

An Introduction to Games and Their Theory

1.1 WHAT IS A GAME?

In ordinary English, a **game** is any pastime or diversion. This definition covers a lot of terrain. Let's try to organize it somewhat. We can start by identifying different varieties of games. First, there are games played on a board, *board games*. These include Chess, Checkers, and Monopoly.¹ Chess and Checkers have exactly two players, whereas Monopoly can have anywhere from two to eight players. In Chess and Checkers, the outcome of the game is win, lose, or draw: either one player wins and the other loses, or both players draw. In Monopoly, the player with the most assets at the end of the game is the winner; all the rest are losers. Second, there are games played with cards, *card games*. These include Solitaire, Poker, and Blackjack. Solitaire is special, since it has only one player. Poker and Blackjack can have two to seven players—you study them in the next two chapters. Third, there are games played on a video screen against a computer, *video games*. These games are rather like Solitaire, in that they have only one human player. Computers can play many games; later in this chapter we see how and why. Finally, there are games played on a field or a court, *field games*. These include baseball, football, basketball, and hockey, the major professional sports in the United States. Field games are played by two aggregate players, the opposing teams, which are composed of individuals, the members of each team.

All these things are called games, so, according to Aristotle's theory of categories, they must have some feature in common that brings them all under the same name.² Let's look for what this feature, or set of features, might be. First, all games have *rules*. The rules specify what a player can and cannot do. A player who breaks the rules is penalized—again, according to the rules—and in extreme cases can be removed from the game altogether. Second, in every game *strategy matters*. There are good and bad strategies, and players can be and are criticized for choosing bad strategies. One of the tasks of game theory is to tell the difference between good and bad strategies. Third, there is an *outcome* to the game, for example, one player wins and the other loses. Fourth, this outcome depends on the strategies chosen by each of the players, a phenomenon we call **strategic interdependence**. Even

¹ Games have proper names, just as people do, so we capitalize them in this book.

² See Aristotle, *Categories and De Interpretatione* ed. J. L. Ackrill (Oxford: Oxford University Press, 1963).

a bad strategy can win if the opponent chooses a worse one. We combine these features to define a game as *any rule-governed situation with a well-defined outcome, characterized by strategic interdependence*. This describes the Aristotelian category to which games belong.

A lot of things that aren't called games in ordinary English satisfy this definition. Consider firms competing in the same business. There are rules governing their competition, including the law of contracts and property and government regulations, which specify what a firm can and cannot do. A firm that breaks the rules can be penalized, and in extreme cases, such as bankruptcy, can even be removed from the game. The outcome of firms' competition is typically something observable, such as the amount of money each firm makes or the amount of market share each firm has. These outcomes are often reported in a firm's financial statements. As we will see, a firm's strategies can include price, quantity, advertising, which markets it operates in, what kinds of contracts it offers its employees—a host of things. Finally—and this is the most important truth in this book—the *outcome for a given firm depends not just on what strategy it chooses, but also on what strategies its competitors choose*. A firm can have the greatest product line in the world and still get clobbered by the competition if its strategies are no good. When firms compete in a market, they operate in a rule-governed situation, with well-defined outcomes, characterized by strategic interdependence. We have a word for this: it's a *game*, in this case a **market game**. Firms competing in a market are just as much players in a game as are Poker players seated around a poker table. Indeed, the stakes in a business are almost always a lot higher. What began as a metaphor at the beginning of this century has become literally and categorically true.

Let's take an example that works on both sides of the literal and metaphorical aspect of *game*: a professional sports franchise, the Chicago Bulls.³ The Bulls play professional basketball in the National Basketball Association (NBA), and in the 1990s they played it very well, winning six NBA championships. More recently, they have been among the worst teams in the league. The Bulls are also a corporation—owned by private investors, mostly from the Chicago area—that is in business to make money each year. The more games the Bulls win on the basketball court, the more money the owners of the Bulls make from ticket sales, broadcasting rights, and concessions. Although winning at basketball isn't the same thing as winning in business, the two are closely related in this case.

Another example, not at all sports related, is the cola industry. Cola was invented in Atlanta in 1886. The inventor later went on to found Coca-Cola Corporation. Other cola drinks entered the market soon thereafter. Since the 1920s, Coke and Pepsi, with the largest market shares and highest profits, have dominated the U.S. market and have battled head-to-head for cola business. Their competition, commonly known as the cola wars,⁴ has included price wars, new product competition (Pepsi Twist is a recent entry), new packaging (the 2-liter plastic bottle), and advertising campaigns (Coke Is It! The Pepsi Generation). How much money Pepsi makes from its cola operations depends on what Coke does, and vice versa. This is strategic interdependence manifested in the market.

Business isn't the only arena in which we see games in the extended sense—they are at the very heart of economics, too. Think of domestic economic negotiations, such as those between the White House and Congress on domestic economic policy issues, for example, the budget, or between the White House and the Fed on monetary policy. Think of international

³ All the examples in this book are drawn from real life. You will never encounter the word widget, to say nothing of the markets in which these mythological objects are bought and sold.

⁴ The title is from J. C. Louis, *The Cola Wars* (New York: Everest House, 1980).

economic negotiations, such as those among the G-8 countries, involving international trade and finance.⁵ **Economic negotiations** such as these have all the makings of a game. There are rules (domestic and international law) governing what each party to the negotiations can do. Strategy matters—if you are absent from the negotiations, you should not be surprised if your interests are not represented. There is a clearly defined outcome—the status quo if negotiations break down, some sort of agreement if negotiations are brought to a successful conclusion. Finally, the outcome of negotiations depends not only on what your side does, but on what the other side does, too. Various forms of bargaining, negotiation, and arbitration—all games with at least an element of cooperation—are studied in the last third of this book.

1.2 WHAT IS GAME THEORY, AND WHY?

Game theory is the science that studies games and takes games seriously enough to solve them. Game theory is a product of the twentieth century, the brainchild of one of the century's greatest minds, John von Neumann.⁶ Von Neumann discovered one of the central regularities of games, the solution for 2-person, zero-sum games, which is covered in Chapter 3. He provided the framework that this book uses to study games in general. Together with Oskar Morgenstern, his fellow refugee from fascism, Von Neumann was the first to solve games in business and economics.⁷

We use game theory to solve games, to help players avoid mistakes and find the right strategy if there is only one. Sometimes the best we can do is narrow a player's choice down to a set of strategies. A solution of a game should tell each player what outcome to expect and how to achieve that outcome. This involves some mathematics; indeed, game theory began as pure mathematics. Von Neumann was first and foremost a mathematician, as was his star student, John Nash. Nash formulated a central idea in solution theory, Nash Equilibrium (NE), which we study from Chapter 3 onward. This is *the* John Nash portrayed in the movie *A Beautiful Mind*. Nash, along with two other game theorists, Reinhard Selten and John Harsanyi, won the Nobel Prize in economics in 1994 for their advances in game theory. You will learn about those advances as you go through this book.

Although some mathematics is unavoidable, we try to make that mathematics as user-friendly as possible, applying the adage that you should tell beginners the truth, but not the whole truth. If you are a beginner and you master this book, then you are ready to go on to more advanced material, which we suggest for further reading throughout the text. At the beginning, you are spared the theoretical niceties and the jargon you don't yet need or can't yet appreciate. Game theory is a subject that gets deep fast enough as it is. Think of it as a swimming pool in which even the shallow end can be over your head.

Although game theory began as applied mathematics, it has become a powerful mode of reasoning in business and economics, and is heavily used in other fields like formal political

⁵ The G-8 countries are the world's seven largest economies (United States, Japan, Germany, France, United Kingdom, Italy, Canada) plus Russia; their representatives meet frequently to negotiate issues that involve economic relations among them.

⁶ Besides creating game theory, von Neumann also provided the mathematical foundation for quantum mechanics, designed the implosion lens for the atomic bomb, and created the architecture of the computer. An excellent biography is Norman MacRae, *John von Neumann* (New York: Pantheon Books, 1992.)

⁷ The title of their magnum opus, John von Neumann and Oskar Morgenstern, *The Theory of Games and Economic Behavior* (Princeton, N.J.: Princeton University Press, 1944), shows how much their work was driven by the possibility of economic applications. Later on, after the creation of game theory, it became apparent that earlier economists, among them Cournot, Bertrand, Edgeworth, and Stackelberg, had also found solutions to games.

theory, evolutionary biology, psychology, and philosophy. The Nobel Prize-winning macroeconomist Robert Lucas argues that the most important contributions to macroeconomics since Keynes have been the result of formulating macroeconomic problems as games and then solving those games.⁸ It is easy to see why this is true. In the heyday of Keynesian economics, economists designed policies for the government to use assuming that the public's expectations, government policies of other countries, and the like would not be affected. This assumption is tantamount to assuming that the public and other governments are not players in a game. We know better now. Expectations held by the public—which show up, for instance, in the investment and portfolio decisions taken in the private sector—make an enormous difference to the outcome of economic policy. When the energy sector was deregulated in the early 1990s, that opened up all sorts of strategies for firms like Enron. Enron was playing a whole new game, and grew (at least on paper) into a \$50-billion firm by playing that game in fairly unique fashion—before it imploded into bankruptcy at the end of 2001. When something is a game—and the entire economy surely is one—you have to treat it as a game if you are going to understand it at all well.

Another example, which arises both in business and in economics, is called *mechanism design*. Suppose you have a certain goal in mind—for example, collecting \$2 trillion in government revenue—and you want to achieve that goal as inexpensively as possible. You need to design a set of mechanisms—foremost among them, the tax system—to achieve the goal. The mechanisms so designed lead to games, since taxpayers have strategies that respond to various kinds of taxation. If you don't include taxpayer strategies in your mechanism, you can be sure that it will fail miserably. Mass transit riders will change their riding strategy to some extent when fares double, and a transit authority that ignores this will not set fares in a strategically sound way. Of course, taxes aren't the only source of government revenue. Until recently, the federal government gave away rights to the electromagnetic spectrum, even though those rights are extremely valuable. Since the mid 1990s, advised and influenced to a large extent by game theorists, the federal government has auctioned off those rights, in the process raising billions of dollars for the U.S. Treasury. We study various kinds of mechanism design problems in Chapters 11 and 12.

There are several purely selfish reasons for studying game theory, above and beyond the desire to know the truth that this science has to offer. These reasons might apply to your job or your career someday, even if you aren't going to be the next John Nash. First and foremost, game theory can improve your strategic decision making. It makes you more aware of when you are in a situation in which strategy matters, to say nothing of making you aware of strategic nuance on the part of your competitors or opponents. Second, it can improve your ability to run a business and to evaluate changes in policy. The phrases "competitive advantage," "everyday low pricing," and "winner's curse" will make a lot more sense to you after they have been strategically explicated (in Chapters 3, 4, and 12, respectively). Finally, game theory can help you become a better economist or a better manager. Game theory is the central paradigm of economics and finance. It contains or informs all the current buzzwords, such as market failure, credibility, incentive contracts, hostile takeovers, and coalition building, to name just a few. And—if you are into that sort of thing—game theory might even make you a better Blackjack player.

Like any theory that involves mathematics, there is no royal road to understanding game theory. You have to follow the arguments, understand the proofs, work the problems

⁸ Lucas is one of the principal architects of the rational expectations revolution in macroeconomics. See his *Models of Business Cycles* (Oxford: Basil Blackwell, 1987).

(there are problems at the end of each chapter), study the examples, and think about it all in terms of your own life and your own work. If you do the work, then you will complete this book knowing how to set up and solve a whole arsenal of games—and even more important, knowing how to think about games in an intelligent way. You won't fall for the wrong answers in strategic situations, if you have analyzed them correctly.

The book is organized as follows. Part One covers the fundamentals, ideas and techniques that are used throughout game theory. Part Two looks at games within games (sub-games), repeated games, and games that are out of equilibrium. Part Three is especially concerned with games of imperfect information, where mechanism design is the central objective. Part Four focuses on games in which cooperation pays off—negotiation, bargaining, and arbitration. Part Five deals with the interaction between games, markets, and politics and introduces the link with micro- and macroeconomics.

This is a lot of material to cover, so let's get started. Theory is something you do, not something you talk about doing. We begin with the simplest possible games, those with a single player and perfect information.

1.3 ONE-PERSON DECISIONS WITH PERFECT INFORMATION


Situations with a single person involved are called **decisions**, to distinguish them from games, which need at least two players. At the same time, some of the same language and machinery that are useful in decisions are useful in games as well. In any science, it makes sense to start with really simple problems, real morale boosters. Solving these prepares you for tackling tougher problems later on. That's precisely why we start with 1-person decisions.

The information available to the players in a game makes a big difference in what they can or should do. A player has **perfect information** if he or she knows exactly what has happened every time a decision needs to be made. A game has perfect information if every player in it has perfect information. Chess is a game with perfect information. As long as both players can see the board, they can see exactly what has happened every time they have to make a move.⁹ If some player does not have perfect information, then the game is one of **imperfect information**. Poker is a game with imperfect information. One player does not know—unless he or she is cheating—what cards have been dealt to another player. Games with imperfect information are very important, but they are more difficult to solve than games with perfect information. In this section we study games with perfect information; later, we tackle games with imperfect information.

As a paradigm for a 1-player decision with perfect information, consider a player about to enter the maze in Figure 1.1. Give this player a number—say, 1. Although a number isn't needed to identify a single player, numbers do become useful to identify the players once there are several of them. Action begins when player 1 enters the maze. Player 1's goal is to get to the pot of gold at the end of the maze without running into a wall first. The minute the player runs into a wall, the call is over. The pot of gold is worth M (for money) dollars, and if the player reaches it, the outcome is M . Player 1 gets paid the amount M , which is positive. If player 1 runs into a wall, the game ends and the outcome is 0. In this book, and for the vast majority of time in real life, a player prefers more money to less, and so prefers the outcome M to the outcome 0. These are the two possible outcomes when the player enters the maze shown in Figure 1.1.

⁹ There is a version of Chess, called Kriegspiel, in which players do not see the board; Kriegspiel thus has imperfect information.

FIGURE 1.1 THE MAZE



Now we are going to walk with player 1 through the maze. Player 1 reaches a *decision point*, marked *a*, soon after entering the maze. At this decision point, player 1 has two choices: go left or go right. If player 1 goes right, she runs into a wall and the game ends. If player 1 goes left, she encounters a second decision point, marked *b*. At this decision point, player 1 can go either left or right. If player 1 goes left, she runs into a wall and the game ends. If player 1 goes right, she eventually reaches the pot of gold.

Now let's see how decision looks at this maze. Decision theory has a precise way of describing any decision, a description called the **extensive form** of the decision. Figure 1.2 shows the extensive form of the maze in Figure 1.1. The decision is dubbed Pot of Gold. We will walk you through the extensive form of Pot of Gold, and then look at some definitions.

Decision theory is all about the decisions the player makes, so the description of Pot of Gold begins with the first decision point, *a*, shown as the circled *a* and 1 at the left of Figure 1.2. A circle means that some player must make a decision at that point. The number of the player whose decision it is—here, player 1—is included inside the circle. Coming out of this decision point are two straight lines. These straight lines represent the two choices player 1 can make at decision point *a*: left or right. The line labeled Right runs into a wall, causing the game to end. The end of the game is represented by an *endpoint*. An outcome is attached to each endpoint. Since running into a wall means getting \$0, the outcome 0 is attached to this endpoint. If player 1 goes left, then she reaches another decision point, *b*. At this decision point, player 1 has two choices, left or right. These choices (also called moves) are again represented by lines coming out of decision point *b*, labeled Left and

Right. If player 1 goes left, she again runs into a wall (an endpoint) and the outcome is 0. If player 1 goes right, she exits from the maze and reaches the pot of gold, the outcome M .

The extensive form shown in Figure 1.2 contains all the information about Pot of Gold that is needed to solve it. Mathematically speaking, the extensive form is a *tree diagram*, so called because it looks like a tree if you are at the starting point, facing right. There are fancy names for the elements of the extensive form. Points in the tree are called *nodes*. Every game begins with an *initial node*. In Pot of Gold, this is the node a . A node with a circle around it and a player number inside is called an **information set**. An information set shows which player has to move and what the player knows when making the move. The lines coming out from a node are called *branches*, again in keeping with the tree metaphor. The outcomes attached to endpoints are called *payoffs*. Pot of Gold has perfect information: whenever a decision has to be made, that is, at nodes a and b , player 1 knows exactly where he or she is in the maze. Later we will see what imperfect information looks like in a tree diagram.

There is a way to solve Pot of Gold and games like it that always works: start from the pot of gold and work backward until the entrance is reached. Working from the end of a game back to the beginning in order to solve it is called **backward induction** and is the procedure used to solve every 1-person game with perfect information. Here's how backward induction is used to solve Pot of Gold (see Figure 1.3).

Start at the last decision point of the game, node b . At node b , player 1 reaches the pot of gold by going right. We denote this by an arrow on the branch labeled Right coming out of node b . Now it is up to player 1 to reach node b , since this puts the player in position to get the gold. The only way to reach node b is by going left at node a . We denote this by an

FIGURE 1.2 POT OF GOLD, EXTENSIVE FORM

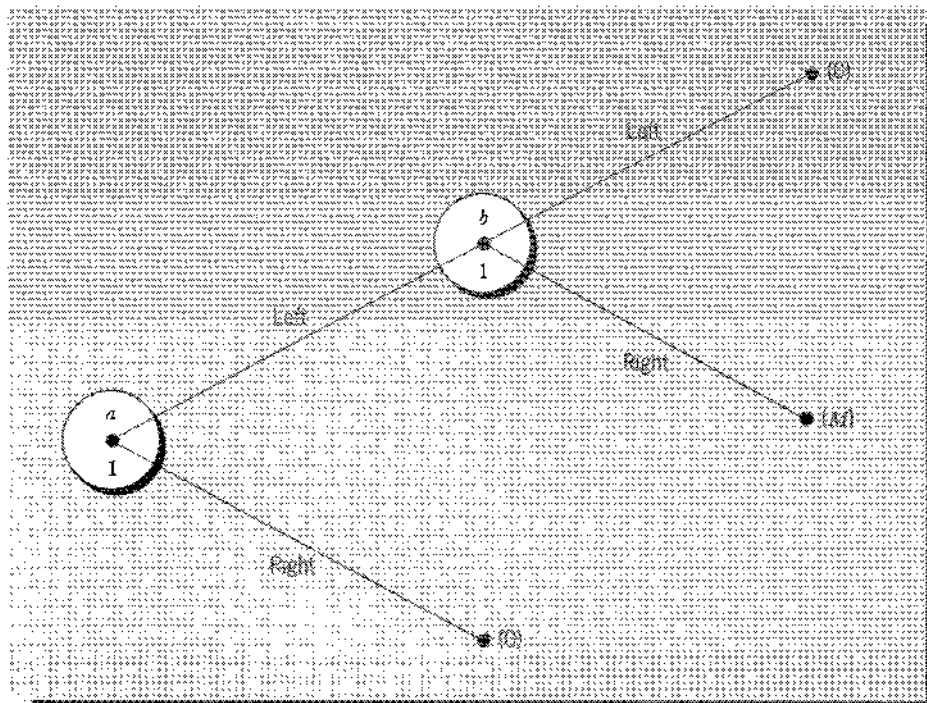



FIGURE 1.3 POT OF GOLD, BACKWARD INDUCTION



arrow on the branch labeled Left coming out of node a . We now have a complete path from the initial node a to the endpoint where the pot of gold is:

go left at a , go right at b

This plan of play solves Pot of Gold and gets the payoff M , the pot of gold. We have just solved your first game.

It is possible to get the payoff 0 in Pot of Gold. All it takes is a bad strategy. A **strategy** is a complete plan of play for a game. The good strategy in Pot of Gold leads to the pot of gold. There are three other strategies for Pot of Gold, all of which are bad because they run into a wall and get payoff 0. To construct a strategy, you first identify every decision that a player might make. Here, player 1 has to make a decision at node a and again at node b . At node a , player 1 has two choices; at node b , two choices. A strategy for Pot of Gold fills in the following blanks:

_____ at a , _____ at b

Since there are two possible ways to fill in the blank at a and two possible ways to fill in the blank at b , there are $(2)(2) = 4$ strategies for Pot of Gold. We have already identified the good strategy. Here are the three bad strategies:

right at a , right at b

left at a , left at b

right at a , left at b

Good play avoids these three strategies.

We have seen that the extensive form of Pot of Gold consists of nodes, branches, endpoints, and payoffs. Game theory has another way to describe Pot of Gold. This description, based only on strategies, is called the **normal form**. The normal form of a 1-player game lists each of the player's strategies and the payoff alongside each. Figure 1.4 shows Pot of Gold in normal form. It lists player 1's four strategies and the payoff, 0 or M , alongside each. The normal form codes all the information of the extensive form into a matrix, here a matrix with four rows and one column. In this code, it is easy to spot the solution to Pot of Gold: pick the strategy (left at a , right at b) that pays M . This is the best strategy, as it is the only one to pay a positive payoff. The three other strategies pay zero.

You ought to be wondering why there are two representations, normal form and extensive form, of the same game. Some purists think that we should study games only in extensive form. In particular, for games where the sequence of moves is crucial, the extensive form has a lot of appeal. Von Neumann and Morgenstern, the originators of the distinction, thought studying games in normal form led to more general insights. For games where all players move simultaneously, the normal form has a lot of appeal. This book takes a pragmatic approach: pick the form that is easier to write down and solve. This is usually the normal form. However, if the game is already written down in extensive form, you can solve it in that form without bothering with the normal form. In 1-player decisions of perfect information, the same solution is derived from both the extensive form and the normal form.¹⁰

FIGURE 1.4 POT OF GOLD, NORMAL FORM

Player 1	
Left at a , left at b	0
Left at a , right at b	M
Right at a , left at b	0
Right at a , right at b	0

¹⁰ Chapter 7 shows to what extent this is true for more complicated games.

1.4 TWO-PERSON GAMES WITH PERFECT INFORMATION

The decision we studied in the last section had only one player. The rest of this chapter is devoted to games with two players. For now, we restrict our attention to 2-person games with perfect information, a simple example of which is shown in Figure 1.5, Two-Person Pot of Gold.

Player 1 moves first, and can go either left or right. If player 1 goes right, the game ends. At the endpoint is attached the payoff vector \mathbf{u} , given by

$$\mathbf{u} = (u_1, u_2) = (0, 0)$$

where u_i is the notation for the payoff to player i . So each player gets 0 at this endpoint. If player 1 goes left, then it is player 2's turn to move. Player 2 can go either left or right. If player 2 goes left, then the game ends with the payoff vector


$$\mathbf{u} = (0, 0)$$

If player 2 goes right, then the game ends as well. In this event, the players have found the pot of gold, whose value they divide evenly:

$$\mathbf{u} = M/2, M/2$$

Solving Pot of Gold with two players is no different from solving Pot of Gold with one player. We use backward induction, starting with player 2, the last player to move. Player 2, who prefers half the pot of gold to none, moves right. This brings us to player 1's move.

FIGURE 1.5 TWO-PERSON POT OF GOLD, EXTENSIVE FORM



Player 1, who also prefers half the pot of gold to none, moves left. This solution is denoted by the arrows through Figure 1.5. We have just solved our first 2-person game.

Pot of Gold with two players is a parable—it tells a large story in a small space. Consider any business deal that has profit potential for both parties. The pot of gold valued M represents the deal. What the two sides have to do is find their way to the deal. The path by which 1 goes left, 2 goes right represents how the two sides find their way to the deal. This path is called the *solution path*. In 2-person Pot of Gold there is only one solution path, the path that gets to the pot of gold. All the other paths get zeros instead. Of course, most business deals are a lot more complicated than 2-person Pot of Gold; this complexity is addressed in later chapters.

Although the extensive form of 2-person Pot of Gold closely resembles the extensive form of the 1-person version (compare Figures 1.2 and 1.5), its normal form looks a lot different (compare Figures 1.4 and 1.6). To construct its normal form, we list the strategies for each player. Player 1 is the row player, and player 2 is the column player. Player 1 has a single information set, with two choices at that set, so player 1 has $(1)(2) = 2$ strategies. These are Left or Right, listed alongside the rows of the matrix. Similarly, player 2 has a single information set, with two choices at that set, so player 2 also has two strategies, Left or Right. These are listed above the columns of the matrix. We array these strategies in a 2×2 matrix (that is, a matrix with two rows and two columns). The rows correspond to player 1's two strategies; the columns, to player 2's two strategies. We move the payoff vectors from the endpoints of the extensive form into the appropriate cell of the normal form matrix. For instance, the pair of strategies

1 goes left, 2 goes right

leads to the payoff vector $(M/2, M/2)$, which appears in row 1 (player 1 goes left) and column 2 (player 2 goes right) in the matrix. In every payoff vector, player 1's payoff comes first, followed by player 2's payoff. Throughout game theory, 2×2 games prove useful as parables of larger questions, a good reason to avail ourselves of them as often as possible. Let's now turn to a class of very special 2-person games with perfect information—games like Chess.

FIGURE 1.6 TWO-PERSON POT OF GOLD, NORMAL FORM

		Player 2	
		Left	Right
Player 1	Left	(0,0)	$(\frac{M}{2}, \frac{M}{2})$
	Right	(0,0)	(0,0)

1.5 GAMES LIKE CHESS


A 2-person game with perfect information is a **game like Chess** if it satisfies the following requirements. First, the players have perfect information. Second, each player has at most a finite number of strategies. Third, the outcomes are limited to win, lose, or draw. Either player 1 wins and player 2 loses (w,l) or both players get a draw (d,d) or player 1 loses and player 2 wins (l,w).

Chess is a game like Chess. The fact that both players can see the board at all times means that it has perfect information. Chess is finite; it can't go on forever, because of a special stopping rule. If the same configuration appears on the board three times, the game is ruled a draw. The two players alternate turns, and the outcomes are limited to win, lose, or draw. So Chess qualifies. The first theorem ever proved about games says:



Theorem on Games like Chess. In games like Chess, exactly one of the following is true: player 1 can guarantee a win, player 2 can guarantee a win, or each player can guarantee a draw.¹¹

FIGURE 1.7 GAMES LIKE CHESS



¹¹ This theorem was essentially stated and proved by the mathematician Zermelo in 1913. See J. Schwalbe and P. Walker, "Zermelo and the early history of game theory," *Games and Economic Behavior* 34 (2001): 123–137, for the latest research on Zermelo's theorem.

We will prove this theorem for the class of extensive form games like Chess shown in Figure 1.7, all of whose normal forms are 2×2 . Player 1 moves first and can move either left or right. If player 1 moves right, the game ends with the payoff vector \mathbf{x} . If player 1 moves left, then it is player 2's turn. Player 2 can move either left or right. If player 2 moves left, the payoff vector is \mathbf{u} ; if player 2 moves right, the payoff vector is \mathbf{v} . The payoff vectors \mathbf{u} , \mathbf{v} , and \mathbf{x} can take any of the following forms:

$$(w,l) \ (d,d) \ (l,w)$$

Since each possible payoff vector can take any of three forms, and there are three payoff vectors, $3^3 = 27$ possible games are represented in Figure 1.7. The following proof applies to each of these 27 games.

The proof works via backward induction. Start with player 2, at the end of the game. Player 2 maximizes payoff, given the payoff vectors \mathbf{u} and \mathbf{v} . There are three possible cases facing player 2.

- *Case 1.* Player 2 can reach a win in at least one of the payoff vectors \mathbf{u} or \mathbf{v} . Then player 2 so moves. In this case, if the play reaches player 2, he can guarantee a win.
- *Case 2.* Player 2 cannot reach a win with either move, but he can reach a draw from at least one of the payoff vectors \mathbf{u} or \mathbf{v} . Then player 2 so moves. In this case, if the play reaches player 2, he can guarantee a draw.
- *Case 3.* Player 2 faces a loss no matter what move he makes: $\mathbf{u} = \mathbf{v} = (w,l)$. In this case, player 1 can guarantee a win by moving left. The play reaches player 2, who is sure to lose.

Now we come back to the start of the game, and player 1's move. If case 2 holds, then the worst player 1 can do is get a draw. If the payoff vector $\mathbf{x} = (w,l)$, then player 1 moves right and guarantees a win. Otherwise, player 1 moves left and guarantees a draw. Finally, suppose that case 1 holds. If $\mathbf{x} = (l,w)$, then player 1 loses no matter what she chooses, left or right. In this case, player 2 guarantees a win. If $\mathbf{x} = (d,d)$, then player 1 guarantees a draw by moving right. Finally, if $\mathbf{x} = (w,l)$, player 1 guarantees a win. This completes the proof.

There are a lot of games like Chess. Tic-Tac-Toe is one popular example. In Tic-Tac-Toe, the two players play on a 3×3 matrix. Player 1 marks an "X" in one of the 9 cells of the matrix. Then player 2 marks an "O" in one of the 8 remaining open cells of the matrix. Play continues until one of the players has a row, column, or diagonal with all of her or his marks in it, in which case that player wins. If the players have filled the matrix and no player has won, the game is declared a draw.

The theorem says something general about games like Chess. For any specific game like Chess, if we understand it well enough, then we can determine which of the three possible conclusions holds. For Tic-Tac-Toe, we know the following: each player in Tic-Tac-Toe can guarantee a draw. You can check this with your classmates. Pair off, and play 5 rounds of Tic-Tac-Toe. In the early rounds, you may observe wins by player 1 or player 2. By the end of the fifth round, however, all you will observe are draws. Experienced Tic-Tac-Toe players know the strategies that prevent defeat for either player 1 or player 2.

Chess is a lot more complicated than Tic-Tac-Toe. In Tic-Tac-Toe, the player who moves first has 9 possible moves; the player who moves second, 8 possible moves. By contrast, in Chess, player 1 (white) has 20 possible opening moves, and player 2 (black) has 20 possible responses to each opening move by white. Already the extensive form of Chess is challenging the capabilities of conventional human graphics. Moreover, Tic-Tac-Toe ends

in at most 9 moves, while in most Chess games it takes 40 moves before it is clear what the outcome will be. It is estimated that the extensive form of Chess would require something on the order of 10^{30} nodes to write down.

Chess is so complicated that it is still not known which of the three possible conclusions of the theorem is true. In Chess played at the world-class level, we observe all three of the possible conclusions of the theorem. However, we observe wins by player 1 more often than wins by player 2, and draws are more common than a win by either player. This suggests that there is some advantage to being player 1. To level the playing field in tournament conditions, such as the World Championship, there are an even number of matches (12), with the two finalists each playing as first mover an equal number of times (6). Whether player 1 can guarantee a win, or whether each player can guarantee a draw, remains an open question. Fame and fortune await the person—or machine—that ultimately decides this question.

■ 1.6 MAN VS. MACHINE: KASPAROV LOSES TO DEEP BLUE

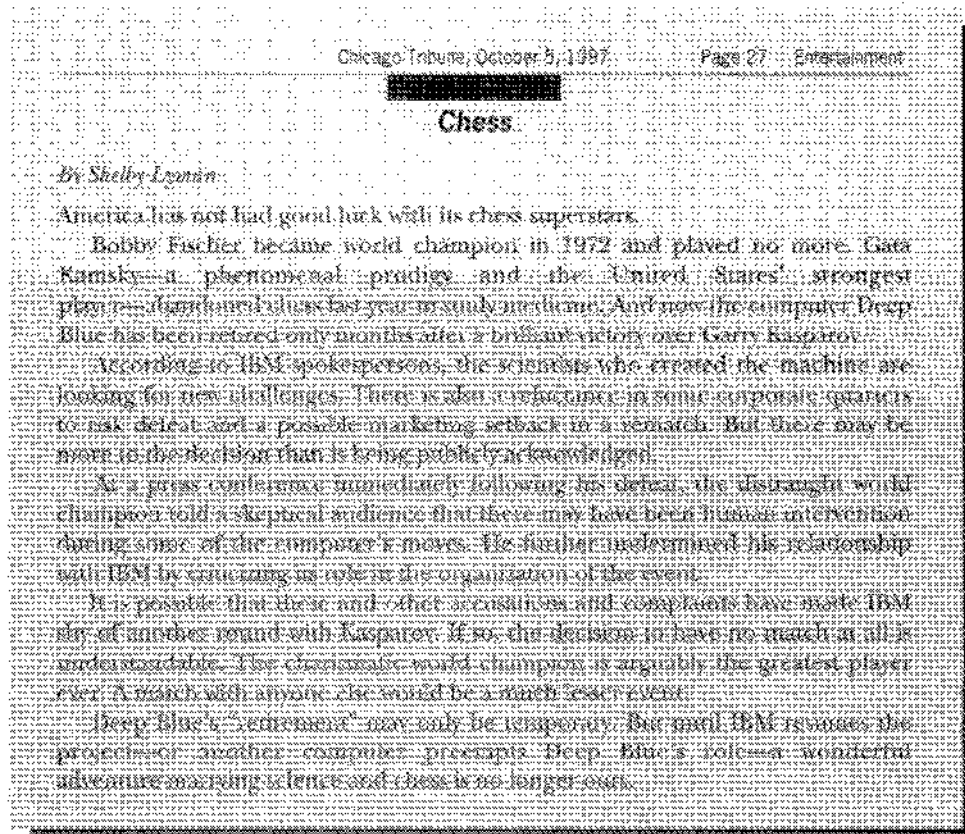
Part of what keeps Chess interesting is that we do not understand it well enough to know which of the three possible conclusions of the theorem is true.¹² If we knew that player 1 can guarantee a win, there would be little or no suspense left in any given match. Player 1 should be the favorite to win, and any other outcome is an upset. But this is precisely what we don't know.

Roughly a hundred million people play Chess, and even more follow the game, making it a business. Important matches, such as the World Championship, are now broadcast over the Internet to large audiences. The most-followed Chess match of all time took place in 1997, pitting the then World Champion, Garry Kasparov, against the highest-ranking Chess machine of all time, Deep Blue. Deep Blue was a creation of IBM, named after the corporation's color. Deep Blue was designed to do nothing else except play. Deep Blue operated by brute force, literally crunching Chess problems to death. Able to calculate the consequences of 200 million moves per second, thanks to its configuration of 32 processors in parallel, Deep Blue was ranked 60th in the world when this match was scheduled. That means Deep Blue was already better than all but 59 human players on the whole planet.

Games like Chess were the first games that machines could play. The only thing a machine needs to play a game like Chess is a program to determine its strategy. Once the machine has a complete plan of play, or a rule for computing the various pieces of a complete plan, it is ready to play. In a game like Chess, strategy is all that matters. Machines have been able to play Tic-Tac-Toe as well as humans for the past 30 years. Chess-playing programs on personal computers can often beat their owners. Still, it came as a shock to most humans when Deep Blue beat Kasparov in their 6-game match. This was the first time in history that a machine beat the World Champion at Chess. Although Kasparov came back to win a subsequent 6-game rematch, it is now established that, on any given day, a sufficiently powerful and well-designed machine can beat any human on the planet at Chess.

Games like Chess, or nearly so, are important not just as amusements. Chapter 7 shows that, with only a slight alteration of the payoffs in Figure 1.7, we get an important game played in business, involving possible entry into a market and attempts by established firms to deter that entry. Games like Chess show with crystalline clarity that strategy matters.

¹² Material in this section is drawn in part from the *Chicago Tribune*, "Deep Blue Association: is a Windfall for IBM," 17 August 1997.

FIGURE 1.8 IS THERE MORE TO THE DECISION BY IBM TO RETIRE DEEP BLUE?

Reprinted with permission.

Suppose you are player 1 in such a game, and it is known that player 1 can guarantee a win. Somehow you manage to lose instead. You have no excuses. You can't blame luck, since there are no chance moves. You can't blame missing information, since you have perfect information. The only reason you lost is that you chose a bad strategy. You snatched defeat from the jaws of victory. You have no one to blame but yourself when you play a bad strategy and get a bad outcome as a result.

1.7 WEAK DOMINANCE OF STRATEGIES AND WEAK DOMINANCE SOLVABILITY

When a player goes about comparing two strategies, the comparison involved is usually not straightforward. Usually the player finds that in some situations, one strategy does better than the other, while in other situations, the reverse is the case. Suppose you are head football coach, your team has the ball, and you are comparing the two strategies Pass and Run. The payoff from each of these strategies, in terms of yards gained, depends on what strategy the defense chooses. If the defense sets up in a 9-man line, then Run pays very little, while Pass stands to gain many yards. Just the opposite is true if the defense sets

up in a prevent formation, with 7 or 8 players defending against the pass. This is the usual situation in games.

However, when one strategy always does at least as well as another in every situation, and does better in at least one situation, then the comparison between them is straightforward. We say that the strategy that does at least as well as another in every situation, and does better in at least one situation, **weakly dominates** the other. You will see why we need the adverb *weakly* in the next section, where we study a stronger kind of domination. If a player has two strategies, and one weakly dominates the other, then she or he has every reason to choose the strategy that weakly dominates; and no good reason to choose the strategy that gets weakly dominated. The strategy that weakly dominates is a good strategy; the strategy that gets dominated is a bad strategy.

To see an example of this, recall the game in Figure 1.6; this is the normal form of 2-person Maze with perfect information. Consider player 1, comparing her two strategies Left and Right. Here is the comparison:

If player 2 chooses Left, then player 1 gets 0 if she chooses Left and 0 if she chooses Right. Since $0 = 0$, the result is a tie in this case.

If player 2 chooses Right, then player 1 gets $M/2$ if she chooses Left and 0 if she chooses Right. Since $M/2 > 0$, Left is better for player 1.

To sum up, Left pays player 1 at least as much as Right in every situation, and Left pays more than Right in one situation, so Left weakly dominates Right.

You can check that for player 2, the strategy Right weakly dominates the strategy Left. So each of the players in this game has a strategy that dominates her or his other strategy. The strategic choice in this game is really straightforward for each of the two players: choose the strategy that weakly dominates. When each player chooses his or her strategy that weakly dominates, player 1 chooses Left and player 2 chooses Right. The result is the payoff $M/2$ to each player in Figure 1.6.

When every player in a game has a strategy that dominates all his or her other strategies, we say the game is **weakly dominance solvable**. Solely by determining which strategy is weakly dominant, and choosing that strategy, we can solve such a game. Although games that are weakly dominance solvable are rare, when they do occur players who understand the game know exactly how to play them.

Often times, games in extreme situations tend to be dominance solvable. Here is an extreme case where this is so. A group of prisoners of war is about to be executed. This is not supposed to happen in war, according to the Geneva Convention of 1925; however, in the 2001 war in Afghanistan, the Geneva Convention was not always followed. Each of the prisoners has a pair of strategies, Do Nothing and Revolt. If a prisoner chooses the strategy Do Nothing, then he will be executed. If a prisoner chooses the strategy Revolt, and if enough other prisoners choose that same strategy, then at least some of the prisoners avoid execution. Thus, Revolt weakly dominates Do Nothing. And Revolt is exactly what happened at the prisoner of war camp outside Mazaar-e-Sharif, Afghanistan. Of some 600 prisoners of war, roughly 90 survived the revolt and fell into the hands of the Red Cross, which does observe the Geneva Convention.

Here is a second example, not involving life-and-death calculations but rather more economics. In a **second-price auction**, each player bids on the item being auctioned. The high bidder wins, but pays a price equal to the second-highest bid. In this kind of auction, bidding what the item is worth to you weakly dominates underbidding. If you bid what the item is

worth and win the auction, then you pay a price lower than what the item is worth to you. In this case, you come out with a positive payoff; you come out ahead. If you underbid and still win the auction, then you pay a price lower than what the item is worth to you—a tie. Finally, if you underbid and lose the auction, you get a zero payoff (the status quo hasn't changed), when you might have won the auction by bidding what the item is worth to you. That's precisely where the weak domination comes in. Auctions are complicated, and second-price auctions are fairly rare; we defer further study of them to Chapter 12.

1.8 STRICT AND WEAK DOMINATION OF STRATEGIES AND STRICT DOMINANCE SOLVABILITY

In the last section, we considered the comparison of two strategies, one of which was at least as good as the other in every situation, and better than the other in at least one situation. If one strategy is better than another in every situation, we say it **strictly dominates**. If every player in a game has a strategy that strictly dominates all other strategies, then playing the strictly dominating strategy is best for the player. We call such a game **strictly dominance solvable**. Solely by determining which strategy is strictly dominant, and choosing that strategy, we can solve such a game. Games that are strictly dominance solvable are even rarer than those that are weakly dominance solvable.

Figure 1.9 shows a game that is strictly dominance solvable. This game represents a certain kind of Teamwork. If each player chooses the strategy Left, then teamwork succeeds, and the payoff is +1 to each player. If the players do not choose the same strategy, then teamwork breaks down, and the payoff is 0 to each. Finally, if each player chooses the strategy Right, then teamwork leads to infighting and lots of other negativity, and each player gets the payoff -1.

Notice that for player 1 in the game of Teamwork, the strategy Left always pays better than the strategy Right: there are no ties. Similarly, for player 2 in Teamwork, the strategy Left always pays better than the strategy Right: again, there are no ties. Each player choosing his or her strictly dominant strategy, Left, leads to the outcome +1 for each player. This is the strictly dominant solution of Teamwork.

FIGURE 1.9 GAME THAT IS STRICTLY DOMINANCE SOLVABLE

		Player 2	
		Left	Right
Player 1	Left	(1,1)	(0,0)
	Right	(0,0)	(-1,-1)

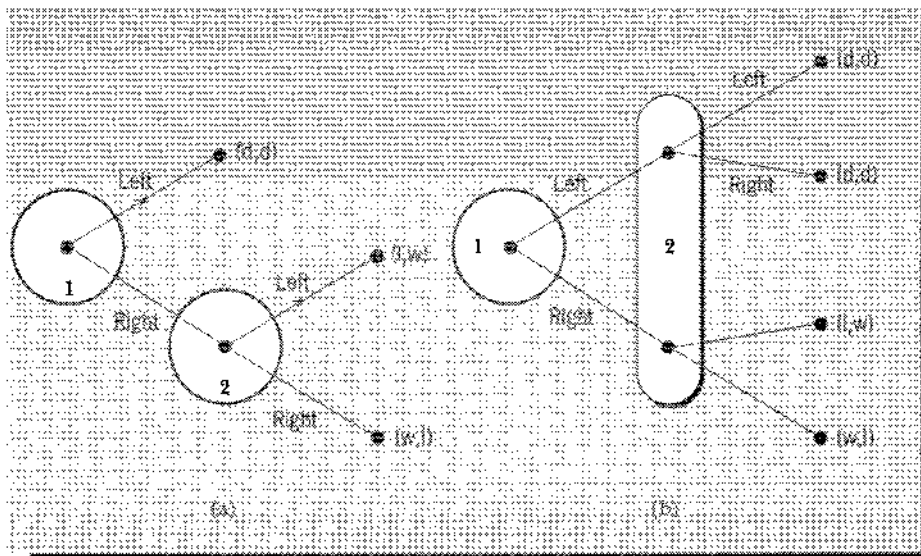
It is instructive to compare Figure 1.9, where the game is strictly dominance solvable, and Figure 1.6, where the game is weakly dominance solvable. In Figure 1.9, there are no ties; the dominance all goes one way, in favor of the teamwork strategy Left. In Figure 1.6, there are ties; the dominance kicks in only when player 1 uses the strategy Left and player 2 uses the strategy Right.

Here is a version of Teamwork that you may not want to play in class. Suppose the payoff represents points toward your final grade. “Left” means “Attend Class”; “Right” means “Skip Class.” If all students skip class, then the professor deducts one grade point from each—that’s harsh. If some students skip class and some attend class, then each student who attends is awarded one grade point, but there is no deduction from any student who does not attend. Finally, if each student attends class, then each student is awarded one grade point. In such a class, attending class is a strictly dominant strategy and attendance in class, the strictly dominant solution.

1.9 EXTENSIVE FORM, NORMAL FORM, AND COALITION FUNCTION FORM

Sometimes the line between a game like Chess and a game that is not like Chess can be very fine indeed. Consider the two extensive form games in Figure 1.10. The game in Figure 1.10a is like Chess, and player 1 can guarantee a draw by going left. The game in Figure 1.10b is not like Chess, since it has imperfect information. This imperfection shows up in the information set of player 2, which has 2 nodes. Player 2 doesn’t know how player 1 has moved, and so does not know at which node play has arrived when it is his turn to move. However, if we look at the normal forms of these two games, we see that they are exactly the same! The normal form of both is shown in Figure 1.11. Each player has one information set, with two choices at that information set, so each player has two strategies. That makes the normal

FIGURE 1.10 GAMES NEARLY LIKE CHESS



form a 2×2 matrix in each case. When both players go left, the outcome is a draw (d,d), and so on; so that the payoffs match up. Even though the game in Figure 1.10b is not like Chess, it is close enough that it has the same normal form as a game that is like Chess.


It often happens that games with different extensive forms have the same normal form, because the normal form suppresses some information that is available in the extensive form. For instance, you can see at once that the game in Figure 1.10a has perfect information, but this inference is no longer available in the normal form. Every extensive form has a unique normal form representation. However, for every game in normal form there are usually several possible games in extensive form that could give rise to it. Thus things that are true for games like Chess may also be true for games nearly like Chess.

In Figure 1.10a, just as in Figure 1.11, player 1 can guarantee a draw by moving left. In the extensive form, we find this out by backward induction. In normal form, we get this outcome by a rather different route. For player 2, choosing the strategy Left weakly dominates choosing the strategy Right. For player 1, neither strategy weakly dominates the other. So this game is not weakly dominance solvable, to say nothing of strictly dominance solvable. However, suppose player 1 takes this weak domination on the part of player 2 into account. Then player 1 does best by choosing the strategy Left and getting a draw. What player 1 has done is appeal to player 2's **rationality**; a rational player, here player 2, has no reason to use a weakly dominated strategy. In such a case, we say a game outcome is **rationalizable**, when it is based on an appeal to rationality. If you can follow an argument like that just given, you are well on your way to thinking like a true game theorist!

In addition to the extensive and normal forms, there is a third form of games, especially useful for games with a largely cooperative character, called the **coalition function form**. This form requires the answers to only two questions: what can each player guarantee for himself or herself, and what can the two players acting together guarantee for themselves? For the win, lose, or draw game, such as that shown in Figure 1.11, the coalition function form simply says that player 1 can guarantee a draw, as can player 2. The players acting together cannot guarantee anything more. For 2-person Pot of Gold, the coalition function form says that player 1 can guarantee 0 (because player 2 goes left), player 2 can guarantee 0 (because player 1 goes left), but together they can guarantee M (because player

FIGURE 1.11 GAMES IN NORMAL FORM

		Player 2	
		Left	Right
Player 1	Left	(d,d)	(d,d)
	Right	(l,w)	(w,l)

FIGURE 1.12 TWO-PERSON POT OF GOLD, COALITION FUNCTION FORM

I goes right and player 2 goes left). This coalition function form is depicted in Figure 1.12: the three points we have just identified are the vertices of a payoff triangle, and the striped area shows various ways of dividing less than M among the two players. The coalition function form is used primarily when studying how the gains of cooperation are divided up among those involved in a deal—a question focused on in Chapter 13 and beyond.

■ SUMMARY

1. A game is any rule-governed strategic situation with a well-defined outcome, characterized by strategic interdependence among the players.
2. There is a great variety of games in the literal sense: board games, card games, video games, and field games. Business and economics exhibit games in the extended sense. The Chicago Bulls provide an example of both senses.
3. Game theory is the science that studies games and takes them seriously enough to solve them. Game theory began as applied mathematics, but is now central to the way we think about business and economics.
4. Games serve as models for market interactions and economic negotiations. For example, the cola wars between Coke and Pepsi are a market game; negotiations among the G-8 countries are a game in international economics.
5. The extensive form is the basic description of a game. An extensive form is a tree diagram, with nodes, branches, an initial node, information sets, and endpoints.
6. A 1-player decision with perfect information can be solved by backward induction, starting at the end and working back to the beginning.
7. A game has perfect information if every information set contains a single node.
8. Two-person games with perfect information can be solved by backward induction, just like 1-person decisions with perfect information.
9. For a game like Chess, exactly one of the following is true: player 1 can guarantee a win, player 2 can guarantee a win, or each player can guarantee a draw.
10. A strategy that pays at least as well as another in every situation, and that pays better than another in at least one situation, weakly dominates. A game where every player has a weakly dominant strategy is weakly dominance solvable.

11. A strategy that pays better than another in every situation strictly dominates. A game where every player has a strictly dominant strategy is strictly dominance solvable.
12. There are three representations of games: extensive form, normal form, and coalition function form. The normal form looks at the implications of strategies, while suppressing some of the detail contained in the extensive form. The coalition function form suppresses even more detail than the normal form and is used mainly for studying cooperation among players.

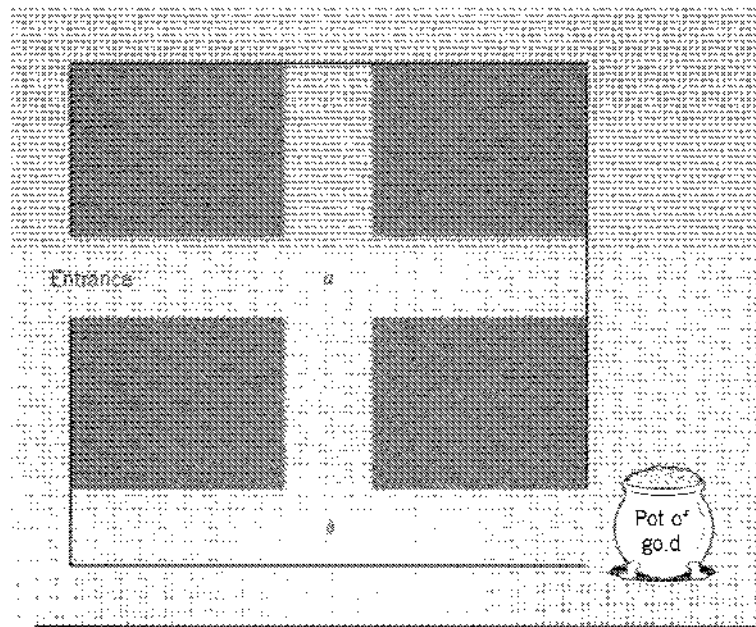
KEY TERMS

game	extensive form	second-price auction
strategic interdependence	information set	strictly dominant
market game	backward induction	strictly dominance solvable
economic negotiation	strategy	rationality
game theory	games like Chess	rationalizable
decisions	weakly dominant	coalition function form
perfect/imperfect information	weakly dominance solvable	

PROBLEMS

1. Give examples of strategic interaction in each of the following industries: entertainment, beer, automobiles, and financial services.
2. Draw the extensive form of the maze in Figure 1.13, and then solve it. Player 1 moves at a; player 2, at b.

FIGURE 1.13 THE MAZE II

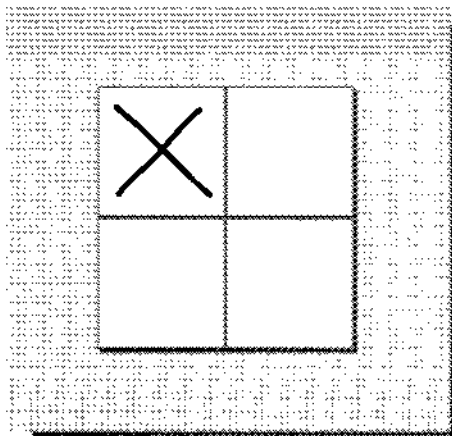


3. Transform the game in problem 2 into normal form. Next, find a weakly dominant strategy for each player. Finally, show that the game is weakly dominance solvable by finding its weak dominance solution.
4. Consider the schema of Figure 1.7. Find payoff vectors \mathbf{u} , \mathbf{v} , and \mathbf{x} such that (a) player 1 can guarantee a win; (b) player 2 can guarantee a win; and (c) either player can guarantee a draw. (There are many right answers.)
5. Suppose you play Tic-Tac-Toe on a 2×2 board, as shown in Figure 1.14. The usual rules apply. Player 1, who goes first, puts an X in any one of the cells. Player 2, who goes second, puts an O in one of the remaining cells. The first player to fill a row, column, or diagonal with his or her marks wins. Show that this game is like Chess. Draw its extensive form. Which player can guarantee a win?
6. In section 1.5, we said experienced Tic-Tac-Toe players on a 3×3 board never lose. Show how the player moving second can always get a draw, by considering how to counter the three kinds of opening moves: (1) player 1 puts X in a corner; (2) player 1 puts X on the side, but not a corner; (3) player 1 puts X in the middle.

Violate any one of the conditions that make a game like Chess and the game may no longer satisfy the conclusion of the theorem for games like Chess. This is the theme of the next four questions.

7. The theorem on games like Chess breaks down if there are three or more players. See if you can find a counterexample. If you can't, here's a big hint: see Chapter 5, section 5.1.
8. The game Pick the Largest Number is a 2-player, win, lose, or draw game with perfect information. Player 1 picks any number. Player 2 hears player 1's number, then picks any number he wants. If player 1's number is bigger than player 2's, the game ends in a win for player 1. If player 1's number equals player 2's number, the game ends in a draw. If player 1's number is smaller than player 2's number, then player 1 gets to choose again. Show that Pick the Largest Number is not like Chess because it is not finite: There is no largest number. In particular, show that player 2 need never lose—and so the game never ends.
9. Now consider the game Pick the Largest Number, as in problem 8, only now the largest number a player is allowed to pick is 3. This game is finite. Show that it is like Chess. In particular, either player can force a draw. For extra credit, show that this game is weakly dominance solvable, and find its weak dominance solution.

FIGURE 1.14 TIC-TAC-TOE ON A 2×2 BOARD, PLAYER 1 HAS JUST MOVED



10. The game Escape and Evasion (shown in Figure 1.15) is a 2-person, win, lose, or draw finite game. Player 1, the escapee, has just escaped from jail, and can go either top or bottom. Player 2, the jailer, can also go either top or bottom, but does not know which way the escapee has gone. If the jailer goes the same way as the escapee, the jailer catches the escapee—a win for the jailer. If the jailer goes a different way from the escapee, the escapee gets away—a win for the escapee. Show that Escape and Evasion has imperfect information. Then show that neither player can guarantee a win. Since draws aren't allowed, you have a counterexample to the theorem on games like Chess.
11. According to the doctrine that "a dollar is a dollar," the same principles apply to the game Pot of Gold whether M is \$1 or \$1 million. Do you agree or disagree? Why?

FIGURE 1.15 ESCAPE AND EVASION

