



Topic
Business
& Economics

Subtopic
Economics

Games People Play: Game Theory in Life, Business, and Beyond

Course Guidebook

Professor Scott P. Stevens
James Madison University



PUBLISHED BY:

THE GREAT COURSES
Corporate Headquarters
4840 Westfields Boulevard, Suite 500
Chantilly, Virginia 20151-2299
Phone: 1-800-832-2412
Fax: 703-378-3819
www.thegreatcourses.com

Copyright © The Teaching Company, 2008

Printed in the United States of America

This book is in copyright. All rights reserved.

Without limiting the rights under copyright reserved above,
no part of this publication may be reproduced, stored in
or introduced into a retrieval system, or transmitted,
in any form, or by any means
(electronic, mechanical, photocopying, recording, or otherwise),
without the prior written permission of
The Teaching Company.



Scott P. Stevens, Ph.D.

Professor of Computer Information Systems
and Management Science
James Madison University

Professor Scott P. Stevens is Professor of Computer Information Systems and Management Science at James Madison University (JMU) in Harrisonburg, Virginia. In 1979, he received B.S. degrees in both Mathematics and Physics from The Pennsylvania State University,

where he was first in his graduating class in the College of Science. In 1987, Stevens received his Ph.D. in Mathematics from The Pennsylvania State University, working under the direction of Torrence Parsons. Parsons himself received his Ph.D. under the direction of Albert W. Tucker, the well-known game theorist. After Parson's unexpected death in 1987, Stevens completed his doctoral work under George E. Andrews, the world's leading expert in the study of integer partitions. Dr. Stevens's doctoral thesis was "Group-Action Graphs and Ramsey Graph Theory: Investigating the Ramsey Numbers $R(K_{1,n}, K_{k,m})$ and $R(K_{1,n}, B_{k,m})$."

Professor Stevens's research interests include combinatorics, game theory, graph theory, statistics, and neural networks. In collaboration with his JMU colleagues, he has published articles on a wide range of topics. These include papers on neural network prediction of survival in blunt-trauma-injured patients; the effect of private school competition on public schools; standards of ethical computer usage in different countries; automatic data collection in business; and optimizing the purchase, transportation, and deliverability of natural gas from the Gulf of Mexico. His publications have appeared in the *European Journal of Operations Research*; *International Journal of Operations & Production Management*; *Political Research Quarterly*; *Omega: The International Journal of Management Science*; *Neural Computing and Applications*; *INFORMS Transactions on Education*; *Decision Sciences Journal of Innovative Education*; and in a number of conference proceedings. Much of his recent research focuses on the more effective delivery of mathematical concepts to students.

Professor Stevens has consulted to a number of firms, including Corning Glass, C&P Telephone, and the Globaltec Corporation. He is a member of the American Mathematical Association, the Institute for Operations Research and the Management Sciences, and Alpha Kappa Psi Business Fraternity.

Professor Stevens's primary professional focus since moving to JMU in 1984 has been his deep commitment to excellence in teaching. He was the 1999 recipient of JMU's Carl Harter Award, the university's highest teaching award. In 2001, he was named Outstanding Graduate Teacher in the JMU MBA program and became the first professor to be named Outstanding Teacher five times by the students of the undergraduate business program. His teaching interests are wide and include game theory, statistics, operations research, physics, calculus, and the history of science.

Table of Contents

INTRODUCTION

Professor Biography	i
Course Scope	1

LECTURE GUIDES

LECTURE 1

The World of Game Theory	4
--------------------------------	---

LECTURE 2

The Nature of the Game	7
------------------------------	---

LECTURE 3

The Real Life Chessboard—Sequential Games	10
---	----

LECTURE 4

Life's Little Games—The 2×2 Classic Games	13
--	----

LECTURE 5

Guessing Right—Simultaneous Move Games	16
--	----

LECTURE 6

Practical Applications of Game Theory	20
---	----

LECTURE 7

A Random Walk—Dealing with Chance Events	23
--	----

LECTURE 8

Pure Competition—Constant-Sum Games	27
---	----

LECTURE 9

Mixed Strategies and Nonzero-Sum Games	30
--	----

LECTURE 10

Threats, Promises, and Commitments	34
--	----

Table of Contents

LECTURE 11

Credibility, Deterrence, and Compellence	38
--	----

LECTURE 12

Incomplete and Imperfect Information	41
--	----

LECTURE 13

Whom Can You Trust?—Signaling and Screening	45
---	----

LECTURE 14

Encouraging Productivity—Incentive Schemes	49
--	----

LECTURE 15

The Persistence of Memory—Repeated Games	53
--	----

LECTURE 16

Does This Stuff Really Work?	56
------------------------------------	----

LECTURE 17

The Tragedy of the Commons	60
----------------------------------	----

LECTURE 18

Games in Motion—Evolutionary Game Theory	63
--	----

LECTURE 19

Game Theory and Economics—Oligopolies	67
---	----

LECTURE 20

Voting—Determining the Will of the People	70
---	----

LECTURE 21

Auctions and the Winner’s Curse	73
---------------------------------------	----

LECTURE 22

Bargaining and Cooperative Games	77
--	----

LECTURE 23

Game Theory and Business—Co-opetition	81
---	----

Table of Contents

LECTURE 24

All the World's a Game	84
------------------------------	----

SUPPLEMENTAL MATERIAL

Timeline	88
Glossary	95
Biographical Notes	105
Bibliography	110

Games People Play: Game Theory in Life, Business, and Beyond

Scope:

To many, modern game theory began in 1944 when John von Neumann and Oskar Morgenstern published their landmark book, *Theory of Games and Economic Behavior*, putting the ideas of neoclassical economics into a more general framework. Earlier economic analysis was capable of describing outcomes in a market that was essentially unaffected by the action of any one individual. Game theory's ability to analyze situations where individual choice *does* matter has resulted in a long line of game theorists receiving Nobel Prizes in Economics.

But what *is* game theory? Simply put, game theory is the study of strategic, interactive decision making among rational individuals. Any time people are making decisions that affect others or in response to the actions—or even the *expected* actions—of others, they're playing a game. That's a broad definition, and it means that much of our lives are spent sailing on a sea of games. On that voyage, game theory can serve as both chart and compass.

At sea, the compass lets the captain set a course and stick to it. Given his current position, what direction takes him where he wants to go? In a similar way, game theory lets us model many real-life situations in which we find ourselves. Its analytical tools then help us to gain insight into where to go from there. Like a compass, it helps us set our course. At least as important as the compass, though, is the chart.

One purpose of the chart is to identify waters—situations—that are inherently dangerous or that must be navigated with great care. Game theory helps us to recognize such personal, professional, or political situations. Recognizing the character of the game that we're playing often allows us to play it better. Moreover, recognizing a dangerous game may allow us to sidestep it entirely, replacing it with one that is more to our liking.

These ideas can be applied to an interaction as trivial as where to meet for lunch or as earthshaking as whether to risk provoking nuclear war. The tools are the same and can be applied with varying degrees of sophistication. This course is intended as a first exploration of the world of strategic decision making. We'll provide the basic tools of the trade, clearly demonstrate how to use them, and then look at some of their applications. And the applications are far-reaching. Many of our examples will be in business, especially when we discuss merging competitive and cooperative business frameworks—the business model of “co-opetition.” That said, we'll also examine applications of game theory to economics, military strategy, politics, and biology. We'll even look at less lofty subjects, such as NASCAR, soccer, traffic jams, and getting your kid to do her homework.

Game theory is a young field—less than a century old. In that time, it has made remarkable advances, but it's far from complete. Traditional game theory assumes that the players of games are rational—that they act in best accordance with their own desires given their knowledge and beliefs. This assumption does not always appear to be a reasonable one. In certain situations, the predictions of game theory and the observed behavior of real people differ dramatically. We will look at why this may be so and discuss such ideas as “bounded rationality” that are intended to address this disparity. Also, we'll take a look at some of the exciting work being done in the areas of behavioral game theory and evolutionary game theory, two promising new branches of the field.

Many of the games studied in this course will be small, owing to the limitations of time. The techniques presented, though, can be applied to much larger problems, especially when the number-crunching power of modern computers can be brought to bear. Occasionally, the results of such efforts are remarkable. Political scientist Bruce Bueno de Mesquita has used game theory to predict international events for the CIA and others. In 1984, he predicted that, after the death of Ayatollah Khomeini, Iran's clerical leadership would devolve to the ayatollah Hojatolislam Khamenei and to a junior cleric named Akbar Hashemi Rafsanjani.

The prediction was surprising, to say the least. Khomeini had already designated a successor, and it wasn't Khamenei. And Rafsanjani was such

a political nonentity to the West that his name had not yet even appeared in *The New York Times*. Yet, when Khomeini died five years later, Khamenei and Rafsanjani took up the reins. One must view any particular anecdote with skepticism, but CIA analysts claim that Bueno de Mesquita's success rate in predictions is over 90%.

The world of game theory is extensive and rich, and the mathematics upon which most of it is built can be formidable. This course is intended as an introduction to that world for the intelligent layperson; thus, I'll keep the mathematical complexity to a minimum. Those with the interest and mathematical acumen will find more sophisticated treatments of the subject to be powerful, far-reaching, and often beautiful fields of study. This booklet includes a supplementary reading list, and I've taken care to indicate those books that are intended for the more mathematically oriented viewer.

Although “game theory” is actually strategic decision making, it shares one characteristic with the more traditional meaning of “game”: It’s fun. I hope that by the time you complete this course, you’ll be able to recognize the games you play every day and exclaim with the same zeal as the legendary Sherlock Holmes, “The game is afoot!”

The World of Game Theory

Lecture 1

Any time people are interacting with one another—responding to the choices of others or what they think those choices will be—they’re playing a game, and that’s what game theory is really about. What’s the best way to play the game that we’re in?

We begin our discussion of game theory by playing a game. In this game, you begin with \$100 and a button that you may push. You are playing with 100 other people whose identities are unknown to you. Pushing the button has two effects: When you push your button, every other player loses \$2. Also, if you lose money because other players push their buttons, pushing your button will cut those losses in half.

This simple game introduces many basic game theory concepts that we’ll examine more closely later: players, strategies, payoffs, rationality, and common knowledge. Actual play of this game shows some surprising results. For example, across groups of strangers who have no training in game theory, percentages of people who push the button vary widely—anywhere from 30% to 70%. On average, 50% of people push the button. This means each person who did not push the button is now broke, and each who did still has \$50. This outcome differs from what many people would predict before playing the game. Also, if we imagine a version of the game in which pushing the button leads to not only monetary gain or loss but also to deployment of armed forces or weapons, we see that this game—leading to 30% to 70% of players opting to attack—may be an inherently dangerous one.

Game theory is the study of interactive, strategic decision making among rational individuals. Games include international conflicts, as well as threats and promises in general. Board games, card games, and sports can be analyzed with game theory, but such games are not the focus of the field.

Game theory helps us determine how to play the game we’re in or how to change it into a game that better suits us. Such lessons are useful even to professional decision makers. For example, Max Bazerman, a professor at

Harvard Business School, demonstrated the error of failing to think ahead by auctioning a \$100 bill to Wall Street investors for \$465!

Game theory begins with simple examples, using them to develop general principles that assist in superior decision making. The predictions of game theory give us a baseline for understanding the decisions we make in everyday life.

With the 1944 publication of the *Theory of Games and Economic Behavior*, game theory came to the attention of the larger world. John von Neumann, the father of modern game theory, teamed up with economist Oskar Morgenstern to write the book. Its goal was to put neoclassical economic theory on a firm scientific footing.

Government licensing of the radio spectrum provides a good example of game theory's utility. Historical approaches to licensing include administrative processes; lotteries; first-come, first-served approaches; and auctions. In the United States, game theorists created a multi-objective auction structure to replace the failed administrative process and lottery approaches. This auction system has been a triumph, raising \$400 billion for the U.S. treasury in its first five years and efficiently distributing licenses.

In this course, we'll look at applications of game theory to a wide spectrum of topics and explore important general ideas, such as strategies, threats, promises, brinkmanship, incomplete information, and chance. We'll consider game theory applications to interesting real-world situations, phenomena, and processes such as global warming, voting, market-entry and price-setting decisions, the evolution of cooperative behavior, and much more.

A game has three major components: players, strategies, and payoffs. A player is a decision maker in a game. A strategy is a specification of a decision for each possible situation in which a player may find himself or herself. A

The name “game theory” may be an unfortunate one. A more descriptive name would be “strategic interaction decision making.” Game theory sounds like child’s play, and it’s not.

payoff is the reward or loss a player experiences when all the players follow their respective strategies.

In the next lecture, we'll develop these ideas more thoroughly. We will also lay the groundwork for the rest of the lecture series. ■

Suggested Reading

Cramton, “Spectrum Auctions.”

McCain, *Game Theory*.

Questions to Consider

1. Consider the button-pushing exercise from the lecture. What conditions would have made you change your mind about whether to push the button? Here are some possibilities: knowing the other players, talking to others before making your decision, making your decision publicly, increasing or decreasing the payoffs by a factor of 100, playing the game again after the first-round results are declared, and not being given the \$100 to start with.
2. In the button-pushing game, we looked at which lines of reasoning, if any, were rational. How would you define rationality? Can two rational people make different decisions?

The Nature of the Game

Lecture 2

Our retail game would have been more realistic if I could have charged any amount for my vase, not just \$10 or \$20. Then, every different price I could charge would correspond to a different strategy for me. You, on the other hand, would still only have one choice to make—buy or not—but you'd have to specify which choice to make for every single price that I could offer you. So you'd have an infinite number of strategies, too.

To examine the idea of strategy in game theory more thoroughly, we begin by revisiting the game in which I can price a product at \$10 or \$20, and you choose whether or not to buy it. In this game, we found that I had two strategies, but you had four. My two strategies were to price the product at \$10 and to price it at \$20. You had four strategies because you had two situations in which you might find yourself and two choices in each situation.

Adding complexity to a game can quickly increase the number of possible strategies. Even for simple interactions, a player can have an infinite number of strategies; for example, most sellers can set an infinite number of prices.

Our strategies thus far have been “pure strategies,” which do not involve randomness and tell us what to do in every situation. An example relating to chess illustrates the concept of a strategy as an instruction book and the extent of a pure strategy. In this example, selecting a book from your strategy library is effectively selecting a strategy. Because a “pure strategy” book specifies a course of action for all possible situations, it often includes pages that will never occur in the game or that never could occur. Pages that are never used can still provide deterrents to other strategies for one’s opponent. Strategies that are not pure—that depend on an element of chance—are called “mixed strategies.”

A player’s payoff represents how much he or she likes the outcome of the game. For our purposes, higher payoffs are better than lower payoffs. The

payoffs for a particular player (Player A) reflect what that player cares about, not what another player thinks Player A should care about. Payoffs must reflect the actual preferences of the players, not preferences anyone else ascribes to them. Game theory can represent such ideas as fairness, but only if they are incorporated into the payoffs of players who care about them. In general, the payoffs for different players cannot be directly compared.

For a class of games called “finite games,” all that matters about a player’s payoffs is the order in which he or she ranks them, not the size of the payoffs themselves. If the payoff scale is only a ranking, the payoffs are called “ordinal payoffs.” If the scale measures how much a player prefers one option to another, the payoffs are called “cardinal payoffs.” A finite game is a game in which any player gets a finite number of moves and has only a finite number of choices at each move. Finite games must also have a finite number of players. Finite games require only ordinal payoffs to solve.

In games that aren’t finite, the units used to measure the payoffs may be more complicated. Some players might be risk averse (or risk loving): \$1 million is worth more (or less) to them than a 10% shot at \$10 million. This valuation is a matter of personal preference, not logic. Game theorists often describe payoffs in terms of utility—the general happiness a player gets from a given outcome. Payoffs on other scales can be converted to utility payoffs, as we see in the example of the lottery ticket.

Common knowledge is our assumption that all players know about the game, all players know that all players know about the game, ad infinitum. The idea of common knowledge can seem complicated when we try to formalize it, but in general, we can rely on our intuitive understanding of the concept and have no problems in working with game theory and its applications. ■

A very useful way to imagine a pure strategy is as ... an instruction book.

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

McCain, *Game Theory*.

Questions to Consider

1. In the movie *The Princess Bride*, Vizzini is offered two goblets of wine and must determine which one holds the poison. He begins this “battle of wits” by saying, “Now, a clever man would put the poison into his own goblet because he would know that only a great fool would reach for what he was given. I am not a great fool, so I can clearly not choose the wine in front of you. But you must have known I was not a great fool—you would have counted on it—so I can clearly not choose the wine in front of me.” Relate this reasoning to the idea of common knowledge.
2. The definition of strategy used in game theory specifies how a player will react in every situation, even situations that will not come to pass when the game is actually played. Consider the proposition that what you choose to do often depends on the consequences that would result if you chose to do something else—in other words, your choice is dependent upon events that never occur.

The Real Life Chessboard—Sequential Games

Lecture 3

As the name suggests, sequential games have events unfolding over time. ... Most of the encounters you think of as interaction with other people are sequential games.

A “sequential game” is a game in which events unfold over time. In sequential games, players have at least some information about the earlier actions of other players. Choices made by one player may influence the choices or options of other players later on. Games in which the players make their decisions without any knowledge of what other players choose are called “simultaneous games.”

We consider an example based on airplane manufacturing, a dynamic game of perfect information. The structure and payoffs of the game are as follows: The European company Airbus is deciding whether to enter a market currently monopolized by the American manufacturer Boeing. To enter the market, the cost to Airbus is \$1 billion. Having a monopoly in either the United States or Europe is worth \$900 million for each competitor. Competing with another company in either market earns both competitors \$300 million. The European Economic Community (EEC) and the United States can both pass protective legislation (PL) that excludes the foreign company from competing in their domestic markets. Payoffs to the EEC and the United States are equal to the profits of their domestic companies plus a \$700 million “competition bonus” if both companies compete in the domestic markets. This bonus represents benefits to the markets’ consumers from lower prices. We assume that the EEC decides whether to pass PL, then the United States decides, then Airbus decides whether to build.

Sequential games are usually represented in “extensive form,” also called a “game tree.” A game tree has a node wherever a player makes a decision. A node has one branch for each decision the player can make there. At the end of any sequence of decision nodes, the game tree gives the outcome as a set of payoffs. A powerful tool for solving sequential games is “rollback.” To solve the game, we begin at the last move and work backward to the root

node. Airbus builds if and only if the EEC passes PL—otherwise, it will lose money. Thus, the United States chooses to pass PL to gain a payoff. The EEC is now indifferent—the same outcome occurs regardless of whether or not it passes PL. Either way, the equilibrium outcome gives Boeing a monopoly in both markets.

Some games have a “first-mover advantage”: Some or all players do better if they move sooner. Let’s look at the same game again, this time with Airbus moving before the United States. The United States, which now moves last, has a weakly dominant strategy in not passing PL. Because the United States won’t retaliate, Airbus should build if the EEC passes PL. The EEC, preferring \$1,200 million to nothing, should then pass PL. Airbus builds in this equilibrium situation.

Let’s solve the game once more, this time with Airbus moving first. Because the EEC now moves after Airbus, it has no incentive to pass PL. Once again, Airbus does not build. As the example shows, there is no such thing as a universal first-mover advantage. The existence of an advantage depends on the circumstances of the game.

Our rollback procedure generates a solution that demonstrates a broader concept: “Nash equilibrium.” A collection of strategies, one for each player, is called a “strategy profile.” A strategy profile is a Nash equilibrium if no player benefits by unilaterally changing his or her strategy in the profile. An important refinement of the Nash equilibrium is the “subgame-perfect equilibrium,” which requires the strategy profile to produce a Nash equilibrium in each of a game’s subgames. Subgame perfection guards against equilibria in which players make silly or irrational decisions off the equilibrium path. Equilibria that rely on hollow threats are not subgame-perfect.

This tree rollback approach can be used in a lot of different games. For it to work, they have to be finite and deterministic ... noncooperative and sequential.

Rollback equilibria are always subgame-perfect. Rollback always shows the best way to play any finite, deterministic game with no chance and no hidden information. This result is called “Kuhn’s theorem.” For example, according to Kuhn’s theorem, a perfectly played game of chess is a guaranteed win for white, a guaranteed win for black, or a guaranteed draw. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Kretschmer, “Game Theory.”

Questions to Consider

1. In a haggling situation, is there usually a first-mover advantage, a second-mover advantage, or neither?
2. A small time crook is being interrogated by a policeman. The crook committed a petty crime, but everyone knows that gathering sufficient evidence to prove it will be time consuming for the police. The policeman tells the crook that if she doesn’t confess, he’s going to spend all of his time making sure that she gets the harshest sentence possible. The crook, in light of the threat, confesses. Verify that this scenario represents a Nash equilibrium, provided that the policeman would actually follow through with his threat. Now verify that the subgame-perfect equilibrium is for the crook not to confess and for the policeman to go back on his threat; that is, the policeman’s threat is actually not credible.

Life's Little Games—The 2×2 Classic Games

Lecture 4

I want to share with you some games that are, in a sense, atomic. You can't break them down into smaller games. Because of their tininess ... they appear again and again at the heart of many other games, and knowing them gives you a place to start in evaluating such games.

Simultaneous games are games with no turn-taking. Many classic simultaneous games are commonly found as building blocks of larger games. In simultaneous games, the players don't have to move at the same time; the only restriction is that no players can know the other players' decisions when they make their own. We'll look at four 2×2 simultaneous games. A 2×2 game is one with only two players and only two moves allowed for each player. The four games we'll look at today are the coordination game, the battle of the sexes, chicken, and the prisoner's dilemma.

Let's examine the coordination game in the context of dressing up for a dinner date. You and Taylor are meeting at your favorite restaurant, L'Amour. Each of you has a choice to dress either casually or formally. Above all, you prefer matching with Taylor. You would prefer that you both dress formally rather than both dressing casually.

We can represent this game (and other simultaneous games) in matrix form. In the matrix, each of your strategies gets a row, and each of Taylor's gets a column.

Despite the fact that both players dressing formally seems like the obvious answer, several things could disrupt this equilibrium. If you and Taylor don't share all levels of common knowledge, one of you could rationalize dressing casually. If you think Taylor is irrational—that Taylor will go against his or her own preferences—you might decide to dress casually. This kind of situation explains why we almost always assume that players are rational.

We switch two payoffs in the coordination game and examine a new game, the battle of the sexes. Taylor's preferences are the same as before. You, however, prefer to dress casually. Above all, you prefer to match with Taylor. You prefer "matching and casual" to "matching and formal." Taylor prefers the reverse. If you could talk to Taylor beforehand, you could agree on what to wear and turn the game into a cooperative game, in which binding agreements are possible. Any line of reasoning that will lead you to choose one style of clothes could be echoed by Taylor to decide on choosing the other.

To solve such games as the battle of the sexes, we need a focal point (usually called a "Schelling point" in game theory), an answer that seems like the obvious choice. An experiment on ABC's *Primetime* showed just how good people are at identifying Schelling points. Six teams of two were dropped off randomly in New York City and were told to meet up with another team

before the end of the day. All six teams ended up at either Times Square or the Empire State Building, and almost every team chose noon as the meeting time.

[The prisoner's dilemma] was discovered ... at the RAND Corporation in 1950 ... where you could see John von Neumann and John Nash walking down the halls.

We modify the story of Taylor and L'Amour to illustrate a game of chicken. You and Taylor have broken up, and each of you is dating someone new. It's Valentine's Day, and you are both going out to dinner with your new partners. Your preferences are symmetric: You'd like to be alone with your new date at L'Amour, and if you can't be there, you'd prefer that your ex isn't either. Your worst

payoff is that both you and Taylor go to L'Amour, which would make for a terrible evening. The "fair" solution to the game—no one goes to L'Amour—is not a Nash equilibrium. As soon as you know that your ex isn't going to your restaurant, you will want to go there, and vice versa. Once again, we need to find or create a Schelling point.

We switch two of your payoffs from chicken to arrive at the prisoner's dilemma. You'd rather spend the evening glaring at Taylor than feel exiled

to another restaurant, and Taylor feels the same way. This new matrix is the much-vaunted prisoner's dilemma, a game as maddening as it is important. Regardless of what Taylor does, going to L'Amour is better for you than not going. You gain an extra unit of payoff either way. Because the game is symmetric, Taylor will reason the same way. Thus, everyone goes to L'Amour and has a worse time than if both couples had stayed away.

The paradox in the prisoner's dilemma is that the cooperative outcome is better for everyone than only the Nash equilibrium. The equilibrium is not "Pareto-optimal" (efficient in an economics sense). A solution is Pareto-optimal if the only way to achieve a better payoff for one player is to give a worse payoff to another player. Achieving cooperation in the prisoner's dilemma proves to be a difficult and relevant problem. ■

Suggested Reading

Poundstone, *Prisoner's Dilemma*.

Rapoport, Guyer, and Gordon, *The 2 × 2 Game*.

Questions to Consider

1. Prisoner's dilemma situations arise frequently in life. It's likely that you have played such games and in some cases, you found a way to cooperate with the other player; in other cases, one of you betrayed the other; and in still others, both of you betrayed each other. Consider examples of each situation and ask yourself: What made the difference?
2. Try this game with a friend or family member to see how Schelling points work. You are identified as the "New York player" and your friend or family member is the "California player." In the game, three states may be "claimed": Florida, Idaho, and Montana. Each of you secretly writes down the names of the state that you claim. If each state is claimed by exactly one person, then you and your partner win. Before revealing your choices, how confident do you feel that the two of you have won? How confident are you that you used the same reasoning in arriving at your answer? Try it again with a "harder" set of states.

Guessing Right—Simultaneous Move Games

Lecture 5

The Allies want to go to Cherbourg; it's better. So the Germans are going to expect an attack there. But the Allies are going to know that the Germans are going to expect an attack there, so they're going to choose Normandy. The Germans can realize that, and they'll—you get the idea. Round and round and round until someone gets tired of second-guessing.

Simultaneous games are useful for modeling situations in which communication can't or won't take place among players. This is often the case in competitive or zero-sum games. Simultaneous games don't have to be truly simultaneous, as long as players can't observe one another's moves. One such class of games includes pursuit-evasion games, those in which one player wants both players to choose the same option and the other wants them to choose different options. Examples of pursuit-evasion games include choosing the landing beaches for D-Day and the Battle of the Bismarck Sea.

Let's analyze the Battle of the Bismarck Sea to explore the concept of dominance. The Japanese are trying to ship a convoy of troops from one island to another three days away, while the Allies are trying to bomb the Japanese fleet. The Japanese have two possible routes, north or south. Because weather in the north is expected to be stormy, the Allies will need a day to find the convoy if it travels north. If the Allied patrol finds the Japanese in the south, it can bomb them the same day. If the Allies are wrong the first day, they still find out which route the Japanese took and can bomb the convoy for two days (if the convoy went south) or one day (if the convoy went north).

This game is an example of a “zero-sum game,” or a game of perfect competition. Whatever advantage one player gains, the other loses. One method for solving such a game is von Neumann's minimax approach, in which both players try to minimize potential losses in their worst-case scenarios. If Japan sails north, the convoy can't be bombed for more than

two days. If the Allies search north, the convoy won't be bombed for less than two days. Thus, the convoy will be bombed for two days. Sailing north "weakly dominates" sailing south. This means that the Japanese always do at least as well sailing north as they do sailing south.

Eliminating dominated strategies can prove useful in simplifying larger simultaneous games. We can reduce a 4×4 game by eliminating dominated strategies. To understand how this works, we will consider a scenario in which two vendors, Hamlet and McGuffin, both seek to maximize their profits in the breakfast-sandwich business. Charging \$2 is strongly dominated by charging \$3 for McGuffin. Moreover, charging \$3 is strongly dominated by charging \$4. No matter what price Hamlet sets, charging \$3 always earns McGuffin strictly more profit than charging \$2, and charging \$4 earns more profit than charging \$3. If McGuffin is rational, he will never play these dominated strategies.

Charging \$4 also dominates charging \$5. Charging \$4 is a dominant strategy because it dominates all of McGuffin's other strategies. A rational player will always play a dominant strategy.

Because Hamlet knows that McGuffin will charge \$4, he simply chooses his best response to the \$4 strategy—in this case, \$3. Hamlet's decision is an example of "iterated elimination of dominated strategies" (IEDS). He eliminated dominated strategies for one player, then eliminated them for himself, whereupon he had only one option left. Any cell in the matrix eliminated by IEDS is guaranteed not to be a Nash equilibrium. If IEDS leaves only one cell, that cell is a Nash equilibrium.

IEDS can even solve some games that don't seem to have dominant strategies. Such is the case in a game that involves setting transportation prices. In the problem of setting prices for tour-boat rides and cable-car rides, no row or column in the matrix consistently outperforms the other three; thus, neither player has a dominant strategy. IEDS can reduce the game by eliminating dominated strategies, even if no dominant strategies exist. For the cable-car company, \$7 dominates \$6 and \$8 dominates \$9. If the cable-car company will always charge \$7 or \$8, charging \$7 is a dominant strategy for the boat

company. The cable-car company can then choose its best response to \$7, which is to charge \$8. Both companies make \$66 profit.

Unfortunately, IEDS cannot solve all simultaneous games, even if it can reduce them. To illustrate, we reduce the size of the customer base in the previous example.

With four cells left, there are no more strongly dominated strategies. One approach is to eliminate weakly dominated strategies, as well. The price of \$7 for the boat ride weakly dominates \$6, and \$8 for the cable-car ride weakly dominates \$7. However, this process can eliminate other viable Nash equilibria, even if they're weakly dominated.

Another tactic, called the “best-response method,” can find all the Nash equilibria in a simultaneous game. Mark the best payoff for the row player in each column. The highlighted cells represent the best responses to each of the column player’s strategies. Then mark the best payoff for the column player in each row. These cells represent the best responses for the column player to each of the row player’s strategies. Any cell with both payoffs marked is a Nash equilibrium. There are no other Nash equilibria in the game.

The best-response method shows the modified transportation game to have three Nash equilibria. The only admissible equilibrium—the only one that uses no weakly dominated strategies—is (60, 60); this strategy is the one we’d expect to see played. Even though the equilibrium is not Pareto-optimal, neither player can change strategies to earn more. The equilibrium (70, 70) is not a Nash equilibrium because either player could unilaterally change strategies and make \$72. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Mehlmann, *The Game’s Afoot! Game Theory in Myth and Paradox*.

Questions to Consider

1. You, along with three strangers, see a burglar smash a store window and take a piece of jewelry. All of you could identify the perpetrator. As you leave the scene, you decide whether to call the police. Assume as common knowledge that the benefit from the criminal being apprehended is greater than the cost of becoming involved for each witness. What are the Nash equilibria of the game? What difficulties exist in reaching an equilibrium? What would you do if this actually happened?
2. You and a friend are selected as contestants on a game show. You each must pick a whole number from 1 to 7. You will pick an odd number, and your friend will pick an even number. If the numbers you pick are consecutive, such as 5 and 6, then you each win \$1,000. If both you and your friend understand weakly dominated strategies, you're guaranteed to win the money. How? If you and your friend both understand IEDS for weakly dominate strategies, you can win the money every time, even if the numbers are selected from 1 to 100, with you choosing odd and your friend choosing even. How?

Practical Applications of Game Theory

Lecture 6

We're going to see why cigarette companies were happy when they were banned from the TV airwaves. We'll look at a situation where it could be in your best interest to let someone blackmail you. We'll see how a bid of \$98 can beat a bid of \$102 for a stock. We'll see how insisting that you lose ties could be the one thing that lets you win.

The movie *A Beautiful Mind* shows John Nash thunderstruck by his epiphany on how to resolve a dating problem in noncooperative game theory. Unfortunately, his solution is not a Nash equilibrium. Game theory can explain why tiered bids in corporate buyouts prove to be such a powerful tactic. For example, in 1988, Macy's wanted to buy out Federated Stores (which included the popular chain Bloomingdale's), whose stock was selling at \$100 per share. Macy's offered \$102 per share to the company's shareholders, conditional on Macy's obtaining a majority of the stock. Players had a weakly dominant strategy of selling to Macy's—they would get \$102 per share if Macy's won its bid but still keep \$100 if the takeover bid failed.

Robert Campeau, also wanting to buy out Federated Stores, offered a two-tiered bid of \$105 and \$90. Campeau would buy up to 50% of the shares in Federated Stores for \$105 per share, adding his price to a "pot" that would eventually be divided evenly among all selling stockholders. After the 50% mark, he would add only \$90 per share to the pot. If Campeau won the takeover bid, he could take the company private and appropriate the remaining shares for fair-market value, or \$90 in this case.

Selling to Campeau strongly dominated selling to Macy's. If Macy's obtained at least 50% of the shares, Campeau's offer of \$105 was better than Macy's of \$102. If Campeau obtained at least 50% of the shares, his offer of \$97.50 (the price offered after the division of the two-tiered bids) was at least better than the \$90 that would be paid if the company went private. If no one obtained a majority, selling to Campeau for \$105 would be better than keeping the \$100 stock.

Selling to Campeau was dominant despite its inefficiency. Shareholders could have sold to Macy's for \$102 but instead sold to Campeau for \$97.50. Campeau created a prisoner's dilemma for shareholders with his two-tiered bid approach.

Voting games often have surprising and counterintuitive results when more than two choices are available. One example of such a game involves a three-person committee, composed of A, B, and C, deciding on who will chair the committee for the next year. In this game, each member would like to chair the committee. Failing that, A prefers B to C, B prefers C to A, and C prefers A to B. As the current chair, A decides how the vote will take place, but A graciously decides that B will win ties and C will be allowed to vote first.

We roll back the game tree as usual to find the equilibrium. A, who votes last, has a weakly dominant strategy in voting for himself. (For A, voting for himself is called "sincere voting" because he is voting for the candidate he likes best.) Voting sincerely is also weakly dominant for B. The only way C can avoid B becoming the chair is to cast a second vote for A, even though she would prefer to vote for herself. This is an example of "strategic voting." B could try to avoid reelecting A by promising to vote for C, but he has no reason to follow through after C has voted. Paradoxically, A would not have been elected if he had won ties.

**Success in one
game doesn't
always translate to
success in another.**

We can also use game theory to explain why, when cigarette companies were banned from television advertising in 1971, the industry received an unintentional boon. A simplification reveals cigarette advertising to be a prisoner's dilemma. Suppose there are two competing cigarette companies in the market, both of which can choose to spend \$500 million on advertising. A company increases the size of its market by 5% if it advertises and captures 80% of the market if the other company does not advertise. If both companies advertise or if neither does, they split the market evenly. For both companies, advertising is a dominant strategy, even though each company makes \$1.15 billion with advertising instead of the \$1.5 billion each could make if neither advertised. The government removed the "defection" options from the game

by forbidding advertising, leaving the cigarette companies with their Pareto-optimal payoff.

Game theory can also elucidate the reasons for a surprising fact: One way to make a promise credible is to allow another player to retaliate if you break the promise—effectively, to encourage blackmail. To see how this works, imagine a scenario in which three candidates, Dennis, Rebecca, and Indira, are running in a mayoral election. Polls show Dennis with 40% of the vote, Rebecca with 45%, and Indira with 15%. Indira likes Rebecca a little more than Dennis, but she cares much more about having influence over zoning—a position Dennis could offer her if he wins with her support. Because Dennis has no reason to keep his promise after Indira’s move, he needs to find a way to make his promise credible. Dennis could give Indira information that could cause a scandal if she revealed it while he was in office. Rolling back the tree shows that Indira will reveal the scandal only if Dennis refuses her the planning position. In equilibrium, Dennis wins and Indira gets the position. Allowing himself to be blackmailed actually won Dennis the election! ■

Suggested Reading

Aliprantis and Chakrabarti, *Games and Decision Making*.

McCain, *Game Theory*.

Questions to Consider

1. In the chairmanship vote, who would have won the chairmanship if A won ties?
2. Consider the button game from Lecture 1: Each of 101 players has a button. Pressing the button has two effects. First, it costs all other players \$2. Second, if other players have pressed their buttons, it cuts the losses you take from them in half. What is the Pareto-optimal solution in this game? What are the two Nash equilibria of the game? Suppose that before the game is played, the players met in a room and all agreed not to push. Pushes are still anonymous. Would you push? Would you expect others to? Why or why not?

A Random Walk—Dealing with Chance Events

Lecture 7

There are really three main ways that unpredictability or chance works its way into game theory: as uncertainty about the outcome of an event within the game, as an uncertainty about the structure of the game itself, or as an uncertainty about the pure strategy that a player will choose.

Chance events can affect games in a number of ways. First, they may occur within a game. We can think of chance events as “decisions” made by nature and add “Nature” as a player in the game. Nature makes each move randomly with a certain probability of each option. Also, the structure of a game may be uncertain, resulting in a game of incomplete information. The United States is playing such a game against North Korea, because we do not know their nuclear capabilities (i.e., their strategies). We also do not know the payoffs for those strategies. John Harsanyi discovered a way to turn games of incomplete information into games of imperfect information, in which players know the structure of the game but not necessarily where they are within that structure. We will look at Harsanyi’s approach, for which he won a Nobel Prize, in a later lecture.

Also, chance can play a role in a player’s choice of which pure strategy to play. All two-player games have at least one Nash equilibrium if mixed strategies are allowed. Mixed strategies are those in which the pure strategy used is randomly selected from the available strategies, with each pure strategy having its own probability of selection. With mixed strategies, players need not play each pure strategy with equal probability. Any game that allows mixed strategies cannot be finite. This is because a player can create an infinite number of different strategies by varying the probability with which he or she plays each pure strategy. Because the game is infinite, cardinal payoffs are needed to solve it, not just ordinal payoffs. Probability plays a crucial role in computing payoffs from mixed strategies.

One way to understand probability and how counterintuitive it can be is to find the answer to the following question: In a group containing you and 40

of your friends, what is the probability that two people will share the same birth date? Despite the common guess of 41/365 (a 1 in 9 chance, or 11% probability), the answer is actually more than 90%.

To find the probability of two or more independent events, multiply the two probabilities together. According to a recent article in a technology magazine, 82% of Americans own cell phones. Suppose that you and I each choose an American at random. How likely is it that the person you pick has a cell phone and the person I pick does not? The odds that your person has a cell phone and my person does not are $0.82 \times 0.18 = 0.1476 = 14.76\%$.

Finding the probability of a combination of dependent events (e.g., that two people in a group share a birthday) is slightly different. First, find the probability of the first event. Then, find the probability of the second event given that the first event happened, and so on. Finally, multiply these results together. In other words, the probability for a combination of dependent events is the product of the probability of each individual event.

Returning to the original group with you and 40 of your friends, imagine that you line up and mark your birthdays off on a calendar, one after another. No two people have the same birth date if and only if each person's birthday is unmarked when his or her turn arrives. The probability of a birthday not being marked is the number of open days on the calendar divided by the total number of days in a year. This quotient is 365/365 for the first person, 364/365 for the second, and so on down to 325/365. Multiplying the probability for each individual event, we get $1 \times 364/365 \times 363/365 \times 362/365 \times \dots \times 325/365 \approx 0.097$, or only about a 9.7% chance of no two people having the same birthday.

Let's now turn to the Monty Hall Paradox and discuss its payoffs in terms of expected value. Suppose you are a contestant on the old TV show *Let's Make a Deal*®, hosted by Monty Hall. Monty shows you three doors. Behind two of the doors is trash; behind one of them is a new car. You choose a door, and Monty then opens one of the other doors, revealing trash (he can always do this). You are then given a chance to switch your choice to the other door. Contrary to intuition, switching doors alters your chances of winning,

increasing them from one-third to two-thirds. If you initially guess right (that is, choose the door concealing the car) but then switch, you lose. This happens one time in three. If you initially guess wrong (that is, choose a door concealing trash), Monty must open the other trash door. By switching, you are guaranteed to pick the winning door. This happens two times in three.

We calculate your payoff in the Monty Hall game using expected value. To find expected value, multiply each payoff by its chance of occurring and sum the results. “Expected” here essentially means “average.” You have a two-thirds chance of winning a \$30,000 car and a one-third chance of receiving a payoff of \$0. Your expected value is $2/3 (\$30,000) + 1/3 (\$0) = \$60,000/3 = \$20,000$.

Finally, we look for the optimal responses to some mixed strategies in the game between even Stephen and odd Maude and compute expected values for each player. In this game, each player can “shoot” one, two, or three fingers. The expected payoffs for each player for a variety of mixed-strategy profiles are as follows: When Maude plays a mixed strategy of shooting one 30% of the time, two 60% of the time, and three 10% of the time (0.3, 0.6, 0.1), Stephen’s best payoff comes from shooting two as a pure strategy. On the other hand, when Stephen plays a mixed strategy of shooting one 40% of the time, two 50% of the time, and three 10% of the time (0.4, 0.5, 0.1), Maude’s best payoff comes from shooting one or three in any proportion and not shooting two at all.

If you know the mixed strategy being played by the other player, you should examine each of your pure strategies in turn. If you find only one best expected payoff, you should play that option as a pure strategy, but if two or more strategies are tied, the best response is to play any of the tied strategies as pure strategies or to play any mix among those tied strategies.

To find a mixed-strategy equilibrium, each player’s mix of strategies must be a best response against the other’s. In our next lecture, we will discuss how von Neumann showed that players can find mixed-strategy equilibria in their games. ■

Suggested Reading

Berensen, Krehbiel, and Levine, *Basic Business Statistics*.

Grinstead and Snell, *Introduction to Probability*.

Questions to Consider

1. Probability can be quite a bit trickier than you might think. Suppose you have two audio cassettes in your car. One has country music on both sides; the other has country on one side and rock on the other. You pick a tape at random, pick a random side, and play it. It's country. How likely is it that the other side is country, too? Most people say the probability is one-half, but in fact, the other side will be country two-thirds of the time. Try the experiment several times to convince yourself!
2. One of the common mistakes with probability involves conditional probabilities—how likely one event is given another. Knowing that a drug test is 95% reliable means that a user is 95% likely to be identified as a user, and a nonuser is 95% likely to be identified as a nonuser. When someone tests positive, how likely is it that he or she is a user? Almost everyone says 99%, but the actual answer depends on the fraction of the population using drugs. If only 1% of the population uses drugs, then a person testing positive will be innocent more than five times out of six!

Pure Competition—Constant-Sum Games

Lecture 8

The security values for the two sides weren't equal, and because of this, it triggered the infinite second-guessing game that we've seen so often in these lectures. That's going to happen any time the two sides have different security values, and this is where the minimax theorem comes in.

According to von Neumann's revolutionary minimax theorem, two-player zero-sum games have sensible "optimal" strategies. That is, there's a sensible way to define a best-strategy profile, as well as a sensible way to find it. We revisit the minimax strategies we found in discussing the Battle of the Bismarck Sea in Lecture 5. The Japanese could sail north and hold the Allies to two days of bombing, while the Allies could search north and guarantee at least two days of bombing. When these minimax "security values" of the game are identical, neither player can do better. To illustrate this, consider that, in the Battle of the Bismarck Sea scenario, if the convoy is bombed for more than two days, the Japanese can do better by reverting to their minimax strategy. However, if the convoy is bombed for less than two days, the Allies can do better by playing their minimax strategy. If the other side is playing its minimax, deviating from your minimax won't help you and might hurt you.

On the other hand, suppose the Allies can't go south again if they initially search north. This new game has no equilibria in pure strategies and degenerates into second-guessing. This is because the Allies' minimax strategy is now to sail south, which ensures them one day of bombing. If both sides play their minimax strategies, the result will not be a Nash equilibrium. The Allies can unilaterally change their strategy and get an extra day of bombing. However, then the Japanese will then want to change *their* strategy, and the process will continue in an infinite game of second-guessing. As in this example and others in which two players have unequal security values, pure minimax strategies do not guarantee an equilibrium; they guarantee an equilibrium only when the two players' security values are equal.

The minimax theorem says that if players are allowed to use mixed strategies, then strategies with the same security values can always be found. An examination of penalty kicks in soccer illustrates the concept of

mixed-strategy equilibrium. We assume the goalie is diving left to block kicks with probability q and calculate the expected payoff for the kicker.

“There could be no theory of games without that theorem. I thought there was nothing worth publishing until the minimax theorem was proved.”

—John von Neumann

hence, the one we would expect the goalie to play. Similarly, we assume the kicker kicks left with probability p and calculate the expected payoff for the goalie. The bottom of the goalie’s payoff valley is at the point $p = 0.39$; thus, the kicker should kick left 39% of the time and right 61% of the time.

We graph the kicker’s and goalie’s best responses for all values of p and q . The goalie is going to be happy only with a strategy profile that corresponds to a point on the line representing the kicker’s worst expected payoff. If it doesn’t, he’ll unilaterally change his play. The only point at which no one benefits from unilaterally changing play—the only Nash equilibrium—is where the lines cross. At any equilibrium point, each player will be indifferent to all the options that he or she is playing.

Although real soccer players of course do not make these calculations or generate random numbers to implement their mixed strategy, they play as though they *were* using game theory for strategy selection. They choose strategies that correspond to within 1% of what game theory identifies as optimal.

We revisit a variant of even Stephen and odd Maude and find an easy, nearly magical way of computing optimal strategies for 2×2 games. In this variant, odd Maude and even Stephen each “shoot” either one finger or two fingers. If the resulting sum is odd, Maude wins \$1 per finger from Stephen. If the resulting sum is even, Stephen wins \$1 per finger from Maude.

We take the positive difference of Stephen’s two payoffs in the first column and write the answer over the second column. We do the same with the two payoffs in the second column, writing the answer over the first. Then we divide each number by the total of the payoff differences; those probabilities correspond to Maude’s optimal strategy. Maude shoots one $7/12$ of the time and two $5/12$ of the time. Repeating the tactic for Stephen, we find that his strategy should also be to shoot one $7/12$ of the time and two $5/12$ of the time. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Palacios-Huerta, “Professionals Play Minimax.”

Question to Consider

1. Finding the mixed-equilibrium strategies for zero-sum games can yield results that seem perplexing at first. In the children’s game Rock-Paper-Scissors, the optimal mixed strategy is for each player to randomly choose each option one-third of the time. But suppose that when rock beats scissors, the winning player scores 2 points, not 1. How would you expect optimal play to change? Interestingly, the players play rock less, not more! Paper is played one-half of the time, and rock and scissors both drop to one-fourth. Why? The increased rewards from rock lead the other player to defend against it by avoiding scissors and playing paper. This additional defense makes the choice of rock less attractive. Verify that the strategies specified result in a mixed-strategy equilibrium for the modified Rock-Paper-Scissors game. (Each row and each column will have an expected payoff of 0.)

Mixed Strategies and Nonzero-Sum Games

Lecture 9

Humans are very good at watching other humans, and some people are very good at anticipating the trains of thoughts of other people's minds. If you're convinced that you can read your opponent with a fair degree of accuracy, then you're essentially playing an asymmetric game.

We begin by revisiting the ideas presented in Lecture 8 from a less mathematical, more intuitive perspective. We will start by reexamining the game of even Stephen and odd Maude to find out what it means to say that a mixed strategy is optimal. In doing so, we discover that the optimal strategy for Maude is $(0.25, 0.5, 0.25)$. That is, she should choose one 25% of the time, two 50% of the time, and three 25% of the time. Surprisingly, these are also the optimal choices for Stephen.

To say that a strategy is optimal doesn't mean that another strategy can't work better against a foolish opponent. One player could predict his or her opponent's future behavior based on past patterns. Then he or she could develop a strategy based on these patterns. For example, one player could decide, based on a foolish series of past choices by an opponent, that future choices by that opponent would be similar. However, for this strategy to outperform the ones we have decided are optimal, the player would need to be able to rely on his or her opponent consistently being foolish in the same way. He or she also might develop an ineffective strategy because of basing it on an opponent who had adopted a deceptively foolish pattern of behavior to be abandoned once the first player had accepted his or her pattern as typical—and developed a strategy based on it. There is not a reliable way to outperform an opponent in cleverness without more information than we have available. It is in this sense that the strategies that we have discussed to this point in the course are optimal. These strategies make one player bulletproof relative to his or her opponent's cleverness.

In relation to our example, saying that $(0.25, 0.5, 0.25)$ is Stephen's best strategy means that he is bulletproof if he plays it. His expected payoff from playing this strategy is 0, even if Maude knows that he's doing it. If Stephen

were to play a strategy other than this equilibrium choice, Maude could take advantage of the situation by playing one of her pure strategies, giving Stephen a lower average payoff. If Maude plays her equilibrium strategy, she is bulletproof in the same sense. She will lose \$0 on average, regardless of what Stephen does.

According to the von Neumann minimax theorem, these bulletproof strategies exist for all two-player zero-sum games. The common security value won't necessarily be 0, but if the security value for one player is gaining \$1, the security value for the other player will be losing \$1.

Let's now turn to a game that involves camping. In this game, two families share a campsite, and each has the option to visit it zero, one, or two times in a month. Going zero times is a dominated strategy for both families. The best-response method gives two Nash equilibria in pure strategies. Each family camping once and each family camping twice are both equilibria.

The camping game is actually a version of another classic 2×2 game, the Stag Hunt. The best payoff for each player is hunting stag together, followed by hunting rabbit alone, hunting rabbit together, and hunting stag alone. Stag Hunts frequently appear in other guises, including the arms race between the United States and the Soviet Union and collaboration on team projects at work.

Intuitively, we may think that both players should cooperate to receive the best payoff, but this isn't always what happens in a Stag Hunt. The cooperative equilibrium is better for both players than the other equilibrium. In game theory, we say that this strategy is "payoff dominant." However, the noncooperative equilibrium is preferable if one player thinks it's likely that the other player won't choose the cooperative option. Ironically, the reason one player might not choose the cooperative equilibrium would be that he or she considers it likely that the other player won't cooperate.

**One side tells
the other the
obvious truth:
that they both do
better by "doing
the right thing."**

In the Stag Hunt and other games that share its structure, communication is critically important. Through communication, players can arrive at a nonbinding commitment to cooperate, which can establish a Schelling point. After the Schelling point is established, both players would lose by unilaterally changing their strategies.

What is the mixed-strategy equilibrium for the Stag Hunt? Amazingly, the same trick of taking the differences of payoffs that we learned in the last lecture works for nonzero-sum games. In the mixed-strategy equilibrium, both families camp twice 80% of the time and camp once 20% of the time. The Smiths get a payoff of 44 and the Joneses get 88.

The idea of mixed-strategy equilibria makes much less sense in nonzero-sum games. In zero-sum games, if one player changes away from his or her equilibrium strategy, the other player could do better and the first player might do worse. In nonzero-sum games, both players might be able to do better.

If one family goes camping once with probability greater than 20%, the two families eventually arrive at the payoff-dominant equilibrium. Suppose the Smiths decide to play a 50/50 mixed strategy instead of 20/80. Then, the Joneses get more payoffs by going once 100% of the time. But if the Joneses always make one trip, then the Smiths will also always go once. This mixed-strategy equilibrium, like many others, is inherently unstable. If either player alters his or her strategy, one of the two pure-strategy equilibria will result. Mixed-strategy equilibria in nonzero-sum games often give inefficient payoffs, especially in games that already have pure-strategy equilibria. ■

Suggested Reading

Davis, *Game Theory*.

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Questions to Consider

1. One interesting property of mixed-strategy equilibria for 2×2 nonzero-sum games is the following: The probabilities with which a player plays each option depend solely on the payoffs to the other player, not the player's own payoffs. As an example, suppose that a motorist must decide whether to put money in the parking meter, while the parking lot owner must decide whether to patrol the lot for parking violators. An increase in the cost of patrolling the lot will not change the fraction of the time the lot owner will patrol. It will, however, increase the fraction of the time that the motorist will park illegally.
2. Consider the game of two people approaching one another on a sidewalk. Each chooses right or left. If they make the same choice, they pass one another without a problem and each gets a payoff of 1. If they make opposite choices, they both get payoffs of 0. Find the three Nash equilibria of the game. (One of them is a mixed equilibrium.) Show that the payoff from the mixed equilibrium is only half as good for either player as either of the two pure equilibria.

Threats, Promises, and Commitments

Lecture 10

Strategic moves aren't acts of desperation; they're moves that are added to an existing game. When they're used properly, they can change the outcomes of those games dramatically.

Strategic moves come in three forms: threats, promises, and commitments. They aren't acts of desperation—they're moves that are added to an existing game and can change the outcome of a game dramatically. Because strategic moves are added on to the beginning of a game, they're often useful in games with a first-mover advantage.

A “commitment” is an unconditional statement that a player will make a certain decision. Commitments effectively allow a player to make a move now instead of at its usual position in the game tree. In previous lectures, we've seen players take advantage of commitments. In the battle of the sexes game, you made a commitment when you left a message on Taylor's answering machine that said you were dressing casually for the date.

In game theory, “promises” are the equivalent of saying, “If you make this choice, I will respond with a choice that you'll like—something that you wouldn't normally expect me to do.” Unlike commitments, promises are conditional: They are triggered only if a particular choice is made.

On the other hand, a “threat” in game theory is the inverse of a promise. Where a promise amounts to saying, “Do what I want and I'll make things better for you than you would otherwise expect,” a threat is the equivalent of saying, “Do what I want, or I'll make things worse for you than you would otherwise expect.”

Credibility is a critical issue when judging the effectiveness of a strategic move. Credibility problems are common in using strategic moves because such moves require a player to do something he or she wouldn't normally do. For example, if you promise your child that you'll go to Disneyland if she gets an A in math, she has no incentive to work if she knows the family

is planning to go to Disneyland anyway. Threats and promises mean that, under certain circumstances, you plan to do something you do not want to do when you reach that situation.

A promise you want to follow through on is called an “assurance,” not a promise. Similarly, a threat you want to follow through on is a “warning.” Assurances and warnings don’t change how people play games. Threats, promises, and commitments can, assuming they’re believed.

Promises can help players avoid situations that aren’t Pareto-optimal, as we see in a simple game involving you and Taylor going to the movies. You and Taylor each have a movie you want to see with the other person. Your movie is playing this week and Taylor’s is playing next week. The only rollback equilibrium is for both of you to stay home. You can’t do any better with a commitment to go to Taylor’s movie or with a threat not to go. You can, however, promise to go to Taylor’s movie if you both go to yours.

It is important to remember three important points about promises. First, your promise allows both you and the other player to get a better payoff than you otherwise would have. Second, your promise makes you do something that you didn’t want to do. After seeing your movie, your best payoff comes from going back on the promise and skipping Taylor’s movie in the second week. And finally, none of this works unless your promise is heard and believed. Strategic moves must be observable and credible in order to work.

Let’s look at a hostage negotiation with terrorists in terms of credible strategic moves. We will assume that the terrorists’ priorities are the following: not submitting to the will of the United States, having their demands met, and not killing the hostages. For the United States, the priorities are having the hostages freed and not ceding to terrorist demands. The rollback equilibrium is for the terrorists to take hostages, the United States to meet their demands,

A threat, promise, or commitment always involves ... saying ... in at least one circumstance, you're going to take an action that is *not* in your best interest at the time.

and the terrorists to release the hostages. One possible strategic move for the United States is to make a commitment not to negotiate with terrorists under any circumstances. Other players must know about a strategic move for it to be effective.

As a player, you have two primary approaches for achieving credibility: altering your payoffs and restricting your strategies. You could change your payoffs so that by the time you have to decide whether to keep your commitment, you want to do so. This approach almost always entails reducing one of your payoffs. One obvious way to do this is by making a contract, preferably with a third-party enforcement mechanism. Without a third party, you can't effectively enforce threats, because a threatened player won't hold you to a contract that would harm his or her interests.

In the Ultimatum Game that we played in Lecture 2, I could give a third-party observer \$5 and tell that observer that she can keep the money if I accept less than \$5 from you. If you offer less than \$5, my best payoff comes from rejecting your offer.

The custom of giving engagement rings has served since the 1930s (albeit to a lesser extent in recent years, as the social sanctions against premarital sex have weakened) to make the promise of an upcoming wedding a credible commitment. Men and women have always engaged in premarital sex, but earlier in our history, sleeping with a man outside of marriage could ruin a woman's reputation and endanger her prospects for marriage if her fiancé's commitment was a sham. An expensive engagement ring did, and does, make it too costly for a man to simply break his commitment and leave his fiancée. In game theory, the ring is a "signal" that establishes a separating equilibrium in the marriage game. Men who are serious about the engagement buy the ring; men who are faking won't buy one. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Hild and Laster, "Reinhard Selten."

Questions to Consider

1. Consider the sequential market-entry game in which entry (development) costs for either firm are \$50 million and the returns are \$100 million for a lone entrant and \$40 million each if both firms enter. Firm B, which moves second, says that it is going to enter the market regardless of what Firm A does. Why is this an incredible threat? Show that if Firm B spends more than \$10 million on development costs before Firm A makes a decision, Firm B's threat is now credible.
2. An effective threat is cheaper than an effective promise. Why?

Credibility, Deterrence, and Compellence

Lecture 11

Boeing's announcement was a strategic commitment. The message was: "We're building this plane, Airbus, whether you do or not. If you do, too, we're all going down the tubes, so stay out." If the commitment was taken seriously, then Airbus should stay out of the market. But was Boeing telling the truth?

Credibility often determines the success or failure of a strategic move. For example, in 1993 Boeing signaled to Airbus that it was committed to building a new generation of planes, the "superjumbos." At the time the market had enough demand for only one superjumbo manufacturer. Four years later, Boeing rescinded its commitment, but its strategic—and credible—move prevented Airbus from developing the superjumbos and competing with Boeing's midsized planes. In another example from the world of the airlines, in 1992 Continental Airlines proposed to raise fares on its flights by 5%, and then within the week rescinded the proposal. Essentially, Continental made a promise to its competitors, stating, "We'll raise our fares if you raise yours, and we'll both be better off." When not enough of Continental's competitors followed suit, it dropped the plan.

There are other ways to gain credibility by changing payoffs. Most of them trade under an idea we've seen earlier—an iterated or repeated game. For example, in the "Hamlet approach" (often called "rational irrationality"), one person convinces another that he or she doesn't play rationally or that he or she doesn't value the payoff that the first person ascribes to him or her. This technique can lend credibility to otherwise incredible threats or commitments.

A "reputation" for adhering to commitments increases the credibility of a strategic move and serves as an incentive in itself to keep commitments. Israel's reputation of not negotiating with terrorists discourages terrorism. If Israel were to give in to terrorist demands, it would realize a substantial negative payoff for its ruined reputation. Dale Earnhardt's reputation for

cooperation allowed him to arrange “drafting agreements” with other drivers in NASCAR races.

Multiple techniques may be combined to establish credibility. For example, a government could agree to terrorist demands but double-cross the terrorist organization at the last minute. A government that follows this course of action will have ruined its reputation for dealing with terrorists, which may discourage terrorists from attacking in the future. This move represents rational irrationality. Extending the reputation factor to games with more than two players is usually called “teamwork.” The team or community serves to change the payoffs of individual members via promises of acceptance or threats of censure.

Cutting the game into small slices can chip away at another player’s position bit by bit. This approach is often called “salami tactics.” In the Cold War, the domino theory of the spread of Communism was one example of salami tactics. Salami tactics work well when the only response the other player has is a disproportionate one—a “big gun.” One way to fight back against salami tactics is to mix strategies, slowly increasing the probability that the big-gun response will be used. President Kennedy used this “brinkmanship” approach in the Cuban Missile Crisis.

When dealing with strategic moves, flexibility is a player’s greatest enemy. The easiest way to convince other players that your strategic move is credible is to arrange the conditions of the game so that when a certain situation arises, you must follow through with prescribed actions. We’ve already encountered an example of such aggressive constraints: the doomsday device. Doomsday devices do occur in real life. For example, in response to a hostile takeover bid in 2003, PeopleSoft added a clause to its contracts that promised refunds to customers of two to five times the licensing fees paid if the company was downsized. Oracle eventually acquired PeopleSoft (for a much larger bid than it originally intended) because the doomsday device wasn’t a sufficient deterrent. In this case, the doomsday device was also revocable, which defeated its purpose. ABC’s *Primetime* also created a doomsday device to encourage volunteers to lose weight. If the contestants didn’t lose 15 pounds in two months, ABC would air pictures of them wearing skimpy bathing suits on the giant screen of their local sports stadium.

You can tie your own hands and, thus, increase your credibility with a number of less ominous approaches, including appointing an agent, burning your bridges, and simply getting the last word. An agent may gain a strategic advantage for a player because agents are not usually empowered to radically change the deal. It's also possible to commit—though not to threaten or promise—by leaving yourself only one choice. This approach is often called burning your bridges and entails a complete loss of control. The most extreme example of getting the last word is dying. Courts place a good deal of authority in wills because they are quite literally the last word on what the deceased desired.

In general, goals of strategic moves fall into one of two categories: “deterrence” (maintaining the status quo) and “compellence” (changing the status quo). Threats are usually better for deterrence, and promises are generally better for compellence. Using a promise for deterrence forces you to keep paying over an extended period of time, but threats are essentially free. In the Cold War, the United States threatened the Soviets instead of making promises and stationed a moderate level of troops in Europe to forestall salami tactics. Though the troops could never stop a full-scale Soviet attack, if they were attacked, their deaths would create an outcry in America and virtually ensure nuclear war. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Fursenko and Naftali, *One Hell of a Gamble*.

Questions to Consider

1. How can increasing the magnitude of a threat or promise actually reduce its credibility?
2. You want your employee to reduce his or her rate of absenteeism. Would you use a threat or promise? What are the advantages and disadvantages of each?

Incomplete and Imperfect Information

Lecture 12

Maybe your game includes random events, or maybe you're not exactly sure what game it is that you're playing. This is an unfortunate but real possibility. Games have players, strategies, and payoffs, and we generally assume that these things are common knowledge. But what if they aren't?

So far in this course, we've focused on uncertainty as it applies to mixed strategies, but it can affect games in other important ways. Uncertainty figures into games of "incomplete information," those in which some players don't know the structure of the game. They may be uncertain about the possible strategies or payoffs of the other players.

Uncertainty also plays a central role in games of "imperfect information," or those in which some players are unsure about the history of the game. A special case of imperfect information is "asymmetric information," which is where one player knows something another player doesn't. Games of symmetric information—in which all players have the same knowledge—can usually be solved with the expected-value approach.

The Harsanyi transformation can convert a game of incomplete information into a game of asymmetric information. A Harsanyi transformation creates a spectrum of possible versions (or types) of players in a game. It then assigns each type a probability representing the likelihood of that type's appearance in the game. For example, a solvent IBM is more likely than an insolvent one. Your optimal strategy is the one that has the best expected value against the range of possible opponents.

We will analyze a game of asymmetric information involving used-car sales to see how hidden information can drastically affect the outcome of a game. In our analysis, we will assume that the used-car lot has two kinds of cars: Half are lemons (bad cars) and half are peaches (good cars). As a customer, you value a peach at \$6,000 and a lemon at \$2,000. As the dealer, I value

a peach at \$4,999 and a lemon at \$999. You make an offer, which I either accept or decline, similar to how we played the Ultimatum Game.

If both players can tell a lemon from a peach, this game involves chance but is still a game of perfect information. Rolling back the tree, you offer \$5,000 for a peach and \$1,000 for a lemon and I accept.

If neither of us can tell a lemon from a peach, we simply move the chance node to the end of the game tree. We can still find a rollback equilibrium: You offer \$3,000 and I accept.

If only the car dealer can distinguish lemons from peaches, the game is one of asymmetric information. The game tree begins with a chance node, but you can't tell if you're at the lemon node or the peach node. For you, these nodes are in the same information set. Given that I'll never sell you a peach for less than \$5,000 and your expected payoff from a car is \$4,000, you should never offer enough to make me sell you a peach. This means you should never offer more than I'd accept for a lemon—\$1,000. The asymmetric information completely eliminates the market for peaches.

The “market for lemons” result has a wide range of applications; it can be seen, for example, in a scenario involving health insurance plans. Buying health insurance is a game of asymmetric information because people know more about their own health than the companies insuring them do.

We can represent the cost of insuring Americans as a probability distribution. Imagine a wooden plank with a dollar scale marked on it. For each American, I put a grain of sand on the plank at whichever dollar mark corresponds to his or her health care cost for the next year.

The expected value of providing health care to an American is the point where the plank exactly balances. In this example, this point is at \$10,000. This amount, \$10,000, would be a reasonable amount to pay for insurance only if no Americans knew their own health care costs. If people can anticipate their health care costs, those who expect lower costs will not buy insurance for \$10,000. With everyone left of the balancing point opting out of health insurance, the average cost to insure anyone who will buy it jumps

to \$75,000. Repeating this logic, insurance eventually costs so much that no one wants to buy it.

“Moral hazard” can arise when a party doesn’t bear the full consequences of his or her actions. That player may play less carefully than we would expect. People take more risks with insured vehicles if they believe that their insurance companies will cover the cost. For the same reason, doctors may charge more for medical procedures than individuals might be willing to pay if those individuals are insured and doctors and patients alike know that their insurers will cover part of the procedure costs. In another example of moral hazard, pitchers in the American League throw 15% more beanballs than those in the National League because American League pitchers don’t bat and, therefore, don’t face retaliation in the batter’s box for hitting a batter on the other team. The idea of moral hazard also plays a role in higher car insurance rates in Philadelphia compared to those in Pittsburgh. Fewer people have auto insurance in Philadelphia, which means that more accidents involve uninsured drivers and insurance companies have to pay more damages if their drivers are involved in accidents. When insurers pass these costs on to consumers, fewer people can afford car insurance, which perpetuates the cycle.

Games of imperfect information can become very complicated, very quickly. As a modest example, we examine a simplified version of poker called Tiny Poker. The rules of the game are as follows: At the beginning of the game, you put \$70 into the kitty and I put in \$40. We are each dealt a card from the deck. The deck contains two 1 cards and one 2 card. I can bet \$70 or fold. Folding means the other player gets all the money in the kitty. You can then bet \$30 or fold. If neither player folds, the high card wins. If we tie, each player wins half the kitty. How often I should bluff depends on how often you think I’m bluffing, which in turn, depends on how much I bluff. This is the same infinite loop we found in simultaneous games, and as we will see, we can solve the problem by treating the game as simultaneous. ■

**Every sequential game
can be expressed as
a simultaneous game.
Saying it another way:
If you can express it in
a game tree, you can
express it in a matrix.**

Suggested Reading

Aliprantis and Chakrabarti, *Games and Decision Making*.

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Questions to Consider

1. In the Tiny Poker game at the end of the lecture, each player has two different information sets and four strategies. What are they? Remember, an information set is the information you have available to you at the time you make your decision, and a strategy must tell you what to do given each information set.
2. How does the issue of moral hazard apply to the bailout of private corporations by the government?

Whom Can You Trust?—Signaling and Screening

Lecture 13

Signaling occurs when I know something and I'm trying to convey what I know to you. The flip side of that is screening. ... You're trying to pick my brain. The methodology of both signaling and screening is the same. You look at the actions that the player takes and hope that they convey something about the information that he or she has.

In games of asymmetric information, a problem often arises. Specifically, how can one player convey information to another player and be believed? The player conveying the information is said to be “signaling,” while the player trying to discover the information is “screening.”

In one real-world case, a company wanted to establish a hazardous waste disposal plant near a residential neighborhood, but homeowners feared that the site would cause property values to plummet. The waste disposal company assured homeowners that the chance of a negative environmental impact was insignificant. In fact, the company added, it was confident that the site would bring economic gains to the community and cause property values to rise. Homeowners were suspicious because they knew that the company would make the same claim if the site would decrease housing values. One resident hit upon an interesting way to screen the company: An independent real estate firm would assess the area’s homes, and in five years, the company would offer to buy out those homeowners who wanted to sell their houses for the price of the earlier assessment. When the company’s spokesperson emphatically declined the offer, the true type of the firm was signaled.

Let’s now return to the game of Tiny Poker that we looked at in the last lecture. As you will recall, to begin the game, I put \$40 in the kitty, and you put in \$70. We are then each dealt one card from a deck that has only three cards: two 1s and one 2. If I decide to stay in the game after the deal, I must add \$70 to the kitty; if you stay in the game after the deal, you add \$30 to the kitty. Whoever has the 2 gets all the money in the kitty; if neither of us has the 2, we split the pot evenly.

How often I should bluff depends on how often you think I bluff and vice versa. By framing this game of asymmetric information as a simultaneous game, we can avoid the infinite cycle of second-guessing. In Tiny Poker, each of us has four pure strategies: always bet, bet only with a 2, bet only with a 1, and never bet. Using the “difference of payoffs” trick, we find that I should bluff 40% of the time with a 1, and you should bluff just over half the time, or about 8/15 of the time. On average, this game favors me by \$2. Game theory can mathematically describe even a quintessentially human behavior: bluffing.

Tiny Poker has a “semiseparating equilibrium,” which means that you can deduce my type (card) some of the time but not all of the time. There are two other kinds of equilibria for signaling: “pooling equilibria” and “separating equilibria.” In a pooling equilibrium, all types take the same actions. In a separating equilibrium, all types take different actions.

Examples of screening and signaling abound in Greek mythology. In the *Iliad*, Odysseus tries to avoid fighting in the Trojan War by feigning insanity; among other questionable activities, he plows his fields in random corkscrews. To determine whether Odysseus is really insane, Palamedes creates a separating equilibrium by throwing Odysseus’s infant son, Telemachus, in front of the plow. In the situation created by Palamedes, a sane Odysseus will respond differently than an insane one. Either type of Odysseus can continue to plow, but the costs are different. An insane Odysseus will have a minimal cost; he won’t even know what he has done. But for the sane Odysseus, killing his son is more costly than going to war. Not surprisingly, Odysseus stops plowing and reveals himself to be sane.

An example of signaling in mythology comes to us in a story of the god Dionysus, who was kidnapped by pirates and bound to a ship’s mast with ropes. Dionysus signals that he is a god by making the ropes untie themselves. When the entire crew except for the helmsman ignores this credible signal from Dionysus, he turns them into dolphins.

Examples of signaling abound in the natural world. For instance, when threatened by a cheetah, gazelles in East Africa sometimes perform a maneuver called “stotting” (jumping straight up into the air before running

away). A likely explanation for this behavior is signaling. Stotting serves as a credible signal that the gazelle is healthy, and even a cheetah cannot catch a healthy gazelle. Cheetahs often stop the chase after a gazelle stotts, and no one has ever seen a cheetah catch a stotting gazelle.

The brightly colored markings found in males of some species, such as peacocks, may also be examples of signaling. Given that colorful creatures are easier for predators to spot, males may signal their strength and good genes by surviving even with the disadvantage of distinctive coloring.

Now, let's turn to a signaling example in business: Suppose you're running a high-quality gardening business for the next two months. Three-quarters of the gardeners in the market perform low-quality work and one-quarter perform high-quality work. Low-quality service costs \$300 per month to provide and is worth \$200 to the customer. High-quality service costs \$500 per month to provide and is worth \$1,200 to the customer. Despite a signaling problem akin to the market for lemons, you can find profitable work using an introductory offer: You can credibly signal you are a high-quality gardener by making an introductory offer of \$290 for the first month of work. A low-quality gardener will be fired after one month and will lose \$10 with the offer, but after you establish your credibility as a high-quality gardener, you can charge up to \$1,190 for the second month to recoup your costs.

The film *The Princess Bride* also includes a surprising amount of game theory, including vignettes that illustrate the topics we have covered in this lecture. In the battle-of-wits scene, our hero, Westley, produces a tube of poison. He takes two goblets of wine, and with his back turned to his opponent, Vizzini, pours the poison into the wine. Westley tells Vizzini to choose which goblet he wants; they will then both drink. Clearly, this game is one of asymmetric information because only Westley knows where the poison is. After trying to screen Westley, Vizzini chooses a glass—and dies. What Vizzini didn't know was that Westley had spent the last few years building up an immunity to the poison and, thus, was able to put the poison into both cups. Vizzini didn't know he was playing a game of incomplete information. In essence, two types of Westley existed, and one of them was immune to the poison.

Later on in the movie, Westley has been temporarily paralyzed, and we are unsure whether or not the paralysis has worn off. The evil Prince Humperdinck tries to screen Westley before deciding to attack. Westley breaks the impasse by standing and resolutely ordering Humperdinck to drop his sword. Humperdinck does so and is captured just as Westley collapses. Humperdinck thought he was dealing with a separating equilibrium (only an unparalyzed Westley would stand) when, in reality, both types of Westley (the paralyzed and the unparalyzed) chose to stand. ■

Suggested Reading

McCain, *Game Theory*.

Spence, “Job Market Signaling.”

Questions to Consider

1. Why would one not expect a separating equilibrium in the Westley/Humperdinck duel game?
2. In light of the material in this lecture, how can the lavish downtown office space of an advertising company be justified?

Encouraging Productivity—Incentive Schemes

Lecture 14

Suppose that the screening and signaling is done and you know the kind of person you're dealing with. They're still going to do what they want to do if they're rational. They're going to factor in everything that they care about and then choose the option that gives them the best overall satisfaction. If you want to get people to do what you want them to do, you're going to have to create an alignment between their desires and yours.

How do you get others to do what you want them to do? This question is obviously an important one in business, politics, international relations, and day-to-day life. The answer is generally to create an alignment between the behavior you desire and the rewards that the other player receives. Such systems are generally called “incentive payments” or “incentive schemes.”

For this lecture, we're going to assume that \$1 and 1 util are the same thing. However, we should remember that this assumption doesn't always hold. For example, Western Electric conducted a famous experiment at its facility in Cicero, Illinois, to find out if lighting affected worker productivity, an experiment that demonstrated that \$1 and 1 util are not the same in all situations. In this experiment, when the experimenters raised the level of lighting in the factory, workers produced more. When the experimenters then progressively lowered the lighting until it was below the original level, the workers produced still more. Although experimenters initially were surprised and confounded by this, they eventually figured out that, with all lighting changes, the workers had incentive to work harder because they believed that management was paying attention to their needs.

Incentive schemes present a possible solution for the “principal-agent problem”—the question of how an employer makes an employee work as hard as possible. If the employer can monitor the employee, the problem doesn't arise. But if the employer pays more for hard work and can't directly observe the employee's behavior, the employee can shirk his or

her responsibilities and earn a better payoff. We've seen this problem of information asymmetry before. To understand the principal-agent problem more thoroughly, consider the following scenario: Your company is bidding on a \$250,000 contract. Your proposal writer, Nathan, can either work hard or work at his normal pace. Nathan's hard work increases the probability of the proposal's acceptance from 50% to 80%. Nathan will work at his normal pace for \$30,000 but wants at least \$60,000 to work hard. If you could monitor Nathan's behavior, you could induce him to work hard by offering him a little over \$60,000 for working hard and nothing for working normally.

If Nathan works on commission, he has an incentive to put in the extra work. The key is to make Nathan's pay contingent on the variable you can observe: the success of the project. In order for a commission arrangement to succeed, Nathan has to earn more than \$60,000 if the project is successful, working hard has to be a better deal than working normally, and the company's best payoff has to come from Nathan working hard. Solving the equations for these conditions yields a bonus of about \$100,000 for Nathan and a base salary of -\$20,000. Although a negative base salary is probably unrealistic, some companies do pay their employees entirely from commissions. Team-based incentive schemes seem to be more effective than incentives tied to overall company performance.

In some cases, the cost of information asymmetry is just too high. If the contract is worth only \$150,000, it's better to let Nathan work at his normal pace and pay him \$30,000.

As another example, suppose you run a company with 10 suppliers, and you need at least 8 of them to deliver in order to function. The suppliers can either deliver on time, in which case they make \$2,000 on the delivery, or they can deliver late, in which case they make \$5,000 on the delivery. Not surprisingly, your suppliers often deliver late. Threatening to fire any supplier who delivers late isn't credible because you can fire, at most, 2

Linking employee pay to observables in company success has found wide acceptance in business.

suppliers. That means a supplier's expected payoff for delivering late is $(8/10)(\$5,000) = \$4,000$.

The solution is to assign each of the suppliers a different number from 1–10. Then you can announce, “I'll stop doing business with the lowest-numbered supplier who is late.” Supplier 1 delivers on time because \$2,000 is better than nothing. Supplier 2 knows that 1 will deliver on time and, facing the same situation as Supplier 1, chooses to deliver on time, as well. The same reasoning forces all 10 suppliers to deliver on time.

Billionaire Warren Buffett proposed an incentive scheme to enact campaign finance reform. He wanted a bill introduced in Congress that would outlaw campaign contributions of more than \$5,000 from any individual or organization. The second part of the plan involved an eccentric billionaire (i.e., Buffett) putting up \$1 billion. If the bill failed to pass, the money would be donated to whichever party cast the most votes for it. According to Buffett, the bill would easily pass because neither party would be willing to give the other \$1 billion. In lecture it was noted that Buffett's idea should have also included the provision that if the bill failed with the same number of votes from each party that together added up to less than the number required for passage, the money would go to the minority party in Congress.

Not all incentive schemes are as well conceived as Buffett's. In the 1994 Shell Caribbean Cup, officials decided to make the soccer game more exciting by counting goals made in overtime as 2 points. At the end of the first round, Barbados needed to beat Grenada by two goals in order to advance. With 10 minutes left in the game, Barbados led 2 to 0. In the 83rd minute, Grenada finally scored. The 2-to-1 lead held by Barbados wasn't enough for that team to advance.

With three minutes left in the game, Barbados intentionally scored on its own goal. If the game ended in a tie, Barbados would have a chance to score in overtime for a 4-to-2 victory, which would allow the team to advance to the next round. After a minute of shock, the Grenada team realized that it could advance to the next round by scoring on either goal. The players from Grenada drove frantically at both goals, with Barbados playing defense on

both ends of the field! Barbados managed to hold out until the sudden-death overtime and scored the first goal to advance to the next round. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

McCain, *Game Theory*.

Questions to Consider

1. What incentive function do co-payments and deductibles play for insurance coverage?
2. How might high overtime pay discourage productivity?

The Persistence of Memory—Repeated Games

Lecture 15

A relationship with someone involves a history of past interactions and an anticipation of more interactions to come in the future. ... Many if not most games are really just a link in a much larger chain.

Because humans generally think in terms of repeated interactions, many if not most games depend on past and future interactions or games. The string of interactions in a repeated game is called the “supergame.” The individual one-shot games are called “stage games.” For now, we’ll assume that the stage games are all the same. Each iteration of the stage game is one round of the supergame.

Strategies for repeated games, like strategies for other games, have to tell players what to do in every situation. Repeated games tend to have a significant number of strategies, as well as multiple equilibria. “Open-loop” strategies don’t depend on history—a player playing an open-loop strategy does the same thing regardless of what another player does. A strategy that does depend on past moves is “closed-loop.” If two people play a stage game with more than one Nash equilibrium, like chicken, then the repeated game will also have more than one Nash equilibrium. Usually it will have a lot more. For example, a 10-round iterated game of chicken has well over 1,000 Nash equilibria, and many of them aren’t easy to characterize.

Even a 100-round prisoner’s dilemma, however, has just one subgame-perfect equilibrium: (defect, defect) in each of the 100 rounds. Suppose each player makes a nonbinding commitment to the grim trigger strategy: “I’ll cooperate, but if you ever betray me, I’ll defect for the rest of the game.” At first blush, it makes sense for the players to cooperate until the last round. At that point, there is no chance for the other player to retaliate, so defecting becomes both players’ dominant strategy. Because both players know that the strategy for the last round is mutual defection, the 99th round becomes the “real” last round. Of course, if the 99th round is the last round, both players will betray based on the same logic we saw for the 100th round. Neither player fears defection in the 100th round because both will defect anyway.

Working back through all of the rounds, both players defect on every round of the game. The only subgame-perfect equilibrium for a prisoner's dilemma of any fixed length is constant defection.

Although this result is depressing, we still have ways to elicit cooperation in the prisoner's dilemma. The problem with the 100-round prisoner's dilemma was that neither player had any reason to cooperate in the final round. But if players don't know when the last round will be, it's possible for them to cooperate. Suppose that after every round, there's a chance of continuing the game that we will signify with δ . In order for the grim trigger strategy to be a Nash equilibrium, neither player should be able to do any better by defecting.

More generally, the payoff for the grim trigger strategy is $3/(1 - \delta)$. For betraying each round, the payoff is $2 + 2/(1 - \delta)$. As we calculated, the payoffs are equal for $\delta = 1/2$. For $\delta > 1/2$, cooperation is the only best response to the grim trigger strategy. For $\delta < 1/2$, defecting is the only best response.

It's also possible to interpret δ in terms of what's called "net present value." Imagine that you could put your payoffs from one round in the bank and let them earn interest at a rate of r per round. Then, getting the payoff in the next round would not be as good as getting it now because you would lose the interest for one round. This model is described by exactly the same mathematics that we've done in this lecture, replacing δ with $1/(1 + r)$. In this model, cooperative equilibrium is possible, provided that the interest rates aren't too high.

Unfortunately, although the cooperative equilibrium exists, both players defecting is still an equilibrium. In fact, there's a subgame-perfect equilibrium that guarantees each player any average payoff between 2 and 3 provided that δ is high enough—and this equilibrium doesn't even have to give the same average payoff to the two players. ■

If the duration of the game is unknown but likely to be sufficiently long, rational players may adopt strategies resulting in sustained mutual cooperation.

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Mailath and Samuelson, *Repeated Games and Reputations*.

Questions to Consider

1. In an iterated prisoner's dilemma game, one may decide to punish an opponent who has defected by responding with defection for a certain number of rounds, then reverting to cooperation. This strategy would result in the punished player losing, say, $\$x$ each round (starting next round) for n rounds. Show that the present value of the money lost is $(\partial - \partial n + 1)x/(1 - \partial)$. [Hint: Call the present value of these losses T and parallel the work above.]
2. For many people, the idea of adding together an infinite number of positive numbers and getting a finite result seems perplexing. But consider adding the numbers 0.1, 0.01, 0.001, 0.0001, and so on, forever. The result is obviously just 0.1111 ..., which is the decimal representation of 1/9. On the other hand, it's not enough that the numbers are getting smaller at each step. The sum of $1/2 + 1/3 + 1/4 + 1/5 + \dots$ becomes arbitrarily large as you continue to add terms. The sum diverges.

Does This Stuff Really Work?

Lecture 16

To heck with what people *should* do if they're rational decision makers. If they don't actually do what game theory says, then the theory has no predictive power. Worse than that, if other players aren't doing what they should do, then my best choice for my real situation is probably different than what theory says, too. That is, the model doesn't have prescriptive power either.

Though game theory has proven remarkably effective in predicting outcomes in some situations, it requires a good deal of information in order to do so. As a field that aspires to be the theory of strategic interaction among rational decision makers, game theory relies on several conditions: the game must be modeled correctly, the payoffs must be accurate, players must play rationally, and the structure of the game must be common knowledge. If one player doesn't play a game rationally, game theory's prescriptions for other players probably aren't their optimal strategies.

In the Ultimatum Game, the game-theoretic equilibrium—the proposer offers a penny and the responder takes it—almost never happens in real life. Proposers tend to offer around 40% of the pot, and responders frequently reject offers of less than 20% to 30%.

Researchers in game theory have ruled out a number of possible explanations for these results. Responders have no reason to “teach the proposer a lesson,” because the two won’t have an opportunity to play again. You might think that responders would reject low offers because the amount of money involved is so small, but increasing the size of the pot has almost no effect on the results.

The proposer might propose a high offer out of fear of rejection. The simpler dictator game was devised to test this hypothesis. In the dictator game, the dictator simply chooses how much to give the other player and keeps the rest. That’s the end of the game. Only about 20% of the players keep all the money, even though that strategy is obviously the only equilibrium.

Although most people share the wealth out of a sense of fairness, “fairness” is shorthand for a remarkably complicated payoff modifier. Additional factors, such as maintaining anonymity or casting the dictator as the seller of an item, had a marked effect on offers in the dictator game. With six such factors in place, nearly two-thirds of dictators kept all the money. In general, the larger the social distance between the players, the smaller the donations.

Variations on the Ultimatum Game have been played in experimental contexts to try to uncover why what players do in these games conflicts with game theory. Experimental evidence from these games shows that the appearance of fairness seems to be more important to players than equal payoff.

Behavioral game theory seeks explanations consistent with game theory for why humans often don’t play equilibria. It is unlikely for an equilibrium to appear out of thin air, but John Nash and others have noted that people often move closer to game-theoretic predictions over time.

Neuroscience shows that during the Ultimatum Game, a part of the brain called the “insula” activates whenever the responder receives a low offer. The insula is responsible for generating emotionally relevant context for sensory experience. With the Ultimatum Game, the greater the number of cells in the insula firing, the quicker the rejection of the proposer’s offer. In other words, the rejection of the offer is visceral, not logical. Experimenters have also found circumstances in which responders will accept lopsided divisions that seem to be out of the proposer’s control.

Experiments with highly social capuchin monkeys prove that the idea of fairness extends beyond humans. In these experiments, researchers trained monkeys to trade a rock for a food reward. Later, when researchers gave more enticing food rewards to some monkeys or rewarded some monkeys who had not earned the rewards, other monkeys refused to play with the human researchers, often even throwing food at them.

Cooperative solutions to the prisoner’s dilemma can develop from kinship or from “reciprocal altruism”—the idea that one good turn deserves (and will get) another. The system of reciprocal altruism will break down, however, unless it incorporates some punishment for cheaters.

Evidence from the Ultimatum Game suggests that people are willing to pay the cost of punishment even under rather remarkable circumstances. To move toward a better understanding of this point, we will revisit the prisoner's dilemma. What happens if we add a punishing player to the prisoner's dilemma? An arbiter (punisher) whose payoff is the sum of the players' payoffs will never punish anyone and, thus, will have no effect. An arbiter whose payoff is the sum of the cooperating players' payoffs has no reason to punish or not to punish. Only an arbiter who receives a positive payoff from punishing defectors can achieve reliable cooperation.

Not all work in behavioral game theory deals with changing payoffs. Another research direction addresses "bounded rationality"—the idea of limits on human rationality. The easiest way to explain this concept is with a game. In a group of 100 people, each person picks a number between 0 and 100. The goal is to pick the closest number to 70% of the average of all the numbers chosen. The only Nash equilibrium for this game is for everyone to pick 0. In studies, the median guess is around 35 and the winning guess is generally around 25.

The kind of behavior exhibited in this experiment is called bounded rationality, a term that describes the limit of how far normal people will carry out a chain of reasoning before terminating it. Generally, people carry chains of reasoning further with repeated play, especially when the previous rounds are common knowledge.

When Robert Axelrod set up an iterated prisoner's dilemma competition, the winning strategy was a surprisingly simple program called tit for tat. With this strategy, one player cooperated on the first round and, after that, did whatever the other player had done in the previous round. In a second competition, more than twice as many programs were submitted, including a few programs designed specifically to beat tit for tat, but the tit-for-tat strategy won again.

**Advantages are
conferred on society
by such incentives
as "fairness" and
"altruistic punishment."**

Axelrod believes that tit for tat works because it elicits cooperation from other players. Tit for tat has four main qualities that make it successful: First, the strategy is nice because the player using it is never the first to betray. Second, it's able to be provoked in that it quickly and reliably punishes the other player's betrayals. Third, it's forgiving because it can easily return to cooperation even after a defection. Lastly, it's straightforward because it's simple enough that other players can see that cooperating is the best choice. ■

Suggested Reading

Axelrod, *The Evolution of Cooperation*.

Bueno de Mesquita, *Predicting Politics*.

Green and Armstrong, “The War in Iraq: Should We Have Expected Better Forecasts?”

Thaler, *The Winner’s Curse*.

Questions to Consider

1. You are player 2 in the Ultimatum Game. Player 1 has offered you \$10. Does it make a difference to you if player 1 is dividing \$12, \$20, or \$100? Why? Does it make a difference to you if player 1 is offering his own money, dividing money that he won in a previous game, or dividing money given to him for division? Why? Does it make a difference if player 1 is an individual, an organization, or a computer? Why?
2. Two lanes of traffic are narrowing to one because of an accident on the road ahead. The drivers of 10 cars have ignored your turn signal and refused to allow you to merge into the single lane. Finally, you get in the single lane. Does this history change your decision on letting other drivers into the single lane? Why? In which direction? Is your behavior rational?

The Tragedy of the Commons

Lecture 17

Repeated exposure to the same situation lets us try different strategies and see how they work. If everyone is doing this, we often *slide* toward the game-theoretic equilibrium.

Collective-action games are those with more than two players. If binding agreements are allowed, games with more than two players can become very difficult to analyze. In this lecture, we will focus on noncooperative games. Because these games can grow very large, often with millions of possible outcomes, we simplify them by assuming a community of players who share the same strategies and payoffs.

A “free rider” is a person who benefits from someone else’s work without paying any cost. Essentially, a free rider is shirking responsibility in a many-player social dilemma. A social dilemma is any game in which the equilibrium isn’t Pareto-optimal. Although all the social dilemmas in this lecture have only two strategies for each player, this limited model can describe many different situations. We will call the two strategies “working” and “shirking.” Despite the negative connotation of “shirking,” we will see that it may sometimes be a good idea to let some members of society shirk.

We will begin by looking at air pollution as a social dilemma; players can work to curb emissions or shirk by producing emissions as they normally do. Payoffs depend on the fraction of the population that is working (p). Air pollution is a multi-person version of a prisoner’s dilemma, and it presents to us a situation known as the “Tragedy of the Commons.” In this situation, many people acting in their own self-interest may, over time, destroy a shared resource. Such “tragedies” occur every day, from overfishing to traffic congestion. The button-pushing game from the first lecture involved a Tragedy of the Commons. In social dilemmas, we measure social good as the sum of payoffs to all players. In a Tragedy of the Commons, “everybody works” does not always maximize social good. This is because, if the benefits from shirking are great enough, having some people shirk can maximize social good. However, determining which members of society should shirk

is often a difficult problem. If shirkers can transfer some of their payoffs to workers, then everyone will receive the average payoff. In effect, shirkers pay for the inconvenience they cause.

Social dilemmas include more situations than just the Tragedy of the Commons. Joining a Neighborhood Watch program is one such case; deciding whether to use the metric system is another.

In a Neighborhood Watch program, becoming a watcher helps to keep your house safe; with 40% of the neighbors watching, both shirkers and workers get a payoff of 4.8. Although 40% watching is better than everyone shirking, it still creates a social dilemma because it doesn't maximize social good. Social good is maximized with about 80% of the neighbors watching; the watchers get a payoff of 5.6 and the shirkers get a payoff of 9.6. The watchers do much worse than the shirkers, but they are still making more than they would get at the 4.8 equilibrium. The Neighborhood Watch program corresponds to a multiplayer game of chicken known as the “volunteer’s dilemma.”

Let’s look at another version of this game that involves the question of why America has not switched to the metric system. This game has three equilibria. From most social good to least, they are: everyone shirks (switches to the metric system), everyone works (retains the system of customary units), and 63% work. The mixed equilibrium isn’t stable and is likely to start shifting to one of the other two equilibria. This game is a multiplayer coordination game.

Kitty Genovese’s horrible murder in 1964 is a tragic example of a volunteer’s dilemma. Genovese was murdered in a half-hour attack witnessed by at least a dozen of her neighbors. Not one of them called the police. Startlingly, the more people eligible to “volunteer”—in this case, call the police—the less likely the event is to be reported by anyone. ■

**What happens when
you take the simplest
of games ... and move
them to their many-
player equivalents? ...
These questions have
remarkable relevance to
real-world situations.**

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Olson, *The Logic of Collective Action*.

Questions to Consider

1. Draw a picture similar to those used in this lecture, but have the payoff functions $w(p)$ and $s(p)$ cross at least two times by making them curves. Analyze the behavior of the resulting society, given that the population originally starts at various points along the horizontal axis.
2. The terms “worker” and “shirker” carry strong connotations, but our analysis doesn’t depend on them. In any of our graphs, the role of “worker” and “shirker” could be reversed by reflecting the graph about a vertical axis (essentially “flipping the page” by exchanging left and right) and reversing the colors of the two lines.

Games in Motion—Evolutionary Game Theory

Lecture 18

This version of game theory rests on a different foundation than the one that we've developed so far. ... And in spite of this, we're going to see that many of the results that we've seen already are results that can be paralleled in the field of evolutionary game theory.

In previous lectures, we've seen that people often need time to gravitate toward an equilibrium. Evolutionary game theory applies this idea to a species. In this lecture, we'll apply game theory to outcomes among different members of the same species. The behavior patterns of a given individual are called its "phenotype." A phenotype corresponds to an individual's strategy in traditional game theory, but it is different in that the phenotype is hardwired; an individual can't change its phenotype any more than you can change your height. In our examples, the phenotypes are as follows: A "hawk" may fight for a prize while a "dove" will not. Individuals never change their phenotype, and successful phenotypes propagate through time.

Populations with only one phenotype are called "monomorphic," from the Greek for "one form." Populations with two or more phenotypes are "polymorphic"—"many forms." It is possible for a single phenotype to adopt a mixed strategy, such as behaving like a hawk 40% of the time and like a dove 60% of the time. "Monomorphic" isn't synonymous with "pure strategy."

A "fitness" function measures how well each phenotype does in interacting with the population as a whole. The fitness function is essentially the expected payoff from an interaction for this phenotype. In a species with 10% hawks and 90% doves, hawks have greater fitness because fights are rare. Hawks win the prize 95% of the time and lose only 5% of the time. Doves win only 45% of the time and have to put on a show to scare off other doves. When one phenotype has greater fitness than another, its proportion increases in the next generation. Evolutionary game theory asks what will eventually happen to a species given a set of phenotypes. We can ask this question in two ways. One way is as follows: If we start with a mix of phenotypes and let the population

evolve, what will be the long-term distribution of phenotypes? The second way to ask the question is this: If we start with a distribution of phenotypes that's stable over time, can a small number of a different phenotype successfully invade this population?

Like the rest of game theory, evolutionary game theory gets quite deep. The mathematics behind it can get rather sophisticated.

Let's look at an iterated prisoner's dilemma from an evolutionary perspective. When the opponents in this game meet, they can cooperate or one player can betray the other. The population in this game has two phenotypes: Grims play the grim trigger strategy, and scrooges always betray. For now, we'll also assume that any two players

play the game only once and then find new partners. To work through this game, we will let p be the proportion of scrooges in the population. The scrooge payoff of $3 - 2p$ is always greater than the grim payoff, $2 - 2p$. We would expect to see scrooges eventually take over the population. Moreover, because scrooges do better against other scrooges than grims do, a few grims can't invade a scrooge population.

For evolutionary biologists, a population that can't be invaded by another phenotype is said to exhibit an “evolutionarily stable strategy,” or ESS. An all-scrooge population is a monomorphic ESS. An all-grim population is not, because scrooges can invade grims.

If the game is played twice, grims can retaliate in the second round. The scrooge payoff of $4 - 2p$ is still always greater than the grims' $4 - 3p$ if there are any scrooges in the population at all. An all-scrooge population is once again an ESS, because it can resist invasion from the grims.

What happens if we extend the game to three days? The scrooge payoff is now $5 - 2p$, and the grims' is $6 - 4p$. If we set the two fitness functions equal to each other, we find that they're the same when $p = 1/2$. When the number of scrooges and grims is the same, they both have a fitness function of 4. In this evolutionary game, whichever population has the majority at the outset ends up with greater fitness. Also, both monomorphic populations are ESSs.

Some games, including, our hawk/dove example, have polymorphic ESSs. We can find ESSs in much the same way that we found equilibria in other games. In a game with hawks and doves competing for a prize, neither monomorphic strategy is an ESS because hawks do better against doves and doves do better against hawks.

The ESSs we found in our scrooge-versus-grim games are exactly the same as the admissible, stable Nash equilibria for the games in standard game theory. We can even find the stable mixed Nash equilibrium in the hawk/dove game at $p = 0.7$. The hawk/dove game has two other Nash equilibria: One player plays hawk every time and the other plays dove every time. These strategies can't be ESSs because the population can't be all hawk and all dove at the same time. This correspondence is really remarkable, because we never required the players to be rational or even to know the payoff matrix.

We can draw members from different species, as well, to handle such games as the battle of the sexes. We can also handle as many different phenotypes as we want. When we introduce a new phenotype, it's possible that some other phenotypes may no longer be ESSs. In the three-day prisoner's dilemma, imagine a third phenotype, the sneak, which plays like a grim except that it always betrays on round three. Grims are no longer an ESS because sneaks can invade them.

In a war of attrition, in which all players pay the same increasing cost until all but one gives up, it's easy to end up paying more than the value of the prize. Imagine an auction for a \$100 bill. A "price clock" starts at \$0 and steadily increases. Each player has a buzzer. Pressing the buzzer stops the clock and means you lose the auction. Both players pay the price on the clock, but only the winner gets the \$100 bill.

There's no monomorphic ESS for the War of Attrition, but there is a polymorphic one. Each time the clock clicks another dollar, give yourself about 1 chance in 100 of buzzing in. The actual chance should be 0.995%. This polymorphic strategy outperforms any monomorphic strategy—any fixed stopping point—by at least \$19. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*, 2nd ed.

Gintis, *Game Theory Evolving*.

Questions to Consider

1. We saw in our hawk/dove game that a polymorphic ESS exists with a population of 70% hawks and 30% doves. But we could consider a mutant phenotype in this game that acts as a hawk 70% of the time and a dove 30% of the time. Show that this phenotype is a monomorphic ESS—it cannot be successfully invaded by either hawks or doves.
2. Games played between two different species (necessary for games that aren't symmetric) require a more complicated graphical representation than what we've seen in this lecture. If each of the species has only two phenotypes, then the situation can be represented within a unit square, in which the x -axis records the fraction of species 1 of its first phenotype, and the y -axis records the fraction of species 2 of its first phenotype. The dynamic evolution from any original state now corresponds to a trajectory through the square, like a ball rolling across a hilly tabletop. Locations where the ball would stop and not be dislodged by a slight breeze are ESSs.

Game Theory and Economics—Oligopolies

Lecture 19

In 2002, the FTC filed charges against two pharmaceutical companies, Andrx in Florida and Hoechst in Germany. Hoechst allegedly paid Andrx \$40 million a year *not* to market the generic equivalent to Hoechst's Cardizem CD. ... The two firms evidently created a bottleneck that made it impossible for any other firm to enter the market.

In this lecture, we look at how game theory is used in economics. Our topic is a classic one from microeconomic theory: What is the optimal production level for a monopolist, and how does the situation change when one or more competitors enter the market? We begin with monopolies, which we will later compare to games involving oligopolies (markets that have only a small number of competitors). In both situations, the seller's goal is to maximize profit.

A monopolist, like other producers, faces a demand curve. Each point on the demand curve tells what quantity of the good is demanded at a certain price. In almost all situations, a demand curve is downward sloping: the higher the price, the lower the quantity demanded.

Maximizing revenue is equivalent to maximizing the area of the rectangle formed by the origin and a point on the demand curve. Maximizing profit, the monopolist's true goal, is equivalent to maximizing the rectangle with its upper-right corner on the demand curve and its base on the variable-cost curve.

We will look at an example of a monopoly, in which the monopolist produces chairs. The monopolist faces a demand curve of $p = 700 - 2q$. In this equation, p represents selling price and q represents quantity to be sold. Given a production cost of \$100 per chair, a demand curve can determine total profit for any value of q . In our case, the monopolist should produce 150 chairs to maximize his or her profit.

In a “von Stackelberg duopoly,” one firm chooses its production level, then the other firm chooses its production level based on this information. Such duopolies offer a significant first-mover advantage. The demand function changes to $p = 700 - (q_1 + q_2)$ because two companies are now producing output. Finding the profit levels in a von Stackelberg duopoly is like solving a sequential game of perfect information. To look forward, you must reason backward. Player 2’s optimal response to player 1’s production of q_1 —her best response function—is to produce $q_2 = 150 - q_1/2$. Player 1 knows this response function and can substitute it for q_2 when calculating optimal q_1 . Solving for player 1’s profit (π), we find $\pi_1 = 300q_1 - q_1/2$. Note that the leading firm can calculate the follower’s response function ahead of time.

Taking the derivative yields a maximum profit at $q = 150$ chairs—the same as the monopoly quantity! If demand curves are linear and variable costs are constant, including a follower firm doesn’t change the leading firm’s quantity and cuts its profits precisely in half. The following firm makes half of the leader’s quantity and earns half its profits—one-quarter of the original monopolist’s profits. The other quarter of the monopolist’s profits stays in the hands of consumers.

“Cournot duopolies” are markets cast as simultaneous games. Because our firms are identical and we move simultaneously, our response functions are symmetric. My response function is $q_2 = 150 - q_1/2$, and yours is $q_1 = 150 - q_2/2$. Solving these equations simultaneously gives q_1 and $q_2 = 100$. Each firm earns \$20,000 in profit. ■

A monopolist making \$45,000 could be tempted to pay a prospective entrant to the marketplace somewhere between \$20,000 and \$25,000 to refrain from entering. Both firms earn more this way than with the Cournot duopoly’s \$20,000. Collusion of this nature is fairly common in the real world, especially in such fields as pharmaceuticals. ■

When you strip a situation down to its essentials, it's surprising how often you'll be looking at a game.

Suggested Reading

Aliprantis and Chakrabarti, *Games and Decision Making*.

Gibbons, *Game Theory for Applied Economists*.

McCain, *Game Theory*.

Questions to Consider

1. Recall the pastry/breakfast sandwich example from Lecture 5. These goods are partial substitutes, and their demand equations are $q_1 = 40 - 8p_1 + 2p_2$ and $q_2 = 40 - 8p_2 + 2p_1$, where the subscript 1 refers to pastries and the subscript 2 refers to breakfast sandwiches. Assume that the variable cost is \$1/pastry and \$2/breakfast sandwich. Using techniques parallel to those in the Cournot duopoly model, show that the equilibrium for this market is approximately \$3.49 for pastries and \$3.94 for breakfast sandwiches.
2. Consider a business in which units take time to produce and storage can be expensive. Would you expect such a business to compete by setting quantity or by setting price?

Voting—Determining the Will of the People

Lecture 20

There can be problems with voting. ... Voter fraud, hanging chads, the Electoral College—all problematic, but in some ways, this is just the tip of the iceberg.

We begin by comparing plurality voting, in which the candidate with the most votes wins, with alternative systems. A good alternative is desirable because plurality voting is vulnerable to the “spoiler effect” when an election involves more than two candidates. We will use an example to consider different voting systems. In our example, a hospital is planned to be built in one of four towns: Easton, Northview, Westlake, or Southville. The citizens will vote on where the hospital should be built; all citizens would like the hospital to be as close as possible to them. According to plurality voting, Southville wins a sincere vote because it has the largest population. But this result means that more than two-thirds of the voters got their last choice.

The “Condorcet method” uses a series of head-to-head matches to determine a winner. A Condorcet winner has to beat every alternative in a one-on-one vote. Westlake wins a head-to-head vote against any other town; thus, it is the Condorcet winner. Unfortunately, finding a Condorcet winner isn’t always possible. In some situations, the group of voters may prefer A to B, B to C, and C to A. Studies suggest that this phenomenon, called the “Condorcet paradox,” doesn’t happen too often in real life, but it’s more likely to occur if players vote strategically.

Instead of focusing on first-place votes, the “Borda count” asks voters to rank their choices from best to worst. The least favorite candidate gets 0 points, the next preferred gets 1 point, and so on. In our hospital vote example, Easton edges out Westlake in the Borda count, 375 to 365. Eliminating Northview, which would lose anyway, causes voter preference to shift from Easton to Westlake.

According to “Arrow’s impossibility theorem,” no voting system can satisfy four reasonable-sounding requirements. Kenneth Arrow came up with five conditions that a good voting system would satisfy. (These conditions can be boiled down to four.) He then proved that no voting system except a dictatorship (in which only one person is allowed to vote) could satisfy the first three conditions. First, Arrow’s system couldn’t result in intransitive preferences, as we saw in the Condorcet method. If the system prefers A to B and B to C, it must prefer A to C. Second, the system had to be independent from irrelevant alternatives. If voters prefer A to B, they can’t prefer B to A when C enters the contest. Finally, if every voter prefers A to B, the system must rank A higher than B. The Arrow impossibility theorem was disturbing because it seemed to say that the idea of a democracy was fundamentally flawed. Some critics believe that Arrow’s condition of independence from irrelevant alternatives is too strict.

A popular alternative to plurality voting called “instant runoff voting” (IRV) can also yield some bizarre results. In IRV, candidates need to reach a certain quota—often a majority—to win. The process simulates a series of runoff elections, with one candidate eliminated each round. In the hospital example, Easton is eliminated first and gives its votes to Northview. Next, Westlake loses and also votes for Northview. Northview wins the IRV election. IRV voters don’t have to worry about “wasting” votes—they will indirectly vote for their second-choice candidates. It’s possible to lose an IRV election because public opinion moved closer to your position. It’s also difficult for moderate candidates to win.

A number of simulations by Ka-Ping Yee illustrate outcomes of the various voting systems. With IRV and plurality voting, moderate candidates have a more difficult time winning; a Borda count makes a win easier for moderate candidates. Additionally, with many candidates, IRV yields confusing and troubling results.

Given [recent election] turnout ... perhaps it's worth pointing out the obvious: that no voting system works if we don't vote.

In a one-dimensional election between two candidates, the game has only one Nash equilibrium: both candidates in the exact center. This equilibrium may partly explain why candidates move closer to the center before a general election. If a candidate adopts a position anywhere else, the other player can play just a tiny bit closer to the center and win a majority of the vote.

Arrow's impossibility theorem applies only to ranked systems, but some unranked systems of voting show real promise. One method is "approval voting," in which each voter can vote for or against each candidate, and the candidate with the most votes wins. The United Nations uses approval voting to elect the secretary-general. In "range voting," the voter gives each candidate a score from 0 to 10. The candidate with the highest average score wins. Range voting gives voters more flexibility than simple approval voting. It also seems to be much harder to disrupt with strategic voting than other systems, such as the Borda count. ■

Suggested Reading

Amy, *Behind the Ballot Box*.

Poundstone, *Gaming the Vote*.

Questions to Consider

1. Suppose that three different versions of a bill appear before a committee of 12 people. Four people prefer version A to version B to version C, 4 prefer version B to version C to version A, and 4 prefer version C to version A to version B. Because there is no majority, the chair of the committee will first hold a vote between two of the alternatives, then the winner of this vote will face the remaining version. Show that if people vote their preferences, the chair effectively chooses which version will be the winner.
2. Do you think that instant runoff voting would increase or decrease the viability of a third party when compared to plurality voting?

Auctions and the Winner's Curse

Lecture 21

Assume that a single object is being auctioned. How much is that object worth? If the answer to that question were the same for everyone and were known to everyone exactly, then there wouldn't be an auction.

Auctions play a more significant role in our lives than many of us realize. The revenue from auctioning radio frequency spectrum licenses equals almost 1% of the federal budget. Online search engines, such as Google, use auctions to determine what ads and links we see when we search. The explosion of online auction sites, particularly eBay, has resulted in millions of transactions for billions of dollars worth of merchandise annually.

Auctions have a number of forms but are generally used only when some uncertainty exists about the value of an object. In a private-value auction, the item being auctioned is worth a different amount to each bidder. Each bidder knows only his or her own valuation of the item. In a common-value auction, in contrast, the item may have the same value to all bidders, but each bidder has only an estimate of what the item is worth.

Open-outcry auctions allow anyone to bid at any time, and all bids are observable by all players. Such auctions are what we usually imagine when we think of auctions in general. These auctions can involve either ascending or descending prices and are called “English” and “Dutch” auctions, respectively. In a descending-price auction, the bid continues to decrease until a player volunteers to pay the bid in exchange for the item.

Sealed-bid auctions allow each player to submit only one bid. The highest bid wins. Many sealed-bid auctions are “second price,” or “Vickrey,” auctions—the winner pays the second-highest bid. If the winner pays his or her own bid, the auction is “first price.” English auctions are strategically equivalent to second-price auctions, and Dutch auctions are equivalent to first-price auctions.

In common-value auctions, overbidding is a constant concern. For example, suppose a common-value auction of a car has only one bidder. The seller sets a reserve price at his valuation of the car, which only he knows. You, the bidder, value the car at half again the reserve price. You have only an estimate of the car's value. This one-bidder auction mirrors the market for lemons. If you bid p and the dealer sells you the car, your expected value is $p(1.5)(0.5) = 3/4p$.

What happens if we generalize the auction to multiple bidders and assume the car is worth \$5,000? Five bidders estimate the car's value at \$4,500, \$4,800, \$5,000, \$5,200, and \$5,500. Even in an English or second-price auction, the car sells for a bit over \$5,200, which means that the winner loses more than \$200.

Because the highest estimates are usually too high, this “winner's curse” frequently occurs in common-value auctions. To avoid the winner's curse, players should “shade” their bids, bidding less than they think the item is worth. The important question is: How much am I willing to pay if I think this item is worth \$5,000 and no one else is willing to bid above me? Calculations of how much to shade one's bid can be quite sophisticated.

In private-value auctions, because each player knows his or her own valuation of the object, the winner's curse doesn't pose a problem. However, players still shade their bids in order to ensure a positive payoff. Consider a first-price auction between two bidders in which the private value of the object to each bidder is equally likely to be any amount from \$0 to \$100. The only Nash equilibrium is for each player to bid exactly half of his or her valuation.

In a second-price auction, bidding your valuation weakly dominates any other strategy. Imagine that the item is worth exactly \$100 to you. Bidding more produces a different outcome only if someone else bids more than \$100. In this case, you're better off losing instead of paying more than the item is worth. Bidding less doesn't allow you to win the item for any less, and someone else may win the auction for less than \$100.

The “revenue equivalence theorem,” a celebrated result of auction theory, addresses the question of which type of auction is best for buyers and sellers. According to this theorem, as long as the valuations of the bidders are independent from one another, first-price, second-price, open-outcry, and sealed-bid auctions each give the same expected revenue for the seller and the same expected payoff for the buyers. The revenue equivalence theorem relies on bidders being risk neutral. Research shows that real-time Dutch auctions yield lower revenues, and Internet Dutch auctions yield higher revenues.

Real-life auctions occasionally defy theoretical predictions, often in interesting ways. The April 1997 spectrum auctions in the United States brought in less than 1% of expected revenues. Game theorists noticed reappearing number sequences at the end of bids. Bidders were signaling to one another using the last few digits of their bids. Some companies used such signals to dissuade competitors from bidding up their frequencies and to threaten retaliation in later rounds of bidding.

New Zealanders were furious with their government after a Vickrey spectrum auction apparently failed to bring in as much revenue as expected. The public wanted to know why a bidder who was willing to bid \$72,000 for a license got it for \$4 and why a bidder who had offered more than \$5 million ended up paying only \$3,600. In English auctions, no one ever sees what the winner would have paid. When citizens saw these numbers, they assumed that the auction had failed.

In spite of their shortcomings, auctions offer a number of important benefits. First, they provide a credible signal indicating which entity can best use a resource. Also, because players can watch other bidders’ behavior, auctions serve to give companies an idea of whether their valuations for resources are reasonable. Google makes most of its revenue by selling sponsored links through second-price auctions. ■

**Buying on eBay,
selling on Google,
or determining
who is going to be
providing you with
cell phone service
next year—auctions
are everywhere.**

Suggested Reading

Harford, *The Undercover Economist*.

Thaler, *The Winner's Curse*.

Questions to Consider

1. Consider an all-pay auction for \$100. Each bidder makes a sealed bid for the \$100. The highest bid wins, but all bidders must pay their bids. You can bid any amount up to \$100. How much would you bid? Your answer almost certainly depends on the number of bidders, n . Imagine a scenario in which there are only 2 bidders and one in which there are 10. According to calculus and probability theory, with 2 bidders, you should randomly choose any number between \$0 and \$100, with each number being equally likely. Your average bid would, therefore, be \$50. With n bidders, your average bid is $\$100/n$, and your particular bid is generated by picking a number in the 0 to 1 range (with each number being equally likely), raising it to the $n - 1$ power, and multiplying the result by \$100. For 10 bidders, half of the time, your bid will be 20 cents or less!
2. Why is the winner's curse not a problem for private-value auctions?

Bargaining and Cooperative Games

Lecture 22

In a cooperative game, each subset of players can work together to create or capture a certain amount of value—a payoff. ... The payoff that these groups can capture isn't any bigger than what the big group could do to begin with. In general, the sum is at least as big as its parts. This raises the possibility that players working as a “grand coalition,” with everybody included, can capture the maximum value. The question then becomes: How much of that value does each player receive?

In cooperative games, binding agreements and coalitions among players are possible. Most of the time, players can acquire more payoff by cooperating than by working independently. A coalition that includes all players is called a “grand coalition.”

As we demonstrate in an airport-building example, cooperative games present the problem of how to distribute that extra payoff. One well-known solution involves the “Shapley value.” Three companies each want a private airport for their company planes. Shortline’s plane requires only a short runway, Medway’s needs a medium runway, and Longfellow’s plane needs a long runway. An airport with only a short runway costs \$90,000 to build. A medium runway costs \$150,000 and a long runway costs \$240,000.

Clearly, the three companies should share an airport, but how much should each company pay? Having each company pay \$80,000, for example, doesn’t work. Shortline and Medway would be better off building their own airport with a medium runway for \$150,000. If no player can do better by leaving the grand coalition under a certain allocation, that solution is part of the “core” of the game.

Sometimes, the core includes an infinite number of possible divisions, as in the airport game. Sometimes, there isn’t any core at all. Suppose three people are deciding how to split \$3,000 by majority vote. This game has no core—no matter how you divide the money, there exists another division that is better for two of the players.

One allocation solution takes into account the Shapley value, which is the average value-added resulting from the addition of a given player to a coalition. When we calculate Shapley values, we find that Shortline should pay \$30,000 for the airport, Medway should pay \$60,000, and Longfellow should pay the remaining \$150,000. Put another way, the companies split the cost of a small airport three ways, then Medway and Longfellow split the extra cost of a medium runway, then Longfellow pays the rest. In the division of the \$3,000, the Shapley value suggests \$1,000 per player—a sensible division for a symmetric game.

The Shapley value is computationally intensive, but it has a number of desirable properties. First, it's efficient: The total of the Shapley values of all the players always equals the total size of the pie. Second, if two players contribute the same amount to the coalition, they have identical Shapley values. Finally, a player contributing nothing to the coalition has a Shapley value of 0, and each player gets as much or more from the Shapley value as he or she would have gotten if acting independently of a coalition.

Shapley values are particularly relevant in legislative bodies, where the idea is usually called the “power index.” When a balance of power exists between large voting blocs, minor blocs may gain substantial voting power.

John Nash outlined the “Nash cooperative bargaining solution” for games with two players. Nash bargaining solutions use “best alternatives to negotiated agreements”—BATNAs—in order to gauge a fair division. Your BATNA is what you walk away with if the negotiation fails.

Nash outlined the conditions a sensible allocation should satisfy. First, a good allocation should be efficient—it should give the entire surplus to the players. Additionally, the correct division has to be independent of units. We should get the same answer whether we use dollars, cents, or pounds. And just as Arrow did with voting, Nash, too, required independence from irrelevant alternatives.

Under the Nash cooperative bargaining solution, the players should all get their BATNAs and the surplus should be divided evenly among the players. Nash's solution coincides with the Shapley value, assuming equal bargaining

power. If bargaining powers are not equal, the surplus is divided in the same ratio as the bargaining powers.

We can model bargaining with a noncooperative sequential game of perfect information. Although a bargaining game could go on forever, delaying the final agreement carries a cost to both players. Suppose you manufacture plastic jugs. You need a supply of plastic beads to run your factory, and your usual supplier isn't available. I represent another supplier of plastic beads. You call me the first day with an offer for three days' worth of beads. I either accept and send you the beads or call tomorrow with a counteroffer on two days' worth of beads and so on. Each day without beads causes you to lose \$11,000 in revenue, but a day's worth of beads costs me only \$1,000. In this example, the best solution is for you to offer me \$13,000 for three days' worth of beads in the first round of negotiations and for me to accept your offer immediately.

For any “shrinking-pie” bargaining game, you can calculate each player's surplus by imagining the game played with refusals at every step until no surplus remains. Add up all surplus destroyed by one player's refusals in this hypothetical game. In our problem, each day's worth of beads has a \$10,000 surplus. If I refuse your first offer on day one, I destroy one day's worth of surplus, \$10,000. If you refuse my counteroffer on day two, you destroy \$10,000 worth of surplus, and if I refuse your final offer, I destroy another \$10,000. I destroyed \$20,000 in surplus through this haggling, and you destroyed \$10,000. Thus, in the actual first-round deal, you should be awarded \$20,000 of surplus, and I should receive \$10,000.

**The practical
upshot of this
analysis is ...
impatience in
bargaining has
serious costs.**

Another type of diminishing returns comes from impatience. Let's look at this problem by determining optimal offers in the division of \$100. Mathematically, we quantify impatience by saying that \$1 now is as good as $(1 + r)$ in the next round. For purposes of this lecture, we'll assume \$1 now is worth $(1 + r)$ to me in the next round and $(1 + s)$ to you in the next round. For you to accept my proposal that I keep \$A out of \$100, I need to make sure that the value of my offer to you now is at least as great as

what you'll get from your next-round offer. The symmetric condition holds for you. Solving these two equations simultaneously for A and B gives the optimal offers. If r and s are relatively small, going first doesn't make much of a difference. Discrepancies in r and s can have a significant impact on the game. In real-world situations, such as real estate deals and international negotiations, impatience has serious costs. ■

Suggested Reading

Dixit and Skeath, *Games of Strategy*.

Muthoo, *Bargaining Theory with Applications*.

Questions to Consider

1. Compare and contrast our analysis of bargaining in this lecture with our analysis of repeated games in Lecture 15.
2. The baseball players' strike in 1980 took place in two parts: spring training and late in the season. Analyze this choice by the players in terms of BATNAs. You should know that players are paid uniformly throughout the season when not on strike but not paid in spring training. Also, game attendance is highest late in the season.

Game Theory and Business—Co-opetition

Lecture 23

The applications of a single model are often spread over a wide range of disciplines, but in my remaining lectures, I'm going to reverse this approach. I'm going to look at one area of application—business—and see how the concepts of game theory can be applied to address its problems.

In their book *Co-opetition*, Adam Brandenburger and Barry Nalebuff outline an integrated, step-by-step approach for applying game theory to business. They take particular care to point out that businesses have to cooperate to create the pie before competing for the pieces. Brandenburger and Nalebuff adopt the PARTS approach to games. The acronym stands for Players, Added Value, Rules, Tactics, and Scope. Players include all the other businesses in your game and fall into one or more categories: customers, suppliers, competitors, and complementors. Just as it did for the Shapley value, added value measures the increase in the size of the pie when you enter the coalition. For Brandenburger and Nalebuff, tactics means “perceptions” and includes our previous work with credibility, signaling, and screening. Scope is a reminder that multiple games may be part of larger game. In particular, if any player thinks two games are linked, then in a sense, they are.

Brandenburger and Nalebuff introduce a useful way to visualize a business game: the Value Net. To paraphrase their book *Co-opetition*, a player is your competitor if customers value your product less when they have the other player's. A player is your complementor if customers value your product more when they have the other player's product. Competitors and complementors don't have to be in the same industry. Videoconferencing software competes with commercial airlines; red wine complements dry cleaning. Competitors and complementors also occur on the supply side. In an information economy with low variable costs, supply-side complementors should become increasingly common.

Brandenburger and Nalebuff also emphasize the folly of pigeonholing other firms into one of the four categories and discourage businesses from focusing too much on any one part of the Value Net. Video rental outlets are both competitors and complementors for movie studios. This is because, although movie rentals diminish the incentive for viewers to go to the theater, they

Brandenburger and Nalebuff list seven ways to get paid to play besides cash. ... It's doubtful that all of these are appropriate for your business, but some of them are likely to be.

bring in extra revenue in rentals, especially for lesser-known titles. Today, video sales and rentals generate more money than theatrical releases. Businesses can find complementary opportunities with other players, even competitors. Early automakers banded together to produce the first stretches of a transcontinental highway, a complementary good for all their products.

Added value, the amount you can add to the coalition, largely determines how much of the pie you get to keep. We can explore some common misconceptions about added value in the context of two card games. In

the first card game, Brandenburger gives each of his 26 students a red card from the deck and keeps the 26 black cards for himself. He announces that a third party will pay \$100 for each red/black pair. Most people assume that Brandenburger has all the bargaining power, but each of the students has \$100 added value from his or her red card. Brandenburger has \$2,600 in added value: exactly half of the whole.

Nalebuff plays the same game but removes three of the black cards, leaving himself with only 23. Although Nalebuff seems to have short-changed himself, each student's value added is now zero. Nalebuff can always negotiate with three other students. By the time he has one card left, he can essentially choose among four bidders.

Nintendo, after only eight years in business, used a strategy similar to Nalebuff's to attain a market value higher than that of Nissan or Sony. The company installed a security chip in its systems to ensure complete control over software development. The firm also strictly limited the number of

games it would approve from any one source. By limiting the value added by suppliers and complementors, Nintendo was free to use Nalebuff's approach to keep the lion's share of the profits. With a limited supply, none of the customers could hold out for a better deal.

The mere fact that you enter a game can change the strategies and payoffs for other players. Competition itself can add value to the marketplace. For example, Holland Sweetener decided to battle Monsanto's NutraSweet product for entry into the aspartame market. Holland and Monsanto fought a fierce battle in the European market. Then, just as the patent for aspartame was about to expire, Coke and Pepsi both announced new long-term contracts with Monsanto. Though neither Coke nor Pepsi wanted to change suppliers, they saved hundreds of millions of dollars in their contracts with Monsanto thanks to Holland Sweetener's competition. Brandenburger and Nalebuff suggest that companies be "paid to play" in circumstances where their competition benefits other players. ■

Suggested Reading

Brandenburger and Nalebuff, *Co-opetition*.

Dixit and Nalebuff, *Thinking Strategically*.

Scheff, *Game Over*.

Questions to Consider

1. Students in the red/black card game had to make individual deals. How would the situation change if they could bargain collectively?
2. Create the Value Net for your own business. Think with special care about complementors; they are opportunities that are often overlooked.

All the World's a Game

Lecture 24

“Rules.” Every game has them, and every business deal does, too. Brandenburger and Nalebuff point out that the rules of business are often viewed as being set in stone. To some extent, that's true. ... But a number of apparently minor rule changes, generally in the “details” of contracts, can have a powerful impact on shaping your business relationships.

Co-opetition, the practical application of game theory to real-life business decision making, is built on the PARTS model of business opportunities: Players, Added Value, Rules, Tactics, and Scope. This lecture focuses on the last three of these components.

The third component of Brandenburger and Nalebuff's PARTS formula is Rules. As we've seen with other applications of game theory, apparently minor rule changes can affect your payoff in surprising ways.

Although “most-favored customer” (MFC) provisions guarantee a customer a good price, they also offer hidden benefits to the seller. An MFC clause means that you agree to charge the customer the best price for your product that you give anyone. Having MFC clauses with many customers allows a supplier to make a credible threat: “I'd rather lose your business than drop my prices for everyone.”

MFC clauses can also make customers less aggressive in negotiation. As a customer, I gain the same benefit of lower prices if someone else does the fighting; thus, I have less incentive to do it myself. In this way, MFC clauses create a volunteer's dilemma for customers and can actually allow a supplier to keep a greater amount of the surplus.

A “meet-the-competition” (MCC) clause gives a supplier valuable information at virtually no cost. According to this clause, if a customer finds a lower price, you get a chance to meet it and keep the customer's business

before he or she switches away from you. Even if you choose to lose the business, you still know exactly what the other offer was. Bidding for a competitor's business has numerous hidden costs, including time and effort. MCCs add yet another hidden cost: If you meet the competitor's price for your customer, the effort of making the bid is lost.

“Low-price guarantees”—if you buy something from me and find a better price elsewhere, I’ll refund the difference—actually allow suppliers to charge higher prices. There’s a cost associated with shopping around for a product you’ve already bought, and most people aren’t willing to pay it.

For Brandenburger and Nalebuff, the category of Tactics relates to shaping perceptions. We already discussed one important way to shape perceptions in our study of credibility and strategic moves. Brandenburger and Nalebuff also use the metaphor of the “fog of business” to describe pieces of information that other players suspect but don’t know for certain. Depending on the situation, you may want to either lift the fog or preserve it. Lifting the fog corresponds to our work with signaling and screening in previous lectures. Preserving the fog is a kind of signal jamming.

The last component of PARTS, Scope, reminds us that most games are linked to other games. When SEGA entered the video-game market with the 16-bit Genesis machine, Nintendo refrained from releasing its own 16-bit system to maximize the return on its effective monopoly in the 8-bit market. When Nintendo finally produced a 16-bit system, the fierce competition removed much of the added value from the 8-bit market.

Many of the games we’ve studied are meaningful because they provide parables for more complicated situations and teach us practical lessons. For example, in order to play a game successfully, you need to be allocentric—you have to see the world from other players’ perspectives. It’s important to see what the other players want, not what you would want in their place. It’s equally important to understand what game other players think they’re playing, as well as how they see you. Unless you have a dominant strategy, then without an allocentric perspective, you lose the ability to encourage the kind of behavior you want in others. Keep in mind that players can behave rationally even if they aren’t conforming to your payoff structure.

The assumption of rationality in game theory sometimes goes too far. Bounded rationality studies seem to suggest that people tend to reason two or three steps before stopping, especially in new situations.

Threats, promises, and commitments are powerful tools, but only if you can make them credible. Strategic moves will force you to do something you don't want to do. One way to make such moves credible is to sacrifice your flexibility so that you'll have to follow through. It's never in your best interest to let someone threaten you.

Cooperation is perhaps the most interesting issue in this course. It's certainly one of the most mystifying. Given the antagonism between the United States and the U.S.S.R. in the late 1940s and early 1950s, game theory legend John von Neumann advocated engaging the Soviets in nuclear war as soon as possible. He believed that such a war was inevitable and that the United States should act before the U.S.S.R. gained strength. Fortunately, the nuclear war between the United States and the U.S.S.R. never came to pass. It will be the work of a new generation of game theorists—perhaps behavioral game theorists—to find good models to explain how we avert such disasters.

Robert Axelrod's work with the prisoner's dilemma provided four clues for eliciting cooperation from others. First, you must be nice—don't be the first to betray another. Second, you must be capable of being provoked—be ready to punish those who betray you and, perhaps, even those who betray others. Third, you must be forgiving—once you have punished a betrayal, be ready to work with other players again. And fourth, you must be straightforward—let others know how the actions they take will influence your actions. ■

If we play the game right, the world of game theory can make this world—the one that we all share—a lot nicer place to live.

Suggested Reading

Brandenburger and Nalebuff, *Co-opetition*.

Dixit and Nalebuff, *Thinking Strategically*.

Questions to Consider

1. Another rule in a contract, generally one with a commodity supplier, is the take-or-pay provision. It's usually used with commodity suppliers that have high fixed costs. As an example, you agree to buy 100 units from a supplier at \$80 per unit. If you buy less, you still have to pay \$60 for each unit not bought. How does this provision, which essentially turns \$60 of variable costs into a fixed cost for you, discourage your competition from stealing your business?
2. Play some of the games you've seen in this course, such as the ultimatum and dictator games, with various people under various conditions. See how the results compare to what game theory predicts, what you expect, and the experimental results discussed in Lecture 16.

Timeline

Timeline

0–500 The Babylonian Talmud suggests ways of dividing properties for marriages and bankruptcies. One confusing prescription was shown by R. J. Aumann and M. Maschler (1985) to correspond to the cooperative game theory concept of the nucleous.

1838 Augustin Cournot publishes *Researches into the Mathematical Principles of the Theory of Wealth*, applying the Nash equilibrium to a specific case of duopoly.

1913 Ernst Zermelo publishes the first “theorem” of game theory, stating that in chess, white can force a win, black can force a win, or both sides can force at least a draw.

1928 John von Neumann proves the minimax theorem: If mixed strategies are permitted, every two-person zero-sum game has a unique individually rational payoff vector. The game may have more than one equilibrium solution, but they all give the same payoff vector; John von Neumann introduces the extensive (tree) form for games.

1944 John von Neumann and Oskar Morgenstern publish the *Theory of Games and Economic Behavior*. The work introduces cooperative games and coalitions, in addition to elaborating

on two-person zero-sum games. The axiomatic utility theory introduced in this book is widely adopted in economics.

1950.....John McDonald publishes the first general introduction to game theory, *Strategy in Poker, Business, and War*; Melvin Dresher and Merrill Flood develop the prisoner's dilemma at the RAND Corporation. A. W. Tucker popularizes the "prisoner" version of the game.

1950–1953.....In a series of four papers, John Nash introduces the Nash equilibrium, proves its existence for noncooperative games, and proposes new methods for studying cooperative games. He also contributes to bargaining theory (including the Nash cooperative equilibrium) during this time.

1952.....The first experimental game theory conference is held in Santa Monica, California, sponsored by the Ford Foundation and the University of Michigan; the first textbook on game theory, *Introduction to the Theory of Games* by John Charles McKinsey, is published.

1952–1953.....Lloyd Shapely develops the Core as a solution to coalition games. He also introduces the Shapely value, a solution concept that assigns to each member of a coalition game a payoff commensurate with that player's contribution to the power of the coalitions of which he or she is a part.

1953	In <i>Extensive Games and the Problem of Information</i> , H. W. Kuhn develops the representation of extensive-form games and information sets still in use today.
1954	Shapley and Martin Shubik begin to explore game-theoretic implications in political science by examining voting power in a committee system. In this context, the Shapley value is the power index of an individual.
1955	R. B. Braithwaite presents one of the first applications of game theory to the field of philosophy in his <i>Theory of Games as a Tool for the Moral Philosopher</i> .
1959	Aumann introduces the concept of strong equilibria in coalitional games. A strong equilibrium is a strategy profile in which any coalition deviating from the specified profile is strictly worse off; Shubik discovers the relation between Shapley's core and the contract curve in economics. He also is the first to take a purely noncooperative view of oligopoly games; Much work is done in iterated games. Many "folk theorems" are developed. For example, let $p = (p_1, p_2, \dots, p_n)$ be any convex combination of payoff vectors possible in the stage game, and let (e_1, e_2, \dots, e_n) be the payoffs from some Nash equilibrium of the stage game. Then, as long as $p_i \geq e_i$ for each i , there exists a subgame-perfect Nash equilibrium of the infinitely iterated game in which the expected payoff to each player is given

by p , provided the discount rate (δ) is sufficiently close to 1.

1959–1960 Cooperative game theory becomes more coherent as NTU (nontransferable utility) games begin to be explored.

1960 Thomas Schelling publishes *The Strategy of Conflict*. Among many important ideas, he proposes a focal point (now called the Schelling point) as a method of choosing among equilibria.

1961 R. C. Lewontin applies game-theoretic equilibria to evolutionary biology for the first time; Aumann extends the idea of the Core to nontransferable utility games.

1964 Aumann and Maschler introduce bargaining sets, a more lenient variation of the Core that is always non-empty.

1965 Reinhard Selten introduces subgame-perfect equilibria for iterated and extensive games; kernels of cooperative games (subsets of bargaining sets) are developed.

1966 Aumann and Maschler apply game theory to the Cold War in their paper “Game-Theoretic Aspects of Gradual Disarmament” and invent infinite games of incomplete information in the process.

1966–1968 John Harsanyi gives the modern definition of cooperative and noncooperative games and constructs the theory for games of incomplete

	information. A game is cooperative if and only if agreements, threats, and contracts are completely enforceable.
1968	William Lucas develops a game with no stable set solution.
1972	The <i>International Journal of Game Theory</i> is founded by Oskar Morgenstern; evolutionarily stable strategies (ESSs) are developed by John Maynard Smith and soon begin to find applications in biology and economics.
1974	Aumann and Shapely examine large games in which individual players do not matter, also known as non-atomic games. Such games are important in economics.
1974–1975	Correlated equilibria and trembling-hand perfect equilibria are introduced.
1976	Aumann formalizes the idea of common knowledge in game theory.
1981	Forward induction begins to find use as a tool for solving games; Aumann introduces the idea of players in repeated games as automata, opening a lively new field of research.
1982	David Kreps and Robert Wilson generalize subgame-perfect equilibria into sequential equilibria for games that begin at a node with imperfect information.
1984	The concept of rationalizability is introduced by B. D. Bernheim and D. G. Pearce; Robert Axelrod publishes <i>The Evolution of Cooperation</i> .

1985–1986 Aumann’s idea of players as automata is used by A. Neyman and A. Rubinstein to develop the idea of bounded rationality.

1986 Elon Kohlberg and Jean-François Mertens introduce refinements to the notion of Nash equilibrium for normal form games. Until this point, such refinements focused on the extensive form.

1988 In *A General Theory of Equilibrium Selection in Games*, Harsanyi and Selten develop a set of criteria to choose among multiple possible equilibria for both cooperative and noncooperative games; game theorists begin to formally discuss the underlying assumptions of Nash equilibria and rationalizability.

1991 D. Fudenberg and J. Tirole present an early discussion of perfect Bayesian equilibrium.

1994 Game theorists John Nash, John Harsanyi, and Reinhard Selten win the Nobel Prize in Economics for “their pioneering analysis of equilibria in the theory of non-cooperative games.”

2005 Aumann and Schelling win the Nobel Prize in Economics for “having enhanced our understanding of conflict and cooperation through game-theory analysis.”

2007 Roger Myerson, Eric Maskin, and Leonid Hurwitz win the Nobel Prize in Economics for “having laid the foundations of mechanism design”

theory.” Mechanism design seeks to add a mechanism to a game so that the outcome of the game is socially desirable, even though each player acts only in his or her own self-interest.

2008.....Development of game theory continues. Much work is being done in evolutionary game theory, mechanism design, rationalizability, repeated games, and cooperative game theory.

Glossary

added value: (The value of the game with you in it) – (the value of the game without you). The greater your added value, the more of that value you can claim.

approval voting: A voting system in which voters may vote for all candidates they find acceptable.

auction, Dutch: A multiple-round, highest-price auction. Prices start high on a “price clock,” then descend until one of the bidders “stops the clock” and pays that price.

auction, English: A multiple-round, highest-price auction. The most familiar kind of auction.

auction, Vickrey: Another name for a sealed-bid, second-price auction.

BATNA: Best alternative to a negotiated agreement. A player’s “disagreement value” in a bargaining game. If no deal is struck, a player still receives his or her BATNA.

battle of the sexes: Any 2×2 nonzero-sum game with the ordinal payoff matrix shown below. The game has two equilibria in pure strategies, both of which are Pareto-optimal.

	Left	Right
Left	3, 4	< 3, < 3
Right	< 3, < 3	4, 3

best response: Simply, a pure strategy that gives a maximum payoff given the strategy choices of the other players. A strategy profile in pure strategies is a Nash equilibrium if and only if each player is playing a best response to the strategies chosen by the others. Identifying best responses is a good way to identify pure Nash equilibria in strategic-form games.

Borda count: A voting system in which a candidate gets the most points for being a voter's first choice, fewer points for being the voter's second choice, and so on. The candidate with the most total points wins.

bounded rationality: An idea developed by Selten and others, rooted in the notion that people have only finite powers of computation, memory, and information processing. Individuals often do not "catch on" until near the end of an interaction.

brinkmanship: The strategic loss of control to pose a probabilistic threat in a game when certain threats are too costly to be credible.

chicken: Any 2×2 nonzero-sum game whose ordinal payoffs correspond to the payoff matrix below. The name comes from two cars approaching each other at high speed in a contest to see which driver will be the first to swerve.

	Straight	Swerve
Straight	1, 1	4, 2
Swerve	2, 4	3, 3

common knowledge: A piece of information X is common knowledge if all players know X , all players know that all players know X , all players know that all players know that all players know X , and so on. Most game theory assumes common knowledge of the game structure and the rationality of all players.

complementor: In business, another business is your complementor if customers value your product *more* when they also have the complementor's product.

Condorcet winner: A candidate who wins all head-to-head votes against all other candidates.

coordination game, pure: Any 2×2 nonzero-sum game with the ordinal payoff matrix shown below. The game has two equilibria in pure strategies, but (first, first) is preferred by both players.

	First	Second
First	4, 4	< 3, < 3
Second	< 3, < 3	3, 3

credible threat, promise, or commitment: By definition, a threat, promise, or commitment claim that under certain circumstances, a player will make a choice that will not give that player his or her best payoff. Such a claim is not believable, or credible, in a rational player unless the player changes the game by either changing his own payoffs or restricting his own future choices.

discount rate (δ): A multiplicative factor used in iterated games. The assumption is that a payoff of δ in the current-stage game is equivalent to a payoff of 1 in the next stage. The value of δ is less than 1, and the larger it is, the more patient the players are. If r is the rate of return (such as interest rate), then $\delta = 1/(1 + r)$.

domination: Strategy A (strongly) dominates strategy B if the payoff for A is better than that for B, regardless of the choices made by the other players. If strategy A is sometimes strictly better than strategy B and sometimes equally good, then A *weakly* dominates B. If A dominates B, then B is dominated by A. Dominated strategies are never used in a Nash equilibrium; weakly dominated strategies may be.

duopoly: A market involving only two suppliers.

equilibrium, dominance-solvable: An equilibrium that can be found by the iterated elimination of strictly dominated strategies.

equilibrium, Nash: A strategy profile in which no player can get a better expected payoff by unilaterally changing his or her strategy. Nash equilibrium is the foundation of the solutions found in almost all noncooperative game theory.

equilibrium, payoff-dominant: An equilibrium that gives a higher payoff to every player than any other equilibrium.

equilibrium, pooling: In a signaling game, an equilibrium in which different player types all respond in the same way and, hence, cannot be distinguished from one another.

equilibrium, rollback: The equilibrium obtained by evaluating a game tree node by node from its end to its beginning. The equilibrium rollback is guaranteed to be a subgame-perfect Nash equilibrium.

equilibrium, semiseparating: In a signaling game, an equilibrium in which one choice made by a player will reveal his or her type, while another choice could be made by multiple types.

equilibrium, separating: In a signaling game, an equilibrium in which different player types distinguish themselves by their choices.

equilibrium, subgame-perfect: An equilibrium that, when restricted to a subgame of the original game, is an equilibrium of that subgame. Equilibria that include senseless choices on off-equilibrium path nodes are not generally subgame-perfect.

equilibrium, trembling-hand perfect: An equilibrium in which no player would change strategies if there were an infinitesimal chance of another player accidentally “trembling” and deviating from his or her equilibrium strategy. For two-player games, this is equivalent to both players playing admissible strategies.

expected value: A term synonymous with “mean” or “average.” Note that your expected payoff is not the payoff you get most often but, rather, the average payoff you would expect by playing your strategy in many, many trials. Expected value is needed when games or strategies involve random chance.

extensive form: The game tree representation of a game, most commonly used with sequential games.

game, constant-sum or zero-sum: A game in which the payoffs received by all players always total to a constant sum. It follows that for some player to do better, another player must do worse. Mathematically, no important difference results if the sum is an arbitrary constant or 0.

game, deterministic: A game that involves no chance events, in contrast to a stochastic game.

game, finite: A game with a finite number of players, each of whom gets only a finite number of moves and has only a finite number of options at each move. Infinitely repeated games are not finite, nor are games allowing mixed strategies because the probabilities allow an infinite number of choices of strategy. Finite games require only ordinal (order-of-preference) payoffs.

game, noncooperative: A game in which binding agreements between players are not possible. This does not mean that the players cannot act for their mutual advantage, merely that nothing constrains their choices. In cooperative games, binding agreements are possible.

game, repeated or iterated: A game consisting of a “stage game” that is played consecutively two or more times, with the choices made in one round of the stage game being known to all players before the next round is played. Also called a “supergame.”

game, sequential-move or dynamic: Players’ choices may be contingent upon earlier decisions made by the players. These games are most often represented by game trees.

game, simultaneous move or static: A game in which players make decisions concurrently, without knowledge of the choice made by the other player. These games are most commonly represented in strategic form.

grim trigger: A strategy for two-player iterated games in which a player initially cooperates but replies to any defection by the other player with an unending series of defect responses. This is the least forgiving of all retaliation strategies.

IEDS: Iterated elimination of dominated strategies. A rational player should not play a dominated strategy. Removing a dominated strategy from the game may result in one or more remaining strategies now being dominated. IEDS continues this process as far as possible. One can use IEDS to eliminate weakly dominated strategies, too, but doing so may eliminate some Nash equilibria.

imperfect information: A player is unaware of the outcome of some earlier chance event or of some contemporaneous decision of another player. All simultaneous games are games of imperfect information.

incomplete information: Also called “asymmetric information.” Information is known to one player but not to the other. Such information often involves the payoffs or choices of a player. Games of incomplete information can be transformed into games of imperfect information.

independence of irrelevant alternatives: The concept that a best solution should not change when some inferior alternatives are removed from consideration. This idea plays a role in both Arrow’s impossibility theorem and the Nash bargaining solution. Some consider it to be too strong a requirement for voting or bargaining.

information set: A collection of nodes in a game tree controlled by a given player that are indistinguishable to that player when any of the nodes in the set is reached. In a game of perfect information, each node is its own information set. Nontrivial information sets generally arise when some information in the game is hidden from the active player.

instant runoff voting (IRV): Also called the “Hare vote.” Voters rank all candidates in order of preference. If no one receives a majority of votes as number 1, the candidate receiving the fewest number-1 votes is eliminated, and voters who had chosen this candidate as their favorite now vote for their second choice. This process continues until some candidate gets a majority.

minimax criterion: In a two-player zero-sum game, maximizing your payoff is equivalent to minimizing the payoff of your opponent. One can define one’s optimal strategy as that which makes the maximum payoff available to one’s opponent as small as possible.

Nash bargaining solution: The solution concept that a surplus should be divided between two players so that their individual surpluses above their BATNAs are in the same proportion as their bargaining power.

Pareto-optimal or efficient: A strategy profile is Pareto-optimal (or efficient) if any other solution that gives one player a higher payoff necessarily gives a different player a lower payoff.

PARTS: Players, added value, rules, tactics, and scope.” The analysis framework used by Brandenburger and Nalebuff in *Co-opetition*.

payoff: The “return” received by a player when a game is complete; higher payoffs are generally assumed to be better.

payoffs, cardinal: A system of payoffs on an interval scale, so that a payoff of 4 is preferred to a payoff of 1 to the same extent that a payoff of 7 is preferred to a payoff of 4. Cardinal payoffs are needed for games allowing mixed strategies or chance nodes.

payoffs, ordinal: A system of payoffs requiring only that if option A is preferred to option B, then A has a higher payoff than B. The magnitude of the payoffs is irrelevant, only their ordering. Ordinal payoffs are sufficient for finite games.

plurality voting: Whichever candidate receives the most votes wins.

present value: The value of a future stream of payoffs, expressed in current payoff units by discounting future payoffs.

prisoner's dilemma: Any 2×2 nonzero-sum game with the ordinal payoff structure shown below. The dominant equilibrium is for both players to defect, but this strategy is not Pareto-optimal.

	Cooperate	Defect
Cooperate	3, 3	1, 4
Defect	4, 1	2, 2

rational: The notion that one's choices are made so as to result in the best possible expected payoff given one's knowledge of the situation at the time the choices were made. Most game theory assumes that players are rational, but this idea has been revisited by Selten and others.

rationalizable: A strategy is rationalizable if it survives the iterated elimination of strategies that are never (weak) best responses. In two-player games, this is equivalent to saying that the strategy survives the iterated elimination of strictly dominated strategies. Strategies can be rationalizable without being part of a Nash equilibrium.

Schelling point or focal point: An equilibrium that, for some reason, is the “natural” one for players to choose among the possible equilibria in the game, such as all drivers using the right side of the road. These equilibria are often culture-specific.

screening: Eliciting from another player a piece of information known to that player but not to you.

signaling: Credibly communicating to another player a piece of information known to you.

signal-jamming: Preventing another player from either signaling or screening.

social dilemma: A game (often involving many players) in which the equilibrium is not Pareto-optimal.

Stag Hunt: Also called the Assurance Game. Any 2×2 nonzero-sum game with the ordinal payoff matrix shown below. Some definitions allow the payoffs labeled 2 and 3 to be equal. The Stag Hunt has two equilibria: (work, work) is payoff dominant and (coast, coast) is risk dominant.

	Work	Coast
Work	4, 4	1, 3
Coast	3, 1	2, 2

strategic form: The matrix representation of a game, also called the “normal form.” For a two-player game, the game is represented by a table. One player picks a row of the table, the other picks a column, and the payoffs are the contents of the cell found by cross-indexing these choices.

strategy, admissible: A strategy that is not weakly or strongly dominated by another strategy.

strategy, dominant: A strategy that outperforms every other strategy for a player, regardless of the choices made by the other players. A strategy is “weakly dominant” if it is *at least as good as* every other strategy for the player, regardless of the choices of the other players.

strategy, dominated: See **domination**.

strategy, mixed: A strategy in which the pure strategy used is randomly selected from the available strategies, with each pure strategy having its own probability of selection.

strategy profile: A collection of strategies, one for each player.

strategy, pure: A specification of the single choice that a player will make in each situation in which she may find herself.

subgame: A game contained within a larger game. To be a subgame, it must start at a single node and include all descendants of that node. Further, if the subgame includes any members of an information set, it must contain all of that set's members.

tit for tat: A strategy for symmetric two-player iterated games in which you treat the other player as he or she treated you in the previous round. Most commonly discussed in relation to the iterated prisoner's dilemma. One generally begins by cooperating.

utility: A measure of the satisfaction obtained from an outcome. Properly, payoffs in games requiring cardinal payoffs should be utility measures to compensate for risk love and risk aversion. For any amount x and probability p , a player will be indifferent between (1) receiving a payoff of xp utils with certainty and (2) receiving either x utils (with probability p) or 0 utils (with probability $1 - p$).

Value Net: Brandenburger and Nalebuff's graphical representation of business relations, including customers, suppliers, competitors, and complementors.

winner's curse: The tendency for the winner of an auction to have overpaid for the common-value item purchased.

Biographical Notes

Kenneth Arrow (1921–): Kenneth Arrow is an American economist and game theorist. He earned his doctorate from Columbia University in 1951 after interrupting his college career to join the Air Force during World War II. Though he has made contributions to many areas of economics, especially those disciplines dealing with social justice in economies, his most famous result is Arrow's impossibility theorem, in which he proved that it is impossible to design a voting system to satisfy a set of fairly straightforward criteria. Many of his students have gone on to become famous game theorists in their own right, including John Harsanyi and Richard Myerson.

Robert Aumann (1930–): Robert Aumann is an Israeli mathematician and economist whose important work in game theory includes his definition of correlated game theory in repeated games, as well as one of the first acknowledgments of common knowledge. Though born in Germany, he grew up in New York City, graduating from City College there in 1950. In his earlier research, he was among the first to focus on games without transferable utility and the first to explore bargaining sets. Aumann won the 2005 Nobel Prize in Economics along with Thomas Schelling for “having enhanced our understanding of conflict and cooperation through game-theory analysis.”

Adam Brandenburger: Adam Brandenburger has published several papers on the application of game theory to various situations in management. He was born and raised in London and attended Queens College at Cambridge before immigrating to the United States. Since then, he has taught as a professor at Harvard University and New York University. He is perhaps best known for his book *Co-opetition* (co-authored with Barry Nalebuff), which further explores the business implications of game theory. In addition to his teaching, Brandenburger has consulted with a number of large firms to do practical work in his field, including Fidelity Investments, IBM, Merck & Co., McKinley & Co., and Xerox.

Marie-Jean-Antoine-Nicolas de Caritat, marquis de Condorcet (1743–1794): Condorcet was a French Enlightenment thinker, political scientist, and mathematician who exposed some of the basic difficulties of voting processes. He was an early proponent of human rights in France, especially those of women and blacks, and was a moderate supporter of the French Revolution. Condorcet's work with voting included Condorcet's paradox, his observation that majority preferences are often not transitive with more than two positions. This insight, in turn, led him to introduce a new criterion for elections, now called the Condorcet winner, who would win in a head-to-head contest against any other candidate. Condorcet died in prison at the age of 51, after having run afoul of the more radical French revolutionaries.

Antoine-Augustin Cournot (1801–1877): Cournot was a French philosopher, mathematician, and economist, as well as one of the first to contribute to what would later become the theory of games. After his primary education, he worked as a clerk in his hometown for four years before beginning more serious studies at Besançon. In addition to his introduction of elasticity in economics, his duopoly model, in which neither firm has any incentive to change its production, stands as one of the first applications of Nash equilibria—more than 100 years before Nash's work. However, Cournot's work was not well reviewed by most of his contemporaries, and in large part, his works passed into obscurity for the rest of the 19th century.

John Harsanyi (1920–2000): John Harsanyi was born in Budapest, Hungary, where he grew up attending the same school as his game-theoretical predecessor, John von Neumann. Despite his mathematical gifts, Harsanyi studied philosophy and sociology. His early work was interrupted twice, first by the Nazi occupation of Hungary during his graduate years, then by his narrow escape from Stalinist Hungary shortly after acquiring his degree. After fleeing to Australia, he began to study economics and statistics, eventually acquiring his second Ph.D. in the former. Harsanyi's most important work in game theory focused on the transformation of games of incomplete information to more manageable games of imperfect information. He has also extended Shapely values for cooperative games and Nash bargaining solutions to new kinds of bargaining problems. He was awarded the 1994 Nobel Prize in Economics along with John Nash and Reinhard Selten.

Barry Nalebuff (1958–): Barry Nalebuff, co-author of *Co-opetition* along with Adam Brandenburger, is the Milton Steinbach Professor of Management at Yale University. He completed his undergraduate studies in economics and mathematics at MIT before completing his Ph.D. at Oxford University. His work encompasses a large number of topics in economics, including the effects of game-theoretic strategies and incentives to optimally run businesses—expertise he also employs as a consultant for such companies as American Express and Citibank. Along with his former student Seth Goldman, Nalebuff is also the cofounder of the Honest Tea Company, a startlingly successful tea firm that trades on the organic nature of its product and its socially responsible business practices. In addition to *Co-opetition*, he has also written another book, *Thinking Strategically* (with Avinash Dixit).

John Forbes Nash, Jr. (1928–): John Nash is perhaps one of the best-known game theorists in the world. His development of the Nash equilibrium (1948–1950) revolutionized the theory of noncooperative games by using best-response arguments to provide an extension of von Neumann and Morgenstern's results to nonzerosum games that can involve some mutual gain or loss for the players. His Nash bargaining solution likewise formed the basis of much of modern bargaining theory. Nash's contributions to mathematics, especially results that facilitated future advances in game theory, have been no less impressive. Despite his battle with paranoid schizophrenia which left him incapacitated and unemployed for much of the 1960s and 1970s, Nash recovered and was awarded the 1994 Nobel Prize in Economics along with John Harsanyi and Reinhard Selten.

Thomas Schelling (1921–): Thomas Schelling is an American game theorist best known for his introduction of the Schelling point (or focal point) and his models of self-sustaining segregation. After graduating from Berkeley in 1944 and receiving his Ph.D. in economics from Harvard in 1951, he served with the Marshall Plan in Europe. He was among the first to open the question of choosing among various equilibria in a game and did pioneering work in developing bargaining theory. Strategic moves—reducing one's payoffs or options in order to increase one's equilibrium payoff—is an idea largely credited to Schelling. In the past few decades, he has also become involved in the global-warming debate, framing it as a bargaining problem

to lower emissions. In 2005, he was awarded the Nobel Prize in Economics along with Robert Aumann.

Reinhard Selten (1930–): Reinhard Selten is a German game theorist, notable for his invention of the subgame-perfect equilibrium and his explorations into bounded rationality. Growing up half-Jewish in 1940s Germany understandably made him attuned to politics during his early years, and this, in conjunction with a love for mathematics, contributed to his fascination with game theory. He published his paper on subgame-perfect equilibria, one of the fundamental concepts of extensive games, in 1962. Thirteen years later, he refined the idea by introducing trembling-hand perfection, as well. In cooperation with John Harsanyi, Selten worked to create bargaining models under imperfect information, in addition to some work on nuclear deterrence during the Cold War. He shared the 1994 Nobel Prize in Economics with Harsanyi and Nash for his work in noncooperative games.

William Vickrey (1914–1996): William Vickrey was born in Victoria, British Columbia, but was educated at Yale and Columbia University. A conscientious objector during World War II, he spent some of his time devising a new inheritance tax system for Puerto Rico. After the war, he toured Japan with a team of economists, recommending reforms of that country's tax system. Vickrey is well known for his proposals on congestion pricing. In such a system, a fee is levied on transportation, with the fee being higher at times of higher congestion. He is also famous for his analysis of the second-price, sealed-bid auction, which is still called the Vickrey auction in his honor. Vickrey received the Nobel Prize in Economics in 1996 and died three days after his selection was announced.

John Louis von Neumann (1903–1957): Born in Budapest, Hungary, John von Neumann was arguably the founding father of modern game theory. As a child, he showed prodigious talents in mathematics, eventually receiving his Ph.D. in 1928. He proved many fundamental results of game theory, including a general minimax solution to 2×2 zero-sum games in 1928 and the use of backward induction in 1937. His work culminated in 1944 with his *Theory of Games and Economic Behavior* with Oskar Morgenstern, which extended the minimax result of zero-sum games to situations with imperfect

information and with more than two players. Apart from his exceptional influence on game theory, von Neumann also made important contributions to logic, set theory, economics, quantum mechanics, and computer science.

Bibliography

Aliprantis, Charalambos D., and Subir K. Chakrabarti. *Games and Decision Making*. Oxford, UK: Oxford University Press, 2000. Somewhat more mathematically rigorous than the sources cited by McCain or Dixit and Skeath, this book is appropriate for the reader who would prefer a presentation that is a bit more technical.

Amy, Douglas J. *Behind the Ballot Box: A Citizen's Guide to Voting Systems*. Westport, CT: Praeger Publishers, 2000. A nice discussion of voting systems used in the United States and Western democracies, along with the advantages and disadvantages of each. As the title suggests, the book is intended for the general reader.

Axelrod, Robert. *The Evolution of Cooperation*. New York: Basic Books, 1984. Axelrod relates the structure and results of his computer tournament, in which submitted programs played one another in a repeated prisoner's dilemma environment. Axelrod makes interesting observations about the characteristics of players who are successful in such environments. He goes on to discuss similar behavior during trench warfare in World War I.

Berensen, Mark, Timothy Krehbiel, and Mark Levine. *Basic Business Statistics*, 10th ed. Upper Saddle River, NJ: Prentice Hall, 2006. A competent summary of the practical application of statistical ideas, focusing on business applications. Many other treatments of basic probability provide similarly lucid explanations.

Brandenburger, Adam M., and Barry J. Nalebuff. *Co-opetition*. New York: Doubleday, 1996. Using the PARTS paradigm and the Value Net to analyze businesses and their relationships with one another, Brandenburger and Nalebuff offer a fresh approach to identifying and using business opportunities.

Bueno de Mesquita, Bruce. *Predicting Politics*. Columbus: Ohio State University Press, 2002. In this interesting book, political scientist Bueno de Mesquita uses policy situations and game theory to analyze historical situations, such as the Cold War, as well as the future of Russia and China.

Camerer, Conlin, and Richard H. Thaler. "Anomalies: Ultimatums, Dictators, and Manners." *Journal of Economic Perspectives* 9, no. 2 (Spring 1995): 209–219. An excellent article summarizing research on the differences between Nash equilibrium play and observed play in some very simple games. The results suggest that a large number of factors interact in determining how "fair" a player chooses to be with another.

Cramton, Peter. "Spectrum Auctions." In *Handbook of Telecommunications Economics*, edited by Martin Cave, Sumit Majumdar, and Ingo Vogelsang, pp. 605–639. Amsterdam: Elsevier Science B.V., 2002. A nice discussion of spectrum auctions: the history of spectrum sales, the advent of the auction structure, and the results of its application.

Dardanoni, Valentino. "A Pedagogical Proof of Arrow's Impossibility Theorem." *Social Choice and Welfare* 18, no. 1 (2001): 107–112. A simple proof of Arrow's impossibility theorem when restricted to three voters with preferences over three alternatives. The principles used could be expanded to a more general proof.

Davis, Morton. *Game Theory: A Nontechnical Introduction*. Mineola, NY: Courier/Dover, 1997. One of the best introductory texts in game theory, requiring no background to be accessible. In avoiding mathematical rigor and proof, it may not suit the more technically inclined.

Dixit, Avinash, and Barry Nalebuff. *Thinking Strategically*. New York: Norton, 1991. An excellent and entertaining application of the ideas of game theory to a wide range of topics. This highly readable book should appeal to anyone interested in game theory, especially those adopting a business perspective.

Dixit, Avinash, and Susan Skeath. *Games of Strategy*, 2nd ed. New York: W.W. Norton and Company, 2004. A readable and well-organized text covering many topics in game theory at a solid introductory level. Highly recommended.

Fink, Evelyn C., Scott Gates, and Brian D. Humes. *Game Theory Topics: Incomplete Information, Repeated Games, and N-Player Games*. Thousand Oaks, CA: Sage Publications, 1998. An example-based approach that covers intermediate-level topics in game theory, particularly games of imperfect information, repeated games, and n -person games. Moderate mathematical rigor.

Fursenko, Aleksandr, and Timothy Naftali. *One Hell of a Gamble: The Secret History of the Cuban Missile Crisis*. New York: Norton, 1997. This book includes information from Khrushchev's own files, as well as those of his leadership circle and the KGB. As a result, it offers important insights into events on the Soviet side during the Cuban Missile Crisis.

Gibbons, Robert. *Game Theory for Applied Economists*. Princeton: Princeton University Press, 1992. A good look at many of the topics discussed in this course from an economic perspective. The text is not mathematically intimidating.

Gintis, Herbert. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Behavior*. Princeton: Princeton University Press, 2000. For the mathematically capable, this book uses a hands-on approach to the field of game theory from the evolutionary perspective. A good text for those seeking a clearer picture of the mathematics behind game theory.

Green, Kesten, and J. Scott Armstrong. "The War in Iraq: Should We Have Expected Better Forecasts?" *Foresight: The International Journal of Applied Forecasting* 2 (2005): 50–52. Green and Armstrong's research suggests that game theory is not especially useful in predicting geopolitical outcomes. Structured analogies and role-playing seem to have greater utility.

Grinstead, Charles, and J. Laurie Snell. *Introduction to Probability*. Providence, RI: American Mathematical Society, 1997. An amazing book for

anyone who wants to develop a solid intuitive understanding of probability theory. Frequent use of computer simulations and interesting problems will excite the thoughtful reader.

Harding, Garrett. "The Tragedy of the Commons." *Science* 162 (December 1968): 1243–1248. Harding's original, influential article on how egoistic choices can result in a disastrous outcome for the society as a whole.

Hartford, Timothy. *The Undercover Economist*. New York: Oxford University Press, 2006. Hartford discusses many everyday experiences in terms of economic theory. Engaging, informative, and fun.

Hild, Matthias, and Tim Læsæter. "Reinhard Selten: The Thought Leader Interview." *strategy+business* (Summer 2005). Also available at <http://www.strategy-business.com/press/16635507/05209> (accessed March 30, 2008). An excellent, short interview with one of the great minds of game theory. Selten's opinions on rationality and the need for the study of game theory are thought-provoking.

Kretschmer, Martin. "Game Theory: The Developer's Dilemma, Boeing vs. Airbus." *strategy+business* (2nd quarter 1998). Also available at <http://www.strategy-business.com/press/16635507/15872> (accessed March 30, 2008). This article discusses the Boeing versus Airbus competition in the superjumbo jet market from a game-theoretic perspective.

Kuhn, Harold W. *Lectures on the Theory of Games*. Princeton and Oxford: Princeton University Press, 2003. Based on a Princeton University lecture course from the 1950s, this book builds basic game theory from the ground up. It focuses strongly on relationships among game theory and other branches of mathematics, especially geometrical interpretations of game-theoretic results. Mathematically sophisticated.

Lowenstein, Roger. "Exuberance Is Rational." *New York Times*, February 11, 2001. An interesting article on Richard Thaler and his ideas concerning bounded rationality.

Mailath, George J., and Larry Samuelson. *Repeated Games and Reputations*. Oxford, UK: Oxford University Press, 2006. Assuming a basic knowledge of game theory, this book explores the effect of repetition, various types of monitoring, and players' resulting reputations on (chiefly economic) games. Comprehensive and clearly written for those comfortable with mathematics.

McCain, Roger A. *Game Theory: A Non-Technical Introduction to the Analysis of Strategy*. Mason, OH: Thompson South-Western, 2004. An excellent nontechnical introduction to game theory, including a large number of interesting examples. McCain provides good sections on n -person games and cooperative games.

Mehlmann, Alexander. *The Game's Afoot! Game Theory in Myth and Paradox*. Translated by David Kramer. Providence, RI: American Mathematical Association, 2000. A unique text, drawing game-theoretic concepts out of literature, myth, and other traditional stories. Mehlmann includes a quick foray into evolutionary game theory.

Muthoo, Abinhay. *Bargaining Theory with Applications*. Cambridge, UK: Cambridge University Press, 1999. This text develops a theory of bargaining, starting with Nash's bargaining solution, and proceeds to examine the effects of bargaining procedures, commitment tactics, and asymmetric information on that model. The last chapters deal with repeated bargaining, the shortcomings of existing theories, and possible avenues of research in the future.

Olson, Mancur. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge: Harvard University Press, 2000. A classic in the field of collective action, originally published in 1965. Mancur argues that individual incentives do not generally lead to maximal social good and that a minority sharing selective incentives can dominate a majority.

Osborne, Martin J., and Ariel Rubenstein. *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994. An excellent jumping-off point for

the more theoretical areas of game theory. Appropriate for advanced undergraduates or graduates interested in the field.

Palacios-Huerta, Ignacio. “Professionals Play Minimax.” *Review of Economic Studies* 70, no. 2 (2003): 395–415. The journal article upon which our discussion of mixed strategies in zero-sum games is based. The article goes on to show that play shows no pattern over time, consistent with an optimal mixed strategy.

Poundstone, William. *Gaming the Vote*. New York: Hill and Wang, 2008. An entertaining and informative book on voting systems, including fascinating historical examples. Surprisingly engaging.

———. *Prisoner’s Dilemma*. New York: Anchor Books, 1993. An excellent and readable look at game theory from a historical perspective. The book provides fascinating insights into John von Neumann, the RAND Corporation, and the American mindset at the time when the world was first confronting, as Poundstone puts it, “the puzzle of the bomb.”

Rapoport, Anatol, Melvin Guyer, and David Gordon. *The 2 × 2 Game*. Ann Arbor, MI: University of Michigan Press, 1976. An exhaustive analysis of 2×2 simultaneous games, including experimental results conducted both in the United States and abroad.

Scheff, David. *Game Over: How Nintendo Conquered the World*. New York: Vintage Books, 1993. A riveting account of Nintendo’s rise to power in the video-game industry.

Schelling, Thomas. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press, 1980. A landmark book on the subject of strategy that still provides fresh insights. Good discussion of the notion of brinkmanship.

Spence, Michael. “Job Market Signaling.” *The Quarterly Journal of Economics* 87, no. 3 (1973): 355–374. Spence’s classic article on signaling. His model shows how a college education could be useful as a signaling device, even if the degree itself provides no useful skills.

Thaler, Richard. *The Winner's Curse*. Princeton: Princeton University Press, 1994. Best for those with a reasonable background in economics. Thaler looks at “anomalies”—behavior seen in the real world that is inconsistent with the principles of neoclassicism.

von Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944. The seminal work on game theory. Good for those who wish to see the original work but a bit dense to read in places because of its mathematics.

Weibull, Jörgen W. *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1997. The text covers many elements of current evolutionary game theory, including evolutionarily stable strategies, replicator dynamics, and multipopulation dynamics. It assumes knowledge of calculus, topology, and set theory. An advanced and abstract text.

Williams, J. D. *The Compleat Strategyst: Being a Primer on the Theory of Games of Strategy*. New York: Dover Publications, 1966. A fun introduction to the theory of two player zero-sum games. Very accessible and entertaining.