordon, David. See Rapaport, Guyer & Gordon (1976).
- Graham, Daniel & Robert Marshall (1987) "Collusive Bidding Behavior at Single-Object Second- Price and English Auctions," *Journal of Political Economy*, 95: 1217-1239.
- Green, Edward (1984) "Continuum and Finite-Player Noncooperative Models of Competition," *Econometrica*, 52: 975-993 (July 1984).
- Green, Jerry. See Mas-Colell, Whinston & Green (1994).
- Greenhut, Melvin & Hiroshi Ohta (1975) *Theory of Spatial Pricing and Market Areas*, Durham, N.C.: Duke University Press, ASIN: 0822303337.
- Grinblatt, Mark & Chuan-Yang Hwang (1989) "Signalling and the Pricing of New Issues," *The Journal of Finance*, 44: 393-420 (June 1989).
- Grossman, Gene & Michael Katz (1983) "Plea Bargaining and Social Welfare," *American Economic Review*, 73: 749-757 (September 1983).
- Grossman, Sanford & Oliver Hart (1980) "Takeover Bids, the Free-Rider Problem, and the Theory of the Corporation," *Bell Journal of Economics*, 11: 42-64 (Spring 1980).
- Grossman, Sanford & Oliver Hart (1983) "An Analysis of the Principal Agent Problem," *Econometrica*, 51: 7-45 (January 1983).
- Groves, Theodore (1973) "Incentives in Teams," *Econometrica*, 41: 617-631 (July 1973).
- Guasch, J. Luis & Andrew Weiss (1980) "Wages as Sorting Mechanisms in Competitive Markets with Asymmetric Information: A Theory of Testing," *Review of Economic Studies*, 47: 653-664 (July 1980).
- Guesnerie, Roger. See Freixas et al. (1985).
- Gul, Faruk (1989) "Bargaining Foundations of Shapley Value," *Econometrica*, 57: 81-96 (January 1989).

- Guth, Wener, Rold Schmittberger, & Bernd Schwarze (1982) "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, 3: 367-388 (December 1982).
- Guyer, Melvin. See Rapaport, Guyer & Gordon (1976).
- Haimanko, Ori. See Dubey, Haimanko & Zapechelnyuk (2002).
- Haldeman, H. R. & Joseph DiMona (1978) *The Ends of Power*, New York: Times Books, ASIN: 0812907248.
- Haller, Hans (1986) "Noncooperative Bargaining of $N \geq 3$ Players," *Economic Letters*, 22: 11-13 (1986).
- Haller, Hans. See Crawford & Haller(1990).
- Halmos, Paul (1970) "How to Write Mathematics," *L'Enseignement Mathématique*, 16: 123-152 (May/June 1970).
- Haltiwanger, John & Michael Waldman (1991) "Responders versus Nonresponders: A New Perspective of Heterogeneity," *Economic Journal*, 101: 1085-102 (September 1991).
- Hamilton, William. See Axelrod & Hamilton (1981).
- Harrington, Joseph (1987) "Collusion in Multiproduct Oligopoly Games under a Finite Horizon," *International Economic Review*, 28: 1-14 (February 1987).
- Hammond, Peter. See Dasgupta et al. (1979).
- Han Fei Tzu (c. 250 B.C.) *Basic Writings*, translated by Burton Watson, New York: Columbia University Press (1964).
- Harris, Milton (1987) *Dynamic Economic Analysis*, Oxford, Oxford University Press, ISBN: 0195044061.
- Harris, Milton & Bengt Holmstrom (1982) "A Theory of Wage Dynamics," *Review of Economic Studies*, 49: 315-334 (July 1982).
- Harris, Milton & Arthur Raviv (1992) "Financial Contracting Theory," in *Advances in Economic Theory: Sixth World Congress*, Jean-Jacques Laffont, ed., Cambridge: Cambridge University Press, ASIN: 0521430194.
- Harris, Milton & Arthur Raviv (1995) "The Role of Games in Security Design," *Review of Financial Studies*, 8: 327-367 (Summer 1995).
- Harris, Robert. See Franks et al. (1988).
- Harsanyi, John (1967) "Games with Incomplete Information Played by 'Bayesian' Players, I: The Basic Model," *Management Science*, 14: 159-182 (November 1967).

- Harsanyi, John (1968a) "Games with Incomplete Information Played by 'Bayesian' Players, II: Bayesian Equilibrium Points," *Management Science*, 14: 320-334 (January 1968).
- Harsanyi, John (1968b) "Games with Incomplete Information Played by 'Bayesian' Players, III: The Basic Probability Distribution of the Game," *Management Science*, 14: 486-502 (March 1968) .
- Harsanyi, John (1973) "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2: 1-23 (1973).
- Harsanyi, John (1977) *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, New York: Cambridge University Press, ISBN: 0521208866.
- Harsanyi, John & Reinhard Selten (1988) *A General Theory of Equilibrium Selection in Games*, Cambridge: MIT Press, ISBN: 0262582384.
- Hart, Oliver. See Grossman & Hart (1980, 1983).
- Hart, Oliver & Bengt Holmstrom (1987) "The Theory of Contracts," in Bewley (1987).
- Hart, Sergiu. See Aumann & Hart (1992).
- Hausman, Jerry & James Poterba (1987) "Household Behavior and the Tax Reform Act of 1986," *Journal of Economic Perspectives*, 1: 101-119 (Summer 1987).
- Haywood, O. (1954) "Military Decisions and Game Theory," *Journal of the Operations Research Society of America*, 2: 365-385 (November 1954).
- Henderson, Dale. See Canzoneri & Henderson (1991).
- Hendricks, Ken, Andrew Weiss & Charles A. Wilson (1988) "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 29, 4: 663-680 (November 1988) .
- Hendricks, Ken & Robert Porter (1988) "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78: 865-883 (December 1988).
- Henry, O. (1945) *Best Stories of O. Henry*, Garden City, NY: The Sun Dial Press (1945).
- Herodotus (c. 429 B.C.) *The Persian Wars*, George Rawlinson, translator, New York: Modern Library (1947).
- Hersh, Reuben. See Davis & Hersh (1981).
- Hines, W. (1987) "Evolutionary Stable Strategies: A Review of Basic Theory," *Theoretical Population Biology*, 31: 195-272 (April 1987).
- Hinkley, David. See Cox & Hinkley (1974).

- Hirshleifer, David (1995) "The Blind Leading the Blind: Social Influence, Fads, and Informational Cascades," pp. 188-215 (chapter 12) of *The New Economics of Human Behavior*, edited by Mariano Tommasi and Kathryn Ierulli, Cambridge: Cambridge University Press, ISBN: 0521479495
- Hirshleifer, David & Eric Rasmusen (1989) "Cooperation in a Repeated Prisoner's Dilemma with Ostracism," *Journal of Economic Behavior and Organization*, 12: 87-106 (August 1989).
- Hirshleifer, David & Sheridan Titman (1990) "Share Tendering Strategies and the Success of Hostile Takeover Bids," *Journal of Political Economy*, 98: 295-324 (April 1990).
- Hirshleifer, David. See Png & Hirshleifer (1987), Bikhchandani, Hirshleifer & Welch (1992).
- Hirshleifer, Jack (1982) "Evolutionary Models in Economics and Law: Cooperation versus Conflict Strategies," *Research in Law and Economics*, 4: 1-60 (1982).
- Hirshleifer, Jack (1987) "On the Emotions as Guarantors of Threats and Promises," In *The Latest on the Best: Essays on Evolution and Optimality*, edited by John Dupre. Cambridge: MIT Press (1987). ASIN: 0262040905
- Hirshleifer, Jack & Juan Martinez-Coll (1988) "What Strategies can Support the Evolutionary Emergence of Cooperation?" *Journal of Conflict Resolution*, 32: 367-98 (June 1988).
- Hirshleifer, Jack & Eric Rasmusen (1992) "Are Equilibrium Strategies Unaffected By Incentives?" *Journal of Theoretical Politics*, 4: 343-57 (July 1992).
- Hirshleifer, Jack & John Riley (1979) "The Analytics of Uncertainty and Information: An Expository Survey," *Journal of Economic Literature*, 17: 1375-421 (December 1979).
- Hirshleifer, Jack & John Riley (1992) *The Analytics of Uncertainty and Information*, Cambridge: Cambridge University Press, ASIN: 0521239567
- Hoffman, Elizabeth & Matthew Spitzer (1985) "Entitlements, Rights and Fairness: An Experimental Examination of Subjects' Concepts of Distributive Justice," *Journal of Legal Studies*, 14: 269-297 (June 1985).
- Hofstadter, Douglas (1983) "Computer Tournaments of the Prisoner's Dilemma Suggest how Cooperation Evolves," *Scientific American*, 248: 16-26 (May 1983).
- Holmes, Oliver (1881) *The Common Law*, Boston: Little, Brown and Co. (1923).
- Holmstrom, Bengt (1979) "Moral Hazard and Observability," *Bell Journal of Economics*, 10: 74-91 (Spring 1979).
- Holmstrom, Bengt (1982) "Moral Hazard in Teams," *Bell Journal of Economics*, 13: 324-40 (Autumn 1982).

- Holmstrom, Bengt & Paul Milgrom (1987) "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55: 303-328 (March 1987).
- Holmstrom, Bengt & Paul Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics and Organization*, 7: 24-52 (Special Issue, 1991).
- Holmstrom, Bengt & Roger Myerson (1983) "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51: 1799-819 (November 1983).
- Holmstrom, Bengt. See Harris & Holmstrom (1982) and Hart & Holmstrom (1987).
- Holt, Charles & David Scheffman (1987) "Facilitating Practices: The Effects of Advance Notice and Best-Price Policies," *Rand Journal of Economics*, 18: 187-97 (Summer 1987).
- Hotelling, Harold (1929) "Stability in Competition," *Economic Journal*, 39: 41-57 (March 1929). Reprinted in Rasmusen (2001).
- Hughes, Patricia (1986) "Signalling by Direct Disclosure Under Asymmetric Information," *Journal of Accounting and Economics*, 8: 119-142 (June 1986).
- Humes, Brian. See Gates & Humes (1996).
- Hwang, Chuan-Yang. See Grinblatt & Hwang (1989).
- Intriligator, Michael (1971) *Mathematical Optimization and Economic Theory*, Englewood Cliffs, NJ: Prentice-Hall, ASIN: 0135617537.
- Isoda, Kazuo. See Nikaido & Isoda (1955).
- Jarrell, Gregg & Sam Peltzman (1985) "The Impact of Product Recalls on the Wealth of Sellers," *Journal of Political Economy*, 93: 512-536 (June 1985).
- Jensen, Michael. See Baker et al. (1988).
- Johnson, Norman & Samuel Kotz (1970) *Distributions in Statistics*, 3 vol., New York: John Wiley and Sons (1970).
- Joskow, Paul (1985) "Vertical Integration and Longterm Contracts: the Case of Coal-Burning Electric Generating Plants," *Journal of Law, Economics and Organization*, 1: 33-80 (Spring 1985).
- Joskow, Paul (1987) "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets," *American Economic Review*, 77: 168-185 (March 1987).
- Kahneman, Daniel, Paul Slovic, & Amos Tversky, eds. (1982) *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press, ISBN: 0521240646.
- Kakutani, Shizuo (1941) "A Generalization of Brouwer's Fixed Point Theorem," *Duke Mathematical Journal*, 8: 457-9 (September 1941).

- Kalai, Ehud, Dov Samet & William Stanford (1988) "Note on Reactive Equilibria in the Discounted Prisoner's Dilemma and Associated Games," *International Journal of Game Theory*, 17: 177-186 (1988).
- Kamien, Morton & Nancy Schwartz (1982) *Market Structure and Innovation*, Cambridge: Cambridge University Press, ASIN: 0521221900
- Kamien, Morton & Nancy Schwartz (1991) *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*, 2nd edn. New York: North Holland (1991, 1st edn. 1981), 0444016090.
- Kandori, M. and H. Matsushima (1998) "Private Observation, Communication, and Collusion," *Econometrica* 66: 627-652.
- Kaplow, Louis & Steven Shavell (1996) "Property Rules Versus Liability Rules: An Economic Analysis," *Harvard Law Review*, 109: 713-89 (February 1996).
- Karlin, Samuel (1959) *Mathematical Methods and Theory in Games, Programming and Economics*, Reading: Addison-Wesley (1959).
- Katz, Lawrence (1986) "Efficiency Wage Theory: A Partial Evaluation," In *NBER Macroeconomics Annual 1986*, edited by Stanley Fischer. Cambridge: MIT Press, ASIN: 0262061058.
- Katz, Michael & Carl Shapiro (1985) "Network Externalities, Competition, and Compatibility," *American Economic Review*, 75: 424-40 (June 1985).
- Katz, Michael. See Grossman & Katz (1983).
- Katz, Michael. See Moskowitz et al. (1980).
- Kennan, John & Robert Wilson (1993) "Bargaining with Private Information," *Journal of Economic Literature*, 31: 45-104 (March 1993).
- Kennedy, Peter (1979) *A Guide to Econometrics*, 1st edn., Cambridge: MIT Press 1979 (3rd edn. 1992), ISBN: 0262110733.
- Keynes, John Maynard (1933) *Essays in Biography*, New York: Harcourt, Brace and Company (1933).
- Keynes, John Maynard (1936) *The General Theory of Employment, Interest and Money*, London: Macmillan (1947).
- Kierkegaard, Søren (1938) *The Journals of Søren Kierkegaard*, translated by Alexander Dru, Oxford: Oxford University Press (1938).
- Kihlstrom, Richard & Michael Riordan (1984) "Advertising as a Signal," *Journal of Political Economy*, 92: 427-50 (June 1984).
- Kilgour, D. Marc. See Brams & Kilgour (1988).
- Kindleberger, Charles (1983) "Standards as Public, Collective and Private Goods," *Kyklos*, 36: 377-96 (1983).

- Klein, Benjamin, Robert Crawford, & Armen Alchian (1978) "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *Journal of Law and Economics*, 21: 297- 326 (October 1978).
- Klein, Benjamin & Keith Leffler (1981) "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89: 615-41 (August 1981).
- Klein, Benjamin & Lester Saft (1985) "The Law and Economics of Franchise Tying Contracts," *Journal of Law and Economics*, 28: 345-61 (May 1985).
- Klemperer, Paul (1987) "The Competitiveness of Markets with Switching Costs," *Rand Journal of Economics*, 18: 138-50 (Spring 1987).
- Klemperer, Paul (1998) "Auctions with Almost Common Values: The 'Wallet Game' and Its Applications," *European Economic Review*, xxx (xxx 1998).
- Klemperer, Paul (1999) "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys*, 13: 227-86 (July 1999).
- Klemperer, Paul, ed. (2000) *The Economic Theory of Auctions*. Cheltenham, England: Edward Elgar (2000).
- Klemperer, Paul. See Bulow, Geanakoplos & Klemperer (1985) and Bulow & Klemperer (1996).
- Kohlberg, Elon & Jean-Francois Mertens (1986) "On the Strategic Stability of Equilibria," *Econometrica*, 54: 1003-7 (September 1986).
- Kotz, Samuel. See Johnson & Kotz (1970).
- Kreps, David (1990a) *A Course in Microeconomic Theory*, Princeton: Princeton University Press (1990). ISBN: 0691042640
- Kreps, David (1990b) *Game Theory and Economic Modeling*, Oxford: Oxford University Press, ISBN: 0198283571
- Kreps, David. See Cho & Kreps (1987).
- Kreps, David, Paul Milgrom, John Roberts, & Robert Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27: 245-52 (August 1982). Reprinted in Rasmusen (2001).
- Kreps, David & Jose Scheinkman (1983) "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14: 326-37 (Autumn 1983).
- Kreps, David & A. Michael Spence (1985) "Modelling the Role of History in Industrial Organization and Competition," in *Issues in Contemporary Microeconomics and Welfare*, edited by George Feiwel. London: Macmillan, ISBN: 0333354826
- Kreps, David & Robert Wilson (1982a) "Reputation and Imperfect Information," *Journal of Economic Theory*, 27: 253-79 (August 1982).

- Kreps, David & Robert Wilson (1982b) "Sequential Equilibria," *Econometrica*, 50: 863-94 (July 1982).
- Krishna, Vijay. See Benoit & Krishna (1985).
- Krouse, Clement (1990) *Theory of Industrial Economics*, Oxford: Blackwell (1990).
- Kuhn, Harold (1953) "Extensive Games and the Problem of Information," In Kuhn & Tucker (1953).
- Kuhn, Harold (ed.) (1997) *Classics in Game Theory*, Princeton: Princeton University Press, ISBN: 0691011923.
- Kuhn, Harold & Albert Tucker, eds. (1950) *Contributions to the Theory of Games, Volume I, Annals of Mathematics Studies, No. 24*, Princeton: Princeton University Press, ISBN: 069107934X.
- Kuhn, Harold & Albert Tucker, eds. (1953) *Contributions to the Theory of Games, Volume II, Annals of Mathematics Studies, No. 28*, Princeton: Princeton University Press (1953).
- Kydland, Finn & Edward Prescott (1977) "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85: 473-491 (June 1977).
- Kyle, Albert (1985) "Continuous Auctions and Insider Trading," *Econometrica* 53: 1315-1336 (November 1985).
- Laffont, Jean-Jacques (1997) "Game Theory and Empirical Economics: The Case of Auction Data," *European Economic Review*, 41: 1-35 (January 1997).
- Laffont, Jean-Jacques & Jean Tirole (1986) "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94: 614-641 (June 1986).
- Laffont, Jean-Jacques & Jean Tirole (1993) *A Theory of Incentives in Procurement and Regulation*, Cambridge: MIT Press, ISBN: 0262121743.
- Lakatos, Imre (1976) *Proofs and Refutations: The Logic of Mathematical Discovery*, Cambridge: Cambridge University Press, ISBN: 0521290384.
- Lane, W. (1980) "Product Differentiation in a Market with Endogenous Sequential Entry," *Bell Journal of Economics*, 11: 237-260 (Spring 1980).
- Layard, Richard & George Psacharopoulos (1974) "The Screening Hypothesis and the Returns to Education," *Journal of Political Economy*, 82: 985-998 (September/October 1974).
- Lazear AER 1999 paper.xxx For 4th ediotn. AS nad MS both happened after a certain job went from hourly to piece rate. The people changed. Aug 99, AER, 90: 1346-1361.
- Lazear, Edward & Sherwin Rosen (1981) "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89: 841-864 (October 1981).

- Leffler, Keith. See Klein & Leffler (1981).
- Leibenstein, Harvey (1950) "Bandwagon, Snob and Veblen Effects in the Theory of Consumers' Demand," *Quarterly Journal of Economics*, 64: 183-207 (May 1950).
- Leland, Hayne & David Pyle (1977) "Informational Asymmetries, Financial Structure, and Financial Intermediation," *Journal of Finance*, 32: 371-387 (May 1977).
- Leshno, Moshe, Haim Levy & Yishay Spector (1997) "A Comment on Rothschild and Stiglitz's 'Increasing Risk I: A Definition'," *Journal of Economic Theory*, 77: 223-228 (November 1997).
- Levering, Robert. See Moskowitz et al. (1980).
- Levine, David. See Fudenberg & Levine (1983, 1986).
- Levitan, Richard & Martin Shubik (1972) "Price Duopoly and Capacity Constraints," *International Economic Review*, 13: 111-122 (February 1972).
- Levmore, Saul (1982) "Self-Assessed Valuation for Tort and Other Law," *Virginia Law Review*, 68: 771-861 (April 1982).
- Levy, Haim. See Leshno, Levy & Spector (1997).
- Lewis, David (1969) *Convention: A Philosophical Study*, Cambridge: Harvard University Press, ASIN: 0674170253.
- Liebowitz, S. & Stephen Margolis (1990) "The Fable of the Keys," *Journal of Political Economy*, 33: 1-25 (April 1990). Reprinted in *Famous Fables of Economics: Myths of Market Failures*, Daniel F. Spulber (Editor), Oxford: Blackwell Publishers (2001), ISBN: 0631226753.
- Lipsey, Richard. See Eaton & Lipsey (1975).
- Locke, E. (1949) "The Finan-Seer," *Astounding Science Fiction*, 44: 132-140 (October 1949).
- Lorberbaum, Jeffrey. See Boyd & Lorberbaum (1987).
- Loury, Glenn (1979) "Market Structure and Innovation," *Quarterly Journal of Economics*, 93: 395-410 (August 1979).
- Lucas, Robert. See Stokey & Lucas (1989).
- Luce, R. Duncan & Howard Raiffa (1957) *Games and Decisions: Introduction and Critical Survey*, New York: Wiley.
- Luce, Duncan & Albert Tucker, eds. (1959) *Contributions to the Theory of Games, Volume IV, Annals of Mathematics Studies, No. 40*, Princeton: Princeton University Press.
- Macaulay, Stewart (1963) "Non-Contractual Relations in Business," *American Sociological Review*, 28: 55-70 (February 1963).

- Macey, Jonathan & Fred McChesney (1985) "A Theoretical Analysis of Corporate Greenmail," *Yale Law Journal*, 95: 13-61 (November 1985).
- Macho-Stadler, Ines & J. David Perez-Castillo (1997) *An Introduction to the Economics of Information: Incentives and Contracts*, Oxford: Oxford University Press, ASIN: 0198774672
- Macrae, Norman (1992) *John von Neumann*, New York: Random House (1992).
- Margolis, Stephen. See Liebowitz & Margolis (1990).
- Marschak, Jacob & Roy Radner (1972) *Economic Theory of Teams*, New Haven: Yale University Press, ASIN: 0300012799.
- Marshall, Alfred (1961) *Principles of Economics*, 9th (variorum) edn. London, Macmillan (1st edn. 1890).
- Martin, Stephen (1993) *Advanced Industrial Economics*, Oxford: Blackwell Publishers, ASIN: 063117852X.
- Martinez-Coll. See J. Hirshleifer & Martinez-Coll (1988).
- Mas-Colell, Andreu, Michael Whinston & Jerry Green (1995) *Microeconomic Theory*, Oxford: Oxford University Press, ISBN: 0195073401.
- Maskin, Eric. See Dasgupta & Maskin (1986a, 1986b), Dasgupta et al. (1979), and Fudenberg & Maskin (1986).
- Maskin, Eric & John Riley (1984) "Optimal Auctions with Risk Averse Buyers," *Econometrica*, 52: 1473-1518 (November 1984).
- Maskin, Eric & John Riley (1985) "Input vs. Output Incentive Schemes," *Journal of Public Economics*, 28: 1-23 (October 1985).
- Maskin, Eric & Jean Tirole (1987) "Correlated Equilibria and Sunspots," *Journal of Economic Theory*, 43: 364-373 (December 1987).
- Masten, Scott & Keith Crocker (1985) "Efficient Adaptation in Long-Term Contracts: Take-or-Pay Provisions for Natural Gas," *American Economic Review*, 75: 1083-1093 (December 1985).
- Matthews, Steven (1987). "Comparing Auctions for Risk-Averse Buyers: A Buyer's Point of View," *Econometrica*, 55: 633-646. xxx MONTH?
- Mathewson, G. Frank & Ralph Winter (1985) "The Economics of Franchise Contracts," *Journal of Law and Economics*, 28: 503-526 (October 1985).
- Matsushima, H. See Kandori & Matsushima (1998).
- Matthews, Steven & John Moore (1987) "Monopoly Provision of Quality and Warranties: An Exploration in the Theory of Multidimensional Screening," *Econometrica*, 55: 441-467 (March 1987).

- Mayer, Colin. See Franks et al. (1988).
- Maynard Smith, John (1974) "The Theory of Games and the Evolution of Animal Conflicts," *Journal of Theoretical Biology*, 47: 209-221 (September 1974).
- Maynard Smith, John (1982) *Evolution and the Theory of Games*, Cambridge: Cambridge University Press, ISBN: 0521288843.
- McAfee, R. Preston & John McMillan (1986) "Bidding for Contracts: A Principal-Agent Analysis," *Rand Journal of Economics*, 17: 326- 338 (Autumn 1986).
- McAfee, R. Preston & John McMillan (1987) "Auctions and Bidding," *Journal of Economic Literature*, 25: 699-754 (June 1987).
- McAfee, R. Preston & John McMillan (1988) *Incentives in Government Contracts*, Toronto: University of Toronto Press, ISBN: 0802066380.
- McAfee, R. Preston & John McMillan (1996) "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10: 159-175 (Winter 1996). Reprinted in Rasmusen (2001).
- McChesney, Fred. See Macey & McChesney (1985).
- McCloskey, Donald (1985) "Economical Writing," *Economic Inquiry*, 24: 187-222 (April 1985).
- McCloskey, Donald (1987) *The Writing of Economics*, New York: Macmillan, ASIN: 0023795204.
- McDonald, John & John Tukey (1949) "Colonel Blotto: A Problem of Military Strategy," *Fortune* p. 102 (June 1949). Reprinted in Rasmusen (2001).
- McGee, John (1958) "Predatory Price Cutting: The Standard Oil (N.J.) Case," *Journal of Law and Economics*, 1: 137-169 (October 1958).
- McMillan, John (1986) *Game Theory in International Economics*, Chur, Switzerland: Harwood Academic Publishers, ISBN: 3718602776.
- McMillan, John (1992) *Games, Strategies, and Managers: How Managers can use Game Theory to Make Better Business Decisions*, Oxford, Oxford University Press, ISBN: 0195074300.
- McMillan, John. See McAfee & McMillan (1986, 1987, 1988, 1996).
- Mead, Walter, Asbjorn Moseidjord, & Philip Sorenson (1984) "Competitive Bidding under Asymmetrical Information: Behavior and Performance in Gulf of Mexico Drainage Lease Sales 1959-1969," *Review of Economics and Statistics*, 66: 505-508 (August 1984).
- Meckling, William. See Jensen & Meckling (1976).
- Meese, G. See Baldwin & Meese (1979).

- Mertens, Jean-Francois & S. Zamir (1985) "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory* 14: 1-29 (1985).
- Milgate, Murray. See Eatwell et al. (1989).
- Milgrom, Paul. See Kreps et al. (1982), Glosten & Milgrom (1985), and Holmstrom & Milgrom (1987, 1991).
- Milgrom, Paul (1981a) "An Axiomatic Characterization of Common Knowledge" *Econometrica*, 49: 219-222 (January 1981).
- Milgrom, Paul (1981b) "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12: 380-391 (Autumn 1981).
- Milgrom, Paul (1981c) "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49: 921-943 (July 1981).
- Milgrom, Paul (1987) "Auction Theory," in Truman, Bewley, ed., *Advances in Economic Theory, Fifth World Congress*, Cambridge: Cambridge University Press, ASIN: 0521340446.
- Milgrom, Paul (1999) *Auction Theory for Privatization*, Cambridge: Cambridge University Press (1999). ISBN: xxx.
- Milgrom, Paul & John Roberts (1982a) "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50: 443-459 (March 1982).
- Milgrom, Paul & John Roberts (1982b) "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27: 280-312 (August 1982).
- Milgrom, Paul & John Roberts (1986) "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94: 796-821 (August 1986).
- Milgrom, Paul & John Roberts (1990) "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58: 1255-1279 (November 1990).
- Milgrom, Paul & John Roberts (1992) *Economics, Organizations, and Management*, Englewood Cliffs, New Jersey: Prentice-Hall, ISBN: 0132246503.
- Milgrom, Paul & Robert Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089-1122 (September 1982).
- Miller, Geoffrey (1986) "An Economic Analysis of Rule 68," *Journal of Legal Studies*, 15: 93-125 (January 1986).
- Mirrlees, James (1971) "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38: 175-208 (April 1971).
- Mirrlees, James (1974) "Notes on Welfare Economics, Information and Uncertainty," In *Essays on Economic Behavior under Uncertainty*, edited by M. Balch, Daniel McFadden, and S. Wu. Amsterdam: North Holland.

- Monteverde, K. & David Teece (1982) "Supplier Switching Costs and Vertical Integration in the Automobile Industry," *Bell Journal of Economics*, 13: 206-213 (Spring 1982).
- Mookherjee, Dilip (1984) "Optimal Incentive Schemes with Many Agents," *Review of Economic Studies*, 51: 433-446 (July 1984).
- Mookherjee, Dilip & Ivan Png (1989) "Optimal Auditing, Insurance, and Redistribution," *The Quarterly Journal of Economics*, 104: 399-415 (May 1989).
- Moore, John. See Matthews & Moore (1987).
- Moreaux, Michel (1985) "Perfect Nash Equilibria in Finite Repeated Game and Uniqueness of Nash Equilibrium in the Constituent Game," *Economics Letters*, 17: 317-320 (1985)
- Morgenstern, Oskar. See von Neumann & Morgenstern (1944).
- Morris, Peter (1994) *Introduction to Game Theory*, Berlin: Springer Verlag (1994).
- Morrison, Clarence (1998) "Cournot, Bertrand, and Modern Game Theory," *Atlantic Economic Review*, 26: 172-174 (June 1998).
- Morrow, James (1994) *Game Theory for Political Scientists*, Princeton: Princeton University Press, ISBN: 0691034303.
- Moseidjord, Asbjorn. See Mead et al. (1984).
- Moulin, Herve (1986) *Game Theory for the Social Sciences*, 2nd and revised edn. New York, NYU Press, ISBN: 0814754317.
- Moulin, Herve (1986) *Eighty-Nine Exercises with Solutions from Game Theory of the Social Sciences*, 2nd and revised edn. New York, NYU Press, ASIN: 0814754325.
- Mulherin, J. Harold (1986) "Complexity in Long-term Contracts: An Analysis of Natural Gas Contractual Provisions," *Journal of Law, Economics, and Organization*, 2: 105-118 (Spring 1986).
- Murphy, Kevin J. (1986) "Incentives, Learning, and Compensation: A Theoretical and Empirical Investigation of Managerial Labor Contracts," *Rand Journal of Economics*, 17: 59-76 (Spring 1986).
- Murphy, Kevin J. See Baker et al. (1988).
- Muthoo, Abhinay (1999) *Bargaining Theory With Applications*, Cambridge: Cambridge University Press, ISBN: 0521572258.
- Muzzio, Douglas (1982) *Watergate Games*, New York: New York University Press, ISBN: 0814753841.
- Myerson, Roger (1979) "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47: 61-73 (January 1979).

- Myerson, Roger (1981). "Optimal Auction Design," *Mathematics of Operations Research*, 6: 58-73.
- Myerson, Roger (1991) *Game Theory: Analysis of Conflict*, Cambridge: Harvard University Press, ISBN: 0674341155.
- Myerson, Roger (1999) "Nash Equilibrium and the History of Game Theory," *Journal of Economic Literature* 37: 1067-1082 (September 1999).
- Myerson, Roger. See Holmstrom & Myerson (1983).
- Myerson, Roger & Mark Satterthwaite (1983) "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29:1-21 (April 1983).
- Nalebuff, Barry. See Ghemawat & Nalebuff (1985), Dixit & Nalebuff (1991).
- Nalebuff, Barry & John Riley (1985) "Asymmetric Equilibria in the War of Attrition," *Journal of Theoretical Biology*, 113: 517-527 (April 1985).
- Nalebuff, Barry & David Scharfstein (1987) "Testing in Models of Asymmetric Information," *Review of Economic Studies*, 54: 265-278 (April 1987).
- Nalebuff, Barry & Joseph Stiglitz (1983) "Prizes and Incentives: Towards a General Theory of Compensation and Competition," *Bell Journal of Economics*, 14: 21-43 (Spring 1983).
- Nasar, Sylvia (1998) *A Beautiful Mind*, New York: Simon and Schuster, ISBN: 0684819066.
- Nash, John (1950a) "The Bargaining Problem," *Econometrica*, 18: 155-162 (January 1950). Reprinted in Rasmusen (2000a).
- Nash, John (1950b) "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences, USA*, 36: 48-49 (January 1950). Reprinted in Rasmusen (2001).
- Nash, John (1951) "Non-Cooperative Games," *Annals of Mathematics*, 54: 286-295 (September 1951). Reprinted in Rasmusen (2001).
- Nelson, Philip (1974) "Advertising as Information," *Journal of Political Economy*, 84: 729-754 (July/August 1974).
- Newbery, David. See Gilbert & Newbery (1982).
- Newman, John. See Eatwell et al. (1989).
- Novshek, William (1985) "On the Existence of Cournot Equilibrium," *Review of Economic Studies*, 52: 85-98 (January 1985).
- Ohta, Hiroshi. See Greenhut & Ohta (1975).
- Olsder, Geert. See Basar & Olsder (1999).

- Ordeshook, Peter (1986) *Game Theory and Political Theory: An Introduction*, Cambridge: Cambridge University Press, ASIN: 0521306124.
- Osborne, Martin (2003) *An Introduction to Game Theory*, Oxford: Oxford University Press, ISBN: 0195128958.
- Osborne, Martin & Carolyn Pitchik (1986) "The Nature of Equilibrium in a Location Model" *International Economic Review*, 27: 223-237 (February 1986).
- Osborne, Martin & Carolyn Pitchik (1987) "Equilibrium in Hotelling's Model of Spatial Competition," *Econometrica*, 55: 911-22 (July 1987).
- Osborne, Martin & Ariel Rubinstein (1994) *A Course in Game Theory*, Cambridge: MIT Press, ISBN: 0262650401.
- Owen, Guillermo (1995) *Game Theory*, 3rd edn., New York: Academic Press (1st edn. 1968), ISBN: 0125311516.
- Palfrey, Thomas & Sanjay Srivastava (1993) *Bayesian Implementation*, New York: Harwood Academic Publishers, ISBN: 3718653141.
- Pearce, David (1984) "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52: 1029-1050 (July 1984).
- Pearce, David. See Abreu et al. (1986, 1990).
- Peleg, Bezalel. See Bernheim et al. (1987).
- Peltzman, Sam (1991) "The Handbook of Industrial Organization: A Review Article," *Journal of Political Economy*, 99: 201-217 (February 1991).
- Peltzman, Sam. See Jarrell & Peltzman (1985).
- Perez-Castillo. See Macho-Stadler & Perez-Castillo (1997).
- Perry, Motty (1986) "An Example of Price Formation in Bilateral Situations: A Bargaining Model with Incomplete Information," *Econometrica*, 54: 313-21 (March 1986).
- Petrakis, Emmanuel. See Rasmusen & Petrakis (1992).
- Phlips, Louis (1983) *The Economics of Price Discrimination*, Cambridge: Cambridge University Press, ISBN: 0521283949.
- Phlips, Louis (1988) *The Economics of Imperfect Information*, Cambridge: Cambridge University Press, ISBN: 0521309204.
- Picker, Randal. See Baird, Gertner & Picker (1994).
- Pigou, A. (1920), *The Economics of Welfare*, 4th edn. (1932), 1st edn. (1920), London: Macmillan and Company (1952, 1st edn. 1920).
- Pitchik, Carolyn. See Osborne & Pitchik (1986, 1987).

- Pliny the Younger (1963) *The Letters of the Younger Pliny*, translated by Betty Radice, Baltimore: Penguin Books (1963).
- Png, Ivan (1983) "Strategic Behaviour in Suit, Settlement, and Trial," *Bell Journal of Economics*, 14: 539-550 (Autumn 1983).
- Png, Ivan. See Mookherjee & Png (1989).
- Png, Ivan & David Hirshleifer (1987) "Price Discrimination through Offers to Match Price," *Journal of Business*, 60: 365-383 (July 1987).
- Polemarchakis, Heraklis. See Geanakoplos & Polemarchakis (1982).
- Polinsky, A. Mitchell & Yeon-Koo Che (1991) "Decoupling Liability: Optimal Incentives for Care and Litigation," *Rand Journal of Economics*, 22: 562-570 (Winter 1991).
- Popkin, Samuel (1979) *The Rational Peasant: The Political Economy of Rural Society in Vietnam*, Berkeley: University of California Press, ISBN: 0520039548.
- Porter, Robert (1983a) "Optimal Cartel Trigger Price Strategies," *Journal of Economic Theory*, 29: 313-338 (April 1983).
- Porter, Robert (1983b) "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, 14: 301-314 (Autumn 1983).
- Porter, Robert. See Hendricks & Porter (1988).
- Posner, Richard (1975) "The Social Costs of Monopoly and Regulation," *Journal of Political Economy*, 83: 807-827 (August 1975).
- Poterba, James. See Hausman & Poterba (1987).
- Prescott, Edward. See Kydland & Prescott (1977).
- Poundstone, William (1992) *Prisoner's Dilemma: John von Neumann, Game Theory, and the Puzzle of the Bomb*, New York: Doubleday, ISBN: 0385415672.
- Psacharopoulos, George. See Layard & Psacharopoulos (1974).
- Pyle, David. See Leland & Pyle (1977).
- Rabin, Matthew. See Farrell & Rabin (1996).
- Radner, Roy (1980) "Collusive Behavior in Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22: 136-156 (April 1980).
- Radner, Roy (1985) "Repeated Principal-Agent Games with Discounting," *Econometrica*, 53: 1173-1198 (September 1985).
- Radner, Roy. See Marschak & Radner (1972).
- Raff, Daniel & Lawrence Summers (1987) "Did Henry Ford Pay Efficiency Wages?" *Journal of Labor Economics*, 5: 57-86 (October 1987).

- Raiffa, Howard (1992) "Game Theory at the University of Michigan, 1948-52," pp. 165-76 of *Toward a History of Game Theory*, edited by E. Roy Weintraub, Durham: Duke University Press, ISBN: 0822312530.
- Raiffa, Howard. See Luce & Raiffa (1957).
- Ramsyer, J. Mark. See Wiley, Rasmusen & Ramseyer (1990).
- Rapoport, Anatol (1960) *Fights, Games and Debates*, Ann Arbor: University of Michigan Press (1960).
- Rapoport, Anatol (1970) *N-Person Game Theory: Concepts and Applications*, Ann Arbor: University of Michigan Press, ASIN: 0472050176.
- Rapoport, Anatol & Albert Chammah (1965) *Prisoner's Dilemma: A Study in Conflict and Cooperation*, Ann Arbor: University of Michigan Press, ISBN: 0472061658.
- Rapoport, Anatol, Melvin Guyer & David Gordon (1976) *The 2x2 Game*, Ann Arbor: University of Michigan Press, ASIN: 0472087428.
- Rappoport, Peter. See Cooter & Rappoport (1984).
- Rasmusen, Eric (1987) "Moral Hazard in Risk-Averse Teams," *Rand Journal of Economics*, 18: 428-435 (Fall 1987).
- Rasmusen, Eric (1988a) "Entry for Buyout," *Journal of Industrial Economics*, 36: 281-300 (March 1988).
- Rasmusen, Eric (1988b) "Stock Banks and Mutual Banks," *Journal of Law and Economics*, 31: 395-422 (October 1988).
- Rasmusen, Eric (1989a) *Games and Information*, Oxford: Basil Blackwell, 1989. Japanese translation by Moriki Hosoe, Shozo Murata, and Yoshinobu Arisada, Kyushu University Press, vol. I (1990), vol. 2 (1991). Italian translation by Alberto Bernardo, Milan: Ulrico Hoepli Editore (1993). Spanish translation (*Juegos e Informacion*) by Roberto Mazzoni, Mexico City: Fondo de Cultura Economica (1997). French translation forthcoming, Editions de Boeck & Larcier. Chinese translation forthcoming, Sanlian Press. xxx add Taiwan.
- Rasmusen, Eric (1989b) "A Simple Model of Product Quality with Elastic Demand," *Economics Letters*, 29: 281-283 (1989).
- Rasmusen, Eric (1992a) "Folk Theorems for the Observable Implications of Repeated Games," *Theory and Decision*, 32: 147-164 (March 1992).
- Rasmusen, Eric (1992b) "Managerial Conservatism and Rational Information Acquisition," *Journal of Economics and Management Strategy*, 1: 175-202 (Spring 1992).
- Rasmusen, Eric (1992c) "An Income-Satiation Model of Efficiency Wages," *Economic Inquiry*, 30: 467-478 (July 1992).

- Rasmusen, Eric (1997a) "Signal Jamming and Limit Pricing: A Unified Approach," in *Public Policy and Economic Analysis*, Moriki Hosoe and Eric Rasmusen, editors, Fukuoka, Japan: Kyushu University Press (1997).
- Rasmusen, Eric (1998) "Nuisance Suits," in *The New Palgrave Dictionary of Economics and the Law*, Peter Newman, editor, London: Macmillan Press, ISBN: 033367667X.
- Rasmusen, Eric, editor (2001) *Readings in Games and Information*, Oxford: Blackwell Publishers, ISBN: 0631215565.
- Rasmusen, Eric (2000b) "Writing, Speaking, and Listening," in Rasmusen (2001).
- Rasmusen, Eric & Emmanuel Petrakis (1992) "Defining the Mean-Preserving Spread: 3-pt versus 4-pt," *Decision Making Under Risk and Uncertainty: New Models and Empirical Findings*, edited by John Geweke. Amsterdam: Kluwer, ISBN: 0792319044.
- Rasmusen, Eric & Todd Zenger (1990) "Diseconomies of Scale in Employment Contracts," *Journal of Law, Economics and Organization*, 6: 65-92 (June 1990).
- Rasmusen Eric. See D. Hirshleifer & Rasmusen (1989), J. Hirshleifer & Rasmusen (1992), and Wiley, Rasmusen & Ramseyer (1990).
- Raviv, Arthur. See M. Harris & Raviv (1992, 1995).
- Ray, D. See Mookherjee & Ray (1992).
- Reeve, Hudson. See Dugatkin & Reeve (1998).
- Reinganum, Jennifer (1985) "Innovation and Industry Evolution," *Quarterly Journal of Economics*, 100: 81-99 (February 1985).
- Reinganum, Jennifer (1988) "Plea Bargaining and Prosecutorial Discretion," *The American Economic Review*, 78: 713-728 (September 1988).
- Reinganum, Jennifer (1989) "The Timing of Innovation: Research, Development and Diffusion," in Schmalensee & Willig (1989).
- Reinganum, Jennifer & Nancy Stokey (1985) "Oligopoly Extraction of a Common Property Natural Resource: the Importance of the Period of Commitment in Dynamic Games," *International Economic Review*, 26: 161-174 (February 1985).
- Reiss, Peter. See Bresnahan & Reiss (1991).
- Reynolds, Robert. See Salant et al. (1983).
- Richerson, Peter. See Boyd & Richerson (1985).
- Riker, William (1986) *The Art of Political Manipulation*, New Haven: Yale University Press, ISBN: 0300035926.
- Riley, John (1980) "Strong Evolutionary Equilibrium and the War of Attrition," *Journal of Theoretical Biology*, 82: 383-400 (February 1980).

- Riley, John. See J. Hirshleifer & Riley (1979,1992) Maskin & Riley (1984,1985), and Nalebuff & Riley (1985).
- Riordan, Michael. See Kihlstrom & Riordan (1984).
- Roberts, John. See Kreps et al. (1982) and Milgrom & Roberts (1982a, 1982b, 1986, 1990, 1992).
- Roberts, John & Hugo Sonnenschein (1976) "On the Existence of Cournot Equilibrium without Concave Profit Functions," *Journal of Economic Theory*, 13: 112-117 (August 1976).
- Robinson, Marc (1985) "Collusion and the Choice of Auction," *Rand Journal of Economics*, 16: 1-5 (Spring 1985).
- Rogerson, William (1982) "The Social Costs of Monopoly and Regulation: a Game-Theoretic Analysis," *Bell Journal of Economics*, 13: 391-401 (Autumn 1982).
- Romp, Graham (1997) *Game Theory: Introduction and Applications*, Oxford: Oxford University Press, ISBN: 0198775024.
- Rosen, Sherwin (1986) "Prizes and Incentives in Elimination Tournaments," *American Economic Review*, 76: 701-715 (September 1986).
- Rosen, Sherwin. See Lazear & Rosen (1981).
- Rosenberg, David & Steven Shavell (1985) "A Model in Which Suits are Brought for their Nuisance Value," *International Review of Law and Economics*, 5: 3-13 (June 1985).
- Ross, Steven (1977) "The Determination of Financial Structure: The Incentive-Signalling Approach," *Bell Journal of Economics*, 8: 23-40 (Spring 1977).
- Roth, Alvin (1984) "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92: 991-1016 (December 1984).
- Roth, Alvin, ed. (1985) *Game Theoretic Models of Bargaining*, Cambridge: Cambridge University Press, ISBN: 0521267579.
- Rothkopf, Michael (1980) "TREES: A Decision-Maker's Lament," *Operations Research*, 28: 3 (January/February 1980). Reprinted in Rasmusen (2001).
- Rothschild, Michael. See Diamond & Rothschild (1978).
- Rothschild, Michael & Joseph Stiglitz (1970) "Increasing Risk I," *Journal of Economic Theory*: 225-243. Reprinted in Diamond & Rothschild (1978).
- Rothschild, Michael & Joseph Stiglitz (1976) "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90: 629-649 (November 1976).

- Rubin, Paul (1978) "The Theory of the Firm and the Structure of the Franchise Contract," *Journal of Law and Economics*, 21: 223-233 (April 1978).
- Rubinfeld, Daniel. See Cooter & Rubinfeld (1989).
- Rubinstein, Ariel (1979) "An Optimal Conviction Policy for Offenses that May Have Been Committed by Accident," pp. 406-13 of *Applied Game Theory*, edited by Steven Brams, A. Schotter, Gerhard Schrodiauer, Physica-Verlag, ISBN: 3790802085. Reprinted in Rasmusen (2001).
- Rubinstein, Ariel (1982) "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50: 97-109 (January 1982). Reprinted in Rasmusen (2001).
- Rubinstein, Ariel (1985a) "A Bargaining Model with Incomplete Information about Time Preferences," *Econometrica*, 53: 1151-1172 (September 1985).
- Rubinstein, Ariel (1985b) "Choice of Conjectures in a Bargaining Game with Incomplete Information," in Alvin Roth, ed. (1985) *Game Theoretic Models of Bargaining*, Cambridge: Cambridge University Press, ISBN: 0521267579.
- Rubinstein, Ariel, ed. (1990) *Game Theory in Economics*, Brookfield, VT: Edward Elgar Publishing Company, ISBN: 1852781696.
- Rubinstein, Ariel. See Binmore et al. (1986), Osborne & Rubinstein (1994).
- Rudin, Walter (1964) *Principles of Mathematical Analysis*, New York: McGraw-Hill (1964).
- Saft, Lester. See Klein & Saft (1985).
- Salanie, Bernard (1997) *The Economics of Contracts: A Primer*, Cambridge: MIT Press, ISBN: 0262193868.
- Salant, Stephen, Sheldon Switzer, & Robert Reynolds (1983) "Losses from Horizontal Merger: The Effects of an Exogenous Change in Industry Structure on Cournot-Nash Equilibrium," *Quarterly Journal of Economics*, 98: 185-199 (May 1983).
- Salant, Stephen. See Gaudet & Salant (1991).
- Saloner, Garth. See Farrell & Saloner (1985).
- Salop, Steven & Joseph Stiglitz (1977) "Bargains and Ripoffs; A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 44: 493-510 (October 1977).
- Samet, Dov. See Kalai, Samet & Stanford (1988).
- Samuelson, Paul (1958) "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 66: 467-482 (December 1958).
- Samuelson, William (1984) "Bargaining under Asymmetric Information," *Econometrica*, 52: 995-1005 (July 1984).

- Samuelson, William. See Chatterjee & Samuelson (1983).
- Satterthwaite, Mark. See Myerson & Satterthwaite (1983).
- Savage, Leonard (1954) *The Foundations of Statistics*, New York: Wiley (1954).
- Scarf, Herbert. See Debreu & Scarf (1963).
- Scharfstein, David. See Nalebuff & Scharfstein (1987).
- Scheffman, David. See Holt & Scheffman (1987).
- Scheinkman, Jose. See Kreps & Scheinkman (1983).
- Schelling, Thomas (1960) *The Strategy of Conflict*, Cambridge: Harvard University Press.
- Schelling, Thomas (1978) *Micromotives and Macrobehavior*, New York: W. W. Norton, ISBN: 0393090094.
- Scherer, Frederick (1980) *Industrial Market Structure and Economic Performance*, 2nd edn. Chicago: Rand McNally (1st edition, 19xxx).
- Schmalensee, Richard (1982) "Product Differentiation Advantages of Pioneering Brands," *American Economic Review*, 72: 349-365 (June 1982).
- Schmalensee, Richard & Robert Willig, eds. (1989) *The Handbook of Industrial Organization*, New York: North-Holland, ISBN: 0444704353.
- Schmittberger, Rold. See Guth et al. (1982).
- Schumpeter, Joseph (1911/1934) *Theory of Economic Development*, translated from the German 3rd edn. by Redvers Opie, Cambridge: Harvard University Press (1st edn., 1911).
- Schwartz, Nancy. See Kamien & Schwartz (1981, 1991).
- Schwarze, Bernd. See Guth et al. (1982).
- Seligman, Daniel (1992) *A Question of Intelligence: The IQ Debate in America*, New York: Carol Publishing, ISBN: 1559721316.
- Selten, Reinhard (1965) "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die gesamte Staatswissenschaft*, 121: 301-24, 667-689 (October 1965).
- Selten, Reinhard (1975) "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4: 25-55 (1975).
- Selten, Reinhard (1978) "The Chain-Store Paradox," *Theory and Decision*, 9: 127-159 (April 1978).
- Selten, Reinhard. See Harsanyi & Selten (1988).

- Shaked, Avner (1982) "Existence and Computation of Mixed Strategy Nash Equilibrium for 3-Firms Location Problem," *Journal of Industrial Economics*, 31: 93-6 (September/December 1982). Reprinted in Rasmusen (2001).
- Shaked, Avner & John Sutton (1983) "Natural Oligopolies," *Econometrica*, 51: 1469-1483 (September 1983).
- Shaked, Avner & John Sutton (1984) "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model," *Econometrica*, 52: 1351-1364 (November 1984).
- Shanley, Mark. See Besanko, Dranove & Shanley (1996).
- Shapiro, Carl (1982) "Consumer Information, Product Quality and Seller Reputation," *Bell Journal of Economics*, 13: 20-35 (Spring 1982).
- Shapiro, Carl (1983) "Premiums for High Quality Products as Returns to Reputation," *Quarterly Journal of Economics*, 98: 659-679 (November 1983).
- Shapiro, Carl (1989) "The Theory of Business Strategy," *Rand Journal of Economics*, 20: 125-137 (Spring 1989).
- Shapiro, Carl. See Farrell & Shapiro (1988) and Katz & Shapiro (1985).
- Shapiro, Carl & Joseph Stiglitz (1984) "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review*, 74: 433-444 (June 1984).
- Shapley, Lloyd (1953a) "Open Questions," p. 15 of *Report of an Informal Conference on the Theory of n-Person Games*, Princeton Mathematics mimeo (1953).
- Shapley, Lloyd (1953b) "A Value for n-Person Games," pp. 307-317 of Kuhn & Tucker (1953).
- Shavell, Steven (1979) "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10: 55-73 (Spring 1979).
- Shavell, Steven. See Rosenberg & Shavell (1985), Kaplow & Shavell (1996).
- Shell, Karl. See Cass & Shell (1983).
- Shleifer, Andrei & Robert Vishny (1986) "Greenmail, White Knights, and Shareholders' Interest," *Rand Journal of Economics*, 17: 293-309 (Autumn 1986).
- Shubik, Martin (1954) "Does the Fittest Necessarily Survive?" pp. 43-6 of *Readings in Game Theory and Political Behavior*, edited by Martin Shubik, Garden City, New York: Doubleday (1954). Also published in Rasmusen (2001).
- Shubik, Martin (1971) "The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation," *Journal of Conflict Resolution*, 15: 109-11 (March 1971). Reprinted in Rasmusen (2001).
- Shubik, Martin (1982) *Game Theory in the Social Sciences: Concepts and Solutions*, Cambridge: MIT Press, ISBN: 0262191954.

Shubik, Martin (1992) "Game Theory at Princeton, 1949-1955: A Personal Reminiscence," pp. 151-64 of *Toward a History of Game Theory*, edited by E.Roy Weintraub, Durham: Duke University Press, ISBN: 0822312530.

Shubik, Martin. See Levitan & Shubik (1972).

Shy, Oz (1996) *Industrial Organization, Theory and Applications*, Cambridge: MIT Press, ISBN: 0262691795.

Simon, Leo (1987) "Games with Discontinuous Payoffs," *Review of Economic Studies*, 54: 569-598 (October 1987).

Skeath, Susan. See Dixit & Skeath (1999).

Slade, Margaret (1987) "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion," *Journal of Industrial Economics*, 35: 499-516 (June 1987).

Slatkin, Montgomery (1980) "Altruism in Theory," review of Scott Boorman & Paul Levitt, *The Genetics of Altruism*. *Science*, 210: 633-647 (November 1980).

Slovic, Paul. See Kahneman, Slovic & Tversky (1982).

Smith, Abbie. See Antle & Smith (1986).

Smith, Adam (1776) *An Inquiry into The Nature and Causes of the Wealth of Nations*, Chicago: University of Chicago Press 1977), ISBN: 0226763749.

Smith, J. See DeBrock & Smith (1983).

Sobel, Joel. See Border & Sobel (1987) and Crawford & Sobel (1982).

Sobel, Joel & Ichiro Takahashi (1983) "A Multi-Stage Model of Bargaining," *Review of Economic Studies*, 50: 411-426 (July 1983).

Sonnenschein, Hugo (1983) "Economics of Incentives: An Introductory Account," in *Technology, Organization, and Economic Structure: Essays in Honor of Prof. Isamu Yamada*, edited by Ryuzo Sato & Martin Beckmann, Berlin: Springer-Verlag, ASIN: 0387119981.

Sonnenschein, Hugo. See Roberts & Sonnenschein (1976).

Sorenson, Philip. See Mead et al. (1984).

Sowden, Lanning. See Campbell & Sowden (1985).

Spector, Yishay. See Leshno, Levy & Spector (1997).

Spence, A. Michael (1973) "Job Market Signalling," *Quarterly Journal of Economics*, 87: 355-374 (August 1973).

Spence, A. Michael (1977) "Entry, Capacity, Investment, and Oligopolistic Pricing," *Bell Journal of Economics*, 8: 534- 544 (Autumn 1977).

- Spence, A. Michael. See Kreps & Spence (1984).
- Spitzer, Matthew. See Hoffman & Spitzer (1985).
- Spulber, Daniel (1989) *Regulation and Markets*, Cambridge: MIT Press, ISBN: 0262192756.
- Srivastava, Sanjay. See Palfrey & Srivastava (1993).
- Stacchetti, Ennio. See Abreu et al. (1986, 1990)
- Stackelberg, Heinrich von (1934) *Marktform und Gleichgewicht*, Berlin: J. Springer. Translated by Alan Peacock as *The Theory of the Market Economy*, London: William Hodge (1952).
- Stahl, Saul (1998) *A Gentle Introduction to Game Theory*, Providence, RI: American Mathematical Society. ISBN: xxx.
- Stanford, William. See Kalai, Samet & Stanford (1988).
- Starmer, Chris (2000) "Developments in Non-Expected Utility Theory, " *Journal of Economic Literature*, 38:332-382 (June 2000).
- Staten, Michael & John Umbeck (1982) "Information Costs and Incentives to Shirk: Disability Compensation of Air Traffic Controllers," *American Economic Review*, 72: 1023-1037 (December 1982).
- Staten, Michael & John Umbeck (1986) "A Study of Signaling Behavior in Occupational Disease Claims," *Journal of Law and Economics*, 29: 263-286 (October 1986).
- Stevenson, Robert (1987) *Island Nights' Entertainments*, London: Hogarth. ISBN: xxx.
- Stigler, George (1964) "A Theory of Oligopoly," *Journal of Political Economy*, 72: 44-61 (February 1964).
- Stigler, George. See Becker & Stigler (1974).
- Stiglitz, Joseph (1982a) "Self-Selection and Pareto-Efficient Taxation," *Journal of Public Economics*, 17: 213-240 (March 1982).
- Stiglitz, Joseph (1982b) "Utilitarianism and Horizontal Equity: The Case of Random Taxation," *Journal of Public Economics*, 18: 1-33 (June 1982).
- Stiglitz, Joseph (1987) "The Causes and Consequences of the Dependence of Quality on Price," *Journal of Economic Literature*, 25: 1-48 (March 1987).
- Stiglitz, Joseph & Andrew Weiss (1981) "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71: 393-410 (June 1981).
- Stiglitz, Joseph & Andrew Weiss (1989) "Sorting out the Differences Between Screening and Signalling Models," in *Papers in Commemoration of the Economic Theory Seminar at Oxford University*, edited by Michael Dempster, Oxford: Oxford University Press. ISBN: xxx.

- Stiglitz, Joseph. See Dasgupta & Stiglitz (1980), Dixit & Stiglitz (1977), Nalebuff & Stiglitz (1983), Rothschild & Stiglitz (1970, 1976), Salop & Stiglitz (1977), and Shapiro & Stiglitz (1984).
- Stokey, Nancy & Robert Lucas (1989) *Recursive Methods in Economic Dynamics*, Cambridge: Harvard University Press, ISBN: 0674750969.
- Stokey, Nancy. See Reinganum & Stokey (1985).
- Straffin, Philip (1980) "The Prisoner's Dilemma," *UMAP Journal*, 1: 101-103 (1980). Reprinted in Rasmusen (2001).
- Strunk, William & E. B. White (1959) *The Elements of Style*, New York: Macmillan.
- Sultan, Ralph (1974) *Pricing in the Electrical Oligopoly, Vol I: Competition or Collusion*, Cambridge: Harvard University Press, ASIN: 0875841104.
- Summers, Larry. See Raff & Summers (1987).
- Sutton, John (1986) "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53: 709-24 (October 1986).
- Sutton, John (1991) *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*, Cambridge: MIT Press, ISBN: 0262193051.
- Sutton, John. See Shaked & Sutton (1983, 1984).
- Switzer, Sheldon. See Salant et al. (1983).
- Szenberg, Michael, editor (1992) *Eminent Economists: Their Life Philosophies*, Cambridge: Cambridge University Press, ISBN: 0521382122.
- Szenberg, Michael, editor (1998) *Passion and Craft: Economists at Work*, Ann Arbor: University of Michigan Press, ASIN: 0472096850.
- Szep, J. & F. Forgo (1985) *Introduction to the Theory of Games*, Dordrecht: D. Reidel (1985).
- Takahashi, Ichiro. See Sobel & Takahashi (1983).
- Takayama, Akira (1985) *Mathematical Economics*, 2nd edn. Cambridge: Cambridge University Press, ISBN: 0521314984.
- Taylor, Alan (1995) *Mathematics and Politics: Strategy, Voting, Power and Proof*, Berlin: Springer Verlag, ISBN: 0387943919.
- Teece, David. See Monteverde & Teece (1982).
- Telser, Lester (1966) "Cutthroat Competition and the Long Purse," *Journal of Law and Economics*, 9: 259-77 (October 1966) .
- Tenorio, Rafael (1993) "Revenue-Equivalence and Bidding Behavior in a Multi-Unit Auction Market: An Empirical Analysis," *Review of Economics and Statistics*, 75: 302-14 (May 1993).

- Thaler, Richard (1991) *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, New York: The Free Press, ASIN: 0029324653.
- Thisse, Jacques. See d'Aspremont et al. (1979).
- Thomas, L. (1984) *Games, Theory and Applications*, Chichester, England: Ellis Horwood (1984).
- Tirole, Jean (1986) "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations," *Journal of Law, Economics, and Organization*, 2: 181-214 (Fall 1986).
- Tirole, Jean (1988) *The Theory of Industrial Organization*, Cambridge: MIT Press, ISBN: 0262200716.
- Tirole, Jean. See Freixas et al. (1985), Fudenberg & Tirole (1983,1986a,1986b,1988, 1991a,1991b), Laffont & Tirole (1986,1993), and Maskin & Tirole (1987).
- Titman, Sheridan. See D. Hirshleifer & Titman (1990).
- Topkis, Donald (1998) *Supermodularity and Complementarity*, Princeton: Princeton University Press, ISBN: 0691032440.
- Tsebelis, George (1989) "The Abuse of Probability in Political Analysis: The Robinson Crusoe Fallacy," *American Political Science Review*, 83: 77-91 (March 1989).
- Tucker, Albert (unpublished) "A Two-Person Dilemma," Stanford University mimeo. May 1950. Reprinted in Straffin (1980). Also published in Rasmusen (2001).
- Tucker, Albert. See Drescher et al. (1957), Kuhn & Tucker (1950,1953), and Luce & Tucker (1959).
- Tukey, John (1949) "A Problem in Strategy," *Econometrica*, (supplement), 17: 73 (abstract) (July 1949).
- Tukey, John. See McDonald & Tukey (1949).
- Tullock, Gordon (1967) "The Welfare Costs of Tariffs, Monopolies, and Theft," *Western Economic Journal*, 5: 224-32 (June 1967).
- Tullock, Gordon (1983) *Economics of Income Redistribution*, Boston: Kluwer-Nijhoff (1983).
- Tullock, Gordon (1985) "Adam Smith and the Prisoners' Dilemma," *The Quarterly Journal of Economics*, 100: 1073-81 (September 1985).
- Tversky, Amon. See Kahneman, Slovic & Tversky (1982).
- Umbeck, John. See Staten & Umbeck (1982, 1986).
- Van Damme, Eric (1983) *Refinements of the Nash Equilibrium Concept*, Berlin: Springer-Verlag, ASIN: 0387126902.

- Van Damme, Eric (1987) *Stability and Perfection of Nash Equilibrium*, Berlin: Springer-Verlag.
- Van Damme, Eric (1989) "Stable Equilibria and Forward Induction," *Journal of Economic Theory*, 48: 476-496 (August 1989).
- Varet, L. Gerard. See D'Aspremont & Varet (1979).
- Varian, Hal (1992) *Microeconomic Analysis*, 3rd edn. New York: W. W. Norton, 1992 (2nd edn. 1984), ISBN: 0393957357.
- Vickrey, William (1961) "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16: 8-37 (March 1961).
- Vickrey, William (1964) *Microstatics*, New York: Harcourt, Brace and World (1964).
- Vishny, Robert. See Shleifer & Vishny (1986).
- Vives, Xavier (1990), "Nash Equilibrium with Strategic Complementarities," *Journal of Mathematical Economics*, 19: 305-321.
- Vives, Xavier (2000) *Oligopoly Pricing*, Cambridge: MIT Press, ISBN: 0262220601.
- von Neumann, John (1928) "Zur Theorie der Gesellschaftspiele," *Mathematische Annalen*, 100: 295-320 (1928). Translated by Sonya Bargmann as "On the Theory of Games of Strategy," pp. 13-42 of Luce & Tucker (1959).
- von Neumann, John & Oskar Morgenstern (1944) *The Theory of Games in Economic Behavior*, New York: Wiley (1944).
- Waldegrave, James (1713) "Excerpt from a Letter," (with a preface by Harold Kuhn), in Baumol & Goldfeld (1968).
- Waldman, Michael (1987) "Noncooperative Entry Deterrence, Uncertainty, and the Free Rider Problem," *Review of Economic Studies*, 54: 301-10 (April 1987).
- Waldman, Michael. See Haltiwanger & Waldman (1991).
- Weber, Robert. See Milgrom & Weber (1982).
- Weibull, Jorgen (1995) *Evolutionary Game Theory*, Cambridge: MIT Press, ISBN: 0262231816.
- Weiner, E. (1984) *The Oxford Guide to the English Language*, Oxford: Oxford University Press, ISBN: 0198691319.
- Weintraub, E. Roy, ed. (1992) *Toward a History of Game Theory*, Durham: Duke University Press, ISBN: 0822312530.
- Weiss, Andrew (1990) *Efficiency Wages*, Princeton: Princeton University Press, ISBN: 0691003882.
- Weiss, Andrew. See Guasch & Weiss (1980) and Stiglitz & Weiss (1981, 1989).

- Weitzman, Martin (1974) "Prices vs. Quantities," *Review of Economic Studies*, 41: 477-91 (October 1974).
- Welch, Ivo. See Bikhchandani, D. Hirshleifer & Welch (1992).
- Weston, J. Fred. See Copeland & Weston (1988).
- Whinston, Michael. See Bernheim et al. (1987), Bernheim & Whinston (1987), and Mas-Colell, Whinston & Green (1994).
- White, E.B. See Strunk & White (1959).
- Wicksteed, Philip (1885) *The Common Sense of Political Economy*, New York: Kelley (1950).
- Wiley, John, Eric Rasmusen & Mark Ramseyer (1990) "The Leasing Monopolist," *UCLA Law Review*, 37: 693-732 (April 1990).
- Williams, J. (1966) *The Compleat Strategyst: Being a Primer on the Theory of Games of Strategy*, New York: McGraw-Hill (1966).
- Williamson, Oliver (1975) *Markets and Hierarchies: Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*, New York: Free Press, ISBN: 0029353602
- Willig, Robert. See Schmalensee & Willig (1989).
- Wilson, Charles (1980) "The Nature of Equilibrium in Markets with Adverse Selection," *Bell Journal of Economics*, 11: 108-30 (Spring 1980).
- Wilson, Robert (1979) "Auctions of Shares," *Quarterly Journal of Economics*, 93: 675-89 (November 1979).
- Wilson, Robert (unpublished) Stanford University 311b course notes.
- Wilson, Robert. See Kennan & Wilson (1993), Kreps & Wilson (1982a, 1982b) and Kreps et al. (1982).
- Winter, Ralph. See Mathewson & Winter (1985).
- Wolfe, Philip. See Drescher et al. (1957).
- Wolfson, M. (1985) "Empirical Evidence of Incentive Problems and their Mitigation in Oil and Tax Shelter Programs," pp. 101-25 of *Principals and Agents: The Structure of Business*, edited by John Pratt & Richard Zeckhauser. Boston: Harvard Business School Press, ISBN: 0875841643
- Wolinsky, Asher. See Binmore et al. (1986).
- Wydick, Richard (1978) "Plain English for Lawyers," *California Law Review*, 66: 727-64 (1978).
- Yellen Janet. See Akerlof & Yellen (1986).

Zahavi, Amotz (1975) "Mate Selection: A Selection for a Handicap," *Journal of Theoretical Biology*, 53: 205-14 (September 1975).

Zamir, S. See Mertens & Zamir (1985).

Zapechelnyuk, Andriy. See Dubey, Haimanko & Zapechelnyuk (2002).

Zenger, Todd. See Rasmusen & Zenger (1990).

Von Zermelo, E. (1913) "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," *Proceedings, Fifth International Congress of Mathematicians*, 2: 501-4 (1913). Also published in Rasmusen (2001).

Zimmerman, Jerold. See Gaver & Zimmerman (1977).

Mathematical Appendix

This appendix has three purposes: to remind some readers of the definitions of terms they have seen before, to give other readers an idea of what the terms mean, and to list a few theorems for reference. In accordance with these limited purposes, some terms such as “boundary point” are left undefined. For fuller exposition, see Rudin (1964) on real analysis, Debreu’s *Theory of Value* (1959), and Chiang (1984) and Takayama (1985) on mathematics for economists. Intriligator (1971) and Varian (1992) both have good mathematical appendices and are strong in discussing optimization, and Kamien & Schwartz (1991) covers maximizing by choice of functions. Border’s 1985 book is entirely about fixed point theorems. Stokey & Lucas (1989) is about dynamic programming. Fudenberg & Tirole (1991a) is the best source of mathematical theorems for use in game theory.

*A.1 Notation

Σ Summation. $\sum_{i=1}^3 x_i = x_1 + x_2 + x_3$.

Π Product. $\prod_{i=1}^3 x_i = x_1 x_2 x_3$.

$|x|$ Absolute value of x . If $x \geq 0$ then $|x| = x$ and if $x < 0$ then $|x| = -x$.

$|$ “Such that,” “given that,” or “conditional upon.” $\{x|x < 3\}$ denotes the set of real numbers less than three. $Prob(x|y < 5)$ denotes the probability of x given that y is less than 5.

$:$ “Such that.” $\{x : x < 3\}$ denotes the set of real numbers less than three. The colon is a synonym for $|$.

$\mathbf{R}^n$ The set of n -dimensional vectors of real numbers (integers, fractions, and the least upper bounds of any subsets thereof).

$\{x, y, z\}$ A set of elements x, y , and z . The set $\{3, 5\}$ consists of two elements, 3 and 5.

$\in$ “Is an element of.” $a \in \{2, 5\}$ means that a takes either the value 2 or 5.

$\subset$ Set inclusion. If $X = \{2, 3, 4\}$ and $Y = \{2, 4\}$, then $Y \subset X$ because Y is a subset of X .

$[x, y]$ The closed interval with endpoints x and y . The interval $[0, 1000]$ is the set $\{x|0 \leq x \leq 1000\}$. Square brackets are also used as delimiters.

(x, y) The open interval with endpoints x and y . The interval $(0, 1000)$ is the set $\{x|0 < x < 1000\}$. $(0, 1000]$ would be a half-open interval, the set $\{x|0 < x \leq 1000\}$. Parentheses are also used as delimiters.

¹xxx Footnotes starting with xxx are the author’s notes to himself. Comments welcomed.

$x!$ x -factorial. $x! = x(x-1)(x-2)\dots(2)(1)$. $4! = 4(3)(2)(1) = 24$.

$\binom{a}{b}$ The number of unordered combinations of b elements from a set with a elements.
 $\binom{a}{b} = \frac{a!}{b!(a-b)!}$, so $\binom{4}{3} = \frac{4!}{3!(4-3)!} = 24/6 = 4$. (See *combination* and *permutation* below.)

$\times$ The Cartesian product. $X \times Y$ is the set of points $\{x, y\}$, where $x \in X$ and $y \in Y$.

ϵ An arbitrarily small positive number. If my payoff from both *Left* and *Right* equals 10, I am indifferent between them; if my payoff from *Left* is changed to $10 + \epsilon$, I prefer *Left*.

$\sim$ We say that $X \sim F$ if the random variable X is distributed according to distribution F .

$\exists$ "There exists..."

$\forall$ "For all..."

$\equiv$ "Equals by definition."

$\rightarrow$ If f maps space X into space Y then $f: X \rightarrow Y$.

$\frac{df}{dx}$, $\frac{d^2f}{dx^2}$ The first and second derivatives of a function. If $f(x) = x^2$ then $\frac{df}{dx} = 2x$ and $\frac{d^2f}{dx^2} = 2$.

f', f'' The first and second derivatives of a function. If $f(x) = x^2$ then $f' = 2x$ and $f'' = 2$.
 Primes are also used on variables (not functions) for other purposes: x' and x'' might denote two particular values of x .

$\frac{\partial f}{\partial x}$, $\frac{\partial^2 f}{\partial x \partial y}$ Partial derivatives of a function. If $f(x, y) = x^2y$ then $\frac{\partial f}{\partial x} = 2xy$ and $\frac{\partial^2 f}{\partial x \partial y} = 2x$.

y_{-i} The set y minus element i . If $y = \{y_1, y_2, y_3\}$, then $y_{-2} = \{y_1, y_3\}$.

$Max(x, y)$ The maximum of two numbers x and y . $Max(8, 24) = 24$.

$Min(x, y)$ The minimum of two numbers x and y . $Min(5, 3) = 3$.

$\lceil x \rceil$ Ceiling (x). A number rounded up to the nearest integer. $\lceil 4.2 \rceil = 5$. This notation is not well known in economics.

$\lfloor x \rfloor$ Floor (x). A number rounded down to the nearest integer. $\lfloor 6.9 \rfloor = 6$. This notation is not well known in economics.

$Sup X$ The supremum (least upper bound) of set X . If $X = \{x | 0 \leq x < 1000\}$, then $sup X = 1000$. The supremum is useful because sometimes, as here, no maximum exists.

$Inf X$ The infimum (greatest lower bound) of set X . If $X = \{x | 0 \leq x < 1000\}$, then $inf X = 0$.

Argmax The argument that maximizes a function. If $e^* = \operatorname{argmax} EU(e)$, then e^* is the value of e that maximizes the function $EU(e)$. The argmax of $f(x) = x - x^2$ is $1/2$.

Maximum The greatest value that a function can take. $Maximum(x - x^2) = 1/4$.

Minimum The lowest value that a function can take. $Minimum(-5 + x^2) = -5$.

*A.2 The Greek Alphabet

A	α	alpha
B	β	beta
Γ	γ	gamma
Δ	δ	delta
E	ϵ or ε	epsilon
Z	ζ	zeta
H	η	eta
Θ	θ	theta
I	ι	iota
K	κ	kappa
Λ	λ	lambda
M	μ	mu
N	ν	nu
Ξ	ξ	xi
O	o	omicron
Π	π	pi
P	ρ	rho
Σ	σ	sigma
T	τ	tau
Υ	υ	upsilon
Φ	ϕ	phi
X	χ	chi
Ψ	ψ	psi
Ω	ω	omega

*A.3 Glossary

almost always See “generically.”

annuity A riskless security paying a constant amount each year for a given period of years, with the amount conventionally paid at the end of each year.

closed A closed set in $\mathbf{R}^n$ includes its boundary points. The set $\{x : 0 \leq x \leq 1000\}$ is closed.

combination The number of unordered sets of b elements from a set with a elements, denoted $\binom{a}{b} = \frac{a!}{b!(a-b)!}$. If we form sets of 2 element from the set $A = \{w, x, y, z\}$, the possibilities are $\{w, x\}, \{w, y\}, \{w, z\}, \{x, y\}, \{x, z\}, \{y, z\}$. Thus, $\binom{4}{2} = \frac{4!}{2!(4-2)!} = 24/6 = 6$. (See *permutation* for the ordered version.)

compact If set X in $\mathbf{R}^n$ is closed and bounded, then X is compact. Outside of Euclidean space, however, a set being closed and bounded does not guarantee compactness.

complete metric space All compact metric spaces and all Euclidean spaces are complete.

concave function The continuous function $f(x)$ defined on interval X is concave if for all elements w and z of X , $f(0.5w + 0.5z) \geq 0.5f(w) + 0.5f(z)$. If f maps $\mathbf{R}$ into $\mathbf{R}$ and f is concave, then $f'' \leq 0$. See figure A.1.




Figure 1 Concavity and Convexity

continuous function Let $d(x, y)$ represent the distance between points x and y . The function f is continuous if for every $\epsilon > 0$ there exists a $\delta(\epsilon) > 0$ such that $d(x, y) < \delta(\epsilon)$ implies $d(f(x), f(y)) < \epsilon$.

continuum A continuum is a closed interval of the real line, or a set that can be mapped one-to-one onto such an interval.

contraction The mapping $f(x)$ is said to be a contraction if there exists a number $c < 1$ such that for the metric d of the space X ,

$$d(f(x), f(y)) \leq cd(x, y), \text{ for all } x, y \in X. \quad (1)$$

convex function The continuous function $f(x)$ is convex if for all elements w and z of X , $f(0.5w + 0.5z) \leq 0.5f(w) + 0.5f(z)$. See Figure 1. Convex functions are only loosely related to convex sets.

convex set If set X is convex, then if you take any two of its elements w and z and a real number $t : 0 \leq t \leq 1$, then $tw + (1 - t)z$ is also in X .

correspondence A correspondence is a mapping that maps each point to one or more other points, as opposed to a function, which only maps to one.

domain The domain of a mapping is the set of elements it maps from—the property over which it reigns and can change as it pleases. (The mapping maps from the domain onto the range.)

function If f maps each point in X to exactly one point in Y , f is called a function. The two mappings in figure A.1 are functions, but the mapping in Figure 2 is not.

generically If a fact is true on set X generically, “except on a set of measure zero,” or “almost always,” then it is false only on a subset of points Z that have the property that if a point is randomly chosen using a density function with support X , a point in Z is chosen with probability zero. This implies that if the fact is false on $z \in \mathbf{R}^n$ and z is perturbed by adding a random amount ϵ , the fact is true on $z + \epsilon$ with probability one. See pp. xxx.

integration by parts This is a technique to rearrange integrals so they can be solved more easily. It uses the formula

$$\int_{z=a}^b g(z)h'(z)dz = g(z)h(z) \Big|_{z=a}^b - \int_{z=a}^b h(z)g'(z)dz. \quad (2)$$

To derive this, differentiate $g(z)h(z)$ using the chain rule, integrate each side of the equation, and rearrange.

Lagrange multiplier The Lagrange multiplier λ is the marginal value of relaxing a constraint in an optimization problem. If the problem is

$$\left\{ \begin{array}{l} \text{Maximize} \\ x \end{array} \right. x^2 \text{ subject to } x \leq 5 \}, \text{ then } \lambda = 2x^* = 10.$$

lattice A lattice is a partially ordered set (the $\geq$ ordering is defined) where for any two elements a and b , the values $\inf(a, b)$ and $\sup(a, b)$ are also in the set.

A lattice is complete if the infimum and supremum of each of its subsets are in the lattice.

lower semicontinuous correspondence The correspondence ϕ is lower semicontinuous at the point x_0 if

$$x_n \rightarrow x_0, y_0 \in \phi(x_0), \text{ implies } \exists y_n \in \phi(x_n) \text{ such that } y_n \rightarrow y_0, \quad (3)$$

which means that associated with every x sequence leading to x_0 is a y sequence leading to its image. See Figure 2. This idea is not as important as upper semicontinuity.

maximand A maximand is what is being maximized. In the problem “Maximize $f(x, \theta)$ by choice of x ”, the maximand is f .

mean-preserving spread See the **Risk** section below.

measure zero See “generically.”

metric The function $d(w, z)$ defined over elements of set X is a metric if (1) $d(w, z) > 0$ if $w \neq z$ and $d(w, z) = 0$ if and only if $w = z$; (2) $d(w, z) = d(z, w)$; and (3) $d(w, z) \leq d(w, y) + d(y, z)$ for points $w, y, z \in X$.

metric space Set X is a metric space if it is associated with a metric that defines the distance between any two of its elements.

one-to-one The mapping $f : X \rightarrow Y$ is one-to-one if every point in set X maps to a different point in Y , so $x_1 \neq x_2$ implies $f(x_1) \neq f(x_2)$.

onto The mapping $f : X \rightarrow Y$ is onto Y if every point in Y is mapped onto by some point in X .

open In the space $\mathbf{R}^n$, an open set is one that does not include all its boundary points. The set $\{x : 0 \leq x < 1000\}$ is open. In more general spaces, an open set is a member of a topology.

permutation The number of ordered sets of b elements from a set with a elements, which equals $\frac{a!}{b!(a-b)!}$. If we form sets of 2 elements from the set $A = \{w, x, y, z\}$, the possibilities are $\{w, x\}, \{x, w\}, \{w, y\}, \{y, w\}, \{w, z\}, \{z, w\}, \{x, y\}, \{y, x\}, \{x, z\}, \{z, x\}, \{y, z\}, \{z, y\}$. The number of these is $\frac{4!}{4-2!} = 24/2 = 12$. (See *combination* for the unordered version.)

perpetuity A riskless security paying a constant amount each year in perpetuity, with the amount conventionally paid at the end of each year.

quasi-concave The continuous function f is quasi-concave if for $w \neq z$, $f(0.5w + 0.5z) > \min[f(w), f(z)]$, or, equivalently, if the set $\{x \in X | f(x) > b\}$ is convex for any number b . Every concave function is quasi-concave, but not every quasi-concave function is concave.

range The range of a mapping is the set of elements to which it maps— the property over which it can spew its output. (The mapping maps from the domain onto the range.)

risk See the **Risk** section below.

stochastic dominance See the **Risk** section below.

strict The word “strict” is used in a variety of contexts to mean that a relationship does not hold with equality or is not arbitrarily close to being violated. If function f is concave and $f' > 0$, then $f'' \leq 0$, but if f is strictly concave, then $f'' < 0$. The opposite of “strictly” is “weakly.” The word “strong” is sometimes used as a synonym for “strict.”

supermodular See the **Supermodularity** section below.

support The support of a probability distribution $F(x)$ is the closure of the set of values of x such that the density is positive. If each output between 0 and 20 has a positive probability density, and no other output does, then the support of the output distribution is $[0,20]$.

topology Besides denoting a field of mathematics, a topology is a collection of subsets of a space called “open sets” that includes (1) the entire space and the empty set, (2) the intersection of any finite number of open sets, and (3) the union of any number of open sets. In a metric space, the metric “induces” a topology by defining an open set. Imposing a topology on a space is something like defining which elements are close to each other, which is easy to do for $\mathbf{R}^n$ but not for every space (e.g., spaces consisting of functions or of game trees).

upper semicontinuous correspondence The correspondence $\phi : X \rightarrow Y$ is upper semicontinuous at point x_0 if

$$x_n \rightarrow x_0, y_n \in \phi(x_n), y_n \rightarrow y_0, \text{ implies } y_0 \in \phi(x_0), \quad (4)$$

which means that every sequence of points in $\phi(x)$ leads to a point also in $\phi(x)$. See Figure 2. An alternative definition, appropriate only if Y is compact, is that ϕ is upper semicontinuous if the set of points $\{x, \phi(x)\}$ is closed.




Figure 2 Upper Semicontinuity

vector A vector is an ordered set of real numbers, a point in $\mathbf{R}^n$. The point $(2.5, 3, -4)$ is a vector in $\mathbf{R}^3$.

weak The word “weak” is used in a variety of contexts to mean that a relationship might hold with equality or be on a borderline. If f is concave and $f' > 0$, then $f'' \leq 0$, but to say that f is weakly concave, while technically adding nothing to the meaning, emphasizes that $f'' = 0$ under some or all parameters. The opposite of “weak” is “strict” or “strong.”

*A.4 Formulas and Functions

$$\log(xy) = \log(x) + \log(y).$$

$$\log(x^2) = 2\log(x).$$

$$a^x = (e^{\log(a)})^x.$$

$$e^{rt} = (e^r)^t.$$

$$e^{a+b} = e^a e^b.$$

$$a > b \Rightarrow ka < kb, \text{ if } k < 0.$$

The Quadratic Formula: Let $ax^2 + bx + c = 0$. Then $x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$.

Derivatives

$f(x)$	$f'(x)$
x^a	ax^{a-1}
$1/x$	$-\frac{1}{x^2}$
$\frac{1}{x^2}$	$-\frac{2}{x^3}$
e^x	e^x
e^{rx}	re^{rx}
$\log(ax)$	$1/x$
$\log(x)$	$1/x$
a^x	$a^x \log(a)$
$f(g(x))$	$f'(g(x))g'(x)$

Determinants

$$\begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} = a_{11}a_{22} - a_{21}a_{12}.$$

$$\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = a_{11}a_{22}a_{33} - a_{11}a_{23}a_{32} + a_{12}a_{23}a_{31} - a_{12}a_{21}a_{33} + a_{13}a_{21}a_{32} - a_{13}a_{22}a_{31}.$$

Table 1 Some Useful Functional Forms

$f(x)$	$f'(x)$	$f''(x)$	<i>Slope</i>	<i>Curvature</i>
$\log(x)$	$\frac{1}{x}$	$-\frac{1}{x^2}$	increasing	concave
$\sqrt{x}$	$\frac{1}{2\sqrt{x}}$	$-\frac{1}{4x^{(3/2)}}$	increasing	concave
x^2	$2x$	2	increasing	convex
$\frac{1}{x}$	$-\frac{1}{x^2}$	$\frac{2}{x^3}$	decreasing	convex
$7 - x^2$	$-2x$	-2	decreasing	concave
$7x - x^2$	$7 - 2x$	-2	increasing/decreasing	concave

The signs of derivatives can be confusing. The function $f(x) = x^2$ is increasing at an *increasing* rate, but the function $f(x) = \frac{1}{x}$ is decreasing at a *decreasing* rate, even though $f'' > 0$ in each case.

***A.5 Probability Distributions** The definitive listing of probability distributions and their characteristics is the three volume series of Johnson & Kotz (1970). A few major distributions are listed here. A **probability distribution** is the same as a **cumulative density function** for a continuous distribution.

The Exponential Distribution

The exponential distribution, which has the set of nonnegative real numbers as its support, has the density function

$$f(x) = \frac{e^{-x/\lambda}}{\lambda}. \quad (5)$$

The cumulative density function is

$$F(x) = 1 - e^{-x/\lambda}. \quad (6)$$

The Uniform Distribution

A variable is uniformly distributed over support X if each point in X has equal probability. The density function for support $X = [\alpha, \beta]$ is

$$f(x) = \begin{cases} 0 & x < \alpha \\ \frac{1}{\beta - \alpha} & \alpha \leq x \leq \beta \\ 0 & x > \beta, \end{cases} \quad (7)$$

and the cumulative density function is

$$F(x) = \begin{cases} 0 & x < \alpha \\ \frac{x - \alpha}{\beta - \alpha} & \alpha \leq x \leq \beta \\ 1 & x > \beta \end{cases} \quad (8)$$

The Normal Distribution

The normal distribution is a two-parameter single-peaked distribution which has as its

support the entire real line. The density function for mean μ and variance σ^2 is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \quad (9)$$

The cumulative density function is the the integral of this, often denoted $\Phi(x)$, which cannot be simplified analytically. Refer to a computer program such as *Mathematica* or to tables in statistics texts for its values.

The Lognormal Distribution

If $\log(x)$ has a normal distribution, it is said that x has a lognormal distribution. This is a skewed distribution which has the set of positive real numbers as its support, since the logarithm of a negative number is not defined.

A.6 Supermodularity

Suppose that there are N players in a game, subscripted by i and j , and that player i has a strategy consisting of $\bar{s}^i$ elements, subscripted by s and t , so his strategy is the vector $y^i = (y_1^i, \dots, y_{\bar{s}^i}^i)$. Let his strategy set be S^i and his payoff function be $\pi^i(y^i, y^{-i}; z)$, where z represents a fixed parameter. We say that the game is a **supermodular game** if the following four conditions are satisfied for every player $i = 1, \dots, N$:

(A1) S^i is a complete lattice.

(A2) $\pi^i : S \rightarrow R \cup \{-\infty\}$ is order semicontinuous in y^i for fixed y^{-i} , and order continuous in y^{-i} for fixed y^i , and has a finite upper bound.

(A3) π^i is supermodular in y^i , for fixed y^{-i} . For all strategy combinations y and y' in S ,

$$\pi^i(y) + \pi^i(y') \leq \pi^i(\sup\{y, y'\}) + \pi^i(\inf\{y, y'\}). \quad (10)$$

(A4) π^i has increasing differences in y^i and y^{-i} . For all $y^i \geq y^{i'}$, the difference $\pi^i(y^i, y^{-i}) - \pi^i(y^{i'}, y^{-i})$ is nondecreasing in y^{-i} .²

In addition, it is sometimes useful to use a fifth assumption:

(A5) π^i has increasing differences in y^i and z for fixed y^{-i} ; for all $y^i \geq y^{i'}$, the difference $\pi^i(y^i, y^{-i}, z) - \pi^i(y^{i'}, y^{-i}, z)$ is nondecreasing with respect to z .

The conditions for **smooth supermodularity** are:

A1' The strategy set is an interval in $R^{\bar{s}^i}$:

$$S^i = [\underline{y}^i, \overline{y}^i]. \quad (11)$$

A2' π^i is twice continuously differentiable on S^i .

²xxx Check on this. Shouldn't y^{-i} be in there too?

A3' (Supermodularity) Increasing one component of player i 's strategy does not decrease the net marginal benefit of any other component: for all i , and all s and t such that $1 \leq s < t \leq \bar{s}^i$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial y_t^i} \geq 0. \quad (12)$$

A4' (Increasing differences in one's own and other strategies) Increasing one component of i 's strategy does not decrease the net marginal benefit of increasing any component of player j 's strategy: for all $i \neq j$, and all s and t such that $1 \leq s \leq \bar{s}^i$ and $1 \leq t \leq \bar{s}^j$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial y_t^j} \geq 0. \quad (13)$$

The fifth assumption becomes

A5': (Increasing differences in one's own strategies and parameters) Increasing parameter z does not decrease the net marginal benefit to player i of any component of his own strategy: for all i , and all s such that $1 \leq s \leq \bar{s}^i$,

$$\frac{\partial^2 \pi^i}{\partial y_s^i \partial z} \geq 0. \quad (14)$$

Theorem 1

If the game is supermodular, there exists a largest and smallest Nash equilibrium in pure strategies.³

Theorem 1 is useful because it shows (a) existence of an equilibrium in pure strategies, and (b) if there are at least two equilibria (note that the largest and smallest equilibria might be the same strategy combination), then two of them can be ranked in the magnitudes of the components of each player's equilibrium strategy.

Theorem 2

If the game is supermodular and assumption (A5') is satisfied, then the largest and smallest equilibria are nondecreasing functions of the parameter z .

Theorem 3

If a game is supermodular, then for each player there is a largest and smallest serially undominated strategy, where both of these strategies are pure.

Theorem 4

Let $\underline{y}^i$ denote the smallest element of player i 's strategy set S^i in a supermodular game. Let y^* and $y^{*'}$ denote two equilibria, with $y^* \geq y^{*'}$, so y is the "big" equilibrium. Then,
1. If $\pi^i(\underline{y}^i, y^{-i})$ is increasing in y^{-i} , then $\pi^i(y^*) \geq \pi^i(y^{*'})$.

³The theorems are taken from Milgrom & Roberts (1990). Theorem 1 is their corollary to Theorem 5. Theorem 2 is their Theorem 6 and corollary. Theorem 3 is their Theorem 5, and 4 is their Theorem 7. For more on supermodularity, see Milgrom & Roberts (1990) or pp. 489-97 of Fudenberg & Tirole (1991). See also Topkis's 1998 book.

2. If $\pi^i(y^i, y^{-i})$ is decreasing in y^{-i} , then $\pi^i(y^*) \leq \pi^i(y^{*'})$.
3. If the condition in (1) holds for a subset N_1 of players, and the condition in (2) holds for the remainder of the players, then the big equilibrium y^* is the best equilibrium for players in N_1 and the worst for the remaining player, and the small equilibrium $y^{*'}$ is the worst equilibrium for players in N_1 and the best for the remaining players.

A.7 Fixed Point Theorems

Fixed points theorems say that various kinds of mappings from one set to another result in at least one point being mapped back onto itself. The most famous fixed point theorem is Brouwer's Theorem, illustrated in Figure 3. I will use a formulation from page 952 of Mas-Colell, Whinston & Green (1994).




Figure 3 A mapping with three fixed points

The Brouwer Fixed Point Theorem. Suppose that set A in R^N is nonempty, compact, and convex; and that $f : A \rightarrow A$ is a continuous function from A into itself. (“Compact” means closed and bounded, in Euclidean space.) Then f has a fixed point; that is, there is an x in A such that $x = f(x)$.

The usefulness of fixed point theorems is that an equilibrium is a fixed point. Consider equilibrium prices. Let p be a point in N -space consisting of one price for each good. Let P be the set of all possible price points. P will be convex and compact if we limit it to finite prices. Agents in the economy look at p and make decisions about consumption and output. These decisions change p into $f(p)$. An equilibrium is a point p^* such that $f(p^*) = p^*$. If you can show that f is continuous, you can show that p^* exists.

This is also true for Nash equilibria. Let s be a point in N – space consisting of one strategy for each player – a strategy combination. Let S be the set of all possible strategy combinations. This will be compact and convex if we allow mixed strategies (for convexity) and if strategy sets are closed and bounded. Each strategy combination s will cause each

player to react by choosing his best response $f(s)$. A Nash equilibrium is s^* such that $f(s^*) = s^*$. If you can show that f is continuous – which you can do if payoff functions are continuous – you can show that s^* exists.

The Brouwer theorem is useful in itself, and conveys the intuition of fixed point theorems, but to prove existence of prices in general equilibrium and existence of Nash equilibrium in game theory requires the Kakutani fixed point theorem. That is because the mappings involved are not one-to-one functions, but one-to-many- point correspondences. In general equilibrium, one firm might be indifferent between producing various amounts of output. In game theory, one player might have two best responses to another player's strategy.

The Kakutani Fixed Point Theorem (Kakutani [1941]) Suppose that set A in R^N is a nonempty, compact, convex set and that $f : A \rightarrow A$ is an upper hemicontinuous correspondence from A into itself, with the property that the set $f(x)$ is nonempty and convex for every x . Then f has a fixed point; that is, there is an x in A such that x is one element of $f(x)$.

Other fixed point theorems exist for other kinds of mappings – for example for a mapping from a set of functions back into itself. In deciding which theorem to use, care must be taken to identify the mathematical nature of the strategy combination set and the smoothness of the best response functions.

*A.8 Genericity

Suppose we have a space X consisting of the interval between 0 and 100 on the real line, $[0, 100]$, and a function f such that $f(x) = 3$ except that $f(15) = 5$. We can then say any of the following:

- 1 $f(x) = 3$ except on a set of measure zero.
- 2 $f(x) = 3$ except on a null set.
- 3 Generically, $f(x) = 3$.
- 4 $f(x) = 3$ almost always.
- 5 The set of x such that $f(x) = 3$ is dense in X .
- 6 The set of x such that $f(x) = 3$ has full measure.

These all convey the idea that if parameters are picked using a continuous random density, $f(x)$ will be 3 with probability one, and any other value of $f(x)$ is very special in that sense. If you start with point x , and add a small random perturbation ϵ , then with probability one, $f(x + \epsilon) = 3$. So unless there is some special reason for x to take the particular value of 15, you can count on observation $f(x) = 3$.

Statements like these always depend on the definition of the space X . If, instead, we took a space Y consisting of the integers between 0 and 100, which is 0,1,2,...,100, then it is not true that “ $f(x) = 3$ except on a set of measure zero.” Instead, if x is chosen randomly,

there is a $1/101$ probability that $x = 15$ and $f(x) = 5$.

The concept of “a set of measure zero” becomes more difficult to implement if the space X is not just a finite interval. I have not defined the concept in these notes; I have just pointed to usage. This, however, is enough to be useful to you. A course in real analysis would teach you the definitions. As with the concepts of “closed” and “bounded”, complications can arise even in economic applications because of infinite spaces and in dealing with spaces of functions, game tree branchings, or other such objects.

Now, let us apply the idea to games. Here is an example of a theorem that uses genericity.

Theorem: “Generically, all finite games of perfect information have a unique subgame perfect equilibrium.”

Proof. A game of perfect information has no simultaneous moves, and consists of a tree in which each player moves in sequence. Since the game is finite, each path through the tree leads to an end node. For each end node, consider the decision node just before it. The player making the decision there has a finite number N of choices, since this is a finite game. Denote the payoffs from these choices as $(P_1, P_2, \dots, P_N)$. This set of payoffs has a unique maximum, because generically no two payoffs will be equal. (If they were, and you perturbed the payoffs a little, with probability one they would no longer be equal, so games with equal payoffs have measure zero.) The player will pick the action with the biggest payoff. Every subgame perfect equilibrium must specify that the players choose those actions, since they are the unique Nash strategies in the subgames at the end of the game.

Next, consider the next-to-last decision nodes. The player making the decision at such a node has a finite number of choices, and using the payoffs determined from the optimal choice of final moves, he will find that some move has the maximum payoff. The subgame perfect equilibrium must specify that move. Continue this procedure until you reach the very first move of the game. The player there will find that some one of his finite moves has the largest payoff, and he will pick that one move. Each player will have one best action choice at each node, and so the equilibrium will be unique. Q.E.D.

Genericity entered this as the condition that we are ignoring special games in the theorem’s statement – games that have tied payoffs. Whether those are really special or not depends on the context.

*A.9 Discounting

A model in which the action takes place in real time must specify whether payments and receipts are valued less if they are made later, i.e., whether they are **discounted**. Discounting is measured by the discount rate or the discount factor.

*The **discount rate**, r , is the extra fraction of a unit of value needed to compensate for delaying receipt by one period.*

*The **discount factor**, δ , is the equivalent in present units of value of one unit to be received*

one period from the present.

The discount rate is analogous to the interest rate, and in some models the interest rate determines the discount rate. The discount factor represents exactly the same idea as the discount rate, and $\delta = \frac{1}{1+r}$. Models use r or δ depending on notational convenience. Not discounting is equivalent to $r = 0$ and $\delta = 1$, so the notation includes zero discounting as a special case.

Whether to put discounting into a model involves two questions. The first is whether the added complexity will be accompanied by a change in the results or by a surprising demonstration of no change in the results. A second, more specific question is whether the events of the model occur in real time, so that discounting is appropriate. The bargaining game of Alternating Offers from Section 12.3 can be interpreted in two ways. One way is that the players make all their offers and counteroffers between dawn and dusk of a single day, so essentially no real time has passed. The other way is that each offer consumes a week of time, so that the delay before the bargain is reached is important to the players. Discounting is appropriate only in the second interpretation.

Discounting has two important sources: time preference and a probability that the game might end, represented by the rate of time preference, ρ , and the probability each period that the game ends, θ . It is usually assumed that ρ and θ are constant. If they both take the value zero, the player does not care whether his payments are scheduled now or ten years from now. Otherwise, a player is indifferent between $\frac{x}{1+\rho}$ now and x guaranteed to be paid one period later. With probability $(1 - \theta)$ the game continues and the later payment is actually made, so the player is indifferent between $(1 - \theta)x/(1 + \rho)$ now and the promise of x to be paid one period later contingent upon the game still continuing. The discount factor is therefore

$$\delta = \frac{1}{1 + r} = \frac{(1 - \theta)}{(1 + \rho)}. \quad (15)$$

Table 2 summarizes the implications of discounting for the value of payment streams of various kinds. We will not go into how these are derived, but they all stem from the basic fact that a dollar paid in the future is worth δ dollars now. Continuous time models usually refer to rates of payment rather than lump sums, so the discount factor is not so useful a concept, but discounting works the same way as in discrete time except that payments are continuously compounded. For a full explanation, see a finance text (e.g., Appendix A of Copeland & Weston [1988]).

Table 2 Discounting

Payoff Stream	Discounted Value	
	r-notation (discount rate)	δ -notation (discount factor)
x at the end of one period	$\frac{x}{1+r}$	δx
x at the end of each period in perpetuity	$\frac{x}{r}$	$\frac{\delta x}{1-\delta}$
x at the start of each period in perpetuity	$x + \frac{x}{r}$	$\frac{x}{1-\delta}$
x at the end of each period up through T (first formula)	$\sum_{t=1}^T \frac{x}{(1+r)^t}$	$\sum_{t=1}^T \delta^t x$
x at the end of each period up through T (second formula)	$\frac{x}{r} \left(1 - \frac{1}{(1+r)^T}\right)$	$\frac{\delta x}{1-\delta} (1 - \delta^T)$
x at time t in continuous time	$x e^{-rt}$	—
Flow of x per period up to time T in continuous time	$\int_0^T x e^{-rt} dt$	—
Flow of x per period in perpetuity, in continuous time	$\frac{x}{r}$	—

The way to remember the formula for an annuity over a period of time is to use the formulas for a payment at a certain time in the future and for a perpetuity. A stream of x paid at the end of each year is worth $\frac{x}{r}$. A payment of Y at the end of period T has a present value of $\frac{Y}{(1+r)^T}$. Thus, if at the start of period T you must pay out a perpetuity of x at the end of each year, the present value of that payment is $\left(\frac{x}{r}\right) \left(\frac{1}{1+r}\right)^T$. One may also view a stream of payments each year from the present until period T as the same thing as owning a perpetuity but having to give away a perpetuity in period T . This leaves a present value of $\left(\frac{x}{r}\right) \left(1 - \left(\frac{1}{1+r}\right)^T\right)$, which is the second formula for an annuity given in Table 2. Figure 4 illustrates this approach to annuities and shows how it can also be used to value a stream of income that starts at period S and ends at period T .




Figure A.4 Discounting

Figure 4: Discounting

Discounting will be left out of most dynamic games in this book, but it is an especially important issue in infinitely repeated games, and is discussed further in Section 5.2.

*A.10 Risk

We say that a player is **risk averse** if his utility function is strictly concave in money, which means that he has diminishing marginal utility of money. He is **risk neutral** if his utility function is linear in money. The qualifier “in money” is used because utility may be a function of other variables too, such as effort.

We say that probability distribution F **dominates** distribution G in the sense of **first-order stochastic dominance** if the cumulative probability that the variable will take a value less than x is greater for G than for F , i.e. if

$$\text{for any } x, F(x) \leq G(x), \quad (16)$$

and (16) is a strong inequality for at least one value of x . The distribution F dominates G in the sense of **second-order stochastic dominance** if the area under the cumulative distribution G up to $G(x)$ is greater than the area under F , i.e. if

$$\text{for any } x, \int_{-\infty}^x F(y)dy \leq \int_{-\infty}^x G(y)dy, \quad (17)$$

and (17) is a strong inequality for some value of x . Equivalently, F dominates G if, limiting U to increasing functions for first-order dominance and increasing concave functions for

second-order dominance,

$$\text{for all functions } U, \int_{-\infty}^{+\infty} U(x)dF(x) > \int_{-\infty}^{+\infty} U(x)dG(x). \quad (18)$$

If F is a first-order dominant gamble, it is preferred by all players; if F is a second-order dominant gamble, it is preferred by all risk-averse players. If F is first-order dominant it is second-order dominant, but not vice versa. See Copeland & Weston (1988) for further details.⁴

Milgrom (1981b) has used stochastic dominance to carefully define what we mean by **good news**. Let θ be a parameter about which the news is received in the form of message x or y , and let utility be increasing in θ . The message x is more favorable than y (is “good news”) if for every possible nondegenerate prior for $F(\theta)$, the posterior $F(\theta|x)$ first-order dominates $F(\theta|y)$.

Rothschild & Stiglitz (1970) shows how two gambles can be related in other ways equivalent to second-order dominance, the most important of which is the **mean preserving spread**. Informally, a mean-preserving spread is a density function which transfers probability mass from the middle of a distribution to its tails. More formally, for discrete distributions placing sufficient probability on the four points a_1, a_2, a_3 , and a_4 ,

a mean-preserving spread is a set of four locations $a_1 < a_2 < a_3 < a_4$ and four probabilities $\gamma_1 \geq 0, \gamma_2 \leq 0, \gamma_3 \leq 0, \gamma_4 \geq 0$ such that $-\gamma_1 = \gamma_2, \gamma_3 = -\gamma_4$, and $\sum_i \gamma_i a_i = 0$.




Figure 5: Mean Preserving Spreads

Figure 5 shows how this works. Panel (a) shows the original distribution with solid

⁴xxx Find chapter or page number.

bars. The mean is $0.1(2) + 0.3(3) + 0.3(4) + 0.2(5) + 0.1(6)$, which is 3.9. The spread has mean $0.1(1) - 0.3(.3) - 0.1(4) + 0.2(6)$, which is 0, so it is mean-preserving. Panel (b) shows the resulting spread-out distribution.

The definition can be extended to continuous distributions, and can be alternatively defined by taking probability mass from one point in the middle and moving it to the sides; panel (c) of figure A.5 shows an example. See Rasmusen & Petrakis (1992), which also, with Leshno, Levy, & Spector (1997), fixes an error in the original Rothschild & Stiglitz proof.

"41d" indicates that the concept is *defined* on page 41, "33n" that the entry is in an *endnote*, and "54p" that the entry is in a homework *problem*. Articles and books are indexed separately, in the bibliography.

This index was prepared by the author, not by a research assistant or by the publisher.

Action combination 13d
 Action set 13d
 Action 13d
 Additively separable 130
 Advantages
 first-mover 29
 last-mover 300
 second-mover 37p
 Adverse selection 161, 211-39
 Advertising 289n, 291p
 Affiliated variables 338n
 Affine transformation 76
 Agent 162d
 Almost always 396, 404-06
 Almost perfect information 63dn
 Alternating Offers game 299, 320
 Annuity 395,
 Anonymity condition 271
 Argmax 394
 Assessment 140
 Assurance game 35n
 Asymmetric information 49d,50
 Auctions 323-339
 categories 324
 common-value 324, 331
 descending 327
 Dutch 327
 English 325
 first-price 325
 open-exit 325
 private-value 324,330
 second-price 327,337
 Auditing Game I, II, III 79, 86n
 Authority 184p
 Axelrod tournament 151, 156n

 Backward induction 110
 Bagehot model 226
 Bank loans 153, 194, 231
 Bankruptcy constraint 186, 209p
 Bargaining 295-322, 165, 169
 Alternating Offers game 299, 320
 bargaining power 165, 169, 196
 Splitting a Pie game 296
 take-it-or-leave-it offer 165
 Battle of the Bismarck Sea game 21, 34n
 Battles of the Sexes game 28, 35n, 89p, 157p
 Bayes's Rule 55d, 64p
 Bayesian equilibrium 54
 Beer-Quiche game 155n
 Behavior strategy 84n
 Beliefs

Out-of-equilibrium 143, 155n
 Posterior 55
 Prior 54
 Updating 55
 Benoit-Krishna game 134p
 Bertrand equilibrium 35n, 343, 367n, 370p
 Best reply 19d
 Best response 19d, 82
 Bid-ask spread 226
 Bilateral Trading games I, II, III, and IV 309, 321n
 Bimatrix game 34n
 Biology 106n, 125, 133n, 289n
 Blackboxing 100
 Boiling in oil contract 175, 182n
 Boundary 32
 Bounded set 70
 Bourgeois strategy 128
 Boxed Pigs game 25, 35
 Branch 40d
 Bridge card game 85n
 Broadway Game
 I 174
 II 178
 Brouwer fixed-point theorem 404
 Budget-balancing constraint 202, 207n, 262
 Buyout, entry for 386, 392p

 Calculus 394, 399, 400
 Capacity constraints 368n, 387, 345
 Cardinal utility 34n, 76
 Certain information 47d
 Chainstore paradox 110, 129
 Cheap talk 75
 Chatterjee-Samuelson mechanism 315
 Chicken game 71, 85n
 Closed set 70, 394, 395
 Civic Duty game 78
 Closed-loop 105n
 Cloud 44
 Coalition 319n
 Coalitionproofness 103
 Coarsening 45d
 Coase Conjecture 369n
 Coase Theorem 170
 Co-insurance 200
 Colonel Blotto games 87n
 Combination 395
 Combination, strategy 17d, 33n
 Commitment 29, 80, 117, 237p
 Common knowledge 47d, 62n, 63nd, 158p
 In Entry Deterrence IV and V
 Common-value auction 324, 331
 Communication 33, 75, 105, 32
 Compact set 395
 Comparative statics 357
 Competition constraint 121, 213
 Competitive fringe 87n
 Complementary inputs 360
 Complete information 50d, 51
 Complete robustness 146
 Completely mixed strategy 67d
 Concave function 395
 Concordant beliefs 47d
 Conditional likelihood 56

Constraints

- bankruptcy 186, 209p
- competition 121, 213
- incentive compatibility 121, 173, 213
- nonpooling 213
- participation 173, 181n, 213, 233n
- self selection 213

Continuous function 395

Continuous strategies 81, 100

Continuous time 85n

Continuum 395

Continuum of players 69, 85, 303

Continuum of equilibria 72, 207n, 296, 306

Contractible variable 170

Contraction mapping 395

Contracts

- boiling-in-oil 175, 182n
- first-best 172d, 246
- forcing 167
- linear 167
- pooling 142, 145, 213d
- second-best 172d
- selling-the-store 178, 183p
- threshold 167

Contribution games 77

Convex function 396

Convex set 173, 396

Cooperative games 21, 319n

Co-opting 205n

Coordination games 29, 35n, 77

Copayment 200

Corner solution 72

Correlated strategies 74

Correlated-value auction 324

Correspondence 396

Cournot model 30, 81, 87n, 156p, 340, 367n, 370p

Cream skimming 224

Credible threat 94

Cumulative density function 400

Customer Switching Costs game 123

Dangerous Coordination game 30

Decision theory 14

Decision tree 14

Deductible 200

Density 400

Derivatives, calculus 394, 399, 400

Descending auction 327

Determinants 399

Deviation 33n

Differential games 87n

Dimensionality condition 114, 131n

Direct incentive-compatible mechanism 315

Disagreement point 297

Discontinuity

Discoordination games 37p, 77

Discount factor 406

Discount rate 406

Discounting 73, 299, 406-08

- in bargaining 301

- in the Folk Theorem 113

Distribution 400

- exponential 400

- lognormal 400

uniform 400

Dollar Auction game 336n

Domain 396

Dominance-solvable 23d

Dominance, stochastic 408

Dominant diagonal condition 361

Dominant strategy 19d

Dominant-strategy equilibrium 20d

Dominant-strategy mechanism 262

Dominated strategy 19d, 61

Double auction mechanism 311

Double-sided moral hazard 179n

Dry Cleaners game 14

Duelling games 74

Duopoly 30, 81, 87n, 156p, 340, 367n, 370p

Durable monopoly 362, 369n

Dutch auction 327, 336

Dynamic consistency 105n

Dynamic game 90

Edgeworth paradox 345

Education Game

- I 268

- II-IV 271

- V 277

- VI 278

- VII 281

Efficiency wage 181n, 185, 205n, 210p, 230

Efficient rationing 367n

Elmer's Appetite 64p

Emotions 101, 132n

End node (end point) 40d

English auction 325

Entry deterrence 372-392

Entry Deterrence game 106p, 134p

- I 93

- II and III 139

- IV and V 147, 158p

Epsilon equilibrium 130n

Equilibrium 18d

- continua 72, 207n, 296, 306

- epsilon 130n

- existence 70, 225, 281

- fully revealing 213, 246

- implausible 91, 143

- multiple 18

- Pareto-dominant 28, 37p

- partially revealing 213

Equilibrium concept 18d

- coalitionproof 103

- dominant-strategy 20d

- evolutionarily stable strategies 125d, 135p

- iterated dominance 23d, 118

- maximin 116

- Nash 26d, 36p, 54, 100

- perfect bayesian 140d, 155

- sequential 140d

- Stackelberg 30, 83, 87n, 106n

- subgame perfect 91d, 137

- trembling-hand-perfect 139

Equilibrium outcome 18

Equilibrium path 89d

Equilibrium refinement 28

- forward induction 158p

- Equilibrium strategy 18
- Evolutionary dynamics 128
- Evolutionarily stable strategy (ESS) 125d, 135p
- Exemplifying theory 2
- Existence of equilibrium 70, 225, 281
- Exotic refinements 155n
- Expected externality mechanism 316, 321n
- Expected utility 48, 63n
- Expensive Talk Game 157p
- Exponential distribution 400
- Extensive form 40d, 94, 155n
- Extrinsic uncertainty 74

- Fairness 101, 298
- Fat-cat effect 390n
- Finer information partition 45d
- First-best
 - contract 172d, 246
 - outcome 182p
- First-mover advantage 29
- First-order condition approach 173, 181n
- First-order stochastic dominance 408
- First-price auction 325
- Fixed-dash theorem 403
- Focal point 31-33, 36n, 89p
- Folk Theorem 112d, 120, 131n
- Follow the Leader game
 - I 39
 - II 49
 - III 51
- Forcing contract 167
- Ford, Henry, wages 23
- Forgiving strategy 152
- Forward induction 158p
- Free rider problem 378
- Fully insured 197
- Fully revealing equilibrium 213, 246
- Function 396
- Functional forms 400

- Game theory
 - history 1-2, 6-7n, 51
 - method 2
- Game tree 15, 40d
- Games--in general
 - cooperative vs. noncooperative 21
 - zero- vs variable-sum 25d, 34n
- Games--specific games
 - Alternating Offers 299, 320
 - Auditing Game I, II, III 79, 86n
 - Battle of the Bismarck Sea 21, 34n
 - Beer-Quiche 155n
 - Benoit-Krishna 134p
 - Bilateral Trading I, II, III, IV 309, 321n
 - Boxed Pigs 25, 35
 - Broadway Game I, II 174, 178
 - Chicken 71, 85n
 - Civic Duty 78
 - Colonel Blotto 87n
 - Cournot Game 81
 - Customer Switching Costs 123
 - Dangerous Coordination 30
 - Discoordination 37p
 - Dollar Auction 336n
 - Durable Monopoly 362
 - Education I, II, III, IV, V, VI, VII 268, 271, 277, 278, 281
 - Entry for Buyout 386
 - Expensive Talk 157p
 - Follow the Leader I, II, III
 - The Free Rider Problem in Takeovers 378
 - General 2x2 game 75
 - Government Procurement 255
 - Grab the Dollar 73, 135p
 - Greenmail to Attract White Knights 381
 - Hawk-Dove 127
 - The Hotelling Location Game 354
 - The Hotelling Pricing Game 350
 - Insurance Game I, II, III 197, 222
 - The Iteration Path Game 25
 - Lemons I, II, III, IV 215, 218
 - Limit Pricing as Signal Jamming 285
 - Matching Pennies 88p
 - The Minimax Illustration Game 116
 - Modeller's Dilemma 27
 - Nuisance Suits I, II, III 96
 - Patent Race for a New Market 374
 - Patent Race for an Old Market 376
 - Phd Admissions 144
 - The Png Settlement Game 59
 - The Police Game 86n
 - Predatory Pricing 383
 - Product Quality 119
 - Production Game I, II, III, IV, V, VI, VIa, VII 164, 169, 211, 242
 - Repossession Game I, II, III 194
 - The Salesman Game 244
 - Splitting a Pie 296
 - The Swiss Cheese Game 24
 - Testing 232
 - Underpricing New Stock Issues 283
 - The Utopian Exchange Economy Game 126
 - The War of Attrition 72
 - The Welfare Game 67
- Gang of Four 149
- Generically 396, 404-06
- Good news 409
- Government Procurement game 255
- Grab the Dollar game 73, 135p
- Greek alphabet 395
- Greenmail 380
- Grim strategy 111d, 135p
- Groves mechanism 261, 264n

- Harsanyi doctrine 53, 62
- Harsanyi transformation 51, 63n
- Hawk-Dove game 127
- Hidden actions 161
- Hidden knowledge 161, 179n, 241
- Hold-up potential 193
- Hotelling models 350

- Imperfect information 47d, 63n
- Imperfectly separating equilibrium 213
- Implicit function theorem 358
- Implementing mechanism 241
- Incentive compatibility 121, 173, 213
- Incomplete information 50d, 51, 97
- Incomplete information folk theorem 151d

Individual rationality constraint 181n
 Infimum 394
 Infinite-horizon model 119, 130n
 Infinitely repeated game 119
 Information partition 45d
 Information set 43d
 Informational rents 248, 260
 Informed player 162
 Innovation 372, 390n
 Insurance 46, 99, 182n, 196, 222, 230, 235n, 237p
 Insurance Game
 I and II 197
 III 222
 Intensity rationing 345d, 367n
 Interbuyer price discrimination 248, 250
 Interquantity price discrimination 248
 Intuitive criterion 145, 155n
 Invariance 297
 Inverse-intensity rationing 345d
 Invertible function 172
 Iterated dominance 21, 105n
 Iterated-dominance equilibrium 23d, 118
 Iteration Path Game 25

 Joint Ventures game 65p

 Kakutani fixed-point theorem 404
 Kyle Model of Market Microstructure 228

 Lagrange multiplier 396
 Last-mover advantage 300
 Lattice 396
 Law 59, 95, 107p, 191, 266p
 Least-cost-avoider
 Lemons Game 265n
 I and II 215, 234n
 III and IV 218, 219, 234n
 Limit pricing 157p, 285
 Linear contract 167
 Linear equilibrium 229, 313
 Location models 349, 369n
 Lognormal distribution 400
 Lower semicontinuous correspondence 397
 Lucky Executive Game 186

 Macroeconomics 36n, 105n
 Malice 101
 Marginal likelihood 55
 Market microstructure 225
 Markov strategy 123d
 Matching Pennies game 88p
 Matrix game 34n
 Maximand 397
 Maximin equilibrium 116
 Maximin strategy 115d
 Mean-preserving spread 409
 Measure zero 397, 404-06
 Mechanism 240-66, 308-18
 Mediation 33, 314
 Menu of contracts 246
 Metric 397
 Metric space 395, 397
 Minimax criterion 116
 Minimax Illustration Game 116
 Minimax strategy 114d
 Minimax Theorem 116
 Minimax value 114d
 Mixed extension 67d
 Mixed strategy 66-81, 237p, 287
 the payoff-equating method 71
 correlated strategies 74
 versus randomizing 80
 Modeller's Dilemma game 27
 Monitoring 86n, 207p, 208p
 Monotone likelihood ratio property 181n
 Monty Hall Problem 64p
 Moral hazard 132n, 161-210
 Moral hazard with hidden knowledge 241, 263n
 Most-favored-customer clause 370n
 Move 13d
 Multiple equilibria 18
 Mutation 129
 Mutual knowledge 47
 Myerson-Satterthwaite model 308, 322p

 Nash bargaining solution 296, 319n, 321p
 Nash equilibrium 26d, 36p, 54, 100
 Nature 13d
 Niceness 152
 Nixon and the Madman Theory 105n
 Node 40d, 62n
 Noisy duel 74
 Nonlinear pricing 249, 252
 Noncooperative game theory 21
 Nonpooling constraint 213
 Nonresponders 368n
 Nonseparating constraint 213
 Nonzero-sum game 25d
 Normal distribution 400
 Normal form of a game 39d, 94
 Nuisance Suits I to III 96

 Oligopoly 342
 One-dash price equilibrium 312
 One-shot game 110
 One-sided Prisoner's Dilemma 118, 132n
 One-to-one mapping 172, 397
 Onto 397
 Open set 394, 397
 Open-exit auction 325
 Open-loop 05n
 Open-set problem 99
 Opportunism 193
 Order of play 40
 Ordinal utility 34n
 Out-of-equilibrium behavior 307
 Out-of-equilibrium beliefs 143, 155n
 forward induction 158p
 intuitive criterion 145, 155n
 passive conjectures 142, 145
 Outcome 14d
 Outcome matrix 39d, 62n
 Overlapping generations model 122, 133p
 Overtaking criterion 130

 Pareto dominance 28, 35nd, 37p, 104
 Pareto frontier in bargaining 297
 Pareto optimality 82

- Pareto Perfection Puzzle game 104
- Partially pooling equilibrium 213
- Partially revealing equilibrium 213
- Partially separating equilibrium 213
- Participation constraint 173, 181n, 213, 233n
- Partition, information 45d
- Passive conjectures 142, 145
- Patent races 373, 390n, 392p
- Path 40d
- Payoff 12d
- Payoff-equating method for mixed strategies 71
- Perfect bayesian equilibrium 140d, 155
- Perfect equilibrium 91, 105n, 108p
- Perfect information 47d, 63n, 196
- Perfect Markov equilibrium 124d
- Perfect price discrimination 248, 250
- Perfect recall assumption 84n
- Permutation 397
- Perpetuity 398
- PhD Admissions game 144
- Player 12d
- Pliny and the Freedman's Trial 107p
- Png Settlement Game 59
- Poker example for information 50
- Police game 86n
- Politics 88p, 106p
- Pooling 142, 145, 213d
- Posterior belief 55
- Precommitment 117
- Predatory pricing 93, 383, 391p
- Predecessor node 40d
- Pre-emption game 73
- Present discounted value 406-08
- Price discrimination 247, 263n, 266p
- Price dispersion 230
- Principal 162d, 179n
- Principal-agent model 161
- Prior belief 54
- Prisoner's Dilemma game 20, 33n, 34n
 - repeated 110, 134p, 150
 - one-sided 118
- Private information 50
- Private-value auction 324, 330
- Probability density 400
- Procurement 255
- Product differentiation 349-57
- Product quality 119, 132n, 134p, 188, 205n, 290p
- Production Game 164
 - I 165
 - II-V 169
 - VI-VIa 211
 - VII 242
- Proportional rationing 345d
- Provokable strategy 152
- Pseudo-player 13d
- Public goods 203, 261
- Pure strategy 66d
- Quadratic formula 399
- Quality 119, 132n, 134p, 188, 205n, 290p
- Quasi-concave 398
- Randomization 80
- Range 398

- Rao-Blackwell Theorem 181n
- Ratchet effect 263n
- Rate-of-return regulation 255
- Rationalizable strategy 35n
- Rationing 345
 - Efficient 367n
 - Intensity 345d, 367n
 - Inverse-Intensity 345d
 - Proportional 345d
- Reaction function 83, 348
- Real time 42
- Realization of a game 13
- Refinements
 - forward induction 158p
 - intuitive criterion 145, 155n
 - of equilibrium 28
 - of information 45d, 179
- Regression to the mean 58, 64n, 204n
- Renegotiation 103, 193
- Rent-seeking 73, 338p, 376, 392p
- Repeated games 109-136
 - Prisoner's Dilemma 110, 134p, 150
- Repossession Game I, II, III 194
- Reputation 117, 132n
 - Gang of Four model 149
 - for irrationality 105n
 - in product quality 119, 132n
- Reservation utility 165
- Reservation wage 165
- Reserve prices in auctions 337n
- Residual claimant 178
- Responder 368n
- Revelation Principle 243, 263n
- Revenue-equivalence in auctions 328
- Risk 46, 65p, 182n, 196, 330, 408-10
- Risk-aversion 207n, 221, 408
- Risk-neutrality 408
- Rules of the game 12
- Running from the Gestapo game 88p
- Salesman Game 244
- Samaritan's Dilemma 67
- Screening 161, 277
- Second-best contract 172d
- Second-mover advantage 37p
- Second-order stochastic dominance 408
- Second-price auction 327, 336
- Security value 114d
- Self-selection constraint 213
- Selling-the-store contract 178, 183p
- Semi-separating equilibrium 213
- Separability 181n, 363
- Separating equilibrium 143, 145, 213d, 246
- Sequential equilibrium 140d
- Sequential-move game 38
- Sequential rationality 92
- Serially undominated 23
- Service flow 362
- Settlement range in litigation 98
- Shapley value 319n
- Share auction 337n
- Shifting supports 177
- Side payments 21
- Signal jamming 285

- two signals 282
- versus screening 161, 277
- Silent duel 74
- Simultaneous-move games 38
- Single-crossing property 249, 270, 280
- Singleton information set 45d, 47
- Sleeping patent 378
- Solution concept 18
- Splitting a Pie game 296
- Stability 83, 87n
- Stackelberg 30, 83, 87n, 106n
- Starting node 40d
- State of the world (state of Nature) 53d
- State-space diagram 196, 223, 237p, 238p
- Stochastic dominance 408
- Strategic complement 360, 368n
- Strategic form of a game 39d, 94
- Strategic substitute 368n
- Strategically equivalent auctions 328
- Strategy 16d
 - equilibrium strategy 18
 - grim strategy 111d, 135p
 - Markov strategy 123d
 - Tit-for-tat 112d, 151
 - trigger strategy 131n
 - versus action 16
- Strategy combination 17d, 33n
- Strategy set 16d
- Strategy space 16d, 87n
- Strict (strong) 398
- Subgame 91d, 105n
- Subgame perfect equilibrium 91d, 137
- Successor node 40d
- Sufficient statistic condition 175, 181n
- Sunk costs 99
- Sunspot model 74
- Supergame 132nd
- Supermodularity 359, 401
- Support of a distribution 398
- Supremum 394
- Swiss Cheese Game 24
- Switching costs 123
- Symmetric information 49d
- Synergism 368n
- Take-it-or-leave-it offer 165
- Takeovers 378
- Team theory 206n
- Team 202d, 209p
- Tender Trap game 35n
- Testing 232, 237n
- Threat point 297
- Threats 94, 108p
- Three-step procedure 173
- Three-Way Duel game 106p
- Threshold contract 167
- Time consistency 105n
- Time line 42
- Time-separable 363
- Tit-for-tat strategy 112d, 151
- Topology 398
- Tournament 188, 205n, 207n
- Tremble 92
- Tremble Game 92
- Trigger strategy 131n
- Truth-telling. constraint 243
- Two-by-two game 20,34
- Two-sided asymmetry 320n
- Type of a player 53d
- Uncertainty of information 48d
- Underpricing New Stock Issues game 283
- Uniform distribution 400
- Uninformed player 162
- Uniqueness of equilibrium 18, 130n
- Unravelling 242, 263n, 265p
- Upper semicontinuous correspondence 397, 398
- Utopian Exchange Economy game 126
- Valuation in auctions 324
- Value in auctions 324
- Variable-sum game 25d
- Vector 398
- Verifiability 170
- Vickrey auction 327
- Von Neumann-Morgenstern utility 48, 63n
- Voting Paradox game 89p
- War of attrition 72, 85n
- Weak 398
- Weak Nash equilibrium 27
- Weakly dominated strategy 22d
- Welfare Game 67
- Winner's curse 331, 337n
- Yardstick competition 188
- Zero-sum game 25d,34n