Learning Goals

In this chapter, you will learn

1. How game theory can help us understand strategic behavior
2. What the Nash equilibrium is in the prisoner's dilemma game
3. Why backward induction is critical for thinking strategically
4. How reputation can play a role in repeated games
5. How game theory can help us understand why threats and promises may not always be credible

Chapter 14

STRATEGIC BEHAVIOR

During the 2000 television season, viewers were enthralled by the first season of the show *Survivor*. Sixteen contestants were marooned on Pulau Tiga, an uninhabited island off the coast of Borneo, and at the end of each episode, the survivors voted to kick out one of their members. The last surviving contestant collected $1 million. By episode 13, the original group had been reduced to just three final contestants: Kelly, Rudy, and Rich. In the course of a trial that involved answering questions about the thirteen previously evicted castaways, Kelly won immunity in the first round of voting that would reduce the group to the final pair. She still had to decide, though, how to cast her own vote. If, as seemed certain, Rudy votes against Rich and Rich votes against Rudy, Kelly's vote would decide who stays and who leaves. How should Kelly vote?

To know whom she should vote out in the first round, Kelly must think about how she was likely to fare in the final round. Kelly must think backward from the end of the contest. Rudy seemed popular with the audience and with the other thirteen contestants who would get to vote for the final winner. So Kelly would reason that if the final contest came down to her and Rudy, Rudy would probably win. On the other hand, Rich seemed to be very unpopular. So Kelly would reason that if it came down to a choice between her and Rich, she would probably win. Her best chance would be to face off against Rich, not Rudy, in the final voting. Even if Kelly really dislikes Rich, her best strategy is to oust Rudy. As it turned out, that is exactly what happened. Rich and Kelly voted against Rudy. Unfortunately for Kelly, the voters ended up picking Rich over her, but her strategy was still the right one.

The participants in *Survivor* had to think strategically. They needed to anticipate how their rivals would respond to the decisions they themselves made. They had to consider how their own position would depend on who else survived and who didn't, and they had to use that information in deciding how to vote.

Thinking strategically doesn't help just in a made-up environment such as a television show. We all face situations that call for strategic thinking. Economists try to understand the choices individuals and firms make, and researchers studying strategic behavior have extended the reach of economics into many new areas.

Economists examine the choices made by rational individuals and profit-maximizing firms. In the basic competitive model presented in the first two parts of this book, individuals and firms do not need to behave strategically. Consumers and firms can buy or sell as much as they want at the market price. A firm does not need to worry about how its rivals will react if it decides to produce more. Nor does a monopoly, but for a different reason: a monopoly has no rivals. In the basic competitive model and in a monopoly model, **strategic behavior**—decisions that take into account the possible reactions of others—plays no role.

Things were different in Chapter 12 when we studied oligopolies. With only a few firms in the industry, each firm needs to be concerned about how its rivals might react whenever it contemplates expanding production or cutting its price. Strategic behavior becomes important. When AMD considers cutting the prices of its various computer processor chips, it must try to assess how Intel will respond. If Intel reacts by also cutting prices, then AMD may not gain much market share and its revenues will decline as a result of the lower prices. But if Intel keeps its prices unchanged, AMD may gain market share and its revenues might rise as it sells more chips.

Because oligopolies engage in strategic behavior, Chapter 12 used game theory—and the simple prisoner's dilemma game—to demonstrate why it may be difficult for firms to collude. In this chapter, we return to the prisoner's dilemma and see how its basic insights can be applied to other areas of economics. The usefulness of game theory goes well beyond this simple model, however, for decisions and choices often must incorporate the potential reactions of others. You will learn more about game theory and how it helps us understand the choices made by individuals, firms, unions, and governments. Game theory provides a framework for studying strategic behavior. Using this framework, economists have found that many instances of strategic behavior can be understood by relying on the core concepts of incentives and information.[1]

Review of the Prisoner's Dilemma

Let's very briefly recall the prisoner's dilemma game introduced in Chapter 12. Two prisoners, A and B, are alleged by the police to be conspirators in a crime. After being taken into custody, the two are separated. A police officer tells each, "Now here's the situation. If your partner confesses and you remain silent, you'll get five years in prison. But if your partner confesses and you confess also, you'll both get three years. On the other hand, if both you and your partner remain silent, we'll be able to convict you only of a lesser charge and you'll get one year in prison. But if your partner remains silent and you confess, we'll let you out in three months." The same deal is offered to both prisoners.

[1]If you would like to learn even more about game theory, an accessible textbook is available: Avinash Dixit and Susan Skeath, *Games of Strategy*, Second Edition (New York: W. W. Norton, 2004).

Figure 14.1, which reproduces a diagram from Chapter 12, shows the results of the deal the police have offered the prisoners. This type of grid, showing the payoffs to each player, is called a **game table.** In Chapter 12, we saw that based on self-interest, each individual prisoner believes that confession is best, whether his partner confesses or not. But by following self-interest and confessing, each ends up worse off than if neither had confessed. The prisoner's dilemma is a simple game in which both parties suffer because they independently act in self-interest. Both would be better off if they could get together to agree on a story and if each could threaten to punish the other if he deviates from the story.

This simple game has been widely applied in economics and in other fields such as international relations and political science. In Chapter 12, we used it to explain why two oligopolists would find it difficult to sustain an agreement to collude. We will discuss some further examples of the prisoner's dilemma and then consider other types of game situations. First, though, we will need to clarify our method of analyzing strategic situations to make predictions about how individuals and firms will behave.

		Prisoner B	
		Confesses	Remains silent
Prisoner A	Confesses	A gets 3 years / B gets 3 years	A gets 3 months / B gets 5 years
	Remains silent	A gets 5 years / B gets 3 months	A gets 1 year / B gets 1 year

FIGURE 14.1

PRISONER'S DILEMMA GAME

Each prisoner's dominant strategy is to confess.

DOMINANT STRATEGIES

Behaving strategically means that each player must try to determine what the other player is likely to do. Will your accomplice confess or keep quiet? Will your rival match price reductions if you cut your prices? The decision of one player depends on how she thinks the other player will respond.

In the basic prisoner's dilemma game, we assume that players reason along the following lines: "For each choice that I might make, what is the best choice for the other player to make?" In analyzing the prisoner's dilemma, we ask, "If prisoner A doesn't confess, what is the best strategy for prisoner B? If prisoner A confesses, what is the best strategy for prisoner B?" In both cases, we conclude that confessing is prisoner B's *best response.* If B's best response is to confess, no matter what prisoner A does, then A will conclude that B will confess, so A needs to decide what his best response is to prisoner B confessing. As we saw, prisoner A's best option is also to confess.

Such a strategy—one that works best no matter what the other player does—is called a **dominant strategy.** Recall that an objective of game theory is to predict what strategy each player will choose. When a player has a dominant strategy, that is the strategy we should predict a rational decision maker will choose.

NASH EQUILIBRIUM

It is easy to predict the outcome of a game—its equilibrium—if each player has a dominant strategy. Each will play his dominant strategy. Thus, in the prisoner's dilemma, the equilibrium is both players confessing. The situation is not quite so simple when only one player has a dominant strategy, or when neither player has one; we will learn about such games later. To predict the equilibrium outcomes in these more complex games, we need to look more closely at why confessing is the equilibrium in the prisoner's dilemma.

In the prisoner's dilemma, each prisoner confesses because doing so leads to the best, or *optimal* payoff—the least amount of time in prison—given what the other prisoner can be expected to do. The outcome is an equilibrium in the sense that neither prisoner would change her chosen strategy if offered the chance to do so at the end of the game. By confessing, both have played their best response. Such an equilibrium is called a **Nash equilibrium,** and it is the most fundamental idea for predicting the actions of players in a strategic game.

John Nash developed the notion that bears his name when he was just twenty-one years old and a graduate student in mathematics at Princeton University. Economists have found the concept of Nash equilibrium extremely useful for predicting the outcomes of games and understanding economic problems. In recognition of its importance, John Nash shared the 1994 Nobel Prize in Economics. Nash is not an economist; he is a mathematician and the only winner of the Nobel Prize in Economics whose life has been the subject of a best-selling biography and a popular movie. Published in 1998, *A Beautiful Mind* by Sylvia Nasar chronicles Nash's early mathematical brilliance, his struggle with mental illness, and his eventual recovery. On being informed he had won the Nobel Prize, Nash commented that he hoped it would improve his credit rating.

The prisoner's dilemma arises in many contexts, both in economics and in other social sciences. The following examples, briefly sketched, provide some illustrations.[2]

Example: Collusion In Chapter 12, we studied an application of the prisoner's dilemma to the problem faced by two rivals who can benefit by colluding to restrict output. Colluding results in a higher price and therefore greater profits for each. The higher price can be maintained only if both firms continue to restrict output. But at the higher price, each firm perceives that it would be even better off if it could expand its output a little and sell more at the high price. Of course, this means that each firm has an incentive to cheat on their agreement to restrict output, just as each prisoner had an incentive to confess. And if both firms fail to stick to their bargain, their expansion causes the price, and their profits, to fall, thereby making each worst off than if they had continued to collude. In this game, each firm has a dominant strategy—cheat on the agreement by expanding output. Here, the outcome in which each firm expands is the Nash equilibrium; each firm is choosing its best response, given the actions of its rival. When both players have a dominant strategy, as in the prisoner's dilemma, both will choose to play it; no other outcome would satisfy the definition of a Nash equilibrium. In this case, there is only one Nash equilibrium, a situation economists describe by saying there is a *unique Nash equilibrium.*

Example: Politicians and Negative Ads Why do politicians engage in negative advertising even though they all promise not to?

[2]You should test your understanding of each example by filling out a game table, identifying the dominant strategy for each player, and finding the Nash equilibrium.

Consider the case of politicians A and B. If neither runs a negative campaign, the public thinks highly of both of them, but neither gains any advantage over the other. If both run negative campaign ads, the public thinks poorly of both, but again neither gains an advantage. Each is tarred by the other's ads. If politician A runs a clean campaign, politician B can gain an advantage by running a negative ad that tarnishes A's reputation. Conversely, A gains by running a negative ad if politician B runs a clean campaign.

Each politician will reason as follows: "If my opponent runs a negative ad, I'm better off if I also run negative ads. And if my opponent doesn't run a negative ad, then I can gain an advantage if I run negative ads. Either way, I'm better off if I run negative ads." Each politician has a dominant strategy, and in this unique Nash equilibrium both run negative campaign ads, despite their promises not to.

Example: Military Spending Two countries, A and B, are locked in a military balance. Each must decide whether to build a new generation of missiles. If neither builds the missile system, the military balance is preserved and each country will remain secure. If one builds the system while the other does not, one will gain a military advantage. If they both build the system, they each spend billions of dollars, but neither gains an advantage: each country now has the new missile system, preserving the military balance.

Each country reasons that if the other country fails to build the system, it can gain an advantage by going ahead and building the system. Both also know that if one country builds the new missiles, the other will be worse off if it fails to build the missile system as well. Each country has a dominant strategy—build the missile system. Both countries spend billions, only to find themselves left in the same military balance as before.

Example: Sports Owners and Player Salaries Sports teams compete to hire the best players. Suppose there are just two teams, the Yankees and the Athletics. If both teams collude and keep salaries low, the owners' profits go up. If the Yankees' owner instead offers high salaries while the Athletics owner does not, the Yankees will attract all the good players and will generate higher profits for the owner. Meanwhile, the Athletics end up with weaker players and have a poor season. Low attendance causes the owner to lose money. Conversely, if the Athletics offer high salaries and the Yankees do not, the Athletics get all the good players and earn the higher profits, while the Yankees lose money. If both offer high salaries, neither team gets all the good players, but increased salary costs cut into the owners' profits.

In a Nash equilibrium, both team owners offer players high salaries, and the owners are worse off than if they been able to collude to keep salaries down. Of course, our simple example ignores all the other factors—TV revenues, cable deals, sales of team paraphernalia, and so on—that support total payrolls and thus affect salaries and profits. But the example does suggest one of the factors behind the desire of owners in most professional sports to institute caps on player salaries. Without such caps, each owner has an incentive to try to outbid the others for the best players. A salary cap enables the owners to collude more effectively.

Wrap-Up

THE PRISONER'S DILEMMA

In the prisoner's dilemma, each player has a dominant strategy. In a unique Nash equilibrium, the players choose their dominant strategy. Yet the players are worse off in this Nash equilibrium than if they had chosen their alternative strategy. Applications of the prisoner's dilemma arise in many branches of economics, as well as in other social sciences and in everyday life.

Strategic Behavior in More General Games

In the prisoner's dilemma, both players have a dominant strategy. In most games, however, this is not the case. What each player finds it best to do depends on what the other player does, and determining the outcome of a game is therefore harder. But we can often make an accurate prediction by thinking through the consequences from each player's perspective, just as we did in the prisoner's dilemma.

GAMES WITH ONLY ONE DOMINANT STRATEGY

Consider the positions of two firms deciding on whether to cut prices. Discounters Delux and Quality Brands compete with each other. Discounters Delux promises its customers the lowest prices; it would suffer a large loss of customers if it failed to deliver on that promise. Quality Prices has higher costs; perhaps it provides better health care benefits to its workers or spends more on the displays in its stores. Because of these higher costs, it would prefer not to cut prices. However, it risks losing some of its business if it does not match Discounters Delux's price reductions. The profits each expects to earn are shown in Figure 14.2. The payoffs to Discounters Delux are below the diagonal lines, while those of Quality Brands are above the diagonals.

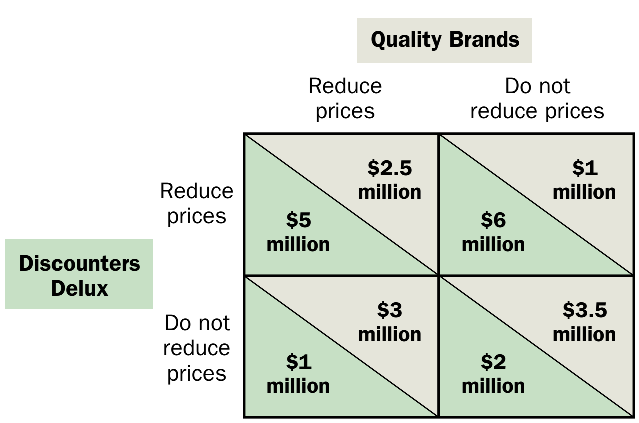

FIGURE 14.2

PRICE-CUTTING GAME

Discounters Delux has a dominant strategy, which is to reduce prices. Thus, the best response for Quality Brands is also to cut prices.

Discounters Delux has a dominant strategy—reduce prices. It makes more in profits under this strategy, regardless of what Quality Brands does. Quality Brands, in contrast, does not have a dominant strategy. If Discounters Delux reduces prices, Quality Brand's best response is to cut prices as well, since otherwise it would lose too many sales. However, if Discounters Delux does not reduce prices, then Quality Brands is better off not reducing its prices.

Even though Quality Brands does not have a dominant strategy, we can predict its move if we reason as follows. Quality Brands

International Perspective

BEGGAR-THY-NEIGHBOR TARIFF POLICIES

During the international depression of the 1930s, many countries debated whether to impose restrictions on international trade. Some argued that restricting imports from other countries would boost the demand for goods produced domestically. Increased demand, in turn, would enable domestic firms to expand employment and production. By reducing the demand for goods produced in other countries, such policies would lead to fewer sales by foreign firms, who would then have to cut production and employment. The gains at home would come at the expense of foreign producers and workers, so these were often called "beggar-thy-neighbor" policies. The trade wars of the 1930s can be understood as another example of the prisoner's dilemma.

In 1930, the United States passed the Smoot-Hawley Tariff Act. This act raised tariffs on imported goods, making them more expensive for American consumers and thereby providing an incentive for Americans to shift their demand to domestically produced goods. Other countries did not stand by idly while American tariffs reduced their market in the United States. They retaliated by raising their tariffs on goods produced in the United States. The result was that all countries suffered from the decline in world trade. By imposing tariffs that reduced trade, all countries lost the benefits of trade.

The outcome in the sort of trade war set off by the Smoot-Hawley Act can be illustrated in terms of a simple game. The diagram shows the payoffs to each of two countries that are deciding whether to impose trade restrictions. Payoffs to country A are shown below the diagonal line in each box; payoffs to country B, above the diagonal. The payoffs are defined as the gains (or losses) to a country's income relative to what that income would be if neither imposed trade restrictions. The numbers are hypothetical and can be thought of as, say, tens of billions of dollars. Each country gains the most if it is the only one to impose restrictions. Both are worse off if both impose restrictions and better off if neither imposes them.

Each country has a dominant strategy—impose trade sanctions. Country A, for example, would reason as follows: "If country B imposes trade restrictions, our country will be better off if we also impose restrictions. If country B does not impose restrictions, we are also better off if we impose restrictions. Therefore, we should impose trade restrictions regardless of what country B does." Country B would reason in exactly the same manner. Unfortunately, when they both impose restrictions, both are left worse off than if no restrictions had been imposed.

Another way to look at this situation is to recognize that both countries would be better off if they could cooperate and mutually agree not to impose trade restrictions. The problem is, our simple example provides no mechanism to ensure that they cooperate. Just as in the collusion example discussed in Chapter 12, each player has an incentive to violate any voluntary agreement not to impose trade restrictions. One of the roles of international organizations such as the World Trade Organization (WTO) is to lay down rules designed to promote international trade and cooperation, enabling nations to make credible commitments not to raise tariffs or other barriers to trade. The WTO can impose sanctions on the countries that violate these rules.

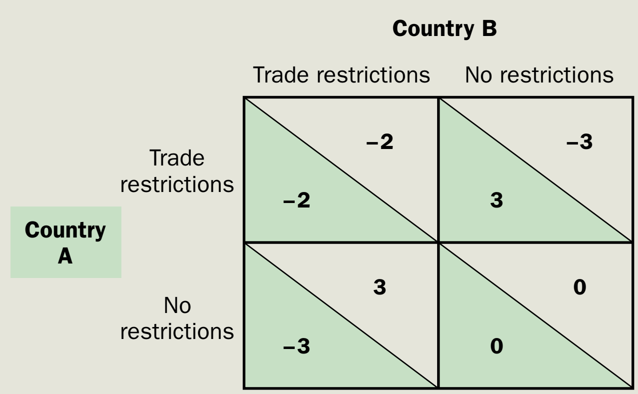

knows that Discounters Delux will reduce prices, since that is Discounters Delux's dominant strategy. So the fact that Quality Brands would find it best to keep prices high if Discounters Delux keeps its prices high is irrelevant. Quality Brands knows that Discounters Delux will cut prices, and therefore its best strategy is also to cut prices. The outcome, or equilibrium, in this game will have both firms reducing their prices.

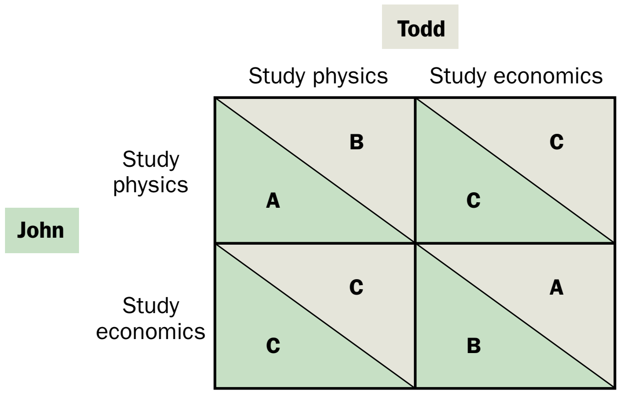

FIGURE 14.3

THE STUDY GAME

Neither player has a dominant strategy, but there are two Nash equilibria: both study physics or both study economics.

GAMES WITHOUT DOMINANT STRATEGIES

Both the prisoner's dilemma game and the price-cutting game have unique Nash equilibria. Often, however, a game will have more than one Nash equilibrium, as the following example illustrates.

Two friends are enrolled in the same physics and economics classes. They both know that studying together is much more efficient than studying alone. By working through problems together, each will greatly improve his performance on upcoming tests in the two classes. However, John is really worried about his physics grade and would prefer that he and Todd focus on physics. Conversely, Todd is most worried about economics and would rather devote their study time to it. The game table for their circumstances is shown in Figure 14.3, with the payoffs expressed in the average grade for the two courses (the letter below the diagonal line in each box is the payoff to John).

Does either John or Todd have a dominant strategy? No. If John insists on studying physics, Todd is better off joining John in hitting the physics books than in going it alone and studying economics. On the other hand, if John is willing to study economics, then Todd's best response is obviously to study economics too. Similarly, John's best strategy is to study physics if that is what Todd also does, while John's best strategy is to study economics if Todd does that instead. Neither player has a dominant strategy, one that is best regardless of what the other does.

Even though there are no dominant strategies in this game, there are two Nash equilibria—either both friends study physics or both study economics. If Todd pulls out Stiglitz and Walsh's economics textbook and starts reviewing the material on trade-offs and incentives, John's best strategy is to join him. Studying economics is John's best choice, given that Todd is studying economics. After all, if he goes off and studies physics by himself, his average grade in the two classes will be a C; while if he instead studies economics with Todd, he will do worse in physics, but he will do so much better in economics that his average grade will be a B. The same is true for Todd—given that John is studying economics, Todd's best strategy is to join him. So the lower right box in the diagram is a Nash equilibrium. But it is not the only one. The upper left box, where both end up studying physics together, is also a Nash equilibrium. While it may not enable us to predict a *unique* equilibrium in a game, the concept of a Nash equilibrium can help eliminate some outcomes. Neither the upper right nor the lower left box in the diagram is a Nash equilibrium. If Todd studies physics, then studying economics alone is not John's best response.

Internet Connection

THE ZERO-SUM GAME SOLVER

Some games have a fixed total payoff: in *Survivor,* the winner receives $1 million and the next few survivors receive smaller amounts. Nothing the contestants do can affect the total prize money available. If one contestant receives more, another receives less. Games like this are called *zero-sum games.* When we think of games, we usually think of sports or chess, or perhaps gambling—all zero-sum games. In sports, one team wins and the other loses. In gambling, every dollar you win is someone else's loss. While many people think that economic exchange is a zero-sum game, you should understand by now that it is not—exchange can leave *both* parties better off. On the Web site of Professor David Levine at UCLA, you can find a program that will find the solution to any two-person zero-sum game. Invent a game yourself and find its solution at http://levine. sscnet.ucla.edu/Games/zerosum.htm.

Wrap-Up

THE BASIC TYPES OF GAMES

1. The prisoner's dilemma—both players have dominant strategies and the game has one Nash equilibrium. Application: understanding why collusion is difficult.
2. The price-cutting game—only one player has a dominant strategy and the game has one Nash equilibrium. Application: understanding competition between duopolists.
3. The study game—neither player has a dominant strategy and there are two Nash equilibria. Application: understanding banking panics (see p. 322).

Repeated Games

In the basic prisoner's dilemma game, each party makes only one decision. The game is played just a single time. The two players could do better if they could somehow cooperate and agree not to pick the dominant strategy. But when the game is actually played, each has an incentive to break any prior agreement and do what is in his own best interest. If the players or parties interact many times, then the strategies for each can become more complicated. There may be additional ways to try to enforce cooperation that would benefit both parties. Games that are played many times over by the same players are called **repeated games.**

To see how the nature of the game is changed when it is repeated, let us consider the actions of two politicians running for the Senate. At the start of an election campaign, suppose each candidate announces that she will refrain from running negative ads as long as her rival does. But if the rival cheats on this agreement and runs a negative ad, the other candidate responds by running her own negative attack ads. Can this threat of retaliation ensure that the two candidates run clean campaigns?

Imagine that there are several weeks remaining until the date of the election. Each candidate will figure that she should run negative ads in the last week of the campaign: this is her best strategy, because the threat of retaliation carries no force after the election. There is no longer any payoff in continuing to cooperate, since the campaign (and the game) ends on election day. The agreement breaks down the week before the election.

Now consider what happens two weeks before the election. Both candidates know that the other will start running negative ads the following week. But if they are not going to honor their agreement during the last week of the campaign anyway, then the threat of future retaliatory attacks is completely meaningless. Hence, each candidate will reason that it pays to cheat on the agreement by running a negative ad. The agreement breaks down two weeks before the election. Reasoning backward through time, we see that the agreement not to run negative ads will break down almost immediately, no matter when the election is held.

This example illustrates an important principle of strategic thinking: think first about the end of the game and work backward from there to identify the best current choice. Making decisions in this way is called **backward induction.** For each decision a player can make, she needs to work out her opponent's optimal response and what her own payoff will be. Then, in the first stage of the game, she can adopt the strategy that gives her the highest payoff.

Backward induction also applies in the various games considered earlier, from Kelly making her decision in *Survivor* to the players in a prisoner's dilemma game. For example, prisoner A reasoned, "Suppose my partner confesses; what is my best strategy? And if my partner does not confess, then what is my best strategy?" Each thought about the consequences of his opponent's choices and worked backward to determine what he should do.

Our analysis of collusion in a repeated game setting may seem too pessimistic about the ability of firms or individuals to cooperate. Certainly we see that individuals often do find ways to cooperate, trading a short-term gain to establish longer-term relationships that yield higher benefits. And firms behave similarly, offering services or providing higher-quality products that lower their immediate profits but contribute to higher profits in the future. In strategic games that do not have a finite end—that always offer the possibility of another round—a variety of strategies may enable players to cooperate to achieve better outcomes.

Wrap-Up

BACKWARD INDUCTION

When strategic interactions occur for a repeated but fixed number of times, the best approach is to start from the end of the game and work backward to determine the optimal strategy. Backward induction helps the player focus on the future consequences of his current decision.

REPUTATIONS

Developing good reputations can be useful when players are engaged in repeated interactions. A firm that relies on local customers for repeat business has more of an incentive to develop a reputation for good service than does one with little repeat business. A car mechanic might have an incentive to pad the bill or otherwise cheat a customer if he never expects to service that particular car again, but in the long run he might profit more by gaining a reputation for good service and relying on repeat business from his customers.

Gaining a reputation is costly in the short run—initially the car mechanic will be unable to charge any more than do garages unconcerned about their reputations. His lower profits in the short run are like an investment: they will pay off in the future when the reputation he has developed enables him to charge more than most mechanics.

TIT FOR TAT

Economists have set up laboratory experiments, much like those used in other sciences, to test how individuals actually behave in these different games. The advantage of this sort of **experimental economics** is that the researcher seeking the crucial determinants of behavior can change one aspect of the experiment at a time. One set of experiments has explored how individuals cooperate in situations like the prisoner's dilemma. These experiments tend to show that participants often evolve simple strategies that, although sometimes apparently irrational in the short run, can be effective in inducing cooperation (collusion) as the game is repeated a number of times. One common strategy is tit for tat. In the case of two oligopolists, one might threaten to increase output if the other does, even if doing so does not maximize its short-term profits. If the rival finds this threat credible—as it may well be, especially after it has been carried out a few times—the firm may decide that it is more profitable to cooperate and keep production low than to cheat. In the real world, such simple strategies may play an important role in ensuring that firms do not compete too vigorously in markets in which there are only three or four dominant firms.

Internet Connection

THE PRISONER'S DILEMMA

You can play a repeated game version of the prisoner's dilemma against a computer at www. princeton.edu/-mdaniels/PD/PD.html. Try a tit-for-tat strategy and see how the computer responds.

INSTITUTIONS

In many situations, institutions ensure that a cooperative outcome is reached. International organizations such as the World Trade Organization (WTO) serve to enforce agreements that promote international trade. Member countries agree to abide by certain rules that forbid the type of trade restrictions and beggar-thy-neighbor policies that proved so disastrous during the 1930s. Professional sports leagues impose salary caps, which limit the ability of teams to boost salaries. Deposit insurance eliminates the incentive for depositors to withdraw funds when there are rumors of financial trouble at their bank, because they know their money is protected even if their bank fails.

Case in Point

BANKING PANICS

Between 1930 and 1933, the United States suffered a massive financial panic that forced about nine thousand banks to fail. This disruption of the financial system contributed to the severity of the Great Depression, when unemployment reached levels as high as 25 percent of the labor force. As banks closed their doors, businesses that relied on bank credit to finance their inventories and investments were forced to cut back production and lay off workers. How can game theory help us understand why so many banks failed?

During the banking panics of the early 1930s, depositors rushed to withdraw their savings, leading to many bank failures.

If you have a deposit in a commercial bank today, the federal government insures it (up to $100,000): if your bank makes bad investments and goes bankrupt, the federal government will make sure that you receive all your money back. Before 1933, however, bank deposits were not insured. If your bank went bankrupt, you could lose everything—that is, unless you acted at the first hint of trouble and withdrew your deposits before the bank ran out of cash. Banks lend out most of the money they receive as deposits, holding only a small fraction as cash to meet daily unpredictable fluctuations in deposits and withdrawals. If all depositors were to suddenly demand their money back, a bank would quickly run out of cash and be forced to close. That's what almost happened to George Bailey's bank in the movie *It's a Wonderful Life,* and it is exactly what happened in real life in the 1930s as depositors raced to be the first to withdraw their deposits. When they all tried to withdraw their deposits at the same time, there simply was not enough money on hand to pay everyone. Thousands of banks had to shut their doors. Everyone would have been better off leaving their deposits alone, since in that case the banks could have remained open.

The concept of Nash equilibrium can help us understand bank runs and financial panics. Consider a simple example of a bank with just two depositors—call them A and B. Each depositor must decide whether to try to withdraw her deposits from the bank or to leave them there. Assume each has deposited $1,000 at the bank. The bank has used these funds to make loans and investments but keeps $200 on hand in its vault. If the bank's loans are repaid, the bank can pay an interest rate of 5 percent to its depositors.

If neither depositor tries to withdraw funds from the bank, let us assume they both will eventually receive the full value of their deposits plus the 5 percent interest (for a total of $1,050). If depositor A withdraws while depositor B does not, A can take out $200, all the cash the bank has on hand. The bank then must shut its doors, and depositor B receives nothing. The reverse happens if depositor B tries to withdraw its funds while A does not. If both try to withdraw their money, the most each can get is $100. The payoffs are shown in Figure 14.4.

Clearly both depositors are better off if neither attempts a withdrawal. In that case, they both would eventually receive $1,050. This is also a Nash equilibrium. If depositor A leaves her money in the bank, depositor B's best strategy is to do the same. Conversely, if depositor B leaves her money in the bank, depositor A's best strategy is to do the same. Each reasons as follows: "If the other leaves her deposits in the bank, my best strategy is also to leave my deposits in the bank." So there is an equilibrium in which neither depositor trys to withdraw funds, and the bank stays open.

Just as in the earlier example of the two friends deciding what to study, there are two Nash equilibria to the deposit withdrawal game. The second Nash equilibrium occurs when both depositors try to withdraw their money and the bank fails. In this case, each reasons as follows: "If the other depositor tries to withdraw, I'm better off if I also try to withdraw. That way, at least I get $100,

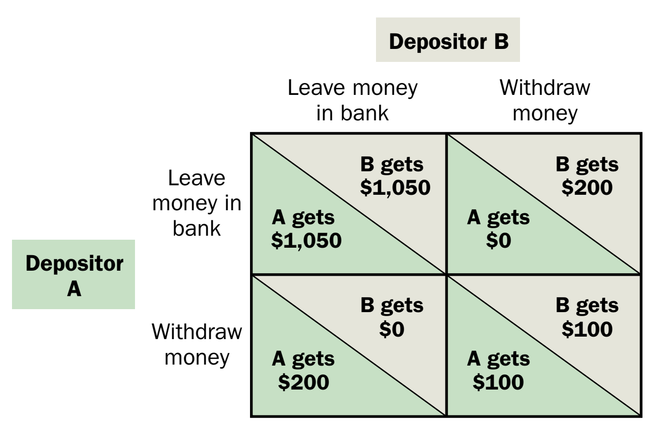

FIGURE 14.4

A BANK PANIC AS A STRATEGIC GAME

Like the study game, this game has two Nash equilibria: both withdraw money or both leave money in the bank. The equilibrium in which both leave their money in the bank leaves both better off than if they both withdraw their money.

which is better than nothing.” So there is an equilibrium in which each rushes to the bank, and the bank is unable to fully meet its obligations to the two depositors. The bank fails.

In this example, there is a good equilibrium—the one in which the bank remains open and the depositors eventually receive all their money plus interest—and a bad equilibrium, when the bank is forced to close. In the prisoner’s dilemma, in contrast, the only equilibrium was inferior to an alternative set of strategies (neither prisoner confesses). The example of John and Todd illustrated another possible outcome: one of the two equilibria was preferred by Todd and the other by John.

Financial panics can be thought of as shifts from the good to the bad equilibrium. Such a shift can occur if depositors start to worry about the financial soundness of the banking sector, even if such fears are unfounded. The simple argument illustrated by this game provides part of the rationale behind federal deposit insurance. Deposit insurance gives each depositor confidence that her money is safe, regardless of what other depositors do. No depositor has an incentive to try to outmaneuver others by getting her money out first.

Sequential Moves

In the prisoner’s dilemma, each player must make a choice without knowing what the other has done. The players move simultaneously. In many situations, however, one player must move first, and the second player then responds directly to what the first has done. This is called a **sequential game,** since players take turns and each can observe what choices were made in earlier moves. In such a game, the player who moves first must consider how the second player will respond to each possible move he can make.

Sports offer many instances of strategic behavior. The baseball manager deciding whether to bring in a relief pitcher is a prime example of someone playing a sequential game. Conventional wisdom in baseball says that a left-handed batter has more success against a right-handed pitcher, and a right-handed batter does better against a left-handed pitcher. Does that mean that if a right-handed player is coming to bat, the manager should bring in a right-handed pitcher? Not necessarily. Once a right-handed pitcher is brought in, the manager of the other team can send a left-handed pinch hitter to the plate. So a manager considering a change of pitchers needs to think about how his counterpart will respond to each of the possible pitching choices he can make.

This example illustrates an important aspect of a sequential game. The player who moves first must anticipate how the second player will respond. Take the case of a firm facing the potential entry of a new rival into its market. Suppose a software firm, call it Redhat, is considering the launch of a new operating system that will compete with Microsoft Windows. Redhat must decide whether to enter the business or stay out. If it enters, Microsoft must decide whether to peacefully compete or to wage a price war. Suppose that Redhat enters and Microsoft competes peacefully. Assume Microsoft will earn profits of $50 billion and Redhat $10 billion on its operating system. If instead Microsoft engages in a price war, assume both firms will lose money, with Microsoft losing $1 billion and Redhat losing $500 million. If Redhat decides not to enter, it earns $0, while Microsoft earns profits of $80 billion. Will Redhat enter? And will Microsoft engage in a price war?

To simplify this complex scenario we can use a **game tree** diagram, which is the standard way to represent a sequential game. Different branches on a game tree indicate the various outcomes that could occur, given all the possible strategies the players could follow. For example, the entry game involving Microsoft and Redhat is represented by the game tree of Figure 14.5. At the end of each branch, the payoffs to Redhat and Microsoft are shown. The first number is Redhat's payoff; the second is Microsoft's.

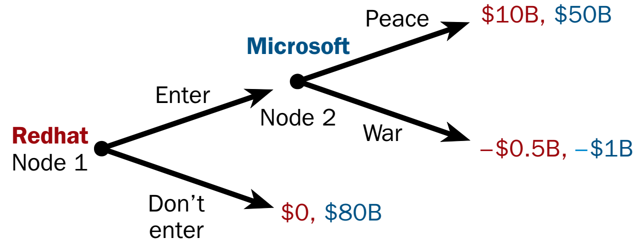

FIGURE 14.5

GAME TREE FOR A SEQUENTIAL GAME

Redhat makes the first move. If Redhat decides to enter the market, Microsoft may then choose between peaceful coexistence or a price war. By reasoning backward from these options, we can see that Redhat will enter the market and Microsoft will choose to compete peacefully.

Each node—the points at which new branches split—represents a decision point for one of the players. In this game, Redhat moves first (node 1). Microsoft moves second, after it has learned whether Redhat has entered or not. If Redhat decides to enter, we move along the upper branch of the game tree from node 1 to node 2. Microsoft must then choose whether to compete peacefully or wage a price war. From node 2, we move along the upper branch if Microsoft competes peacefully, and the payoffs are $10 billion to Redhat and $50 billion to Microsoft; or we move along the lower branch if Microsoft wages a price war, in which case the payoffs are –$500 million and –$1 billion, respectively. If Redhat decides not to enter, we move along the lower branch from node 1. Here, there is nothing further Microsoft must decide, so the game ends and the players receive the payoffs shown ($0 for Redhat and $80 billion for Microsoft).

In deciding whether to enter, Redhat's managers will reason as follows: "If we enter, Microsoft can either wage a price war or compete peacefully. In the former case, it loses $1 billion, and in the latter case it makes $50 billion. Clearly, once we enter, it will be in Microsoft's best interest to compete peacefully, so we should enter."

This example again illustrates backward induction. "Thinking strategically" requires that one think first about the end of the game, and work backward from there to determine the best current choice. Redhat asks itself what Microsoft will do if Redhat has entered. It works backward from there to determine if it should enter.

Using backward induction is easy when a game tree is employed . At each node that leads to an end to the game, determine the best strategy of the player who makes a decision at that node. Then work backward. At node 2, Microsoft's best strategy is to compete peacefully. It gets $50 billion that way, while it would lose $1 billion in a price war. Now work back to the previous node—node 1, where Redhat makes its decision. From its analysis of Microsoft's options at node 2, Redhat knows that if it enters the market, Microsoft will compete peacefully, leaving Redhat with a $10 billion profit. Redhat's other option at node 1 is not to enter, which would leave it with profits of $0. Clearly, Redhat's best strategy is to enter.

Thinking Like an Economist

INFORMATION AND THINKING STRATEGICALLY

Information plays a critical role in strategic behavior, but sometimes in surprising ways. Take the prisoner's dilemma, for example. Each prisoner must make a decision about whether to confess without knowing what the other prisoner has done. It might seem that the outcome would change if one prisoner could know the other's choice in advance, but it turns out that providing this extra information makes no difference. It is still best to confess, because confessing is a dominant strategy: it is the best strategy for each prisoner, regardless of what the other did.

In the absence of a dominant strategy, changing the information that the player has *is* likely to alter the player's best strategy. Consider the case of an insurance company that offers health insurance. Suppose the insurance company can obtain information such as whether an individual smokes. Since smoking is associated with many health problems, the insurance company will offer different policies to smokers and nonsmokers; smokers will pay higher insurance premiums to reflect the likelihood they will incur higher medical bills. Now suppose a law is passed that forbids the insurance company from collecting such information. Governments frequently pass such laws on the grounds that collecting certain types of information is an invasion of privacy. If the company offers just a single policy to everyone, it faces an adverse selection problem of the type we will consider in Chapter 15—those with the poorest health are the ones most likely to buy the insurance. But if the firm thinks strategically, it might reason along the following lines: "If we offer only one type of policy, we run the risk that only those with health problems will buy it and we will lose money. Instead, let's offer two different policies. One will have a high deductible, so that patients themselves have to pay a large amount for doctor visits and other services before insurance kicks in. The other policy will have a low deductible, with insurance paying for most medical services. We will charge more for the policy with the low deductible. The low-deductible policy will be more attractive to individuals who think they will need lots of medical services. It will be attractive to smokers. The policy with the high deductible will be more attractive to individuals who think they are less likely to need lots of medical care. It will be attractive to nonsmokers. By offering these separate policies, we can get individuals to reveal information about their health risks."

By thinking strategically, the insurance company is able to overcome some of the information problems it faces. By offering different policies, it is able to separate individuals into high- and low-risk groups.

Time Inconsistency

Threats and promises are common components of strategic behavior. We have already discussed the case of a monopolist who threatens to wage a price war if another firm enters its market (our Redhat and Microsoft example). In that case, the firm that is considering entering the market understands that once it enters, the existing firm's best strategy is to compete peacefully. Because the potential entrant can use backward induction to make this determination, the initial threat is ineffective in deterring entry—it is not credible.

Now instead of looking at this problem from the perspective of the firm contemplating entrance, consider the situation of the monopolist trying to protect its market. It makes sense for the monopolist to try to scare off potential rivals by threatening a price war if any business attempts to enter the industry. But are these threats *time consistent;* that is, will it be in the monopolist's best interest to act in a manner consistent with its statements by actually carrying out the threat? The answer is no; as we have seen, the monopolist's best strategy if a rival enters the market is simply to compete peaceably. We can describe the threat as *time inconsistent;* it will not make sense to actually carry out the threat when it comes time to do so. The rival, knowing that the monopolist will not do as it threatened, knows that the threats can be ignored.

Time inconsistency arises in many contexts, usually in situations in which one player's promise or threat is designed to influence the other player's actions. Consider the case of Sarah, who has just graduated from high school. Her parents believe that holding down a summer job will help Sarah learn responsibility, so they offer to help pay her college tuition in the fall if she works during the summer. The implicit threat is that they will not help with her tuition if Sarah loafs around over the summer and does not work.[3] But Sarah can see that such a threat is time inconsistent. If she hangs out at the beach with her friends all summer, the only choice her parents face in the fall is whether to help her pay for college. Because Sarah knows her parents want her to receive a college education, regardless of how lazy she is, she knows they will help pay for college whatever she does. Since she would rather loaf than work, she does not get a job. Her parents' threat, designed to force her to find a job, was ineffective because it was time inconsistent. Sarah knew they would never actually carry out the threat, because it would not be their best strategy come the fall.

COMMITMENT

Consider again the situation faced by Sarah's parents at the end of the summer. At that point, it was too late to affect Sarah's summer activities, so the threat to not fund her college is no longer worth carrying out. Things would have been different if her parents could have somehow tied their hands in a way that would have prevented them from paying for Sarah's tuition unless she had actually worked during the summer. A demonstrated commitment to undertake a future action may be necessary for threats or promises to be credible.

Military strategists commonly face the problem of making threats credible. During the cold war, the United States stated that it would not rule out being the first to use nuclear weapons. This policy raised the possibility that if the Soviet Union invaded Western Europe, the United States would retaliate against the Soviet Union with nuclear weapons if necessary. Such retaliation would then lead the Soviet Union to launch a nuclear strike against the United States. A Soviet military planner using backward induction might reason that if the Soviet Union invaded Europe, the U.S. government would be faced with the choice of either launching a nuclear attack on the Soviet Union, and having millions of Americans killed in the ensuing nuclear war, or accepting a Soviet victory in Europe. Faced with that choice, the United States might decide to accept a Soviet victory. The U.S. threat to retaliate would not be credible. One argument for maintaining thousands of U.S. troops in Europe was that their loss in a Soviet invasion would force the U.S. military to launch a strike against the Soviet Union. By committing troops to Europe, the Americans made their threat more credible.

Imperfectly competitive markets provide a firm many opportunities to take concrete actions that deter rivals when threats alone would not be credible. A variant of the earlier example of Microsoft and Redhat can illustrate this point. Consider a coffee company, which we'll call Northwest Coffee, that opens a coffee store on every

[3]This example is from Herb Taylor, "Time Inconsistency: A Potential Problem for Policymakers," *Economic Review* (Federal Reserve Bank of Philadelphia), March/April 1985, pp. 3–12.

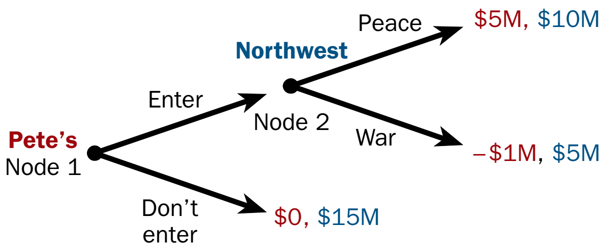

FIGURE 14.6

MARKET ENTRY GAME WITHOUT COMMITMENT

Northwest may threaten to wage a price war if Pete's Coffee enters the market, but such a threat lacks credibility.

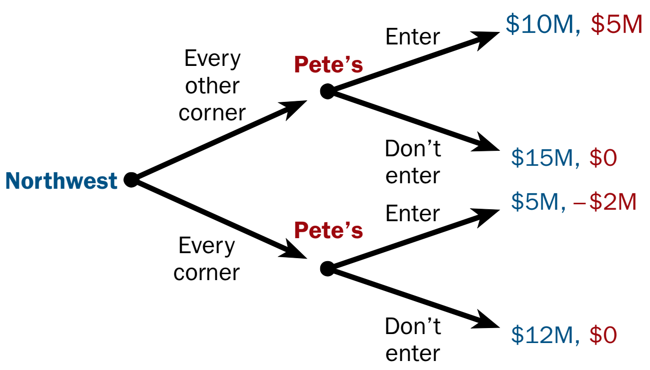

FIGURE 14.7

MARKET ENTRY GAME WITH COMMITMENT

Even though it costs more to operate stores on every street corner, Northwest will do so in order to prevent Pete's Coffee from entering the market. Opening stores on every street corner serves as a commitment mechanism.

other corner of a city's downtown. To deter a potential rival, Pete's Coffee, from opening stores on the empty corners, Northwest Coffee might threaten a price war if a rival opens up competing stores. The game tree and payoffs are shown in Figure 14.6.

Just as in our earlier example, Pete's Coffee can use backward induction to determine that if it enters, Northwest will find it more profitable to compete peacefully. Any threat by Northwest Coffee would not be credible.

Now let's change the game so that Northwest first decides whether to open on every *other* corner, or on *every* corner. The order of play now starts with Northwest's choice about how many stores to open. In the second stage of the game, Pete's Coffee must decide if it will enter. When it makes that decision, it knows whether Northwest has opened on every corner or only on every other corner. The new game tree appears in Figure 14.7.

First, note that if Northwest opens on every other corner and Pete's Coffee enters, the payoffs—$10 million for Northwest Coffee and $5 million for Pete's Coffee— reflect the previous finding that Northwest will not wage a price war. Second, the new branches of the tree show the possible outcomes if Northwest opens on every corner. Because of its added costs, Northwest's profit is smaller when it opens the additional stores and Pete's Coffee stays out of the market ($12 million) than if it opens only on every other corner and Pete's stays out ($15 million). If Northwest has a store on every corner and Pete's Coffee decides to enter the market, Northwest's profits fall to $5 million, while Pete's Coffee ends up losing $2 million.

We again use backward induction to determine the Nash equilibrium. Start from the Pete's Coffee decision node along the top branch (the branch followed if Northwest opens on every other corner). If it finds itself at this node, the best strategy is for Pete's Coffee to enter the market, earning $5 million. Now look at the Pete's Coffee decision node along the bottom branch (the branch followed if Northwest opens on every corner). Here, the best strategy is to not enter. That way, at least, it does not lose any money.

By applying backward induction, we can now analyze Northwest Coffee's decision at the start of the game. Northwest knows that if it puts a store on every other corner, Pete's will enter and Northwest will earn $10 million. If it opens on every corner, Pete's will not enter, and Northwest will earn $12 million. Northwest's best strategy therefore is to open a store on every corner. Even though the decision may appear to be counterintuitive, lowering Northwest's profits from $15 million to $12 million, the extra stores are worth opening because they deter a potential rival. Northwest's extra stores are a more credible threat to a potential rival than is a promise to engage in a price war. Stores that are open and in place serve as a commitment mechanism that deters potential rivals from entering. Northwest is worse off than if it opened fewer stores and Pete's stayed out of the market, but it is better off than if Pete's had entered.

Review and Practice

SUMMARY

1. In perfectly competitive markets, firms and consumers can decide how much to produce and how much to consume without taking into account how others might react. In imperfectly competitive markets, firms must take into account how their rivals will respond to the firms' production or pricing decisions. Firms must behave strategically in such situations. Individuals also face many situations in which they must behave strategically. Economists use game theory to predict how individuals and firms will behave.
2. In a Nash equilibrium, each player in a game is following a strategy that is best, given the strategies followed by the other players. A game may have a unique Nash equilibrium, or it may have several equilibria.
3. A dominant strategy is one that is best regardless of what the other player chooses to do. Looking for dominant strategies can help analysts predict behavior.
4. Backward induction is crucial for strategic behavior. Thinking strategically means looking into the future to predict how others will behave, and then using that information to make decisions.
5. Strategic choices are often designed to influence the choice of others. Once others have made their choices, however, carrying out the initial strategy may no longer be best. When this is the case, the original strategy was time inconsistent.

KEY TERMS

strategic behavior
game table
dominant strategy
Nash equilibrium
repeated game
backward induction
experimental economics
sequential game
game tree
time inconsistency

REVIEW QUESTIONS

1. Firms in perfectly competitive markets do not need to behave strategically. Why not? Why do oligopolists need to behave strategically? Does a monopolist need to behave strategically?
2. Professional sports leagues often have salary caps that limit the amount individual teams can pay players. Using the prisoner's dilemma game, explain why such a restriction might make the team owners better off.
3. What is a dominant strategy? Explain why each player in the prisoner's dilemma has a dominant strategy.
4. What is a Nash equilibrium? What is the unique Nash equilibrium in the prisoner's dilemma game? Can a game have more than one Nash equilibrium? Give an example to illustrate your answer.
5. In the prisoner's dilemma, each player has a dominant strategy and there is a unique Nash equilibrium in which the players choose their dominant strategy. Give an example of a game in which only one player has a dominant strategy. What is the Nash equilibrium?
6. This chapter opened with a discussion of the television show *Survivor*. What principle of strategic behavior did Kelly need to use?
7. What is a sequential game? Why does the player who moves first need to use backward induction?
8. An old saying, when a parent punishes a child, is "This hurts me more than it hurts you." Drawing on the idea of a repeated game, explain why a parent might still punish the child even if doing so really did hurt her more than the child.
9. Why might threats and promises not be credible?

PROBLEMS

1. Consider two oligopolists, each choosing between a "high" and a "low" level of production. Given their

choices of how much to produce, their profits will be as follows:

Firm B

Firm A	High production	Low production
High production	A gets $2 million profit; B gets $2 million profit	A gets $5 million profit; B gets $1 million profit
Low production	A gets $1 million profit; B gets $5 million profit	A gets $4 million profit; B gets $4 million profit

Explain how firm B will reason that it makes sense to produce the high amount, regardless of what firm A chooses. Then explain how firm A will reason that it makes sense to produce the high amount, regardless of what firm B chooses. How might collusion assist the two firms in this case?

2. Use the prisoner's dilemma analysis to describe what happens in the following two situations:
(a) Consider two rivals—say, producers of cigarettes. If Benson and Hedges alone advertises, it diverts customers from Marlboro. If Marlboro alone advertises, it diverts customers from Benson and Hedges. If they both advertise, each retains its customer base.

    Are they genuinely unhappy with government regulations prohibiting advertising? In practice, cigarette firms have complained quite bitterly about such government restrictions, including those aimed at children. Why?

(b) Consider two rivals—say, producers of camera film, Fuji and Kodak. Consumers want a film that accurately reproduces colors and is not grainy. Assume that initially, the two firms had products that were comparable in quality. If one does research and improves its product, it will steal customers from its rivals. If they both do research and develop comparable products, then they will continue to share the market as before. Thus, the hypothetical payoff matrix (in millions of dollars) appears below (the profits in the case of research take into account the expenditures on research):

Fuji

Kodak	No research	Research
No research	$100 Kodak profit; $100 Fuji profit	$0 Kodak profit; $150 Fuji profit
Research	$150 Kodak profit; $0 Fuji profit	$80 Kodak profit; $80 Fuji profit

Explain why both will engage in research, even though doing so reduces their profits. Could this choice make society better off, even though their profits are lower?

3. Draw the game tree for the game discussed on pages 324–325 in which Microsoft moves first and decides whether to threaten a price war; Redhat then decides whether to enter; and finally, Microsoft decides whether to compete peacefully or wage the price war. Verify that Microsoft's decision at the first stage of the game has no effect on Redhat's strategy.
4. Suppose firm A is a monopolist. Firm A threatens a price war if any potential rival enters its market. Suppose firm B is contemplating such a move. The game tree is as follows (firm B's payoffs are shown first; firm A's are shown second):

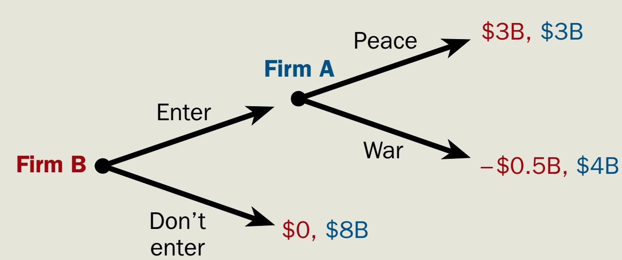

Should firm B enter? Is firm A's threat credible? Why? What makes this example different from the outcome in the Redhat-Microsoft example given in the text?

5. Draw the game tree for the sequential game between Sarah and her parents that was described on page 327. The first move is Sarah's; she decides whether to work during the summer or not work. At the end of the

summer her parents decide whether or not to pay for Sarah's tuition. Using backward induction, what is the equilibrium for this game? Now add a new first stage in which Sarah's parents announce that they will only pay for the tuition if Sarah works. Is this announcement credible? Use this example to explain what is meant by time inconsistency.

6. "Tying one's hands" can be a way to commit credibly to a certain course of action. In the 1960s film *Dr. Strangelove,* the Soviet Union deployed a doomsday device that could destroy the world and would be automatically triggered if the United States attacked. Explain how such a device could serve as a credible threat to deter a U.S. attack. In the movie, the Soviet Union did not inform the United States that the device had been deployed. Why was this a really bad strategy?
7. Suppose Quality Brands and Discounters Delux are involved in a repeated price-cutting game. Explain what a tit-for-tat strategy would be. Is a promise to "match any available price" a way for one firm to signal that it is playing tit for tat? Explain.
8. Restaurants often locate along major highways. Since most customers at such restaurants will not return, does the restaurant have an incentive to develop a reputation for good food? If reputations are important, which restaurant will have a greater incentive to offer good service—one that is part of a national chain (such as McDonald's or Burger King) or one that is locally owned?
9. How might a cultural or group norm or expectations about "correct" behavior, such as that summarized in the old saying "honor among thieves," help enforce cooperation in the prisoner's dilemma?