13 Game Theory and Competitive Strategy

CHAPTER OUTLINE

LIST OF EXAMPLES

In Chapter 12, we began to explore some of the strategic output and pricing decisions that firms must often make. We saw how a firm can take into account the likely responses of its competitors when it makes these decisions. However, there are many questions about market structure and firm behavior that we have not yet addressed. For example, why do firms tend to collude in some markets and to compete aggressively in others? How do some firms manage to deter entry by potential competitors? And how should firms make pricing decisions when demand or cost conditions are changing or new competitors are entering the market?

To answer these questions, we will use game theory to extend our analysis of strategic decision making. The application of game theory has been an important development in microeconomics. This chapter explains some key aspects of this theory and shows how it can be used to understand how markets evolve and operate, and how managers should think about the strategic decisions they continually face. We will see, for example, what happens when oligopolistic firms must set and adjust prices strategically over time, so that the prisoners' dilemma, which we discussed in Chapter 12, is repeated over and over. We will show how firms can make strategic moves that give them advantages over competitors or an edge in bargaining situations, and how they can use threats, promises, or more concrete actions to deter entry. Finally, we will turn to auctions and see how game theory can be applied to auction design and bidding strategies.

13.1 GAMING AND STRATEGIC DECISIONS

First, we should clarify what gaming and strategic decision making are all about. A **game** is any situation in which *players* (the participants) make *strategic decisions*—i.e., decisions that take into account each other's actions and responses. Examples of games include firms competing with each other by setting prices, or a group of consumers bidding against each other at an auction for a work of art. Strategic decisions result in **payoffs** to the players: outcomes that generate rewards or benefits. For the price-setting firms, the payoffs are profits; for the bidders at the auction, the winner's payoff is her consumer surplus—i.e., the value she places on the artwork less the amount she must pay.

A key objective of game theory is to determine the optimal strategy for each player. A **strategy** is a rule or plan of action for playing the

• **game** Situation in which players (participants) make strategic decisions that take into account each other's actions and responses.

• **payoff** Value associated with a possible outcome.

• **strategy** Rule or plan of action for playing a game.

• **optimal strategy** Strategy that maximizes a player's expected payoff.

game. For our price-setting firms, a strategy might be: "I'll keep my price high as long as my competitors do the same, but once a competitor lowers his price, I'll lower mine even more." For a bidder at an auction, a strategy might be: "I'll make a first bid of $2000 to convince the other bidders that I'm serious about winning, but I'll drop out if other bidders push the price above $5000." The **optimal strategy** for a player is the one that maximizes her expected payoff.

We will focus on games involving players who are *rational*, in the sense that they think through the consequences of their actions. In essence, we are concerned with the following question: *If I believe that my competitors are rational and act to maximize their own payoffs, how should I take their behavior into account when making my decisions?* In real life, of course, you may encounter competitors who are irrational, or are less capable than you of thinking through the consequences of their actions. Nonetheless, a good place to start is by assuming that your competitors are just as rational and just as smart as you are.[1] As we will see, taking competitors' behavior into account is not as simple as it might seem. Determining optimal strategies can be difficult, even under conditions of complete symmetry and perfect information (i.e., my competitors and I have the same cost structure and are fully informed about each others' costs, about demand, etc.). Moreover, we will be concerned with more complex situations in which firms face different costs, different types of information, and various degrees and forms of competitive "advantage" and "disadvantage."

Noncooperative versus Cooperative Games

• **cooperative game** Game in which participants can negotiate binding contracts that allow them to plan joint strategies.

• **noncooperative game** Game in which negotiation and enforcement of binding contracts are not possible.

The economic games that firms play can be either *cooperative* or *noncooperative*. In a **cooperative game**, players can negotiate binding contracts that allow them to plan joint strategies. In a **noncooperative game**, negotiation and enforcement of binding contracts are not possible.

An example of a cooperative game is the bargaining between a buyer and a seller over the price of a rug. If the rug costs $100 to produce and the buyer values the rug at $200, a cooperative solution to the game is possible: An agreement to sell the rug at any price between $101 and $199 will maximize the sum of the buyer's consumer surplus and the seller's profit, while making both parties better off. Another cooperative game would involve two firms negotiating a joint investment to develop a new technology (assuming that neither firm would have enough know-how to succeed on its own). If the firms can sign a binding contract to divide the profits from their joint investment, a cooperative outcome that makes both parties better off is possible.[2]

An example of a noncooperative game is a situation in which two competing firms take each other's likely behavior into account when independently setting their prices. Each firm knows that by undercutting its competitor, it can capture more market share. But it also knows that in doing so, it risks setting off a price war. Another noncooperative game is the auction mentioned above: Each bidder must take the likely behavior of the other bidders into account when determining an optimal bidding strategy.

---

[1]When we asked, 80 percent of our students told us that they were smarter and more capable than most of their classmates. We hope that you don't find it too much of a strain to imagine competing against people who are as smart and capable as you are.

[2]Bargaining over a rug is called a *constant sum* game because no matter what the selling price, the sum of consumer surplus and profit will be the same. Negotiating over a joint venture is a *nonconstant sum* game: The total profit that results from the venture will depend on the outcome of the negotiations (e.g., the resources that each firm devotes to the venture).

Note that the fundamental difference between cooperative and noncooperative games lies in the *contracting possibilities*. In cooperative games, binding contracts are possible; in noncooperative games, they are not.

We will be concerned mostly with noncooperative games. Whatever the game, however, keep in mind the following key point about strategic decision making:

> *It is essential to understand your opponent's point of view and to deduce his or her likely responses to your actions.*

This point may seem obvious—of course, one must understand an opponent's point of view. Yet even in simple gaming situations, people often ignore or misjudge opponents' positions and the rational responses that those positions imply.

How to Buy a Dollar Bill Consider the following game devised by Martin Shubik.[3] A dollar bill is auctioned, but in an unusual way. The highest bidder receives the dollar in return for the amount bid. However, the second-highest bidder must also hand over the amount that he or she bid—and get nothing in return. *If you were playing this game, how much would you bid for the dollar bill?*

Classroom experience shows that students often end up bidding more than a dollar for the dollar. In a typical scenario, one player bids 20 cents and another 30 cents. The lower bidder now stands to lose 20 cents but figures he can earn a dollar by raising his bid, and so bids 40 cents. The escalation continues until two players carry the bidding to a dollar against 90 cents. Now the 90-cent bidder has to choose between bidding \$1.10 for the dollar or paying 90 cents to get nothing. Most often, he raises his bid, and the bidding escalates further. In some experiments, the "winning" bidder has ended up paying more than \$3 for the dollar!

How could intelligent students put themselves in this position? By failing to think through the likely response of the other players and the sequence of events it implies.

In the rest of this chapter, we will examine simple games that involve pricing, advertising, and investment decisions. The games are simple in that, *given some behavioral assumptions*, we can determine the best strategy for each firm. But even for these simple games, we will find that the correct behavioral assumptions are not always easy to make. Often they will depend on how the game is played (e.g., how long the firms stay in business, their reputations, etc.). Therefore, when reading this chapter, you should try to understand the basic issues involved in making strategic decisions. You should also keep in mind the importance of carefully assessing your opponent's position and rational response to your actions, as Example 13.1 illustrates.

EXAMPLE 13.1 Acquiring a Company

You represent Company *A* (the acquirer), which is considering acquiring Company *T* (the target).[4] You plan to offer cash for all of Company *T*'s shares, but you are unsure what price to offer. The complication is this: The value of Company *T*—indeed, its viability—depends on the outcome of a major oil exploration project. If the project fails, Company *T* under current management will be worth nothing. But if it succeeds, Company *T*'s value under current management

[3] Martin Shubik, *Game Theory in the Social Sciences* (Cambridge, MA: MIT Press, 1982).

[4] This is a revised version of an example designed by Max Bazerman for a course at MIT.

could be as high as \$100/share. All share values between \$0 and \$100 are considered equally likely.

It is well known, however, that Company *T* will be worth much more under the progressive management of Company *A* than under current management. In fact, whatever the ultimate value under current management, *Company T will be worth 50 percent more under the management of Company A*. If the project fails, Company *T* is worth \$0/share under either management. If the exploration project generates a \$50/share value under current management, the value under Company *A* will be \$75/share. Similarly, a \$100/share value under Company *T* implies a \$150/share value under Company *A*, and so on.

You must determine what price Company *A* should offer for Company *T*'s shares. This offer must be made *now—before* the outcome of the exploration project is known. From all indications, Company *T* would be happy to be acquired by Company *A—for the right price*. You expect Company *T* to delay a decision on your bid until the exploration results are in and then accept or reject your offer before news of the drilling results reaches the press.

Thus, *you (Company A) will not know the results of the exploration project when submitting your price offer, but Company T will know the results when deciding whether to accept your offer. Also, Company T will accept any offer by Company A that is greater than the (per share) value of the company under current management*. As the representative of Company *A*, you are considering price offers in the range \$0/share (i.e., making no offer at all) to \$150/share. *What price per share should you offer for Company T's stock?*

Note: The typical response—to offer between \$50 and \$75 per share—is wrong. The correct answer to this problem appears at the end of this chapter, but we urge you to try to answer it on your own.

13.2 DOMINANT STRATEGIES

How can we decide on the best strategy for playing a game? How can we determine a game's likely outcome? We need something to help us determine how the rational behavior of each player will lead to an equilibrium solution. Some strategies may be successful if competitors make certain choices but fail if they make other choices. Other strategies, however, may be successful regardless of what competitors do. We begin with the concept of a **dominant strategy**—*one that is optimal no matter what an opponent does.*

• **dominant strategy** Strategy that is optimal no matter what an opponent does.

The following example illustrates this in a duopoly setting. Suppose Firms *A* and *B* sell competing products and are deciding whether to undertake advertising campaigns. Each firm will be affected by its competitor's decision. The possible outcomes of the game are illustrated by the payoff matrix in Table 13.1. (Recall that the payoff matrix summarizes the possible outcomes of the game; the first number in each cell is the payoff to *A* and the second is the payoff to *B*.) Observe that if both firms advertise, Firm *A* will earn a profit of 10 and Firm *B* a profit of 5. If Firm *A* advertises and Firm *B* does not, Firm *A* will earn 15 and Firm *B* zero. The table also shows the outcomes for the other two possibilities.

In §12.4, we explain that a payoff matrix is a table showing the payoffs to each player given her decision and the decision of her competitor.

What strategy should each firm choose? First consider Firm *A*. It should clearly advertise because no matter what firm *B* does, Firm *A* does best by advertising. If Firm *B* advertises, *A* earns a profit of 10 if it advertises but only 6 if it doesn't. If *B* does not advertise, *A* earns 15 if it advertises but only 10 if it

TABLE 13.1 Payoff Matrix for Advertising Game

		Firm B	
		Advertise	Don't advertise
Firm A	Advertise	10, 5	15, 0
	Don't advertise	6, 8	10, 2

doesn't. Thus advertising is a dominant strategy for Firm *A*. The same is true for Firm *B*: No matter what firm *A* does, Firm *B* does best by advertising. Therefore, assuming that both firms are rational, we know that the outcome for this game is that both *firms will advertise*. This outcome is easy to determine because both firms have dominant strategies.

When every player has a dominant strategy, we call the outcome of the game an **equilibrium in dominant strategies**. Such games are straightforward to analyze because each player's optimal strategy can be determined without worrying about the actions of the other players.

• **equilibrium in dominant strategies** Outcome of a game in which each firm is doing the best it can regardless of what its competitors are doing.

Unfortunately, not every game has a dominant strategy for each player. To see this, let's change our advertising example slightly. The payoff matrix in Table 13.2 is the same as in Table 13.1 except for the bottom right-hand corner—if neither firm advertises, Firm *B* will again earn a profit of 2, but Firm *A* will earn a profit of 20. (Perhaps Firm *A*'s ads are expensive and largely designed to refute Firm *B*'s claims, so by not advertising, Firm *A* can reduce its expenses considerably.)

Now Firm *A* has no dominant strategy. *Its optimal decision depends on what Firm B does.* If Firm *B* advertises, Firm *A* does best by advertising; but if Firm *B* does not advertise, Firm *A* also does best by not advertising. Now suppose both firms must make their decisions at the same time. What should Firm *A* do?

To answer this, Firm *A* must put itself in Firm *B*'s shoes. What decision is best from Firm *B*'s point of view, and what is Firm *B* likely to do? The answer is clear: Firm *B* has a dominant strategy—advertise, no matter what Firm *A* does. (If Firm *A* advertises, *B* earns 5 by advertising and 0 by not advertising; if *A* doesn't advertise, *B* earns 8 if it advertises and 2 if it doesn't.) Therefore, Firm *A* can conclude that Firm *B* will advertise. This means that Firm *A* should advertise (and thereby earn 10 instead of 6). The logical outcome of the game is that both firms will advertise because Firm *A* is doing the best it can given Firm *B*'s decision; and Firm *B* is doing the best it can given Firm *A*'s decision.

TABLE 13.2 Modified Advertising Game

		Firm B	
		Advertise	Don't advertise
Firm A	Advertise	10, 5	15, 0
	Don't advertise	6, 8	20, 2

13.3 THE NASH EQUILIBRIUM REVISITED

To determine the likely outcome of a game, we have been seeking "self-enforcing," or "stable" strategies. Dominant strategies are stable, but in many games, one or more players do not have a dominant strategy. We therefore need a more general equilibrium concept. In Chapter 12, we introduced the concept of a *Nash equilibrium* and saw that it is widely applicable and intuitively appealing.[5]

Recall that a Nash equilibrium is a set of strategies (or actions) such that *each player is doing the best it can given the actions of its opponents*. Because each player has no incentive to deviate from its Nash strategy, the strategies are stable. In the example shown in Table 13.2, the Nash equilibrium is that both firms advertise: Given the decision of its competitor, each firm is satisfied that it has made the best decision possible, and so has no incentive to change its decision.

In Chapter 12, we used the Nash equilibrium to study output and pricing by oligopolistic firms. In the Cournot model, for example, each firm sets its own output while taking the outputs of its competitors as fixed. We saw that in a Cournot equilibrium, no firm has an incentive to change its output unilaterally because each firm is doing the best it can given the decisions of its competitors. Thus a Cournot equilibrium is a Nash equilibrium.[6] We also examined models in which firms choose price, taking the prices of their competitors as fixed. Again, in the Nash equilibrium, each firm is earning the largest profit it can given the prices of its competitors, and thus has no incentive to change its price.

In §12.2, we explain that the Cournot equilibrium is a Nash equilibrium in which each firm correctly assumes how much its competitor will produce.

It is helpful to compare the concept of a Nash equilibrium with that of an equilibrium in dominant strategies:

Dominant Strategies:	I'm doing the best I can *no matter what you do.* You're doing the best you can *no matter what I do.*
Nash Equilibrium:	I'm doing the best I can *given what you are doing.* You're doing the best you can *given what I am doing.*

Note that a dominant strategy equilibrium is a special case of a Nash equilibrium.

In the advertising game of Table 13.2, there is a single Nash equilibrium—both firms advertise. In general, a game need not have a single Nash equilibrium. Sometimes there is no Nash equilibrium, and sometimes there are several (i.e., several sets of strategies are stable and self-enforcing). A few more examples will help to clarify this.

The Product Choice Problem Consider the following "product choice" problem. Two breakfast cereal companies face a market in which two new variations of cereal can be successfully introduced—provided that each variation is introduced by only one firm. There is a market for a new "crispy" cereal and a market for a new "sweet" cereal, but each firm has the resources to introduce only one new product. The payoff matrix for the two firms might look like the one in Table 13.3.

---

[5]Our discussion of the Nash equilibrium, and of game theory in general, is at an introductory level. For a more in-depth discussion of game theory and its applications, see James W. Friedman, *Game Theory with Applications to Economics* (New York: Oxford University Press, 1990); Drew Fudenberg and Jean Tirole, *Game Theory* (Cambridge, MA: MIT Press, 1991); and Avinash Dixit and Susan Skeath, *Games of Strategy*, 2nd ed. (New York: Norton, 2004).

[6]A *Stackelberg equilibrium* is also a Nash equilibrium. In the Stackelberg model, however, the rules of the game are different: One firm makes its output decision before its competitor does. Under these rules, each firm is doing the best it can given the decision of its competitor.

TABLE 13.3 Product Choice Problem

		Firm 2	
		Crispy	Sweet
Firm 1	Crispy	−5, −5	10, 10
	Sweet	10, 10	−5, −5

In this game, each firm is indifferent about which product it produces—so long as it does not introduce the same product as its competitor. If coordination were possible, the firms would probably agree to divide the market. But what if the firms must behave *noncooperatively*? Suppose that somehow—perhaps through a news release—Firm 1 indicates that it is about to introduce the sweet cereal, and that Firm 2 (after hearing this) announces its plan to introduce the crispy one. Given the action that it believes its opponent to be taking, neither firm has an incentive to deviate from its proposed action. If it takes the proposed action, its payoff is 10, but if it deviates—and its opponent's action remains unchanged—its payoff will be −5. Therefore, the strategy set given by the bottom left-hand corner of the payoff matrix is stable and constitutes a Nash equilibrium: Given the strategy of its opponent, each firm is doing the best it can and has no incentive to deviate.

Note that the upper right-hand corner of the payoff matrix is also a Nash equilibrium, which might occur if Firm 1 indicated that it was about to produce the crispy cereal. Each Nash equilibrium is stable because *once the strategies are chosen*, no player will unilaterally deviate from them. However, without more information, we have no way of knowing *which* equilibrium (crispy/sweet vs. sweet/crispy) is likely to result—or if *either* will result. Of course, both firms have a strong incentive to reach *one* of the two Nash equilibria—if they both introduce the same type of cereal, they will both lose money. The fact that the two firms are not allowed to collude does not mean that they will not reach a Nash equilibrium. As an industry develops, understandings often evolve as firms "signal" each other about the paths the industry is to take.

The Beach Location Game Suppose that you (Y) and a competitor (C) plan to sell soft drinks on a beach this summer. The beach is 200 yards long, and sunbathers are spread evenly across its length. You and your competitor sell the same soft drinks at the same prices, so customers will walk to the closest vendor. Where on the beach will you locate, and where do you think your competitor will locate?

If you think about this for a minute, you will see that the only Nash equilibrium calls for both you and your competitor to locate at the same spot in the center of the beach (see Figure 13.1). To see why, suppose your competitor located at some other point (A), which is three quarters of the way to the end of the beach. In that case, you would no longer want to locate in the center; you would locate near your competitor, just to the left. You would thus capture nearly three-fourths of all sales, while your competitor got only the remaining fourth. This outcome is not an equilibrium because your competitor would then want to move to the center of the beach, and you would do the same.

The "beach location game" can help us understand a variety of phenomena. Have you ever noticed how, along a two- or three-mile stretch of road, two or three gas stations or several car dealerships will be located close to each other?

*Mixed Strategies

In all of the games that we have examined so far, we have considered strategies in which players make a specific choice or take a specific action: advertise or don't advertise, set a price of $4 or a price of $6, and so on. Strategies of this kind are called **pure strategies**. There are games, however, in which a pure strategy is not the best way to play.

• **pure strategy** Strategy in which a player makes a specific choice or takes a specific action.

Matching Pennies An example is the game of "Matching Pennies." In this game, each player chooses heads or tails and the two players reveal their coins at the same time. If the coins match (i.e., both are heads or both are tails), Player *A* wins and receives a dollar from Player *B*. If the coins do not match, Player *B* wins and receives a dollar from Player *A*. The payoff matrix is shown in Table 13.6.

Note that there is no Nash equilibrium in pure strategies for this game. Suppose, for example, that Player *A* chose the strategy of playing heads. Then Player *B* would want to play tails. But if Player *B* plays tails, Player *A* would also want to play tails. No combination of heads or tails leaves both players satisfied—one player or the other will always want to change strategies.

Although there is no Nash equilibrium in pure strategies, there is a Nash equilibrium in **mixed strategies**: *strategies in which players make random choices among two or more possible actions, based on sets of chosen probabilities*. In this game, for example, Player *A* might simply flip the coin, thereby playing heads with probability 1/2 and playing tails with probability 1/2. In fact, if Player *A* follows this strategy and Player *B* does the same, we will have a Nash equilibrium: Both players will be doing the best they can given what the opponent is doing. Note that although the outcome is random, the *expected payoff* is 0 for each player.

• **mixed strategy** Strategy in which a player makes a random choice among two or more possible actions, based on a set of chosen probabilities.

It may seem strange to play a game by choosing actions randomly. But put yourself in the position of Player *A* and think what would happen if you followed a strategy *other* than just flipping the coin. Suppose you decided to play heads. If Player *B* knows this, she would play tails and you would lose. Even if Player *B* didn't know your strategy, if the game were played repeatedly, she could eventually discern your pattern of play and choose a strategy that countered it. Of course, you would then want to change your strategy—which is why this would not be a Nash equilibrium. Only if you and your opponent both choose heads or tails randomly with probability 1/2 would neither of you have any incentive to change strategies. (You can check that the use of different probabilities, say 3/4 for heads and 1/4 for tails, does not generate a Nash equilibrium.)

One reason to consider mixed strategies is that some games (such as "Matching Pennies") do not have any Nash equilibria in pure strategies. It can be shown, however, that once we allow for mixed strategies, *every* game has

TABLE 13.6 Matching Pennies

		Player B	
		Heads	**Tails**
Player A	Heads	1, –1	–1, 1
	Tails	–1, 1	1, –1

at least one Nash equilibrium.[7] Mixed strategies, therefore, provide solutions to games when pure strategies fail. Of course, whether solutions involving mixed strategies are reasonable will depend on the particular game and players. Mixed strategies are likely to be very reasonable for "Matching Pennies," poker, and other such games. A firm, on the other hand, might not find it reasonable to believe that its competitor will set its price randomly.

The Battle of the Sexes Some games have Nash equilibria both in pure strategies and in mixed strategies. An example is "The Battle of the Sexes," a game that you might find familiar. It goes like this. Jim and Joan would like to spend Saturday night together but have different tastes in entertainment. Jim would like to go to the opera, but Joan prefers mud wrestling. As the payoff matrix in Table 13.7 shows, Jim would most prefer to go to the opera with Joan, but prefers watching mud wrestling with Joan to going to the opera alone, and similarly for Joan.

First, note that there are two Nash equilibria in pure strategies for this game—the one in which Jim and Joan both watch mud wrestling, and the one in which they both go to the opera. Joan, of course, would prefer the first of these outcomes and Jim the second, but both outcomes are equilibria—neither Jim nor Joan would want to change his or her decision, given the decision of the other.

This game also has an equilibrium in mixed strategies: Joan chooses wrestling with probability 2/3 and opera with probability 1/3, and Jim chooses wrestling with probability 1/3 and opera with probability 2/3. You can check that if Joan uses this strategy, Joan cannot do better with any other strategy, and vice versa.[8] The outcome is random, and Jim and Joan will each have an expected payoff of 2/3.

Should we expect Jim and Joan to use these mixed strategies? Unless they're very risk loving or in some other way a strange couple, probably not. By agreeing to either form of entertainment, each will have a payoff of at least 1, which exceeds the expected payoff of 2/3 from randomizing. In this game as in many others, mixed strategies provide another solution, but not a very realistic one. Hence, for the remainder of this chapter we will focus on pure strategies.

TABLE 13.7 The Battle of the Sexes

		Jim	
		Wrestling	**Opera**
Joan	Wrestling	2, 1	0, 0
	Opera	0, 0	1, 2

[7]More precisely, every game with a finite number of players and a finite number of actions has at least one Nash equilibrium. For a proof, see David M. Kreps, *A Course in Microeconomic Theory* (Princeton, NJ: Princeton University Press, 1990), p. 409.

[8]Suppose Joan randomizes, letting p be the probability of wrestling and $(1 - p)$ the probability of opera. Because Jim is using probabilities of 1/3 for wrestling and 2/3 for opera, the probability that both will choose wrestling is $(1/3)p$, and the probability that both will choose opera is $(2/3)(1 - p)$. Thus, Joan's expected payoff is $2(1/3)p + 1(2/3)(1 - p) = (2/3)p + 2/3 - (2/3)p = 2/3$. This payoff is independent of p, so Joan cannot do better in terms of expected payoff no matter what she chooses.

13.4 REPEATED GAMES

We saw in Chapter 12 that in oligopolistic markets, firms often find themselves in a prisoners' dilemma when making output or pricing decisions. Can firms find a way out of this dilemma, so that oligopolistic coordination and cooperation (whether explicit or implicit) could prevail?

To answer this question, we must recognize that the prisoners' dilemma, as we have described it so far, is limited: Although some prisoners may have only one opportunity in life to confess or not, most firms set output and price over and over again. In real life, firms play **repeated games**: Actions are taken and payoffs received over and over again. In repeated games, strategies can become more complex. For example, with each repetition of the prisoners' dilemma, each firm can develop a reputation about its own behavior and can study the behavior of its competitors.

• **repeated game** Game in which actions are taken and payoffs received over and over again.

How does repetition change the likely outcome of the game? Suppose you are Firm 1 in the prisoners' dilemma illustrated by the payoff matrix in Table 13.8. If you and your competitor both charge a high price, you will both make a higher profit than if you both charged a low price. However, you are afraid to charge a high price because if your competitor charges a low price, you will lose money and, to add insult to injury, your competitor will get rich. But suppose this game is repeated over and over again—for example, you and your competitor simultaneously announce your prices on the first day of every month. Should you then play the game differently, perhaps changing your price over time in response to your competitor's behavior?

In an interesting study, Robert Axelrod asked game theorists to come up with the best strategy they could think of to play this game in a repeated manner.[9] (A possible strategy might be: "I'll start off with a high price, then lower my price. But then if my competitor lowers his price, I'll raise mine for a while before lowering it again, etc.") Then, in a computer simulation, Axelrod played these strategies off against one another to see which worked best.

Tit-for-Tat Strategy As you would expect, any given strategy would work better against some strategies than it would against others. The objective, however, was to find the strategy that was most robust—that would work best on average against *all*, or almost all, other strategies. The result was surprising. The strategy that worked best was an extremely simple **tit-for-tat strategy**: I start out with a high price, which I maintain so long as you continue to "cooperate" and also charge a high price. As soon as you lower your price, however, I follow suit and lower mine. If you later decide to cooperate and raise your price again, I'll immediately raise my price as well.

• **tit-for-tat strategy** Repeated-game strategy in which a player responds in kind to an opponent's previous play, cooperating with cooperative opponents and retaliating against uncooperative ones.

TABLE 13.8 Pricing Problem

		Firm 2	
		Low price	High price
Firm 1	Low price	10, 10	100, –50
	High price	–50, 100	50, 50

[9]See Robert Axelrod, *The Evolution of Cooperation* (New York: Basic Books, 1984).

Why does this tit-for-tat strategy work best? In particular, can I expect that using the tit-for-tat strategy will induce my competitor to behave cooperatively (and charge a high price)?

Infinitely Repeated Game Suppose the game is *infinitely repeated*. In other words, my competitor and I repeatedly set prices month after month, *forever*. Cooperative behavior (i.e., charging a high price) is then the rational response to a tit-for-tat strategy. (This assumes that my competitor knows, or can figure out, that I am using a tit-for-tat strategy.) To see why, suppose that in one month my competitor sets a low price and undercuts me. In that month he will make a large profit. But my competitor knows that the following month I will set a low price, so that his profit will fall and will remain low as long as we both continue to charge a low price. Because the game is infinitely repeated, the cumulative loss of profits that results must outweigh any short-term gain that accrued during the first month of undercutting. Thus, it is not rational to undercut.

In fact, with an infinitely repeated game, my competitor need not even be sure that I am playing tit-for-tat to make cooperation its own rational strategy. Even if my competitor believes there is only *some* chance that I am playing tit-for-tat, he will still find it rational to start by charging a high price and maintain it as long as I do. Why? With infinite repetition of the game, the *expected* gains from cooperation will outweigh those from undercutting. This will be true even if the probability that I am playing tit-for-tat (and so will continue cooperating) is small.

Finite Number of Repetitions Now suppose the game is repeated a *finite* number of times—say, N months. (N can be large as long as it is finite.) If my competitor (Firm 2) is rational *and believes that I am rational*, he will reason as follows: "Because Firm 1 is playing tit-for-tat, I (Firm 2) cannot undercut—that is, *until the last month*. I *should* undercut the last month because then I can make a large profit that month, and afterward the game is over, so Firm 1 cannot retaliate. Therefore, I will charge a high price until the last month, and then I will charge a low price."

However, since I (Firm 1) have also figured this out, I also plan to charge a low price in the last month. Of course, Firm 2 can figure this out as well, and therefore *knows* that I will charge a low price in the last month. But then what about the next-to-last month? Because there will be no cooperation in the last month, anyway, Firm 2 figures that it should undercut and charge a low price in the next-to-last month. But, of course, I have figured this out too, so I *also* plan to charge a low price in the next-to-last month. And because the same reasoning applies to each preceding month, the game *unravels*: The only rational outcome is for both of us to charge a low price every month.

Tit-for-Tat in Practice Since most of us do not expect to live forever, the unravelling argument would seem to make the tit-for-tat strategy of little value, leaving us stuck in the prisoners' dilemma. In practice, however, tit-for-tat can sometimes work and cooperation can prevail. There are two primary reasons.

First, most managers don't know how long they will be competing with their rivals, and this also serves to make cooperative behavior a good strategy. If the end point of the repeated game is unknown, the unravelling argument that begins with a clear expectation of undercutting in the last month no longer applies. As with an infinitely repeated game, it will be rational to play tit-for-tat.

Second, my competitor might have some doubt about the extent of my rationality. Suppose my competitor *thinks* (and he need not be certain) that I am playing tit-for-tat. He also thinks that *perhaps* I am playing tit-for-tat "blindly," or with limited rationality, in the sense that I have failed to work out the logical implications of a finite time horizon as discussed above. My competitor thinks, for example, that perhaps I have not figured out that he will undercut me in the last month, so that I should also charge a low price in the last month, and so on. "*Perhaps*," thinks my competitor, "Firm 1 will play tit-for-tat blindly, charging a high price as long as I charge a high price." Then (if the time horizon is long enough), it is rational for my competitor to maintain a high price until the last month (when he will undercut me).

Note that we have stressed the word *perhaps*. My competitor need not be sure that I am playing tit-for-tat "blindly," or even that I am playing tit-for-tat at all. Just the *possibility* can make cooperative behavior a good strategy (until near the end) if the time horizon is long enough. Although my competitor's conjecture about how I am playing the game might be wrong, cooperative behavior is profitable *in expected value terms*. With a long time horizon, the sum of current and future profits, weighted by the probability that the conjecture is correct, can exceed the sum of profits from price competition, even if my competitor is the first to undercut. After all, if I am wrong and my competitor charges a low price, I can shift my strategy at the cost of only one period's profit—a minor cost in light of the substantial profit that I can make if we both choose to set a high price.

Thus, in a repeated game, the prisoners' dilemma can have a cooperative outcome. In most markets, the game is in fact repeated over a long and uncertain length of time, and managers have doubts about how "perfectly rationally" they and their competitors operate. As a result, in some industries, particularly those in which only a few firms compete over a long period under stable demand and cost conditions, cooperation prevails, even though no contractual arrangements are made. (The water meter industry, discussed below, is an example.) In many other industries, however, there is little or no cooperative behavior.

Sometimes cooperation breaks down or never begins because there are too many firms. More often, failure to cooperate is the result of rapidly shifting demand or cost conditions. Uncertainties about demand or costs make it difficult for the firms to reach an implicit understanding of what cooperation should entail. (Remember that an *explicit* understanding, arrived at through meetings and discussions, could lead to an antitrust violation.) Suppose, for example, that cost differences or different beliefs about demand lead one firm to conclude that cooperation means charging \$50 while a second firm thinks it means \$40. If the second firm charges \$40, the first firm might view that as a grab for market share and respond in tit-for-tat fashion with a \$35 price. A price war could then develop.

EXAMPLE 13.2 Oligopolistic Cooperation in the Water Meter Industry

For some four decades, almost all the water meters sold in the United States have been produced by four American companies: Rockwell International, Badger Meter, Neptune Water Meter Company, and Hersey Products. Rockwell has had about a 35-percent share of the market, and the other three firms have together had about a 50- to 55-percent share.[10]

[10]This example is based in part on Nancy Taubenslag, "Rockwell International," Harvard Business School Case No. 9-383-019, July 1983. In 1979, Neptune Water Meter Company was acquired by Wheelabrator-Frye. Hersey Products is a small privately held company.

Most buyers of water meters are municipal water utilities, who install the meters in residential and commercial establishments in order to measure water consumption and bill consumers accordingly. Because the cost of meters is a small part of the total cost of providing water, utilities are concerned mainly that the meters be accurate and reliable. Price is not a primary issue, and demand is very inelastic. Demand is also very stable; because every residence or business must have a water meter, demand grows slowly along with the population.

In addition, utilities tend to have long-standing relationships with suppliers and are reluctant to shift from one to another. Because any new entrant will find it difficult to lure customers from existing firms, this creates a barrier to entry. Substantial economies of scale create a second barrier to entry: To capture a significant share of the market, a new entrant must invest in a large factory. This requirement virtually precludes entry by new firms.

With inelastic and stable demand and little threat of entry by new firms, the existing four firms could earn substantial monopoly profits if they set prices cooperatively. If, on the other hand, they compete aggressively, with each firm cutting price to increase its own share of the market, profits would fall to nearly competitive levels. The firms thus face a prisoners' dilemma. Can cooperation prevail?

It can and *has* prevailed. Remember that the same four firms have been playing a *repeated game* for decades. Demand has been stable and predictable, and over the years, the firms have been able to assess their own and each other's costs. In this situation, tit-for-tat strategies work well: It pays each firm to cooperate as long as its competitors are cooperating.

As a result, the four firms operate as though they were members of a country club. There is rarely an attempt to undercut price, and each firm appears satisfied with its share of the market. While the business may appear dull, it is certainly profitable. All four firms have been earning returns on their investments that far exceed those in more competitive industries.

EXAMPLE 13.3 Competition and Collusion in the Airline Industry

In March 1983, American Airlines, whose president, Robert Crandall, had become notable for his use of the telephone (see Example 10.5—page 384), proposed that all airlines adopt a uniform fare schedule based on mileage. The rate per mile would depend on the length of the trip, with the lowest rate of 15 cents per mile for trips over 2500 miles, higher rates for shorter trips, and the highest rate, 53 cents per mile, for trips under 250 miles. For example, a one-way coach ticket from Boston to Chicago, a distance of 932 miles, would cost $233 (based on a rate of 25 cents per mile for trips between 751 and 1000 miles).

This proposal would have done away with the many different fares (some heavily discounted) then available. The cost of a ticket from one city to another would depend only on the number of miles between those cities. As a senior vice-president of American Airlines said, "The new streamlined fare structure will help reduce fare confusion." Most other major airlines reacted favorably to the plan and began to adopt it. A vice-president of TWA said, "It's a good move. It's very businesslike." United Airlines quickly announced that it would adopt the plan on routes where it competes with American, which included most of its system, and TWA and Continental said that they would adopt it for all their routes.[11]

Why did American propose this plan, and what made it so attractive to the other airlines? Was it really to "help reduce fare confusion"? No, the aim was to reduce price competition and achieve a collusive pricing arrangement. Prices had been driven down by competitive undercutting, as airlines competed for market share. And as Robert Crandall had learned less than a year earlier, fixing prices over the telephone is illegal. Instead, the companies would implicitly fix prices by agreeing to use the same fare-setting formula.

The plan failed, a victim of the prisoners' dilemma. Only two weeks after the plan was announced and adopted by most airlines, Pan Am, which was dissatisfied with its small share of the U.S. market, dropped its fares. American, United, and TWA, afraid of losing their own shares of the market, quickly dropped their fares to match Pan Am. The price-cutting continued, and fortunately for consumers, the plan was soon dead.

American Airlines introduced another simplified, four-tier fare structure in April 1992, which was quickly adopted by most major carriers. But it, too, soon fell victim to competitive discounts. In May 1992, Northwest Airlines announced a "kids fly free" program, and American responded with a summer half-price sale, which other carriers matched. As a result, the airline industry lost billions.

Why is airline pricing so intensively competitive? Airlines plan route capacities two or more years into the future, but they make pricing decisions over short horizons—month by month or even week by week. In the short run, the marginal cost of adding passengers to a flight is very low—essentially the cost of a soft drink and a bag of peanuts. Each airline, therefore, has an incentive to lower fares in order to capture passengers from its competitors. In addition, the demand for air travel often fluctuates unpredictably. Such factors as these stand in the way of implicit price cooperation.

Thus, aggressive competition has continued to be the rule in the airline industry. In 2002, for example, both American Airlines and US Airways introduced price increases, only to abandon them when other carriers refused to cooperate. In fact, for several reasons, pricing has become even more competitive in recent years. First, discount airlines—such as Southwest and JetBlue—have attracted millions of price-conscious consumers and forced the major carriers to cut fares. Second, during periods of sluggish demand, airlines are compelled to reduce prices in order to attract consumers. Finally, Internet services such as Expedia, Orbitz, and Travelocity have promoted "fare shopping" by online consumers and encouraged more competitive pricing. These developments have forced several major airlines into bankruptcy and resulted in record losses for the industry.

13.5 SEQUENTIAL GAMES

In most of the games we have discussed so far, both players move at the same time. In the Cournot model of duopoly, for example, both firms set output at the same time. In **sequential games**, players move in turn. The Stackelberg model

• **sequential game** Game in which players move in turn, responding to each other's actions and reactions.

[11] "American to Base Fares on Mileage," *New York Times*, March 15, 1983; "Most Big Airlines Back American's Fare Plan," *New York Times*, March 17, 1983.

TABLE 13.9 Modified Product Choice Problem

		Firm 2	
		Crispy	**Sweet**
Firm 1	Crispy	−5, −5	10, 20
	Sweet	20, 10	−5, −5

discussed in Chapter 12 is an example of a sequential game; one firm sets output before the other does. There are many other examples: an advertising decision by one firm and the response by its competitor; entry-deterring investment by an incumbent firm and the decision whether to enter the market by a potential competitor; or a new government regulatory policy and the investment and output response of the regulated firms.

We will look at a variety of sequential games in the remainder of this chapter. As we will see, they are often easier to analyze than games in which the players move at the same time. In a sequential game, the key is to think through the possible actions and rational reactions of each player.

As a simple example, let's return to the product choice problem first discussed in Section 13.3. This problem involves two companies facing a market in which two new variations of breakfast cereal can be successfully introduced as long as each firm introduces only one variation. This time, let's change the payoff matrix slightly. As Table 13.9 shows, the new sweet cereal will inevitably be a better seller than the new crispy cereal, earning a profit of 20 rather than 10 (perhaps because consumers prefer sweet things to crispy things). Both new cereals will still be profitable, however, as long as each is introduced by only one firm. (Compare Table 13.9 with Table 13.3—page 485.)

Suppose that both firms, in ignorance of each other's intentions, must announce their decisions independently and simultaneously. In that case, both will probably introduce the sweet cereal—and both will lose money.

Now suppose that Firm 1 can gear up its production faster and introduce its new cereal first. We now have a sequential game: Firm 1 introduces a new cereal, and *then* Firm 2 introduces one. What will be the outcome of this game? When making its decision, Firm 1 must consider the rational response of its competitor. It knows that whichever cereal it introduces, Firm 2 will introduce the other kind. Thus it will introduce the sweet cereal, knowing that Firm 2 will respond by introducing the crispy one.

The Extensive Form of a Game

Although this outcome can be deduced from the payoff matrix in Table 13.9, sequential games are sometimes easier to visualize if we represent the possible moves in the form of a decision tree. This representation is called the **extensive form of a game** and is shown in Figure 13.2. The figure shows the possible choices of Firm 1 (introduce a crispy or a sweet cereal) and the possible responses of Firm 2 to each of those choices. The resulting payoffs are given at the end of each branch. For example, if Firm 1 produces a crispy cereal and Firm 2 responds by also producing a crispy cereal, each firm will have a payoff of −5.

• **extensive form of a game** Representation of possible moves in a game in the form of a decision tree.

To find the solution to the extensive form game, work backward from the end. For Firm 1, the best sequence of moves is the one in which it earns 20 and

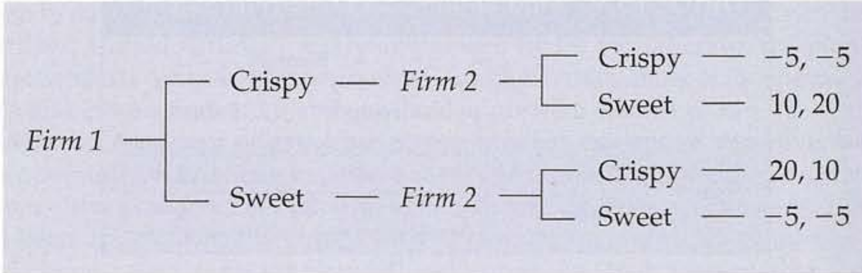

FIGURE 13.2 Product Choice Game in Extensive Form

Firm 2 earns 10. Thus it can deduce that it should produce the sweet cereal because Firm 2's best response is then to produce the crispy cereal.

The Advantage of Moving First

In this product-choice game, there is a clear advantage to moving first: By introducing the sweet cereal, Firm 1 leaves Firm 2 little choice but to introduce the crispy one. This is much like the first-mover advantage that we saw in the Stackelberg model in Chapter 12. In that model, the firm that moves first can choose a large level of output, thereby giving its competitor little choice but to choose a small level.

In §12.2, we explain that the Stackelberg model is an oligopoly model in which one firm sets its output before other firms do.

Recall that in §12.2, we explain that in the Cournot model each firm treats the output of its competitors as fixed, and that all firms simultaneously decide how much to produce.

To clarify the nature of this first-mover advantage, it will be useful to review the Stackelberg model and compare it to the Cournot model in which both firms choose their outputs simultaneously. As in Chapter 12, we will use the example in which two duopolists face the market demand curve

$$P = 30 - Q$$

where Q is the total production, i.e., $Q = Q_1 + Q_2$. As before, we will also assume that both firms have zero marginal cost. Recall that the Cournot equilibrium is then $Q_1 = Q_2 = 10$, so that $P = 10$ and each firm earns a profit of 100. Recall also that if the two firms colluded, they would set $Q_1 = Q_2 = 7.5$, so that $P = 15$ and each firm earns a profit of 112.50. Finally, recall from Section 12.3 that in the Stackelberg model, in which Firm 1 moves first, the outcome is $Q_1 = 15$ and $Q_2 = 7.5$, so that $P = 7.50$ and the firms' profits are 112.50 and 56.25, respectively.

These and a few other possible outcomes are summarized in the payoff matrix in Table 13.10. If both firms move simultaneously, the only solution to the

TABLE 13.10 Choosing Output

		Firm 2		
		7.5	10	15
Firm 1	7.5	112.50, 112.50	93.75, 125	56.25, 112.50
	10	125, 93.75	100, 100	50, 75
	15	112.50, 56.25	75, 50	0, 0

game is that both produce 10 and earn 100. In this Cournot equilibrium each firm is doing the best it can given what its competitor is doing. If Firm 1 moves first, however, it knows that its decision will constrain Firm 2's choice. Observe from the payoff matrix that if Firm 1 sets $Q_1 = 7.5$, Firm 2's best response will be to set $Q_2 = 10$. This will give Firm 1 a profit of 93.75 and Firm 2 a profit of 125. If Firm 1 sets $Q_1 = 10$, Firm 2 will set $Q_2 = 10$, and both firms will earn 100. But if Firm 1 sets $Q_1 = 15$, Firm 2 will set $Q_2 = 7.5$, so that Firm 1 earns 112.50, and Firm 2 earns 56.25. Therefore, the most that Firm 1 can earn is 112.50, and it does so by setting $Q_1 = 15$. Compared to the Cournot outcome, when Firm 1 moves first, it does better—and Firm 2 does much worse.

13.6 THREATS, COMMITMENTS, AND CREDIBILITY

The product choice problem and the Stackelberg model are two examples of how a firm that moves first can create a *fait accompli* that gives it an advantage over its competitor. In this section, we'll take a broader look at the advantage that a firm can have by moving first. We'll also consider what determines *which* firm goes first. We will focus on the following question: *What actions can a firm take to gain advantage in the marketplace?* For example, how might a firm deter entry by potential competitors, or induce existing competitors to raise prices, reduce output, or leave the market altogether?

Recall that in the Stackelberg model, the firm that moved first gained an advantage by *committing itself to a large output*. Making a commitment—constraining its future behavior—is crucial. To see why, suppose that the first mover (Firm 1) could later change its mind in response to what Firm 2 does. What would happen? Clearly, Firm 2 would produce a large output. Why? Because it knows that Firm 1 will respond by reducing the output that it first announced. The only way that Firm 1 can gain a first-mover advantage is by committing itself. In effect, *Firm 1 constrains Firm 2's behavior by constraining its own behavior*.

The idea of constraining your own behavior to gain an advantage may seem paradoxical, but we'll soon see that it is not. Let's consider a few examples.

First, let's return once more to the product-choice problem shown in Table 13.9. The firm that introduces its new breakfast cereal first will do best. *But which firm will introduce its cereal first*? Even if both firms require the same amount of time to gear up production, each has an incentive to *commit itself first to the sweet cereal*. The key word is *commit*. If Firm 1 simply announces it will produce the sweet cereal, Firm 2 will have little reason to believe it. After all, Firm 2, knowing the incentives, can make the same announcement louder and more vociferously. Firm 1 must constrain its own behavior in some way that convinces Firm 2 that Firm 1 has *no choice* but to produce the sweet cereal. Firm 1 might launch an expensive advertising campaign describing the new sweet cereal well before its introduction, thereby putting its reputation on the line. Firm 1 might also sign a contract for the forward delivery of a large quantity of sugar (and make the contract public, or at least send a copy to Firm 2). The idea is for Firm 1 to *commit itself* to produce the sweet cereal. Commitment is a strategic move that will induce Firm 2 to make the decision that Firm 1 wants it to make—namely, to produce the crispy cereal.

Why can't Firm 1 simply *threaten* Firm 2, vowing to produce the sweet cereal even if Firm 2 does the same? Because Firm 2 has little reason to believe the threat—and can make the same threat itself. A threat is useful only if it is credible. The following example should help make this clear.

TABLE 13.11 Pricing of Computers and Word Processors

		Firm 2	
		High price	**Low price**
Firm 1	High price	100, 80	80, 100
	Low price	20, 0	10, 20

Empty Threats

Suppose Firm 1 produces personal computers that can be used both as word processors and to do other tasks. Firm 2 produces only dedicated word processors. As the payoff matrix in Table 13.11 shows, as long as Firm 1 charges a high price for its computers, both firms can make a good deal of money. Even if Firm 2 charges a low price for its word processors, many people will still buy Firm 1's computers (because they can do so many other things), although some buyers will be induced by the price differential to buy the dedicated word processor instead. However, if Firm 1 charges a low price, Firm 2 will also have to charge a low price (or else make zero profit), and the profit of both firms will be significantly reduced.

Firm 1 would prefer the outcome in the upper left-hand corner of the matrix. For Firm 2, however, charging a low price is clearly a dominant strategy. Thus the outcome in the upper right-hand corner will prevail (no matter which firm sets its price first).

Firm 1 would probably be viewed as the "dominant" firm in this industry because its pricing actions will have the greatest impact on overall industry profits. Can Firm 1 induce Firm 2 to charge a high price by *threatening* to charge a low price if Firm 2 charges a low price? No, as the payoff matrix in Table 13.11 makes clear: *Whatever* Firm 2 does, Firm 1 will be much worse off if it charges a low price. As a result, its threat is not credible.

Commitment and Credibility

Sometimes firms can make threats credible. To see how, consider the following example. Race Car Motors, Inc., produces cars, and Far Out Engines, Ltd., produces specialty car engines. Far Out Engines sells most of its engines to Race Car Motors, and a few to a limited outside market. Nonetheless, it depends heavily on Race Car Motors and makes its production decisions in response to Race Car's production plans.

We thus have a sequential game in which Race Car is the "leader." It will decide what kind of cars to build, and Far Out Engines will then decide what kind of engines to produce. The payoff matrix in Table 13.12(a) shows the possible outcomes of this game. (Profits are in millions of dollars.) Observe that Race Car will do best by deciding to produce small cars. It knows that in response to this decision, Far Out will produce small engines, most of which Race Car will then buy. As a result, Far Out will make \$3 million and Race Car \$6 million.

Far Out, however, would much prefer the outcome in the lower right-hand corner of the payoff matrix. If it could produce big engines, *and* if Race Car produced big cars and thus bought the big engines, it would make \$8 million. (Race Car,

TABLE 13.12(a) Production Choice Problem

		Race Car Motors	
		Small cars	**Big cars**
Far Out Engines	Small engines	3, 6	3, 0
	Big engines	1, 1	8, 3

however, would make only $3 million.) Can Far Out induce Race Car to produce big cars instead of small ones?

Suppose Far Out *threatens* to produce big engines no matter what Race Car does; suppose, too, that no other engine producer can easily satisfy the needs of Race Car. If Race Car believed Far Out's threat, it would produce big cars: Otherwise, it would have trouble finding engines for its small cars and would earn only $1 million instead of $3 million. But the threat is not credible: Once Race Car responded by announcing its intentions to produce small cars, Far Out would have no incentive to carry out its threat.

Far Out can make its threat credible by visibly and irreversibly reducing some of its own payoffs in the matrix, thereby constraining its own choices. In particular, Far Out must reduce its profits from small engines (the payoffs in the top row of the matrix). It might do this by *shutting down or destroying some of its small engine production capacity*. This would result in the payoff matrix shown in Table 13.12(b). Now Race Car *knows* that whatever kind of car it produces, Far Out will produce big engines. If Race Car produces the small cars, Far Out will sell the big engines as best it can to other car producers and settle for making only $1 million. But this is better than making no profits by producing small engines. Because Race Car will have to look elsewhere for engines, its profit will also be lower ($1 million). Now it is clearly in Race Car's interest to produce large cars. By taking an action that *seemingly puts itself at a disadvantage*, Far Out has improved its outcome in the game.

Although strategic commitments of this kind can be effective, they are risky and depend heavily on having accurate knowledge of the payoff matrix and the industry. Suppose, for example, that Far Out commits itself to producing big engines but is surprised to find that another firm can produce small engines at a low cost. The commitment may then lead Far Out to bankruptcy rather than continued high profits.

The Role of Reputation Developing the right kind of *reputation* can also give one a strategic advantage. Again, consider Far Out Engines' desire to produce big engines for Race Car Motors' big cars. Suppose that the managers of Far Out

TABLE 13.12(b) Modified Production Choice Problem

		Race Car Motors	
		Small cars	**Big cars**
Far Out Engines	Small engines	0, 6	0, 0
	Big engines	1, 1	8, 3

Engines develop a reputation for being irrational—perhaps downright crazy. They threaten to produce big engines no matter what Race Car Motors does (refer to Table 13.12a). Now the threat might be credible without any further action; after all, you can't be sure that an irrational manager will always make a profit-maximizing decision. In gaming situations, the party that is known (or thought) to be a little crazy can have a significant advantage.

Developing a reputation can be an especially important strategy in a repeated game. A firm might find it advantageous to behave irrationally for several plays of the game. This might give it a reputation that will allow it to increase its long-run profits substantially.

Bargaining Strategy

Our discussion of commitment and credibility also applies to bargaining problems. The outcome of a bargaining situation can depend on the ability of either side to take an action that alters its relative bargaining position.

For example, consider two firms that are each planning to introduce one of two products which are complementary goods. As the payoff matrix in Table 13.13 shows, Firm 1 has a cost advantage over Firm 2 in producing *A*. Therefore, if both firms produce *A*, Firm 1 can maintain a lower price and earn a higher profit. Similarly, Firm 2 has a cost advantage over Firm 1 in producing product *B*. If the two firms could agree about who will produce what, the rational outcome would be the one in the upper right-hand corner: Firm 1 produces *A*, Firm 2 produces *B*, and both firms make profits of 50. Indeed, *even without cooperation*, this outcome will result whether Firm 1 or Firm 2 moves first or both firms move simultaneously. Why? Because producing *B* is a dominant strategy for Firm 2, so (*A*, *B*) is the only Nash equilibrium.

Firm 1, of course, would prefer the outcome in the lower left-hand corner of the payoff matrix. But in the context of this limited set of decisions, it cannot achieve that outcome. Suppose, however, that Firms 1 and 2 are also bargaining over a second issue—whether to join a research consortium that a third firm is trying to form. Table 13.14 shows the payoff matrix for this decision problem. Clearly, the dominant strategy is for both firms to enter the consortium, thereby increasing profits to 40.

Now suppose that Firm 1 *links the two bargaining problems* by announcing that it will join the consortium *only* if Firm 2 agrees to produce product *A*. In this case, it is indeed in Firm 2's interest to produce *A* (with Firm 1 producing *B*) in return for Firm 1's participation in the consortium. This example illustrates how combining issues in a bargaining agenda can sometimes benefit one side at the other's expense.

As another example, consider bargaining over the price of a house. Suppose I, as a potential buyer, do not want to pay more than $200,000 for a house that is actually worth $250,000 to me. The seller is willing to part with the house at any

TABLE 13.13 Production Decision

		Firm 2	
		Produce *A*	Produce *B*
Firm 1	Produce A	40, 5	50, 50
	Produce B	60, 40	5, 45

TABLE 13.14 Decision to Join Consortium

		Firm 2: Work alone	Firm 2: Enter consortium
Firm 1	Work alone	10, 10	10, 20
	Enter consortium	20, 10	40, 40

price above $180,000 but would like to receive the highest price she can. If I am the only bidder for the house, how can I make the seller think that I will walk away rather than pay more than $200,000?

I might declare that I will never, ever pay more than $200,000 for the house. But is such a promise credible? It may be if the seller knows that I have a reputation for toughness and that I have never reneged on a promise of this sort. But suppose I have no such reputation. Then the seller knows that I have every incentive to make the promise (making it costs nothing) but little incentive to keep it. (This will probably be our only business transaction together.) As a result, this promise by itself is not likely to improve my bargaining position.

The promise can work, however, if it is combined with an action that gives it credibility. Such an action must reduce my flexibility—limit my options—so that I have no choice but to keep the promise. One possibility would be to make an enforceable bet with a third party—for example, "If I pay more than $200,000 for that house, I'll pay you $60,000." Alternatively, if I am buying the house on behalf of my company, the company might insist on authorization by the Board of Directors for a price above $200,000, and announce that the board will not meet again for several months. In both cases, my promise becomes credible because I have destroyed my ability to break it. The result is less flexibility—and more bargaining power.

EXAMPLE 13.4 Wal-Mart Stores' Preemptive Investment Strategy

Wal-Mart Stores, Inc., is an enormously successful chain of discount retail stores started by Sam Walton in 1969.[12] Its success was unusual in the industry. During the 1960s and 1970s, rapid expansion by existing firms and the entry and expansion of new firms made discount retailing increasingly competitive. During the 1970s and 1980s, industry-wide profits fell, and large discount chains—including such giants as King's, Korvette's, Mammoth Mart, W. T. Grant, and Woolco—went bankrupt. Wal-Mart Stores, however, kept on growing and became even more profitable. By the end of 1985, Sam Walton was one of the richest people in the United States.

[12]This example is based in part on information in Pankaj Ghemawat, "Wal-Mart Stores' Discount Operations," Harvard Business School, 1986.

How did Wal-Mart Stores succeed where others failed? The key was Wal-Mart's expansion strategy. To charge less than ordinary department stores and small retail stores, discount stores rely on size, no frills, and high inventory turnover. Through the 1960s, the conventional wisdom held that a discount store could succeed only in a city with a population of 100,000 or more. Sam Walton disagreed and decided to open his stores in small Southwestern towns; by 1970, there were 30 Wal-Mart stores in small towns in Arkansas, Missouri, and Oklahoma. The stores succeeded because Wal-Mart had created 30 "local monopolies." Discount stores that had opened in larger towns and cities were competing with other discount stores, which drove down prices and profit margins. These small towns, however, had room for only one discount operation. Wal-Mart could undercut the nondiscount retailers and never had to worry that another discount store would open and compete with it.

By the mid-1970s, other discount chains realized that Wal-Mart had a profitable strategy: Open a store in a small town that could support only one discount store and enjoy a local monopoly. There are a lot of small towns in the United States, so the issue became who would get to each town first. Wal-Mart now found itself in a *preemption game* of the sort illustrated by the payoff matrix in Table 13.15. As the matrix shows, if Wal-Mart enters a town but Company X does not, Wal-Mart will make 20 and Company X will make 0. Similarly, if Wal-Mart doesn't enter but Company X does, Wal-Mart makes 0 and Company X makes 20. But if Wal-Mart and Company X *both* enter, *they both lose 10*.

This game has two Nash equilibria—the lower left-hand corner and the upper right-hand corner. Which equilibrium results depends on *who moves first*. If Wal-Mart moves first, it can enter, knowing that the rational response of Company X will be not to enter, so that Wal-Mart will be assured of earning 20. *The trick, therefore, is to preempt*—to set up stores in other small towns quickly, before Company X (or Company Y or Z) can do so. That is exactly what Wal-Mart did. By 1986, it had 1009 stores in operation and was earning an annual profit of $450 million. And while other discount chains were going under, Wal-Mart continued to grow. By 1999, Wal-Mart had become the world's largest retailer, with 2454 stores in the United States and another 729 stores in the rest of the world, and had annual sales of $138 billion.

In recent years, Wal-Mart has continued to preempt other retailers by opening new discount stores, warehouse stores (such as Sam's Club), and combination discount and grocery stores (Wal-Mart Supercenters) all over the world. Wal-Mart has been especially aggressive in applying its preemption strategy in other countries. As of 2007, Wal-Mart had about 3800 stores in the United States and about 2800 stores throughout Europe, Latin America, and Asia. Wal-Mart had also become the world's largest private employer, employing more than 1.6 million people worldwide.

TABLE 13.15 The Discount Store Preemption Game

		Company X	
		Enter	Don't enter
Wal-Mart	Enter	−10, −10	20, 0
	Don't enter	0, 20	0, 0

13.7 ENTRY DETERRENCE

Barriers to entry, which are an important source of monopoly power and profits, sometimes arise naturally. For example, economies of scale, patents and licenses, or access to critical inputs can create entry barriers. However, firms themselves can sometimes deter entry by potential competitors.

To deter entry, *the incumbent firm must convince any potential competitor that entry will be unprofitable.* To see how this might be done, put yourself in the position of an incumbent monopolist facing a prospective entrant, Firm X. Suppose that to enter the industry, Firm X will have to pay a (sunk) cost of $80 million to build a plant. You, of course, would like to induce Firm X to stay out of the industry. If X stays out, you can continue to charge a high price and enjoy monopoly profits. As shown in the upper right-hand corner of the payoff matrix in Table 13.16 (a), you would earn $200 million in profits.

In §7.1, we explain that a sunk cost is an expenditure that has been made and cannot be recovered.

If Firm X does enter the market, you must make a decision. You can be "accommodating," maintaining a high price in the hope that X will do the same. In that case, you will earn only $100 million in profit because you will have to share the market. New entrant X will earn a *net* profit of $20 million: $100 million minus the $80 million cost of constructing a plant. (This outcome is shown in the upper left-hand corner of the payoff matrix.) Alternatively, you can increase your production capacity, produce more, and lower your price. The lower price will give you a greater market share and a $20 million increase in revenues. Increasing production capacity, however, will cost $50 million, reducing your net profit to $70 million. Because warfare will also reduce the entrant's revenue by $30 million, it will have a net loss of $10 million. (This outcome is shown in the lower left-hand corner of the payoff matrix.) Finally, if Firm X stays out but you expand capacity and lower price nonetheless, your net profit will fall by $70 million (from $200 million to $130 million): the $50 million cost of the extra capacity and a $20 million reduction in revenue from the lower price with no gain in market share. Clearly this choice, shown in the lower right-hand corner of the matrix, would make no sense.

If Firm X thinks you will be accommodating and maintain a high price after it has entered, it will find it profitable to enter and will do so. Suppose you threaten to expand output and wage a price war in order to keep X out. If X takes the threat seriously, it will not enter the market because it can expect to lose $10 million. The threat, however, is not credible. As Table 13.16(a) shows (and as the potential competitor knows), *once entry has occurred, it will be in your best interest to accommodate and maintain a high price.* Firm X's rational move is to enter the market; the outcome will be the upper left-hand corner of the matrix.

TABLE 13.16(a) Entry Possibilities

		Potential Entrant	
		Enter	**Stay out**
Incumbent	High price (accommodation)	100, 20	200, 0
	Low price (warfare)	70, −10	130, 0

TABLE 13.16(b) Entry Deterrence

		Potential Entrant: Enter	Potential Entrant: Stay out
Incumbent	High price (accommodation)	50, 20	150, 0
	Low price (warfare)	70, –10	130, 0

But what if you can make an irrevocable commitment that will alter your incentives once entry occurs—a commitment that will give you little choice but to charge a low price if entry occurs? In particular, suppose you invest the \$50 million *now*, rather than later, in the extra capacity needed to increase output and engage in competitive warfare should entry occur. Of course, if you later maintain a high price (whether or not X enters), this added cost will reduce your payoff.

We now have a new payoff matrix, as shown in Table 13.16 (b). As a result of your decision to invest in additional capacity, your threat to engage in competitive warfare is *completely credible*. Because you already have the additional capacity with which to wage war, you will do better in competitive warfare than you would by maintaining a high price. Because the potential competitor now knows that entry will result in warfare, it is rational for it to stay out of the market. Meanwhile, having deterred entry, you can maintain a high price and earn a profit of \$150 million.

Can an incumbent monopolist deter entry without making the costly move of installing additional production capacity? Earlier we saw that a reputation for irrationality can bestow a strategic advantage. Suppose the incumbent firm has such a reputation. Suppose also that by means of vicious price-cutting, this firm has eventually driven out every entrant in the past, even though it incurred losses in doing so. Its threat might then be credible: The incumbent's irrationality suggests to the potential competitor that it might be better off staying away.

Of course, if the game described above were to be *indefinitely repeated*, then the incumbent might have a *rational* incentive to engage in warfare whenever entry actually occurs. Why? Because short-term losses from warfare might be outweighed by longer-term gains from preventing entry. Understanding this, the potential competitor might find the incumbent's threat of warfare credible and decide to stay out. Now the incumbent relies on its reputation for being rational—and far-sighted—to provide the credibility needed to deter entry. The success of this strategy depends on the time horizon and the relative gains and losses associated with accommodation and warfare.

We have seen that the attractiveness of entry depends largely on the way incumbents can be expected to react. In general, once entry has occurred, incumbents cannot be expected to maintain output at their pre-entry levels. Eventually, they may back off and reduce output, raising price to a new joint profit-maximizing level. Because potential entrants know this, incumbent firms must create a credible threat of warfare to deter entry. A reputation for irrationality can help. Indeed, this seems to be the basis for much of the entry-preventing behavior that goes on in actual markets. The potential entrant must consider that *rational* industry discipline can break down after entry occurs. By fostering an image of irrationality and belligerence, an incumbent firm might convince potential entrants that the risk of warfare is too high.[13]

Strategic Trade Policy and International Competition

We have seen how a preemptive investment can give a firm an advantage by creating a credible threat to potential competitors. In some situations, a preemptive investment—subsidized or otherwise encouraged by the government—can give a *country* an advantage in international markets and so be an important instrument of trade policy.

Does this conflict with what you have learned about the benefits of free trade? In Chapter 9, for example, we saw how trade restrictions such as tariffs or quotas lead to deadweight losses. In Chapter 16 we go further and show how, in a general way, free trade between people (or between countries) is mutually beneficial. Given the virtues of free trade, how can government intervention in an international market ever be warranted? In certain situations, a country can benefit by adopting policies that give its domestic industries a competitive advantage.

To see how this might occur, consider an industry with substantial economies of scale—one in which a few large firms can produce much more efficiently than many small ones. Suppose that by granting subsidies or tax breaks, the government can encourage domestic firms to expand faster than they would otherwise. This might prevent firms in other countries from entering the world market, so that the domestic industry can enjoy higher prices and greater sales. Such a policy works by creating a credible threat to potential entrants. Large domestic firms, taking advantage of scale economies, would be able to satisfy world demand at a low price; if other firms entered, price would be driven below the point at which they could make a profit.

The Commercial Aircraft Market As an example, consider the international market for commercial aircraft. The development and production of a new line of aircraft are subject to substantial economies of scale; it would not pay to develop a new aircraft unless a firm expected to sell many of them. Suppose that Boeing and Airbus (a European consortium that includes France, Germany, Britain, and Spain) are each considering developing a new aircraft. The ultimate payoff to each firm depends in part on what the other firm does. Suppose it is only economical for one firm to produce the new aircraft. Then the payoffs might look like those in Table 13.17(a).[14]

If Boeing has a head start in the development process, the outcome of the game is the upper right-hand corner of the payoff matrix. Boeing will produce a

[13]There is an analogy here to *nuclear deterrence*. Consider the use of a nuclear threat to deter the former Soviet Union from invading Western Europe during the Cold War. If it invaded, would the United States actually react with nuclear weapons, knowing that the Soviets would then respond in kind? Because it is not rational for the United States to react this way, a nuclear threat might not seem credible. But this assumes that everyone is rational; there is a reason to fear an *irrational* response by the United States. Even if an irrational response is viewed as very improbable, it can be a deterrent, given the costliness of an error. The United States can thus gain by promoting the idea that it might act irrationally, or that events might get out of control once an invasion occurs. This is the "rationality of irrationality." See Thomas Schelling, *The Strategy of Conflict* (Harvard Univ. Press, 1980).

[14]This example is drawn from Paul R. Krugman, "Is Free Trade Passé?" *Journal of Economic Perspectives* 1 (Fall 1987): 131–44.

TABLE 13.17(a) Development of a New Aircraft

		Airbus: Produce	Airbus: Don't produce
Boeing	Produce	−10, −10	100, 0
	Don't produce	0, 100	0, 0

new aircraft, and Airbus, realizing that it will lose money if it does the same, will not. Boeing will then earn a profit of 100.

European governments, of course, would prefer that Airbus produce the new aircraft. Can they change the outcome of this game? Suppose they commit to subsidizing Airbus and make this commitment before Boeing has committed itself to produce. If the European governments commit to a subsidy of 20 to Airbus if it produces the plane *regardless of what Boeing does*, the payoff matrix would change to the one in Table 13.17 (b).

Now Airbus will make money from a new aircraft whether or not Boeing produces one. Boeing knows that even if it commits to producing, Airbus will produce as well, and Boeing will lose money. Thus Boeing will decide not to produce, and the outcome will be the one in the lower left-hand corner of Table 13.17(b). A subsidy of 20, then, changes the outcome from one in which Airbus does not produce and earns 0, to one in which it does produce and earns 120. Of this, 100 is a transfer of profit from the United States to Europe. From the European point of view, subsidizing Airbus yields a high return.

European governments *did* commit to subsidizing Airbus, and during the 1980s, Airbus successfully introduced several new airplanes. The result, however, was not quite the one reflected in our simplified example. Boeing also introduced new airplanes (the 757 and 767 models) that were quite profitable. As commercial air travel grew, it became clear that both companies could profitably develop and sell new airplanes. Nonetheless, Boeing's market share would have been much larger without the European subsidies to Airbus. One study estimated that those subsidies totalled $25.9 billion during the 1980s and found that Airbus would not have entered the market without them.[15]

TABLE 13.17(b) Development of Aircraft after European Subsidy

		Airbus: Produce	Airbus: Don't produce
Boeing	Produce	−10, 10	100, 0
	Don't produce	0, 120	0, 0

[15] "Aid to Airbus Called Unfair in U.S. Study," *New York Times*, September 8, 1990.

This example shows how strategic trade policy can transfer profits from one country to another. Bear in mind, however, that a country that uses such a policy may provoke retaliation from its trading partners. If a trade war results, all countries can end up much worse off. The possibility of such an outcome must be considered before a nation adopts a strategic trade policy.

EXAMPLE 13.5 DuPont Deters Entry in the Titanium Dioxide Industry

Titanium dioxide is a whitener used in paints, paper, and other products. In the early 1970s, DuPont and National Lead each accounted for about a third of U.S. titanium dioxide sales; another seven firms produced the remainder. In 1972, DuPont was considering whether to expand capacity. The industry was changing, and with the right strategy, those changes might enable DuPont to capture more of the market and dominate the industry.[16]

Three factors had to be considered. First, although future demand for titanium dioxide was uncertain, it was expected to grow substantially. Second, the government had announced that new environmental regulations would be imposed. Third, the prices of raw materials used to make titanium dioxide were rising. The new regulations and the higher input prices would have a major effect on production cost and give DuPont a cost advantage, both because its production technology was less sensitive to the change in input prices and because its plants were in areas that made disposal of corrosive wastes much less difficult than for other producers. Because of these cost changes, DuPont anticipated that National Lead and some other producers would have to shut down part of their capacity. DuPont's competitors would in effect have to "re-enter" the market by building new plants. Could DuPont deter them from taking this step?

DuPont considered the following strategy: invest nearly $400 million in increased production capacity to try to capture 64 percent of the market by 1985. The production capacity that would be put on line would be much more than what was actually needed. The idea was to *deter competitors from investing*. Scale economies and movement down the learning curve would give DuPont a cost advantage. This would not only make it hard for other firms to compete, but would make credible the implicit threat that in the future, DuPont would fight rather than accommodate.

The strategy was sensible and seemed to work for a few years. By 1975, however, things began to go awry. First, because demand grew by much less than expected, there was excess capacity industrywide. Second, because the environmental regulations were only weakly enforced, competitors did not have to shut down capacity as expected. Finally, DuPont's strategy led to antitrust action by the Federal Trade Commission in 1978. The FTC claimed that DuPont was attempting to monopolize the market. DuPont won the case, but the decline in demand made its victory moot.

[16]This example is based on Pankaj Ghemawat, "Capacity Expansion in the Titanium Dioxide Industry," *Journal of Industrial Economics* 33 (December 1984): 145–63; and P. Ghemawat, "DuPont in Titanium Dioxide," Harvard Business School, Case No. 9-385-140, June 1986.

EXAMPLE 13.6 Diaper Wars

For more than two decades, the disposable diaper industry in the United States has been dominated by two firms: Procter & Gamble, with an approximately 50-percent market share, and Kimberly-Clark, with another 30–40 percent.[17] How do these firms compete? And why haven't other firms been able to enter and take a significant share of this \$5-billion-per-year market?

Even though there are only two major firms, competition is intense. The competition occurs mostly in the form of *cost-reducing innovation*. The key to success is to perfect the manufacturing process so that a plant can manufacture diapers in high volume and at low cost. This is not as simple as it might seem. Packing cellulose fluff for absorbency, adding an elastic gatherer, and binding, folding, and packaging the diapers—at a rate of about 3000 diapers per minute and at a cost of about 10 cents per diaper—requires an innovative, carefully designed, and finely tuned process. Furthermore, small technological improvements in the manufacturing process can result in a significant competitive advantage. If a firm can shave its production cost even slightly, it can reduce price and capture market share. As a result, both firms are forced to spend heavily on research and development (R&D) in a race to reduce cost.

The payoff matrix in Table 13.18 illustrates this. If both firms spend aggressively on R&D, they can expect to maintain their current market shares. P&G will earn a profit of 40, and Kimberly-Clark (with a smaller market share) will earn 20. If neither firm spends money on R&D, their costs and prices will remain constant and the money saved will become part of profits. P&G's profit will increase to 60 and Kimberly-Clark's to 40. However, if one firm continues to do R&D and the other doesn't, the innovating firm will eventually capture most of its competitor's market share. For example, if Kimberly-Clark does R&D and P&G does not, P&G can expect to lose 20 while Kimberly-Clark's profit increases to 60. The two firms are therefore in a prisoners' dilemma: Spending money on R&D is a dominant strategy for each firm.

Why hasn't cooperative behavior evolved? After all, the two firms have been competing in this market for years, and the demand for diapers is fairly stable. For several reasons, a prisoners' dilemma involving R&D is particularly hard to resolve. First, it is difficult for a firm to monitor its competitor's R&D activities

TABLE 13.18 Competing through R&D

		Kimberly-Clark	
		R&D	No R&D
P&G	R&D	40, 20	80, –20
	No R&D	–20, 60	60, 40

[17]Procter & Gamble makes Pampers, Ultra Pampers, and Luvs. Kimberly-Clark has only one major brand, Huggies.

the way it can monitor price. Second, it can take several years to complete an R&D program that leads to a major product improvement. As a result, tit-for-tat strategies, in which both firms cooperate until one of them "cheats," are less likely to work. A firm may not find out that its competitor has been secretly doing R&D until the competitor announces a new and improved product. By then it may be too late to gear up an R&D program of its own.

The ongoing R&D expenditures by P&G and Kimberly-Clark also serve to deter entry. In addition to brand name recognition, these two firms have accumulated so much technological know-how and manufacturing proficiency that they would have a considerable cost advantage over any firm just entering the market. Besides building new factories, an entrant would have to make a large investment in R&D to capture even a small share of the market. After it began producing, a new firm would have to continue to spend heavily on R&D to reduce its costs over time. Entry would be profitable only if P&G and Kimberly-Clark stop doing R&D, so that the entrant could catch up and eventually gain a cost advantage. But as we have seen, no rational firm would expect this to happen.[18]

*13.8 AUCTIONS

In this section, we examine **auction markets**—markets in which products are bought and sold through formal bidding processes.[19] Auctions come in all sizes and shapes. They are often used for differentiated products, especially unique items such as art, antiques, and the rights to extract oil from a piece of land. In recent years, for example, the U.S. Treasury has relied on auctions to sell Treasury bills, the Federal Communications Commission has used auctions for the sale of portions of the electromagnetic spectrum for cellular telephone services, the International Olympic Committee has auctioned television rights, and the Department of Defense has used auctions to procure military equipment. Auctions like these have important advantages: They are likely to be less time-consuming than one-on-one bargaining, and they encourage competition among buyers in a way that increases the seller's revenue.

• **auction market** Market in which products are bought and sold through formal bidding processes.

Why have auctions become so popular and so successful? The low cost of transacting is only part of the answer. Unlike sales in retail stores, auctions are inherently interactive, with many buyers competing to obtain an item of interest. This interaction can be particularly valuable for the sale of items such as artwork or sports memorabilia that are unique, and therefore do not have established market values. It can also be helpful for the sale of items that are not unique but whose value fluctuates over time.

An example is the daily auctioning of fresh tuna at a Tokyo fish market.[20] Each tuna is unique in size, shape, and quality, and consequently in value. If each transaction were carried out through rounds of bargaining and negotiation with potential buyers, it would be extremely time-consuming. Instead, sales

[18]Example 15.3 in Chapter 15 examines in more detail the profitability of capital investment by a new entrant in the diaper market.

[19]There is a vast literature on auctions; for example, see Paul Milgrom, "Auctions and Bidding: A Primer," *Journal of Economic Perspectives* (Summer 1989): 3–22; Avinash Dixit and Susan Skeath, *Games of Strategy*, 2nd ed. (New York: Norton, 2004); and Preston McAfee, *Competitive Solutions: The Strategist's Toolkit*, Princeton University Press (2002): ch. 12.

[20]John McMillan, *Reinventing the Bazaar: A Natural History of Markets* (New York, Norton, 2002).

occur every morning by means of an auction in which each tuna is sold to the highest bidder. This format creates large savings in transaction costs and thereby increases the efficiency of the market.

The design of an auction, which involves choosing the rules under which it operates, greatly affects its outcome. A seller will usually want an auction format that maximizes the revenue from the sale of the product. On the other hand, a buyer collecting bids from a group of potential sellers will want an auction that minimizes the expected cost of the product.

Auction Formats

We will see that the choice of auction format can affect the seller's auction revenue. Several different kinds of auction formats are widely used:

• **English (**or **oral) auction** Auction in which a seller actively solicits progressively higher bids from a group of potential buyers.

• **Dutch auction** Auction in which a seller begins by offering an item at a relatively high price, then reduces it by fixed amounts until the item is sold.

• **sealed-bid auction** Auction in which all bids are made simultaneously in sealed envelopes, the winning bidder being the individual who has submitted the highest bid.

• **first-price auction** Auction in which the sales price is equal to the highest bid.

• **second-price auction** Auction in which the sales price is equal to the second-highest bid.

1. **English (**or **oral) auction:** The seller actively solicits progressively higher bids from a group of potential buyers. At each point, all participants are aware of the current high bid. The auction stops when no bidder is willing to surpass the current high bid; the item is then sold to the highest bidder at a price equal to the amount of the high bid.
2. **Dutch auction:** The seller begins by offering the item at a relatively high price. If no potential buyer agrees to that price, the seller reduces the price by fixed amounts. The first buyer who accepts an offered price can buy the item at that price.
3. **Sealed-bid auction:** All bids are made simultaneously in sealed envelopes, and the winning bidder is the individual who has submitted the highest bid. The price paid by the winning bidder will vary, however, depending on the rules of the auction. In a **first-price auction**, the sales price is equal to the highest bid. In a **second-price auction**, the sales price is equal to the second-highest bid.

Valuation and Information

Suppose you want to sell a distinctive and valuable product such as a painting or a rare coin. Which type of auction is best for you? The answer depends on the preferences of the bidders and the information available to them. We consider two cases:

• **private-value auction** Auction in which each bidder knows his or her individual valuation of the object up for bid, with valuations differing from bidder to bidder.

Recall from §11.2 that the reservation price is the maximum amount of money that an individual will pay for a product.

• **common-value auction** Auction in which the item has the same value to all bidders, but bidders do not know that value precisely and their estimates of it vary.

1. In **private-value auctions,** each bidder knows his or her individual valuation or *reservation price,* and valuations differ from bidder to bidder. In addition, each bidder is uncertain about the value that other bidders place on the product. For example, I might value a signed Barry Bonds home run baseball very highly but not know that you value it less highly.
2. In **common-value auctions,** the item to be auctioned has approximately the same value to all bidders. Bidders, however, do not know precisely what that value is—they can only estimate it, and bidders' estimates will vary. For example, in an auction of an offshore oil reserve, the value of the reserve is the price of oil minus the extraction cost, times the amount of oil in the reserve. As a result, the value should be about the same for all bidders. However, bidders will not know the amount of oil or the extraction cost—they can only estimate these numbers. Because their estimates will differ, they might bid very different amounts to get the reserve.

In reality, auctions can have both private-value and common-value elements. In the oil reserve auction, for example, there may be some private-value elements because different oil reserves may entail different extraction costs. However, to simplify matters we will separate the two. We begin our discussion with private-value auctions and then move on to common-value auctions.

Private-Value Auctions

In private-value auctions, bidders have different reservation prices for the offered item. We might suppose, for example, that in an auction for a signed Barry Bonds baseball, individuals' reservation prices range from $1 (someone who doesn't like baseball but is bidding just for fun) to $600 (a San Francisco Giants fan). Of course, if you are bidding for the baseball, you don't know how many people will bid against you or what their bids will be.

Whatever the auction format, each bidder must choose his or her bidding strategy. For an open English auction, this strategy is a choice of a price at which to stop bidding. For a Dutch auction, the strategy is the price at which the individual expects to make his or her only bid. For a sealed-bid auction, the strategy is the choice of bid to place in a sealed envelope.

What are the payoffs in this bidding game? The payoff for winning is the difference between the winner's reservation price and the price paid; the payoff for losing is zero. Given these payoffs, let's examine bidding strategies and outcomes for different auction formats.

We will begin by showing that English oral auctions and second-price sealed-bid auctions generate nearly identical outcomes. Let's begin with the second-price sealed-bid auction. In this auction, bidding truthfully is a *dominant strategy*—there is no advantage to bidding below your reservation price. Why? Because the price you pay is based on the valuation of the *second highest bidder*, not on your own valuation. Suppose that your reservation price is $100. If you bid below your reservation price—say, $80—you risk losing to the second-highest bidder, who bids $85, when winning (at, say, $87) would have given you a positive payoff. If you bid above your reservation price—say $105—you risk winning but receiving a negative payoff.

Similarly, in an English auction the dominant strategy is to continue bidding until the second person is unwilling to make a bid. Then the winning bid will be approximately equal to the reservation price of the second person. In any case, you should stop bidding *when the bidding reaches your reservation price*. Why? Because if you stop bidding at a point below your reservation price, you risk losing a positive payoff; if you continue beyond your reservation price, you will be guaranteed a negative payoff. How high will the bidding go? It will continue until the winning bid is approximately equal to the reservation price of the second-highest bidder. Likewise, in the sealed-bid auction the winning bid will equal the reservation price of the second-highest bidder. Thus, both auction formats generate nearly identical outcomes. (The outcomes should differ in theory only by a dollar or two.) To illustrate, suppose that there are three bidders whose valuations are $50, $40, and $30, respectively, and furthermore the auctioneer and the bidders have complete information about these valuations. In an English auction, if your valuation was $50 you would offer a winning bid of $40.01 in order to win the bidding from the individual whose reservation price was $40.00. You would make the identical bid in a sealed-bid auction.

Even in a world of incomplete information, we would expect similar results. Indeed, you know that as a seller, you should be indifferent between an oral English auction and a second-price sealed-bid auction, because bidders in each case have private values. Suppose that you plan to sell an item using a sealed-bid auction. Which should you choose, a first-price or a second-price auction? You might think that the first-price auction is better because the payment is given by the highest rather than the second-highest bid. Bidders, however, are aware of this reasoning and will alter their bidding strategies accordingly: They will bid less in anticipation of paying the winning bid if they are successful.

The second-price sealed-bid auction generates revenue equal to the second-highest reservation price. However, the revenue implications of a first-price sealed-bid auction for the seller are more complicated because the optimal strategy of bidders is more complex. The best strategy is to choose a bid that you believe will be equal to or slightly above the reservation price of the individual with the second-highest reservation price.[21] Why? Because the winner must pay his or her bid, and it is never worth paying more than the second-highest reservation price. Thus, we see that the first-price and second-price sealed-bid auctions generate the same expected revenue.

Common-Value Auctions

Suppose that you and four other people participate in an oral auction to purchase a large jar of pennies, which will go to the winning bidder at a price equal to the highest bid. Each bidder can examine the jar but cannot open it and count the pennies. Once you have estimated the number of pennies in the jar, what is your optimal bidding strategy? This is a classic common-value auction, because the jar of pennies has the same value for all bidders. The problem for you and other bidders is the fact that the value is unknown.

You might be tempted to do what many novices would do in this situation—bid up to your own estimate of the number of pennies in the jar, and no higher. This, however, is not the best way to bid. Remember that neither you nor the other bidders knows the number of pennies for certain. All of you have independently made estimates of the number, and those estimates are subject to error—some will be too high and some too low. Who, then, will be the winning bidder? If each bidder bids up to his or her estimate, *the winning bidder is likely to be the person with the largest positive error*—i.e., the person with the largest overestimate of the number of pennies.

The Winner's Curse To appreciate this possibility, suppose that there are actually 620 pennies in the jar. Let's say the bidders' estimates are 540, 590, 615, 650, and 690. Finally, suppose that you are the bidder whose estimate is 690 and that you win the auction with a bid of \$6.80. Should you be happy about winning? No—you will have paid \$6.80 for \$6.20 worth of pennies. You will have fallen prey to the **winner's curse**: The winner of a common-value auction is often worse off than those who did not win because the winner was overly optimistic and, as a consequence, bid more for the item than it was actually worth.

▪ **winner's curse** Situation in which the winner of a common-value auction is worse off as a consequence of overestimating the value of the item and thereby overbidding.

The winner's curse can arise in any common-value auction, and bidders often fail to take it into account. Suppose, for example, that your house needs to be painted. You ask five companies to give you cost estimates for the job, telling each that you will accept the lowest estimate. Who will win the job? It will probably be the painter who has most seriously underestimated the amount of work

[21]To be more exact, the best strategy is to choose a bid that you believe will be equal to or slightly above the second-highest expected reservation price *conditional on your value being the highest.*

involved. At first, that painter might be happy to have won the job, only later to realize that much more work is required than was anticipated. The same problem can arise for oil companies bidding for offshore oil reserves when the size of the reserve and cost of extraction are uncertain (so that the value of the reserve is uncertain). Unless the companies take the winner's curse into account, the winning bidder is likely to win by overestimating the value of the reserve and will thus pay more than the reserve is worth.

How should you take the winner's curse into account when bidding for an item in a common-value auction? You must not only estimate the value of the item that you are bidding for, but also account for the fact that your estimate—and the estimates of the other bidders—are subject to error. To avoid the winner's curse, you must reduce your maximum bid below your value estimate by an amount equal to the expected error of the winning bidder. The more precise your estimate, the less you need to reduce your bid. If you can't assess the precision of your estimate directly, you can estimate the variation in the estimates of the other bidders. If there is a lot of disagreement among these bidders, it is likely that your estimate will be similarly imprecise. To measure the variation in bids, you can use the standard deviation of the estimates, which can be calculated using statistical methods.

Oil companies have been bidding for oil reserves for years, and thus are able to estimate this standard deviation quite well. They can thereby take the winner's curse into account by reducing their maximum bids below their value estimates by an amount equal to the expected error of the winning bidder. As a result, oil companies rarely feel they have made a mistake after winning an auction. House painters, on the other hand, are often less sophisticated in their bidding decisions and suffer from the winner's curse.

The winner's curse is more likely to be a problem in a sealed-bid auction than in a traditional English auction. In a traditional auction, if you are the only bidder who is overly optimistic, you can still win the bidding by offering only slightly more than the second-highest bidder. Therefore, for the winner's curse to be a problem, at least two bidders must be overly optimistic. By contrast, in a sealed-bid auction, your optimism could encourage you to outbid everyone else by a substantial margin.

Maximizing Auction Revenue

Now let's return to the question of auction design from the viewpoint of the seller. Here are some useful tips for choosing the best auction format.

1. In a private-value auction, you should encourage as many bidders as possible: Additional bidders increase the expected bid of the winner and the expected valuation of the second-highest bidder as well.
2. In a common-value auction, you should (a) use an open rather than a sealed-bid auction because, as a general rule, an English (open) common-value auction will generate greater expected revenue than a sealed-bid auction; and (b) reveal information about the true value of the object being auctioned, thereby reducing concern about the winner's curse and, consequently, encouraging more bidding.
3. In a private-value auction, set a minimum bid equal to or even somewhat higher than the value to you of keeping the good for future sale. This will protect against a loss if there are relatively few bidders who do not value the good very highly. Moreover, it could increase the size of the bids by signaling to buyers that the object is valuable. Having the opportunity to

try again to sell the good if there is no minimum bid is obviously an advantage; however, it can be a disadvantage if failure to sell the good the first time is seen as a signal of low quality to bidders in future auctions.

Why use an open auction? Recall that in order to avoid the winner's curse, each bidder in a common value auction will bid below his individual valuation. The greater the uncertainty about the true value of the object, the greater the likelihood of an overbid, and therefore the greater the incentive for the bidder to reduce his bid. (If the bidder is risk-averse, this effect will be magnified.) However, the bidder faces less uncertainty in an English auction than in a sealed-bid auction because he can observe the prices at which other bidders drop out of the competition—an advantage that provides information about their valuations. In short, when you provide more information to bidders, risk-averse bidders will be encouraged to bid more because they will be more confident that they can account for the possibility of a winner's curse.

Bidding and Collusion

We have seen that sellers at auctions can obtain a significant share of the gains from trade by encouraging competition among buyers. It follows, therefore, that buyers can increase their bargaining power by reducing the number of bidders or the frequency of bidding. In some cases this can be accomplished legally through the formation of buying groups, but it may also be accomplished illegally through collusive agreements that violate the antitrust laws. Collusion among buyers is not easy, because even if an "agreement" is reached, individual buyers will have an incentive to cheat by increasing their bids at the last minute in order to obtain the desired item. However, repeated auctions allow for participants to penalize those that break from the agreement by outbidding the "cheater" again and again. Buyer collusion is more of a problem in open-bid auctions than in the case of sealed bids because open auctions offer the best opportunity for colluding bidders to detect and punish cheating.

A well-known case of buyer collusion was the agreement in the mid-1980s among baseball owners to limit their bidding for free-agent players. The fact that such bidding was repeated and open made it possible for owners to retaliate against those that bid too often and too aggressively. Collusion, however, is not limited to buyers. In 2001, two of the world's most successful auction houses, Sotheby's and Christie's, were found guilty of agreeing to fix the price of commissions offered to sellers of auctioned items. Former Sotheby's chairman Alfred Taubman was sentenced to a year in jail for his involvement in the scheme.

EXAMPLE 13.7 Auctioning Legal Services

After Sotheby's and Christie's auction houses were found guilty in 2001 of fixing commission prices, a federal class-action lawsuit followed, on behalf of those who paid too much in commissions. The lawsuit was administrated by Judge Kaplan of the Southern District of New York. When federal courts manage class-action suits, they are responsible for awarding attorney's fees. In this case, Judge Kaplan decided to hold an auction to select the law firm that would represent the plaintiff class.

Judge Kaplan entertained secret sealed bids from 20 law firms. Each firm was told to offer a fee arrangement consisting of a base and a percentage. A settlement or trial award at or below the base would be given entirely to the plaintiffs, with the law firm receiving nothing. If the settlement or award was higher than the base, the law firm would receive the stated percentage of the amount over the base. Many attorneys operate under a contingent fee system in which they offer a base of zero and expect to receive one-third of the award.

The winning bidder was the law firm of Boies, Schiller, & Flexner, which bid a base of \$405 million and a percentage of 25 percent to be earned on any award above \$405 million. Some of the losing bidders were outraged that Boies had bid so high to get the business. Indeed, some suggested that the firm might not work hard in the plaintiffs' interest because the minimum might be unachievable. Prior to the bidding, observers expected the case to generate a settlement of \$130 million. In the end, it appears that Judge Kaplan, the plaintiffs' class, and the Boies law firm were all winners. Months after taking on the case, Boies settled with defendants for \$512 million, earning the attorneys a \$26.75 million fee (25 percent of the \$107 excess over the base guarantee of \$405 million) and generating just over \$475 million for the class members.[22]

EXAMPLE 13.8 Internet Auctions

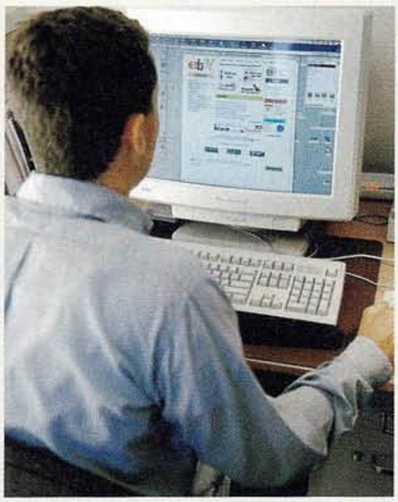

The popularity of auctions has skyrocketed in recent years with the growth of the Internet. Indeed, the Internet has lowered transaction costs by so much that individuals anywhere in the world can now trade relatively low-value items without leaving the comfort of home. Many Internet sites are now devoted to auctions at which participants can buy and sell a wide variety of items. Let's see how these Internet auctions work.[23]

The most popular Internet auction site is **www.ebay.com**. It conducts auctions each day for items ranging from antiques and automobiles to Beanie Babies and Pokémon cards. Founded in 1995 by Pierre Omidyar in an effort to sell a broken laser pointer, eBay dominates the online person-to-person auction industry. It recently listed millions of products for sale, including such unusual items as a Caribbean island, 154 acres in the Catskills, and a ghost town in Nevada. In 2005, eBay accounted for about 85 percent of all online auction sales, totalling over \$50 billion. On average, over 14 million items are listed for sale at any given time.

[22]Some experts have speculated that Boies was successful in obtaining a higher settlement because the defendant was aware that any settlement less than \$405 million would generate no legal fees.

[23]For more information on Internet auctions, see Patrick Bajari and Ali Hortaçsu, "Economic Insights from Internet Auctions," *Journal of Economic Literature* 42 (June 2004): 457–86.

In §4.5 we explain how network externalities affect sales of a product.

How has eBay come to dominate the Internet auction market? Why haven't other Internet auction sites (such as Yahoo and Amazon) succeeded in taking market share from eBay? The answer is that Internet auctions are subject to very strong *network externalities*. If you wanted to auction off some rare coins or Pokémon cards, which auction site would you choose? The one that had the largest number of potential bidders. Likewise, if you wanted to bid for rare coins or Pokémon cards, you would choose the auction site with the largest number of sellers. Thus, both sellers and buyers gravitate to the auction site with the largest market share. Because eBay was the first major Internet auction site, it began with a large market share, and its share grew thanks to the network externality.

Two auction formats are used on eBay: (1) an increasing-bid auction for a single item, in which the highest bidder at the close of the auction wins and pays the seller a price equal to the second-highest bid; and (2) an increasing-bid auction for several identical items, in which the highest n bidders win the n items sold. In both auctions, ties are broken by awarding the item to the buyer who bid first. Notice that neither of these auctions corresponds precisely to any of the four auction formats discussed previously. The first approximates the standard English auction, but the existence of a fixed and known stopping time can cause bidders to place bids strategically at the end of the auction. The second format differs from a conventional Dutch auction in two respects: Bids are increasing rather than decreasing and the auction has a fixed and known stopping time. In both auction formats, sellers can impose a minimum acceptable bid—called a *reserve price*—and although buyers know that a reserve price exists, they are generally not told what it is.

Many Internet auctions are dominated by private-value items. (However, because anyone can put an item up for sale, there are common-value issues—how reliable is the seller, and are there possibilities for resale?) The private-value emphasis of these auctions is especially true of unique antiques that may have considerable value to particular bidders. With private-value auctions, you needn't worry so much about the prior history of bidding: The bids of others tell you about their preferences, but the value that you place on the object is personal to you. Although you want to win the bidding at a price as far below your valuation as possible, the winner's curse needn't be a concern: You can't be disappointed if your value for the object is more than what you paid for it.

Finally, a few caveats are in order when buying items via Internet auctions. Unlike traditional auction houses, low-end auction sites like eBay provide only a forum for buyers and sellers to interact; they provide no quality-control functions. Although many sites, including eBay, make available feedback from buyers for each seller, this is usually the only evidence of a seller's reliability that buyers receive. Furthermore, there is obviously no feedback available for first-time sellers (or for sellers who have recently changed their eBay user names). In addition, the possibility of bid manipulation looms large in Internet auctions. At eBay, for example, a valid e-mail address is the only thing you need to bid on an item. Given the relative ease of obtaining e-mail addresses, sellers may file spurious bids in order to manipulate the bidding process. Thus, caveat emptor is a sound philosophy when buying items on the Internet.

SUMMARY

1. A game is cooperative if the players can communicate and arrange binding contracts; otherwise, it is noncooperative. In either kind of game, the most important aspect of strategy design is understanding your opponent's position, and (if your opponent is rational) correctly deducing the likely response to your actions. Misjudging an opponent's position is a common mistake, as Example 13.1 "Acquiring a Company" (page 481) illustrates.[24]
2. A Nash equilibrium is a set of strategies such that all players are doing their best given the strategies of the other players. An equilibrium in dominant strategies is a special case of a Nash equilibrium; a dominant strategy is optimal no matter what the other players do. A Nash equilibrium relies on the rationality of each player. A maximin strategy is more conservative because it maximizes the minimum possible outcome.
3. Some games have no Nash equilibria in pure strategies but have one or more equilibria in mixed strategies. A mixed strategy is one in which the player makes a random choice among two or more possible actions, based on a set of chosen probabilities.
4. Strategies that are not optimal for a one-shot game may be optimal for a repeated game. Depending on the number of repetitions, a "tit-for-tat" strategy, in which you play cooperatively as long as your competitor does the same, may be optimal for the repeated prisoners' dilemma.
5. In a sequential game, the players move in turn. In some cases, the player who moves first has an advantage. Players may then have an incentive to try to precommit themselves to particular actions before their competitors can do the same.
6. An empty threat is a threat that one has no incentive to carry out. If one's competitors are rational, empty threats are of no value. To make a threat credible, it is sometimes necessary to make a strategic move to constrain one's later behavior, thereby creating an incentive to carry out the threat.
7. Bargaining situations are examples of cooperative games. As in noncooperative games, in bargaining, players can sometimes gain a strategic advantage by limiting their own flexibility.
8. To deter entry, an incumbent firm must convince any potential competitor that entry will be unprofitable. This may be done by investing, and thereby giving credibility to the threat that entry will be met by price warfare. Strategic trade policies by governments sometimes have this objective.
9. Auctions can be conducted in a number of formats, including English (oral with increasing bids), Dutch (oral with decreasing bids), and sealed bid. The opportunity for a seller to raise revenue and for a buyer to obtain an object at a reasonable price depends on the auction format, and on whether the items being auctioned have the same value to all bidders (as in a common-value auction) or different values to different bidders (as in a private-value auction).

QUESTIONS FOR REVIEW

1. What is the difference between a cooperative and a noncooperative game? Give an example of each.
2. What is a dominant strategy? Why is an equilibrium stable in dominant strategies?
3. Explain the meaning of a Nash equilibrium. How does it differ from an equilibrium in dominant strategies?
4. How does a Nash equilibrium differ from a game's maximin solution? When is a maximin solution a more likely outcome than a Nash equilibrium?
5. What is a "tit-for-tat" strategy? Why is it a rational strategy for the infinitely repeated prisoners' dilemma?
6. Consider a game in which the prisoners' dilemma is repeated 10 times and both players are rational and fully informed. Is a tit-for-tat strategy optimal in this case? Under what conditions would such a strategy be optimal?
7. Suppose you and your competitor are playing the pricing game shown in Table 13.8 (page 490). Both of you must announce your prices at the same time. Can you improve your outcome by promising your competitor that you will announce a high price?
8. What is meant by "first-mover advantage"? Give an example of a gaming situation with a first-mover advantage.

[24] Here is the solution to Company A's problem: *It should offer nothing for Company T's stock.* Remember that Company T will accept an offer only if it is greater than the per-share value under current management. Suppose you offer \$50. Thus Company T will accept this offer only if the outcome of the exploration project results in a per-share value under current management of \$50 or less. Any values between \$0 and \$100 are equally likely. Therefore, the *expected value* of Company T's stock, *given that it accepts the offer*—i.e., given that the outcome of the exploration project leads to a value less than \$50—is \$25. Under the management of Company A, therefore, the value would be (1.5)(\$25) = \$37.5, which is less than \$50. In fact, for any price P, if the offer is accepted, Company A can expect a value of only $(3/4)P$.

9. What is a "strategic move"? How can the development of a certain kind of reputation be a strategic move?
10. Can the threat of a price war deter entry by potential competitors? What actions might a firm take to make this threat credible?
11. A strategic move limits one's flexibility and yet gives one an advantage. Why? How might a strategic move give one an advantage in bargaining?
12. Why is the winner's curse potentially a problem for a bidder in a common-value auction but not in a private-value auction?

EXERCISES

1. In many oligopolistic industries, the same firms compete over a long period of time, setting prices and observing each other's behavior repeatedly. Given the large number of repetitions, why don't collusive outcomes typically result?
2. Many industries are often plagued by overcapacity: Firms simultaneously invest in capacity expansion, so that total capacity far exceeds demand. This happens not only in industries in which demand is highly volatile and unpredictable, but also in industries in which demand is fairly stable. What factors lead to overcapacity? Explain each briefly.
3. Two computer firms, *A* and *B*, are planning to market network systems for office information management. Each firm can develop either a fast, high-quality system (High), or a slower, low-quality system (Low). Market research indicates that the resulting profits to each firm for the alternative strategies are given by the following payoff matrix:

		Firm B	
		High	Low
Firm A	High	50, 40	60, 45
	Low	55, 55	15, 20

   a. If both firms make their decisions at the same time and follow *maximin* (low-risk) strategies, what will the outcome be?
   b. Suppose that both firms try to maximize profits, but that Firm *A* has a head start in planning and can commit first. Now what will be the outcome? What will be the outcome if Firm *B* has the head start in planning and can commit first?
   c. Getting a head start costs money. (You have to gear up a large engineering team.) Now consider the *two-stage* game in which, *first*, each firm decides how much money to spend to speed up its planning, and, *second*, it announces which product (*H* or *L*) it will produce. Which firm will spend more to speed up its planning? How much will it spend? Should the other firm spend *anything* to speed up its planning? Explain.
4. Two firms are in the chocolate market. Each can choose to go for the high end of the market (high quality) or the low end (low quality). Resulting profits are given by the following payoff matrix:

		Firm 2	
		Low	High
Firm 1	Low	−20, −30	900, 600
	High	100, 800	50, 50

   a. What outcomes, if any, are Nash equilibria?
   b. If the managers of both firms are conservative and each follows a maximin strategy, what will be the outcome?
   c. What is the cooperative outcome?
   d. Which firm benefits most from the cooperative outcome? How much would that firm need to offer the other to persuade it to collude?
5. Two major networks are competing for viewer ratings in the 8:00–9:00 P.M. and 9:00–10:00 P.M. slots on a given weeknight. Each has two shows to fill these time periods and is juggling its lineup. Each can choose to put its "bigger" show first or to place it second in the 9:00–10:00 P.M. slot. The combination of decisions leads to the following "ratings points" results:

		Network 2	
		First	Second
Network 1	First	20, 30	18, 18
	Second	15, 15	30, 10

   a. Find the Nash equilibria for this game, assuming that both networks make their decisions at the same time.
   b. If each network is risk-averse and uses a maximin strategy, what will be the resulting equilibrium?
   c. What will be the equilibrium if Network 1 makes its selection first? If Network 2 goes first?

d. Suppose the network managers meet to coordinate schedules and Network 1 promises to schedule its big show first. Is this promise credible? What would be the likely outcome?

6. Two competing firms are each planning to introduce a new product. Each will decide whether to produce Product *A*, Product *B*, or Product *C*. They will make their choices at the same time. The resulting payoffs are shown below.

		Firm 2		
		A	*B*	*C*
Firm 1	*A*	−10, −10	0, 10	10, 20
	B	10, 0	−20, −20	−5, 15
	C	20, 10	15, −5	−30, −30

a. Are there any Nash equilibria in pure strategies? If so, what are they?

b. If both firms use maximin strategies, what outcome will result?

c. If Firm 1 uses a maximin strategy and Firm 2 knows this, what will Firm 2 do?

7. We can think of U.S. and Japanese trade policies as a prisoners' dilemma. The two countries are considering policies to open or close their import markets. The payoff matrix is shown below.

		Japan	
		Open	Close
U.S.	Open	10, 10	5, 5
	Close	−100, 5	1, 1

a. Assume that each country knows the payoff matrix and believes that the other country will act in its own interest. Does either country have a dominant strategy? What will be the equilibrium policies if each country acts rationally to maximize its welfare?

b. Now assume that Japan is not certain that the United States will behave rationally. In particular, Japan is concerned that U.S. politicians may want to penalize Japan even if that does not maximize U.S. welfare. How might this concern affect Japan's choice of strategy? How might this change the equilibrium?

8. You are a duopolist producer of a homogeneous good. Both you and your competitor have zero marginal costs. The market demand curve is

$$P = 30 - Q$$

where $Q = Q_1 + Q_2$. Q_1 is your output and Q_2 your competitor's output. Your competitor has also read this book.

a. Suppose you will play this game only once. If you and your competitor must announce your outputs at the same time, how much will you choose to produce? What do you expect your profit to be? Explain.

b. Suppose you are told that you must announce your output *before* your competitor does. How much will you produce in this case, and how much do you think your competitor will produce? What do you expect your profit to be? Is announcing first an advantage or a disadvantage? Explain briefly. How much would you pay for the option of announcing either first or second?

c. Suppose instead that you are to play the first round of a series of 10 rounds (with the same competitor). In each round, you and your competitor announce your outputs at the same time. You want to maximize the sum of your profits over the 10 rounds. How much will you produce in the first round? How much do you expect to produce in the tenth round? In the ninth round? Explain briefly.

d. Once again you will play a series of 10 rounds. This time, however, in each round your competitor will announce its output before you announce yours. How will your answers to (c) change in this case?

9. You play the following bargaining game. Player *A* moves first and makes Player *B* an offer for the division of \$100. (For example, Player *A* could suggest that she take \$60 and Player *B* take \$40.) Player *B* can accept or reject the offer. If he rejects it, the amount of money available drops to \$90, and he then makes an offer for the division of this amount. If Player *A* rejects this offer, the amount of money drops to \$80 and Player *A* makes an offer for its division. If Player *B* rejects this offer, the amount of money drops to 0. Both players are rational, fully informed, and want to maximize their payoffs. Which player will do best in this game?

***10.** Defendo has decided to introduce a revolutionary video game. As the first firm in the market, it will have a monopoly position for at least some time. In deciding what type of manufacturing plant to build, it has the choice of two technologies. Technology *A* is publicly available and will result in annual costs of

$$C^A(q) = 10 + 8q$$

Technology *B* is a proprietary technology developed in Defendo's research labs. It involves a higher fixed cost of production but lower marginal costs:

$$C^B(q) = 60 + 2q$$

Defendo must decide which technology to adopt. Market demand for the new product is $P = 20 - Q$, where Q is total industry output.

a. Suppose Defendo were certain that it would maintain its monopoly position in the market for the entire product lifespan (about five years) without threat of entry. Which technology would you advise Defendo to adopt? What would be Defendo's profit given this choice?

b. Suppose Defendo expects its archrival, Offendo, to consider entering the market shortly after Defendo introduces its new product. Offendo will have access only to Technology *A*. If Offendo does enter the market, the two firms will play a Cournot game (in quantities) and arrive at the Cournot-Nash equilibrium.
   i. If Defendo adopts Technology *A* and Offendo enters the market, what will be the profit of each firm? Would Offendo choose to enter the market given these profits?
   ii. If Defendo adopts Technology *B* and Offendo enters the market, what will be the profit of each firm? Would Offendo choose to enter the market given these profits?
   iii. Which technology would you advise Defendo to adopt given the threat of possible entry? What will be Defendo's profit given this choice? What will be consumer surplus given this choice?

c. What happens to social welfare (the sum of consumer surplus and producer profit) as a result of the threat of entry in this market? What happens to equilibrium price? What might this imply about the role of *potential* competition in limiting market power?

11. Three contestants, *A*, *B*, and *C*, each have a balloon and a pistol. From fixed positions, they fire at each other's balloons. When a balloon is hit, its owner is out. When only one balloon remains, its owner gets a $1000 prize. At the outset, the players decide by lot the order in which they will fire, and each player can choose any remaining balloon as his target. Everyone knows that *A* is the best shot and always hits the target, that *B* hits the target with probability .9, and that C hits the target with probability .8. Which contestant has the highest probability of winning the $1000? Explain why.

12. An antique dealer regularly buys objects at hometown auctions whose bidders are limited to other dealers. Most of her successful bids turn out to be financially worthwhile because she is able to resell the antiques for a profit. On occasion, however, she travels to a nearby town to bid in an auction that is open to the public. She often finds that on the rare occasions in which she does bid successfully, she is disappointed—the antique cannot be sold at a profit. Can you explain the difference in her success between the two sets of circumstances?

13. You are in the market for a new house and have decided to bid for a house at auction. You believe that the value of the house is between $125,000 and $150,000, but you are uncertain as to where in the range it might be. You do know, however, that the seller has reserved the right to withdraw the house from the market if the winning bid is not satisfactory.

a. Should you bid in this auction? Why or why not?

b. Suppose you are a building contractor. You plan to improve the house and then to resell it at a profit. How does this situation affect your answer to (a)? Does it depend on the extent to which your skills are uniquely suitable to improving this particular house?