CHAPTER 14

Game Theory and the Economics of Information

How can we model choice under conditions of strategic interaction and what happens when decisionmakers have less than perfect information?

Chapter Outline

Learning Objectives

- Understand the basics of game theory: a mathematical technique to study choice under conditions of strategic interaction.
- Describe the prisoner's dilemma and its applicability to oligopoly theory as well as many other situations.
- Explore how the outcome in the case of a prisoner's dilemma differs in a repeated-game versus a single-period setting.
- Analyze asymmetric information and market outcomes in the case where consumers have less information than sellers.
- Explain how insurance markets may function when information is imperfect and there is the possibility of either adverse selection or moral hazard.
- Show how limited price information affects price dispersion for a product.
- Investigate advertising and the extent to which it serves to artificially differentiate products versus provide information to consumers about the availability of products and their prices and qualities.

This chapter covers two topics that have been intensively studied by economists in recent years: game theory and the economics of information. Game theory is a mathematical technique developed to study choice under conditions of strategic interaction. It has become the major approach to the study of oligopoly, and our discussion looks at how it has been adapted to give insight into problems of oligopolistic interdependence. In particular, we will see that it helps us understand the difficulties firms face in colluding to raise prices.

In most economic analysis, whether of competition or oligopoly, it is assumed that consumers have complete knowledge of prices and product characteristics. It is clear that this assumption is often violated in the real world, and recent research has begun to analyze what effect this has on the way markets function. When information is costly, consumers are not fully informed and lack either knowledge of the prices different firms charge or knowledge of the qualities of the products they sell, or both. Under these circumstances, the prices, quantities, and qualities of goods traded can be quite different than when consumers have "perfect information." Two especially interesting applications of the effects of costly information involve insurance markets and advertising. We discuss them later in the chapter.

14.1 GAME THEORY

Game theory is a method of analyzing situations in which the outcomes of your choices depend on others' choices, and vice versa. In these strategic situations, there is mutual interdependence among the choices made by the decisionmakers: each decisionmaker needs to account for how the others are affected by the choices made and how the other decisionmakers are likely to respond since their responses may affect what is the best choice to make. Among market structures, oligopoly is the setting in which such interactions are likely to be most important, and game theory has become widely used to analyze oligopolies.

GAME THEORY
a method of analyzing situations in which the outcomes of your choices depend on others' choices, and vice versa

Game theory is a general approach to analyzing strategic interactions, however, and it can be applied to issues other than oligopoly. For example, in determining the defense budget, the U.S. government must not only consider the impact of its budgetary decisions on the military decisions of potential enemies, but also recognize that these potential enemies are trying to predict how the United States will respond to their budgetary decisions. Similarly, a politician thinking about conducting negative campaign tactics must recognize that his or her opponent might respond in a variety of possible ways. Whether it is wise to pursue a negative campaign strategy depends on how the opponent is likely to respond. Game theory can be applied to a wide range of phenomena like these, but our main concern is with how it helps us understand the functioning of oligopolistic markets.

All game theory models have at least three elements in common: players, strategies, and payoffs. The **players** are the decisionmakers whose behavior we are trying to predict. In the case of oligopoly, the players are the firms. The **strategies** are the possible choices of the players. Outputs produced and prices charged are strategies in this sense, but so too are advertising budgets, new product introductions, and product differentiation. For oligopolistic firms, all of these actions can affect rivals. The **payoffs** are the outcomes, or consequences, of the strategies chosen. For firms, it is natural to express the outcomes as the profits or losses realized. It is important to remember, however, that a specific strategy (say, producing 300,000 microwave ovens) does not uniquely determine the profit (payoff) because that will depend on the strategies followed by the other players.

PLAYERS
in game theory, decisionmakers whose behavior we are trying to predict and/or explain

STRATEGIES
in game theory, the possible choices of the players

PAYOFFS
in game theory, the outcomes or consequences of the strategies chosen

Determination of Equilibrium

In addition to the players, strategies, and payoffs, other elements may play a role in determining the outcome in a game theory model, but for the moment, we'll use some simple examples to illustrate how the model links the various elements. Consider a market with just

two firms, A and B. Each firm must choose either a low or a high output; these are the only possible "strategies." However, each firm's profit depends not only on its own output, but also on the output of the other firm.

To go further, we have to be precise about exactly how profits are affected by both firms' choices. To show the possibilities with some illustrative numbers, we will use a payoff matrix. A **payoff matrix** is a simple way of representing how each combination of choices affects firms' profits. Table 14.1 is the payoff matrix for our example. There are four cells in this case. The upper-left cell shows that profit is 10 for firm A and 20 for firm B when both firms choose the low-output strategy. The upper-right cell shows that profit is 9 for firm A and 30 for firm B when A chooses a low output and B chooses a high output. (Thus, when B increases its output from low to high and A holds its output constant at low, B's profit increases from 20 to 30, but A's falls from 10 to 9—perhaps because B's higher output reduces the demand confronting firm A.) Similarly, the lower-left cell shows what happens to profits when A's output is high and B's output is low. Finally, the lower-right cell shows the outcome when both firms choose a high output.

PAYOFF MATRIX
a simple way of representing how each combination of choices affects players' payoffs in a game theory setting

In interpreting this payoff matrix, we assume that firms act independently. Note that this means that firm A has only two possible strategies: It can choose one of the two horizontal rows in the payoff matrix. If it chooses the top row (low output), its profit will be either 10 or 9, depending on what B chooses. Firm A cannot choose which cell it will end up in because it doesn't control B's choice. In the same way, firm B can choose only between the two vertical columns. If it chooses the right column, for example, its profit will be either 30 or 25, but it doesn't know which unless it can predict A's output choice. (We assume that the firms select their outputs simultaneously so neither knows with certainty what the other's choice will be.)

The purpose of this exercise is, of course, to predict what the market equilibrium will be. So let's see if we can figure out what each firm will do. Consider firm A. If firm B chooses a low output, firm A is better off selecting a high output because its profit would then be 20 rather than 10. On the other hand, if B chooses a high output, A is also better off selecting a high output because its profit would then be 18 rather than 9. Consider carefully what this means: A is better off with a high output *regardless of what B does*. In this case A has a *dominant strategy* because the high-output choice is best no matter what B does. In this situation it is easy to predict that firm A will produce a high output.

TABLE 14.1

DOMINANT-STRATEGY EQUILIBRIUM: A SIMPLE OLIGOPOLY GAME

		Firm B	
		Low Output	High Output
Firm A	Low Output	10, 20	9, 30
	High Output	20, 17	18, 25

DOMINANT-STRATEGY EQUILIBRIUM
the simplest game theory outcome, resulting from both players having dominant strategies

Now apply the same reasoning to firm B. If firm A chooses a low output, B is better off selecting a high output because its profit would then be 30 instead of 20. Alternatively, if A chooses a high output, B is still better off selecting a high output because its profit would then be 25 rather than 17. Thus, B also has a dominant strategy: its high-output strategy is best regardless of what A does. We would predict that firm B will produce a high output.

In this example both firms have dominant strategies, and so we expect the equilibrium outcome to include high output from both firms. This is called a **dominant-strategy equilibrium** and is perhaps the simplest game that can be imagined. As you might expect, however, not all situations are so easily analyzed.

APPLICATION 14.1 DOMINANT STRATEGIES IN BASEBALL

In baseball, it is advantageous for any forced base runner (a runner who will be forced to advance to the next base if the batter gets on base) to run on the pitch when there are two outs and the count on the batter stands at three balls and two strikes.[1]

This is so because, if the batter does not hit the pitch, either the pitch is a third strike and the inning ends or it is a fourth ball and the runner and batter advance. If the batter fouls off the pitch, the runner simply returns to the initial base (if the foul isn't caught) or the inning ends (if the foul is caught). But if the batter hits a fair ball, a runner who left base on the pitch has a better chance of advancing or scoring. Thus, running on the pitch is a dominant strategy for a forced runner when there are two outs and the count is three balls and two strikes on the batter.

[1]This application is based on Avinash K. Dixit and Barry J. Nalebuff, *Thinking Strategically* (New York: W. W. Norton, 1991).

A slightly more involved situation is shown in Table 14.2. Here, we have changed only one number from the previous table. When both firms choose a low output (upper-left cell), A's profit is now 22 (rather than 10 as before). Now reconsider A's output choice. If B chooses a low output, A's best choice is a low output (profit of 22 rather than 20), but if B chooses a high output, A's best choice is a high output (profit of 18 rather than 9). Firm A now does not have a dominant strategy: its best choice depends on what firm B does. *What output should firm A choose?*

If firm A can predict what B will do, it would then know which choice was best. So A has to try to figure out what B's choice will be before making its own decision. This is characteristic of most game-theory applications: players have to evaluate and predict what their rivals will do since their own decisions depend on the rivals' choices. Note that this was not necessary in our first example, when both firms had dominant strategies. Each firm's best choice was a high output regardless of the actions of the other. So let's consider how firm A might try to figure out what B will do. In this case it is fairly simple. Firm B still has a dominant strategy. For B, a high output is best regardless of what A does. Thus, A might reasonably predict that B will choose a high output. And if B selects a high output, A's profit is higher when it also selects a high output. Consequently, our analysis suggests that two rational, profit-maximizing firms would both, in the Table 14.2 setting, select a high output, and that that would be the game's equilibrium.

TABLE 14.2

ANOTHER OLIGOPOLY GAME

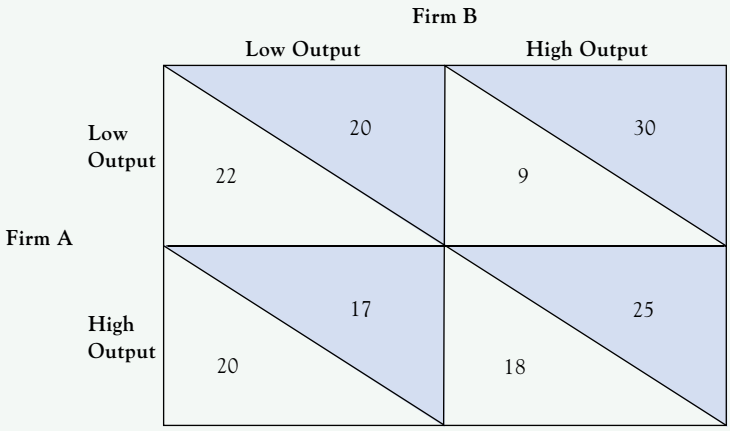

NASH EQUILIBRIUM
a set of strategies such that each player's choice is the best one possible given the strategy chosen by the other player(s)

The equilibrium we have identified in Table 14.2 is not a dominant-strategy equilibrium, however, because that term refers to a case where both firms have dominant strategies, and here only one does. To analyze the outcomes in many games, we need a more general concept of equilibrium than a dominant-strategy equilibrium. The concept most widely used is called a **Nash equilibrium**, after mathematician John Nash, who formalized the notion in 1951 and subsequently won the 1994 Nobel Prize in economics for his contribution. (Nash's genius and lengthy struggle with mental illness were the subjects of a recent award-winning book and movie, both entitled *A Beautiful Mind*.)

A Nash equilibrium is a set of strategies such that each player's choice is the best one possible *given the strategy chosen by the other player(s)*. To see that the high output–high output set of strategies in Table 14.2 is a Nash equilibrium, we evaluate it as follows. What we want to see is whether firm A's best choice is high output *when firm B chooses high output*, and simultaneously whether firm B's best choice is high output *when firm A chooses high output*. If firm B chooses high output, then firm A is best off by choosing high output (profit of 18 rather than 9), so the lower-right cell passes the first part of the test. If A chooses high output, B is best off by choosing high output (profit of 25 rather than 17), so the lower-right cell passes the second part of the test also. The two strategies shown in the lower-right cell then, represent a Nash equilibrium.

If you apply this same reasoning to the equilibrium we identified in Table 14.1, you will see that the high output combination of strategies there is also a Nash equilibrium. A dominant-strategy equilibrium (as in Table 14.1) is a special case of a Nash equilibrium. That is, all dominant-strategy equilibria are also Nash equilibria, though not all Nash equilibria are dominant-strategy equilibria (as the example in Table 14.2 illustrates). Unfortunately, not all games have a Nash equilibrium. Thus, while the concept works in many cases, it does not let us determine the outcome in all strategic situations.

The Nash equilibrium is closely related to the analysis of the Cournot oligopoly model discussed in Chapter 13. Recall that the Cournot equilibrium is one in which each firm is producing its best output *given the outputs of other firms*. This is exactly the description of a Nash equilibrium, and sometimes the equilibrium in such a market situation is described as a Cournot–Nash equilibrium. More importantly, the application of game theory to the strategic interaction in that sort of model has provided a reason why firms might rationally choose to behave in the way Cournot described.

14.2 THE PRISONER'S DILEMMA GAME

PRISONER'S DILEMMA
the most famous game theory model in which self-interest on the part of each player leads to a result in which all players are worse off than they could be if different choices were made

The most famous game theory model is the **prisoner's dilemma**. We will first describe the game in the form that made it famous, thereby illustrating how it got its name. Although it may not be immediately apparent, the prisoner's dilemma is widely applicable—to oligopoly theory as well as many other situations.

Two individuals, Nancy and Sid, are picked up on a public nuisance charge. While questioning the suspects, the district attorney begins to suspect that they may be two key players in an international drug ring. The district attorney, however, does not have enough evidence to convict them of the more serious charge, so she comes up with the following ploy in an attempt to extract a confession. She separates the two prisoners so they cannot communicate with one another and tells each the following:

1. If both confess to drug trafficking, they will both go to jail for 10 years.
2. If neither confesses, they will be charged and convicted of the nuisance offense, and each will receive a 2-year sentence.
3. If one confesses (turns state's evidence) and the other does not, the one who confesses will get a reduced sentence of 1 year; the one who doesn't will be convicted and go to jail for 15 years.

Table 14.3 shows the relevant payoff matrix for this game. Now we can apply our knowledge of game theory to determine what each suspect will do. Let's look at the situation confronting Sid first. If Nancy confesses, Sid's best strategy is to confess also because he then gets 10 instead of 15 years. If Nancy doesn't confess, Sid's best strategy is still to confess because he then gets only 1 year instead of the 2 years he gets if he doesn't confess. *To confess is a dominant strategy for Sid:* Regardless of what Nancy does, Sid does better by confessing.

The situation is exactly symmetrical for Nancy. If Sid confesses, Nancy's best strategy is to confess (10 versus 15 years), and if Sid doesn't confess, Nancy still does better to confess (1 versus 2 years). Confession is a dominant strategy for Nancy also.

Thus, the dominant-strategy equilibrium (and a Nash equilibrium) is for both parties to confess, and that is the expected outcome. This may not seem surprising until you realize fully what it means. The predicted outcome is one where *both* suspects are worse off

TABLE 14.3 THE PRISONER'S DILEMMA

		Nancy: Confess	Nancy: Don't Confess
Sid	Confess	10 years (Sid) / 10 years (Nancy)	15 years (Sid) / 1 year (Nancy)
Sid	Don't Confess	1 year (Sid) / 15 years (Nancy)	2 years (Sid) / 2 years (Nancy)

than they would be if neither confessed (in which case they would get only 2 years each, the lower-right cell). Since they both know this, why do they confess? Because it is in each one's self-interest to confess, even though the *collective* outcome of each pursuing their self-interest is inferior for both. By the way, do not mistakenly think that the reason each confesses is that they believe the other will also. The reason for the predicted outcome is stronger than that: it is in each suspect's interest to confess, regardless of the other's actions.

When they first encounter the prisoner's dilemma, many people try to figure out some way in which the prisoners could realize the best all-around outcome, the 2-year sentence drawn if neither one confesses. By adding some additional elements to the scenario, it is indeed possible to spin a game-theoretic tale where both refuse to confess. For example, if Nancy and Sid are lovers, such that each feels as much pain if the other goes to jail as if they go to jail themselves, they would not confess. Another possibility might be that each suspect believes that if he or she is the only one to confess, the suspect who subsequently does 15 years would be willing to commit murder on release in revenge. In that case, there would probably be no confessions. (Note that the payoffs would be different than just the jail sentences shown in the table.) Taken on its own terms, however, the prisoner's dilemma does show how the individual pursuit of self-interest can, in certain situations, produce results that are inferior for all players.

The prisoner's dilemma has wide-ranging applicability. It has been used to model the interactions between the United States and the Soviet Union in the days of the Cold War. It has also been employed to explain political ticket-splitting (when voters cast ballots for candidates from different parties in various races in the same election) and why the U.S. health care system is so often characterized as being both overly expensive and bureaucratic. In this chapter we explore how the prisoner's dilemma can be applied to explain cheating by members of a cartel, the effects of a "curve"-based grading system on student study effort, and certain aspects of World War I trench warfare.

APPLICATION 14.2 THE CONGRESSIONAL PRISONER'S DILEMMA

Economist Russell Roberts of Washington University in St. Louis likens the difficulty the federal government has encountered in restraining spending over the last three decades to the case in which one's bill at a restaurant is spread evenly across the other 100 diners in the restaurant. If you were responsible for your entire bill, you would usually spend only $6 on a meal and wouldn't splurge on a second drink and dessert that would add $4 to the tab. On the other hand, adding the $4 drink and dessert costs only 4 cents when the bill is spread out evenly over all 100 patrons of the restaurant:

> Splurging is easy to justify now. In fact, you won't just add a drink and dessert, you'll upgrade to the steak and add a bottle of wine. Suppose you and everybody else orders $40 worth of food. The tab for the entire restaurant will be $4,000. Divided by the 100 diners, your bill will be $40. [While you'll get your "fair share"] this outcome is a disaster. When you dined alone, you spent $6. The extra $34 of steak and other treats were not worth it. But in competition with the others, you chose a meal far out of your price range whose enjoyment fell far short of its cost—self-restraint goes unrewarded. If you go back to ordering your $6 meal in hopes of saving money, your tab will be close to $40 anyway, unless the other 99 diners cut back also. The good citizen starts to feel like a chump. And so we read of the freshman Congressman eager to cut pork out of the budget but in trouble back home because local projects will also come under the knife. Instead of being proud to lead the

way, he is forced to fight for the projects, to make sure his district gets its "fair share."[2]

In Roberts' judgment, therefore, the average representative confronts a prisoner's dilemma wherein promoting spending on projects in one's district is a dominant strategy. The resulting equilibrium—an overall high level of government spending (and the taxes that must support such spending) across districts—is inferior to the outcome that would emerge if representatives exercised more restraint in their pursuit of spending on government projects in their individual districts.

[2]Russell Roberts, "If You're Paying, I'll Have Top Sirloin," *Wall Street Journal*, May 18, 1995, p. A18.

The Prisoner's Dilemma and Cheating by Cartel Members

Firms in oligopolistic industries often find themselves in a prisoner's dilemma where each firm acting in its self-interest produces an outcome in which all firms are worse off. Consider a very simple setting, with just two firms in an industry. We will use the example developed in Chapter 13 to illustrate the Cournot model. Utopia and Artesia are the firms, and they sell a homogeneous product. We will consider the possibility that the two firms will coordinate their activities to increase their profits. In other words, the firms try to form a cartel. The cartel agreement would have each firm restrict its output so that the market price will be high. However, each firm might consider raising its output beyond the cartel quota (cheating on the agreement). Thus, for each firm the two possible strategies are to comply and to cheat. How will the firms decide which strategy to choose?

The firms' various alternatives are shown in Table 14.4. The numbers in the cells of the payoff matrix represent the firms' profits in each of the possible outcomes. Importantly, the numbers are not just arbitrarily chosen, but reflect the actual economic environment being investigated. (Or at least the relationship between the numbers, their rank ordering, reflects the problem being studied.) Let's see how we can apply economics to explain each firm's profit for each outcome.

Let's assume that Artesia and Utopia have the same cost curves so when they collude, they agree to produce the same output. Together they will produce the monopoly output for the industry because that maximizes their combined profit. With both producing the same amount, each receives half the monopoly profit. This situation is shown in the lower-right

TABLE 14.4

CHEATING IN A CARTEL

		Artesia: Cheat	Artesia: Comply
Utopia: Cheat		10 / 10	5 / 25
Utopia: Comply		25 / 5	20 / 20

cell where both firms restrict their output, thereby complying with the cartel agreement, and each realizes a profit of 20. Now consider the upper-left cell, where both firms cheat. Because cheating means producing a higher output, price will be lower and, more significantly, profits will be lower. Because the firms are identical, let's suppose that both firms have the same "cheating output" (which we might assume to be the Cournot output), and so will realize the same profit, which we take to be 10.

Now consider the lower-left cell, where Artesia cheats and Utopia complies. The easiest way to see how this affects the profit of each is to imagine starting with the perfect cartel outcome at the lower right. Then Artesia cheats. This increases Artesia's profit from 20 to 25 for the reason we discussed in Chapter 13 (the demand curve confronting a firm is more elastic when it changes output than when all cartel members simultaneously change output). This action, however, reduces Utopia's profit because its output is unchanged but the price of the product is reduced by Artesia's output expansion. Moreover, the reduction in profit for Utopia is greater than the increase for Artesia because we know their combined profit has to be less than 40, the maximum (monopoly) profit shown in the lower-right cell. Here, we assume profit is 5 for Utopia when it complies and Artesia cheats, although the important point is that Utopia's profit falls by more than Artesia's profit increases. Finally, by analogous reasoning, we obtain the figures in the upper-right cell, which is the mirror image of the lower-left cell.

Understanding why the payoffs are as shown, we can analyze the behavior of the two firms. For Artesia, cheating is better if Utopia either complies (profit of 25 versus 20) or cheats (profit of 10 versus 5). Thus, cheating is a dominant strategy for Artesia. For Utopia, cheating is better if Artesia either complies (profit of 25 versus 20) or cheats (profit of 10 versus 5), so cheating is a dominant strategy for Utopia also. Thus, we expect both firms to cheat, and they end up with a profit of 10 each, even though if they both complied they would realize a profit of 20 each. As you might suspect from the outcome, this situation is a version of the prisoner's dilemma game, so the outcome is not surprising. The model shows, in very clear fashion, why firms have an incentive to cheat on a cartel agreement.[3]

It would be a mistake, however, to conclude from this analysis that firms will never successfully form a cartel. Other factors not incorporated here can increase the likelihood of collusion. For example, suppose that cartel agreements are legal and will be enforced. If Artesia and Utopia enter into a contract stipulating that both will comply, each firm will have an incentive to sign since the cartel outcome is better than the outcome realized when both cheat. What this model makes clear, however, is that the agreement must be enforceable because each firm individually still has an incentive to cheat on the agreement. Of course, in the United States cartels are illegal, so if firms illegally enter into a collusive agreement they must have some way of enforcing the agreement on their own to prevent cheating. In other words, for a cartel to be stable, there must be some way to punish firms that cheat. We will consider how this goal might be accomplished in the next section.

This game theory model of cartel behavior can be extended to explain why cheating is more likely the higher the number of firms in the cartel. Suppose there are ten firms initially complying, with each firm making a profit of 20. Then one firm cheats. In our two-firm model, this action increases the cheater's profit from 20 to 25, but with ten firms involved the increase in profit will be greater. This is simply because the cheater will confront a more elastic demand curve when it is one among ten firms than when it is one of two firms. (If

[3]Note that this example appears very similar to Table 14.1: in both cases there are two firms, each choosing between a "low" and a "high" output. Table 14.1 is not, however, a prisoner's dilemma, while Table 14.4 is. What accounts for the difference is that the interdependence is not as pronounced between the firms in Table 14.1. For example, it may be that they are producing goods that are not very good substitutes for one another, while in Table 14.4 the goods are perfect substitutes. Can you see why that would account for the difference in the payoffs shown in the two tables?

one firm among ten increases its output by, say, 20 percent, price will fall by less than after an increase of 20 percent by one of two firms.) In addition, the damage the cheater does to the complying firms will be shared among the remaining nine, and so less noticeable for each individual, noncheating firm. For this reason, the more firms in the cartel, the more common we expect cheating to be. With more firms, cheating produces a greater increase in profit for the cheater while the loss for each complying firm is smaller and harder to detect. This game theory illustration supports our earlier explanation as to why collusion is likely to be less common among a larger number of competitors. Each member has a greater incentive to cheat and so undermine the cartel the greater the number of firms involved.

A Prisoner's Dilemma Game You May Play

To see an application of the prisoner's dilemma that may be relevant in your academic life, consider students competing for grades in a class. The professor grades "on the curve" and assigns a certain distribution of grades in the class regardless of what the absolute grades are. For example, the professor may assign 20 percent of the class a grade of "A" regardless of whether 20 percent or only 5 percent score above 90. The professor may have found that in the past about 20 percent of her students do "A" quality work. If on a particular test only 5 percent get grades above 90, she may reason that the test was unfairly difficult, or that she graded more harshly than usual. Thus, she sets the "A" range to include the top 20 percent even though that means giving an "A" to students with absolute scores below 90. If students' motivation for studying is to get a good grade, they will find themselves in a prisoner's dilemma in this class.

Consider two students, Kaitlyn and Scott—taken as representative of the entire class. In this class, grades are curved so that the average is a "B." If they each study four hours per week, they will both get an absolute score of 85 and a grade of "B," as shown in the lower-right cell of the Table 14.5 payoff matrix. However, if they each study only one hour per week, they both get an absolute score of 60, but still receive a grade of "B" because 60 is now the average score for the class and the average receives a "B." This outcome is shown in the upper-left cell. Although their letter grades are the same in both these cells, the students are better off in the upper-left cell: they have the same grades at a lower cost in study time. However, if Kaitlyn studies one hour and Scott studies four hours, he will receive an "A" and she will receive a "C"—the lower-left cell. (Kaitlyn gets the same absolute score for the

TABLE 14.5

BEATING THE SYSTEM

		Kaitlyn: 1 Hour	Kaitlyn: 4 Hours
Scott: 1 Hour		B (60) / B (60)	A (85) / C (60)
Scott: 4 Hours		C (60) / A (85)	B (85) / B (85)

course, but ranks lower in the class now.) We will assume that a student prefers to earn an "A" with four hours of studying over a "B" with one hour. The same reasoning establishes the grades in the upper-right cell.

The way the students rank the Table 14.5 payoffs gives rise to the prisoner's dilemma. Each student acting independently has a dominant strategy: studying four hours per week. Thus, the outcome is the lower-right cell, a dominant-strategy equilibrium. However, the two students would be better off if they both studied one hour per week and attained the same "B" at a lower cost in time. Students might try to collude and attain the upper-left cell, but they would find enforcing the collusive agreement difficult because each student has an incentive to cheat (that is, studying four hours and "spoiling the curve"), and it is difficult to detect cheating and enforce sanctions. As a result, the students are caught in a prisoner's dilemma.

Does this model explain why students study as much as they do? Actually, it is clear that other factors play a role in studying decisions. For example, if a professor has strict absolute standards (even if you have the highest score in the class, you don't get an "A" unless it is above 90), there will be no prisoner's dilemma. A more serious objection to the analysis concerns the assumption that students care only about the grades they receive. If they care more about learning than about the grades, there is no prisoner's dilemma. Because this is undoubtedly true in your economics classes, the model may not be applicable there. However, it may give you some insight into student behavior in some of your other classes.

14.3 REPEATED GAMES

Oligopolists often find themselves in a prisoner's dilemma if they attempt to collude so as to increase their profits. Because the prisoner's dilemma game has a dominant-strategy equilibrium where all firms cheat (that is, don't collude), it appears that successful collusion never occurs unless binding contracts are permitted and enforced by an external authority. Such a conclusion, however, is overreaching. As explained so far, the implicit assumption is that the prisoner's dilemma game is played only once: the decision to cheat or comply is made just once. This is an appropriate assumption in some settings (such as the case in which Sid and Nancy found themselves), but it is often inappropriate when applied to firms' pricing and output decisions. Firms in an oligopoly play against one another repeatedly as they make decisions week after week. In the jargon of game theorists, the appropriate model for these market conditions is a **repeated-game model** rather than a single-period model.

REPEATED-GAME MODEL
a game theory model in which the "game" is played more than once

Let us now reconsider the prospects for effective collusion in a duopoly. We will use the data in Table 14.4 explained in the last section. Now, however, we want to imagine that the payoffs refer to the profits for Artesia and Utopia each week, and that each week the firms make a choice of either complying (low output) or cheating (high output). In the repeated-game setting, the range of strategies available to the firm increases enormously. It would be impossible to consider all the permutations, but one point should be clear: each firm now has a way to punish its rival for any past transgressions. Because the game is repeated, if Artesia cheats in the fourth week, then it is possible for Utopia to cheat in the fifth week, a tactic that imposes a cost on Artesia. In other words, the firms have a way of enforcing the cartel agreement by punishing one another for cheating.

TIT-FOR-TAT
a strategy in which each player mimics the action (e.g., cheat, comply) taken by the other player in the preceding period

What strategy will firms adopt in this setting, and what is the equilibrium likely to be? These are difficult questions without definitive answers, but we can easily see that collusion is more likely in the repeated-game setting. Suppose that Utopia adopts a tit-for-tat strategy. Under a **tit-for-tat** strategy, Utopia will comply in a given week as long as Artesia complied in the previous week. However, if Artesia cheats in one week, the following week Utopia will cheat (tit-for-tat), and will continue cheating in each successive week until Artesia

complies, after which Utopia will revert to complying. There is nothing sacrosanct about the tit-for-tat strategy, but it is simple and in a computer simulation turns out to be quite effective.[4] We can see why by considering the consequences of adopting this strategy.

Let's examine some of the options open to Artesia when Utopia plays the tit-for-tat strategy and Artesia is aware of it. Table 14.6 shows two possible scenarios that might unfold over a succession of weeks. In the first scenario, both firms comply in the first week and realize a profit of 20 (see Table 14.4 and the discussion there for the source of the profit numbers). In the second week, Utopia complies but Artesia cheats and obtains a higher profit of 25. In this scenario, we investigate the effects if Artesia continues cheating in the following weeks. In the third week, Utopia plays tit-for-tat and cheats. Artesia continues to cheat so both firms realize a profit of 10. As long as Artesia continues to cheat, so will Utopia, and they will continue to obtain a profit of 10 each.

Now imagine you are Artesia, and consider whether you would be better off following the strategy of Scenario 1 (comply in the first week and cheat thereafter) or complying every week. In Scenario 1, you get profits of 20, 25, 10, 10, and so on. By complying, however, you get profits of 20, 20, 20, 20, and so on, because Utopia will comply as long as you do. Clearly, you are better off by complying because the short-term gain you get by cheating in the second week is more than offset by the lower profits you suffer every week thereafter.

As an alternative, however, you might consider cheating one week and then complying in the next. Scenario 2 shows the evolution of the game when Artesia cheats in the second week and then complies in the third week, inducing Utopia to revert to complying in the fourth week. Once again, the sum of profits over time is higher for Artesia if it complies in every period than if it follows this strategy.[5]

This example shows why it is rational for Artesia to comply in every period when it knows that Utopia is playing a tit-for-tat strategy. This does not prove that the collusive outcome results because we have just assumed that Utopia adopted this particular strategy and that Artesia is aware of it. It does, however, suggest why in a repeated-game setting the collusive outcome is more likely: firms realize that if they cheat in the current period they are likely to suffer losses in subsequent periods, and that realization diminishes the incentive to cheat.

TABLE 14.6

CHEATING IN A REPEATED CARTEL GAME

	Scenario 1		Scenario 2	
Period	Artesia	Utopia	Artesia	Utopia
Week 1	Comply (20)	Comply (20)	Comply (20)	Comply (20)
Week 2	Cheat (25)	Comply (5)	Cheat (25)	Comply (5)
Week 3	Cheat (10)	Cheat (10)	Comply (5)	Cheat (25)
Week 4	Cheat (10)	Cheat (10)	Comply (20)	Comply (20)
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.

[4]Of course, the effectiveness of any strategy depends on what strategy opposing players are using. Robert Axelrod, in *The Evolution of Cooperation* (New York: Basic Books, 1984), finds that, on average, tit-for-tat works better than any of the other strategies investigated.

[5]This conclusion depends on the sum of the profits in weeks 2 and 3 (cheat-comply and comply-cheat), here 30, being less than the sum of the profits when Artesia complies in both periods, here 40. Refer back to the discussion of the payoffs in Table 14.4 to see why the relationship between the payoffs must be this way.

APPLICATION 14.3 COOPERATION IN THE TRENCHES OF WORLD WAR I

Cooperation between rivals can emerge in some unlikely places. A poignant example involves the "live-and-let-live" system that surfaced in World War I trench warfare.[6] The system emerged despite the passions of battle and the military logic of "kill or be killed."

The situation along the Western Front of World War I can be represented as a repeated-game prisoner's dilemma. In any given locality, opposing units could either "cheat" (shoot to kill) or "cooperate" (withhold fire or shoot in such a way as to miss). Cheating was the dominant one-period-game strategy for both sides. This is so because weakening the enemy through cheating increased the cheating side's chances of survival. Cheating by both sides, however, resulted in an outcome—heavy losses inflicted on both sides for little or no gain—that was inferior to the one produced by cooperation. And opposing units interacted with each other for what appeared, at least to them, indefinite periods of time.

The diaries, letters, and reminiscences of the trench fighters testify to the "live-and-let-live" (that is, cooperative) equilibrium that eventually emerged. One British staff officer touring the trenches was "astonished to observe German soldiers walking about within rifle range behind their own lines. Our men appear to take no notice." A soldier commented: "It would be child's play to shell the road behind the enemy's trenches, crowded as it must be with ration wagons and water carts, into a bloodstained wilderness . . . but on the whole there is silence. After all, if you prevent your enemy from drawing his rations, his remedy is simple: he will prevent you from drawing yours." Another British officer recounted: "I was having tea with A Company when . . . suddenly a [German] salvo arrived but did no damage. Naturally, both sides got down and our men started swearing at the Germans, when all at once a brave German got on to his parapet and shouted out 'We are very sorry about that; we hope no one was hurt. It is not our fault, it is that damned Prussian artillery [behind the front lines].'"

Believing that tacit truces would undermine troop morale, the high commands of both sides began rotating troops and ordering raids (whose success or failure could be monitored by headquarters staff) in an effort to destroy the "live-and-let-live" system.

[6]Robert Axelrod, *The Evolution of Cooperation* (New York: Basic Books, 1984). The quotes that follow come from Chapter 4.

APPLICATION 14.4 COOPERATIVE LAWYERS: AN OXYMORON?

Because lawyers can face each other frequently and participate in a legal community, their interactions may be repeated game versus one-period game in nature. Professors Jason Scott Johnston and Joel Waldfogel of the University of Pennsylvania examined the extent to which legal disputes are resolved more quickly and are more likely to settle where the lawyers involved interact more frequently, either as individuals or through their firms.[7] Their examination of 2,000 federal civil cases filed in the Eastern District of Pennsylvania in 1994 supports the basic game theory prediction that cooperation is more likely in repeated-game settings. Namely, legal cases involving attorney pairs who interact more regularly are resolved more quickly and are more likely to settle.

[7]Jason Scott Johnston and Joel Waldfogel, "Does Repeat Play Elicit Cooperation? Evidence From Federal Civil Litigation," *Journal of Legal Studies*, 31 No. 1 (January 2002), pp. 39–60.

Do Oligopolistic Firms Always Collude?

Our analysis of the repeated prisoner's dilemma game seems to turn our earlier conclusions on their heads. The initial analysis suggested that firms would inevitably cheat, and now we have an analysis suggesting that over time collusion is likely. But the real world is, alas, even more complicated than the repeated-game model may suggest. It is well to consider the restrictive assumptions (not all explicit) underlying our repeated prisoner's dilemma game: there are only two firms, no entry into the market occurs, the firms have identical costs and produce the same product, each firm has complete knowledge of both firms' payoffs for all strategy combinations, demand and cost conditions do not vary over time, and the game is repeated indefinitely. Relaxing almost any of these assumptions makes it less likely that collusion will be a stable outcome in an oligopolistic industry.

To see how changing one assumption can affect the outcome, consider the number of time periods the game is played. Our previous analysis implicitly assumed that the game was repeated forever. To contrast that, suppose that the two firms know that if they collude for ten weeks, new firms will have time to enter the market, and with new entry in the eleventh week, collusion will become ineffective. Artesia and Utopia then know that the game will last for ten weeks; there are ten time periods in which the data in Table 14.4 are relevant. The easiest way to see how this affects the analysis is to work backward by considering first the decisions of the firms in the last (tenth) week. In that week, it is rational for both firms to cheat because neither can be punished by the other after that, when the market will effectively be competitive. In the last period, the one-period prisoner's dilemma model that we initially discussed is relevant.

Both firms realize that it is in the other's interest to cheat in week 10. Now consider the choice in the ninth week. Because Artesia knows that Utopia will cheat in the tenth week, Artesia has no reason to comply in the ninth week. (The only reason for complying is to avoid its rival cheating in the next period, and it is going to do that anyway.) The same holds true for Utopia. Both firms will cheat in the ninth week, and both know this. Working backward in this way, we find that the firms' incentive is now to cheat in every period! This outcome occurs because of the common knowledge that the repeated game comes to an end at some point in the future.

Modifying the other assumptions in the repeated-game analysis also makes the collusive outcome more difficult to achieve. When demand conditions are variable, for example, the price each firm receives can change either because of cheating by its rival or because of falling market demand. How can the firm pinpoint the cause? If it assumes cheating whenever the price falls, the collusive arrangement can fall apart as the firms retaliate even though cheating may not have occurred.

Reaching a general conclusion about the repeated-game model of cartels is difficult because many factors can clearly influence the outcome. The analysis does, however, suggest how collusion might emerge and be enforced for at least some period of time. Still, the factors we emphasized in our discussion of cartels in Chapter 13—the possibility of entry, cheating, and the difficulty of reaching agreement, especially with heterogeneous firms and products—are forces that tend to undermine collusive arrangements.

Game Theory and Oligopoly: A Summary

We have given only a brief introduction to game theory as it relates to the study of oligopoly. Our purpose has been to convey the nature of this approach to studying markets in which strategic interactions are important. As is apparent from our discussion of collusion as a possible outcome in a repeated game, the strategic interactions become very complicated. Game theory provides a technique that is suited to investigate such interactions, but

as applied in the research literature it is a far more mathematical treatment than might be suggested by our attempt to explain its nature in a simple way.

Because the game-theoretic approach has become the dominant approach to the analysis of oligopoly over the last few decades, let's examine its lessons. Unfortunately, the use of game theory has not produced a general theory of oligopoly—a theory that would tell us, for example, that a market characterized by factors *A* and *B* will operate as in the Cournot model, while a market characterized by factors *C* and *D* will operate as a dominant firm model. A large number of possible outcomes have been enumerated in game theory models, but we do not know when or whether these outcomes will occur in real-world markets. In this sense, the state of oligopoly theory is much as it was before the application of game-theoretic techniques.

Progress has been made, however, in illuminating more clearly the complexity of the whole topic of oligopoly analysis, and many theorists expect further research to provide the basis for a better understanding of the way oligopolistic markets function. For now, we have to be content with the recognition that the outcome in such a market may fall anywhere between the monopoly (the perfect cartel case) and the competitive outcome. And whether the outcome is closer to monopoly or to competition depends on the interplay of a wide variety of factors in ways that are not yet fully understood.

14.4 Asymmetric Information

All of the models presented so far—the competitive model as well as the imperfectly competitive models—have been based on the assumption that market participants have all the information needed to make informed choices. For firms, this means knowing technology, input costs, and the prices consumers will pay for different products. For consumers, this means knowing product characteristics and prices. Although this assumption regarding knowledge is sometimes referred to as the *perfect information* assumption, the term is an exaggeration. Consumers and firms do not have to know everything for the analyses to be valid. Nonetheless, the assumption places meaningful restrictions on the models, and it is important to consider how markets function when market participants lack some information relevant to their decisions. We will begin by considering a common feature in many markets: when consumers have difficulty determining the quality of products prior to purchase.

The "Lemons" Model[8]

We are all familiar with "lemons": products that repeatedly break down or perform unsatisfactorily relative to what we expected. The model we will examine suggests that a high proportion of goods may be lemons in a market where buyers are less well informed about product quality than sellers. The basic assumption is one of **asymmetric information**: participants on one side of the market (in our case, sellers) know more about a good's quality than do participants on the other side (buyers). One market where this characteristic seems prevalent is the one for used cars. A seller of a used car normally has extensive experience using the car and can be expected to know its defects. It is often difficult for a prospective buyer to determine how good the car is until after having purchased and driven it for a while. At the time of the transaction, buyers are likely to have less information than sellers about product quality.

ASYMMETRIC INFORMATION
a case in which participants on one side of the market know more about a good's quality than do participants on the other side

Before examining how asymmetric information affects the functioning of used car markets, let's first consider, for purposes of contrast, a market where all parties are fully informed. Suppose that there are only two types of used cars, "good" (high-quality) cars and "bad" (low-quality) cars. Consumers are willing to pay $12,000 for a good car and $6,000 for a bad car. To simplify, assume that the market demand curves are perfectly elastic at these

[8]The model explained in this section is based on George A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84 No. 3 (August 1970), pp. 488–500.

prices. In this situation, there will be two markets, one for good cars and one for bad cars, and the prices will be $12,000 and $6,000, respectively. Assume also that sales of each type of car are 50,000. These markets will function as the competitive model predicts.

Now assume that buyers cannot distinguish between good and bad cars, but the sellers know the difference. How will this affect the market for used cars? Consider the buyers. Only one price will now prevail in the used car market because buyers cannot distinguish between good and bad cars at the time of purchase. Buyers do know, however, that the car may turn out to be either good or bad. How much the buyer is willing to pay for a used car will then depend on how likely it is that the car will be good or bad. To contrast this market with the full-information outcome given previously, let's initially suppose that consumers believe there is a 50 percent chance that the car they purchase will be a good one and a 50 percent chance it will be a bad one (because 50 percent of the cars sold in the full-information model were of each type). Then we expect the typical buyer to be willing to pay about $9,000 for a used car—because this is the average value of a used car when half turn out to be good (worth $12,000) and half turn out to be bad (worth $6,000).

From this discussion, we might expect that the used car market would now be in equilibrium at a single price of $9,000. But that is premature. We have not considered the response of the used car sellers. The sellers know the qualities of their cars. When sellers of high-quality used cars confronted a price of $12,000 (in the full-information model), we assumed they would choose to sell 50,000 units. Now these sellers will get only $9,000, and so we expect them to offer fewer units for sale. With an upward-sloping supply curve, fewer good cars will be sold when the price is lower. Similarly, if 50,000 bad cars would be sold at a price of $6,000, when the price is $9,000 we expect more owners of lemons to unload their cars. Suppose that at a price of $9,000, sales of good cars would be 25,000 and sales of bad cars would be 75,000.

If the price were $9,000, we would expect consumers to become aware that it is more likely they will get a low-quality than a high-quality car in this market. That will affect their willingness to pay. If consumers perceive correctly that three-fourths of the time they get a bad car, they would be willing to pay only $7,500 for a used car (a weighted average of the $12,000 and $6,000 values of good and bad cars, with a weight of three-fourths on the $6,000 figure). But if the price is $7,500, this will decrease the quantities of both types of used cars. However, it is likely to increase the *proportion* of bad used cars in the market, leading to a further revaluation downward in the price consumers will pay.

Exactly where this process ends—where equilibrium will be reached—depends on the supply elasticities of good and bad used cars. It is possible that the process continues until only low-quality used cars are sold at a price of $6,000. It is also possible that there will be an equilibrium in which both good and bad cars are sold. Which outcome results is not as important as recognizing that *the proportion of used cars sold that are high quality will be lower than when consumers know the qualities before making the purchases*. Low-quality products tend to drive out high-quality products when there is asymmetric information.

Our use of used cars as an example should not be interpreted to mean that the analysis applies only to used products or that it applies to all used products. This analysis may apply in any market in which consumers have difficulty determining product quality. This is often true of the markets for technologically sophisticated products like personal computers, cellular phones or VCRs. It is also true of some services: Did you know the quality of instruction at your college when you enrolled? The problem also often arises in purchasing the services of plumbers, carpenters, doctors, and dentists, to give only a few examples.

Market Responses to Asymmetric Information

Do low-quality products really drive out high-quality products? Certainly this is not always the case and may never be the case. The preceding analysis is intended to draw out the im-

plications of the assumption that consumers have no way of judging quality. Consumers often do have ways of distinguishing low- from high-quality products. Before discussing a few of them, let's consider *why* the preceding analysis is incomplete. It is incomplete because it suggests there are substantial mutual gains to be realized if consumers seeking high-quality products can be paired with sellers of high-quality products. (In our example, sellers of good used cars would be better off selling them at, say, $10,000 than not selling any, and buyers would also be better off.) This means that people have the incentive to acquire (or disseminate) information that allows consumers to know when they are getting a high-quality product. Of course, information itself is a scarce good, and there are costs to acquiring and disseminating information. That is one reason why people are not fully informed: the benefits from acquiring information about product quality will not always be worth its costs.

The way in which information is acquired and used by consumers depends on many factors, including the nature of the product and its price, and therefore will differ from market to market. For low-priced products that are frequently purchased (for example, ballpoint pens), personal experience may be the most economical source of information. When goods are higher priced and are purchased infrequently (automobiles, stereos, and so on), it becomes more important to not be stuck with a lemon, and consumers take some care before making a purchase. In the case of a used car, for example, they are likely to want a test drive and may take it to a mechanic whose quality they already know. For many products, consumers can consult publications like *Consumer Reports* for disinterested evaluations of product quality. Similarly, prospective college students may consult several of the available college guides before choosing a college. The opinions of people you know and trust can also provide useful information.

In many cases, it is more efficient for sellers to take the initiative in providing information about product quality. There are several ways to accomplish this goal. For example, sellers of high-quality products may offer guarantees or warranties. This step communicates to consumers that the products are high quality, and firms are willing to incur the costs of the guarantees because they can charge more when consumers believe they are getting a high-quality product. A key issue here is whether the information provided by the seller is believable. A guarantee is more believable if offered by an established firm than if offered by a stranger peddling "gold" watches on a street corner. In a similar fashion, many firms take actions to develop a reputation, or a brand name, for selling high-quality products. The brand name can provide reliable information about the quality of a firm's products. Finally, liability laws give firms an incentive to avoid at least the most serious quality defects because firms can be bankrupted by suits from consumers.

Do these various ways of acquiring and disseminating information imply that the lemons model tells us nothing about real-world markets? It would be a mistake, we think, to reach that conclusion. The provision and acquisition of information has its costs, and for products where these costs are high we would expect the lemons model to provide some insight. There are, for example, still many purveyors of "effortless weight loss" programs and wrinkle-removal creams, even though their "products" have been discredited (to the satisfaction of most people who have investigated them—but investigation has costs). Even where the cost of product quality information is relatively low and markets function well, it is worthwhile to understand that the markets do function differently than they would if consumers were fully informed without cost.

A final point is important. When a real-world market functions differently than it would if consumers were fully informed, it does not mean that the market is necessarily inefficient. Informing consumers is costly, and that usually means that it is efficient for consumers to be something less than fully informed. That is, the costs of informing consumers may be greater than the benefits produced.

APPLICATION 14.5 IS THERE A LEMONS PROBLEM IN USED CAR MARKETS?

For reasons we have explained, it is doubtful that bad cars will totally drive out good cars in used car markets. However, it is possible that asymmetric information has an effect on the way markets for used cars operate. One study examined this issue through the use of survey data.[9] The author reasoned that to the extent the lemons problem was relevant, the quality of used cars should vary by type of seller. Cars purchased from used car dealers (who may provide warranties and have reputations at stake) and from friends or relatives should be of higher quality than those purchased from unknown individuals through a newspaper ad.

Quality is difficult to measure, of course; three different measures were designed for this study. One was based simply on a buyer's own evaluations of a car's mechanical condition using a 10-point scale, with 1 being a lemon and 10 being a "gem." On this scale, the average used car's condition was rated at 6.65. After controlling for various factors (such as age and mileage), the researcher found that for cars between 1 and 7 years old there were few differences among the various types of sellers. This runs contrary to the lemons model. On the other hand, for older cars (8 to 15 years), cars purchased from dealers and friends were rated higher than those purchased through a newspaper ad from a stranger. Cars purchased from a used car dealer, for example, were rated 0.91 points higher than those purchased through an ad.

This study suggests that asymmetric information about the quality of used cars has no effect for cars less than 8 years old, and a limited effect for older cars. Apparently, consumers do obtain enough information through the sorts of channels we discussed to avoid the extreme outcome predicted by the lemons model.

[9]James M. Lacko, *Product Quality and Information in the Used Car Market*, Bureau of Economics Staff Report, Federal Trade Commission (Washington, D.C.: U.S. Government Printing Office, 1986).

14.5 ADVERSE SELECTION AND MORAL HAZARD

Our analysis of asymmetric information has emphasized markets in which consumers have less information than sellers. In some important instances, for example, insurance markets, it is the firms that are less well informed.[10] Our discussion in Chapter 5 of insurance markets was based on the assumption of full information (that is, both firms and consumers know the risks). Now let's see how insurance markets may function when this assumption is modified.

Adverse Selection

The most important information affecting the operation of insurance markets is the probability that the insured-against event will occur. Insurance companies have amassed statistical data that enable them to estimate these probabilities. They may find, for example, that one out of a thousand houses they insure burns down each year. What they often don't know is how the probability varies from one homeowner to another. Thus, it is quite possible that some homeowners have better information than the insurance companies. In the case of a potential arsonist, this is certain to be the case, but even in less extreme circumstances (people who store flammables alongside electrical wiring in their attics) consumers may know whether their risks are much higher than average. Thus, it is quite possible that

[10]Another example is labor markets: when a firm hires workers, it is less well informed about the quality of the workers' labor services than the workers.

ADVERSE SELECTION
a situation in which asymmetric information causes higher-risk customers to be more likely to purchase or sellers to be more likely to supply low-quality goods

at least some consumers have better information than the firms do. This asymmetric information can have profound effects on the operation of the market.

Assume that some homeowners are much more at risk of suffering fire damage to their homes, and that these homeowners know it. Insurance companies know the average risks based on their experience, and they have to charge premiums based on that. What consumers will find this an attractive deal? Clearly, the high-risk homeowners will find the price attractive; think of how much insurance an arsonist would buy when $100 in coverage can be purchased for $1, where this fee reflects the average risk. On the other hand, the insurance is much less attractive to low-risk homeowners. So the insurance companies will find most of their customers coming from the high-risk group; they get an **adverse selection** from the pool of potential customers. The "undesirable" customers, the high-risk homeowners, are more likely to appear in the market (this is the adverse selection), and the insurance companies cannot distinguish high-risk from low-risk homeowners when they sell policies.

Imagine where this process can lead. As mostly high-risk persons buy insurance, the insurance company finds that it has to pay off on a larger share of policies than initially predicted. The average riskiness of its customers is thus higher than for the population as a whole, and this causes the price of insurance to rise. The higher cost of insurance drives away more low-risk customers and further raises the cost of insurance. In the end, it is possible that only high-risk customers are served, and low-risk customers must go without insurance that could potentially benefit them if it reflected their true risk status. (In a full-information world, high- and low-risk customers would simply be charged different prices, reflecting the difference in risk.) This analysis should sound familiar, for it is essentially the lemons problem in a new setting. Adverse selection was the driving force in our used car example, also. There, sellers of low-quality used cars were adversely selected because buyers could not distinguish between good and bad cars. Here, high-risk customers are adversely selected because firms cannot distinguish between high- and low-risk customers.

The adverse selection problem may be important in many insurance markets. In life and medical insurance markets, customers often have a better idea of their risk status (from their own medical or family history) than do insurance companies. Similarly, automobile drivers who speed or drink and drive are more likely to have accidents, and they know more about their driving behavior than an insurance company does when providing them coverage. Doctors who purchase medical malpractice insurance may also be better informed of their lawsuit risks than their insurance companies. In all these markets, the adverse selection problem possibly leads to a situation where mostly high-risk customers are insured and many low-risk customers choose to remain uninsured.

Of course, the outcome is not likely to be as dire as the analysis so far suggests, and the reason is the same as in the lemons model: there are potential gains to market participants from adjusting their behavior to account for the adverse selection problem. For example, homeowners' insurance covers only the market value of structures and contents. By placing an upper limit on the potential losses, insurance firms reduce the costs imposed by high-risk customers, and this lowers the cost of insurance. (Imagine if potential arsonists could insure the family photo album for its "sentimental value" of $100,000.)

Other insurance company practices also make more sense when the adverse selection problem is understood. For health and life insurance policies, companies often require physical exams (to help distinguish high- from low-risk people) and a waiting period before a policy is in force (some maladies may not be apparent in a physical, even though the consumer is aware of them). In some cases, insurance companies use indirect measures to help identify the riskiness of customers. For instance, men aged 15 to 24 have car accidents with about twice the frequency of women the same age, so gender can be used as an indicator of riskiness. Similarly, women live longer than men and so are at lower risk of dying at any age; hence, life insurance is less costly for them.

Finally, group health plans have been developed partly in response to the adverse selection problem. These plans offer policies covering all of a firm's employees. Because all employees must be enrolled, the likelihood that high-risk people will be overrepresented is smaller; the insured workers are more likely to be representative of the average population. Adverse selection can still occur (high-risk workers choose to work for firms with extensive coverage), but that is less likely to be a problem than in the sale of individual policies, so group plans can provide insurance coverage at lower costs. (There are also some other reasons for the lower cost of group plans, such as lower administrative costs.)

APPLICATION 14.6 ADVERSE SELECTION AND THE AMERICAN RED CROSS

Fewer than 4 percent of all blood donors are paid for their donation. A key reason why so-called *commercial blood* is not more common is adverse selection.[11] Critics of for-profit giving argue that when people are motivated to donate blood for the sake of a financial reward, blood banks are more likely to attract donors who are desperate for the money because they are addicted to drugs or alcohol or have a serious infectious disease. Indeed, numerous studies have found that hepatitis, a disease that can be transmitted through blood transfusion and that inflames the liver and can occasionally be fatal, is much more likely to be present in commercially collected blood than in blood that is donated on a nonprofit basis.

[11]This application is based on: Michael L. Katz and Harvey S. Rosen, *Microeconomics* (Boston: Irwin/McGraw Hill, 1998); and Alvin W. Drake, Stan N. Finkelstein, and Harvey M. Sapolsky, *The American Blood Supply* (Cambridge, MA: MIT Press, 1982).

Moral Hazard

MORAL HAZARD
a situation that occurs when, as a result of having insurance, an individual becomes more likely to engage in risky behavior

Moral hazard is another problem endemic to insurance markets. It occurs when, as a result of having insurance, an individual's behavior changes in such a way that the probability of the unfavorable outcome increases or its cost is greater when it does occur. An uninsured homeowner, for example, would take great precautions against fire. The homeowner might, for example, avoid the use of kerosene heaters, install smoke detectors, prohibit smoking, and have an electrician inspect the wiring frequently. The incentives are altered when the costs of replacing the house are covered by insurance. The benefits from avoiding a fire are now smaller, and so the homeowner is likely to devote fewer resources to that use. What this behavior means is that fires become more likely when the parties are insured, and this leads to higher prices of insurance coverage.

The extent of the moral hazard problem is likely to vary across individuals and types of insurable events. People with car insurance may be more careless about locking their cars or may park them in riskier neighborhoods. But does health insurance coverage lead people to exercise less, eat poorly, and smoke more? One reason the problem may be less severe in this instance is that health insurance does not cover all the costs of illness. Although it may cover all the medical costs, for many illnesses the costs to the patient of the pain, suffering, and disfigurement can be as much or greater. This gives people an incentive to take actions to avoid the illness even when the medical costs are covered. But the moral hazard problem is an important reason why you can't get insurance to compensate you fully for *all* (medical and other) costs you bear from illness: that sort of insurance would give you no reason for taking care of yourself.

The moral hazard problem arises when insurance companies lack knowledge of the actions people take that may affect the occurrence of unfavorable events. If the actions are observable, the policies can be made contingent on performance of those actions. This can work to the advantage of insurance companies as well as insured parties who receive more favorable rates. For instance, homes with security systems and smoke detectors may receive coverage at lower rates, giving homeowners incentives to take actions that lower the probability of theft and fire. Similarly, smokers are generally charged higher health insurance premiums than nonsmokers, making smokers bear the higher cost of medical care that results from their behavior, and at the same time giving them an incentive to quit smoking. "Good driver" policies reward automobile owners with unblemished records through lower rates.

Another type of moral hazard problem is particularly significant to health insurance. Consider a policy that covers all hospital costs. When a person is hospitalized, he or she will, in effect, face a zero price for treatment. This is likely to lead the patient to consume hospital services beyond the point where they are worth what they cost. Any form of medical care that has any benefit, no matter how small, will seem worthwhile if the insurance company is absorbing the expense. In this setting, doctors are likely to prescribe expensive tests and sophisticated treatments, knowing that financial responsibility falls to a faceless third party. But when patients and doctors behave this way, the cost of hospitalization goes up and insurance premiums rise to cover the cost.

A number of practices have evolved in the medical insurance market to deal with moral hazard. One is limitations on the services covered by insurance (no more than three days in the hospital for an appendectomy, for example). Although patients may believe this reflects stinginess on the insurer's part, it actually leads to lower insurance premiums and can be in the long-run interest of the insured. A second approach is to require the insured person to pay part of the costs. Patients might be required to pay 20 percent of the hospital bill; the share of the cost borne by the patient is called the *coinsurance rate*. Use of coinsurance reduces the cost of insurance directly, but also indirectly lowers it by giving patients incentive to be more economical in their use of hospital services. A third approach is the use of *deductibles*. Using a deductible means that a patient must pay, for example, the first $500 of hospital costs before insurance coverage is effective. This gives patients incentive to take account of all costs in the case of minor medical treatments, and also leads to lower insurance premiums.

Insurance markets are profoundly affected by asymmetric information, as this discussion indicates. There is no doubt that these markets function very differently from the way they would if all parties had perfect information. Understanding the problems created by asymmetric information helps us see why certain practices have emerged to mitigate them.

APPLICATION 14.7 MORAL HAZARD IN THE S&L INDUSTRY

The federal government's bailout of savings and loan institutions (S&Ls) ended up costing taxpayers over $150 billion.[12] While the bailout's monetary consequences were sizable, the cause of the widespread bankruptcy in the S&L industry in the late 1980s is straightforward: moral hazard. During the late 1970s and early 1980s the availability of insurance for deposits at S&Ls was expanded by the government through institutions such as the Federal Deposit Insurance Corporation and

[12]Edward J. Kane, *The S&L Insurance Mess* (Washington, D.C.: The Urban Institute Press, 1989).

the Federal Savings and Loan Insurance Corporation. The expanded availability of insurance, coupled with relaxations in other regulations concerning S&Ls, implied that the government would insure virtually all accounts regardless of their size and no matter how risky. Predictably enough, S&Ls increased the riskiness of their loan portfolios during the early 1980s. The increased riskiness of the loan portfolios came back to haunt the S&Ls and ultimately the federal government when the economy slid into recession in the late 1980s.

APPLICATION 14.8 MORAL HAZARD ON THE ROAD AND AT THE PLATE

Moral hazard is not limited to insurance markets. Consider the case of four-wheel-drive motor vehicles. Contrary to the claims of car makers and dealers about the safety of four-wheel drive vehicles, the drivers of such vehicles appear to be more likely to have an accident (holding other factors constant including the recent well-publicized failures of Bridgestone/Firestone tires on Ford sport utility vehicles).[13] The problem appears to be the behavior of drivers once they are behind the wheel of a four-wheel drive automobile. According to one Denver-based traffic investigator: "These people have it in their heads that they're driving Sherman tanks. As soon as snow hits the ground, they go speeding up the highway, and then two minutes later they're rolled over or skidding like crazy." A computer technician from Boston who recently wrecked his four-wheel-drive car notes, "When you're in a four-wheel drive, there is this sense of power and overconfidence that makes you forget everything else. You really don't realize your limitations until you hit something."

Another noninsurance example of moral hazard is the 1973 introduction of the Designated Hitter (DH) rule in the American League but not the National League of Major League Baseball. The DH rule relieved American League pitchers of their responsibility to appear at the plate as a hitter (another player could be designated to hit in their place). Consequently, American League pitchers can throw at opposing hitters with greater impunity (lower cost) than National League pitchers, who must still take their turn at bat. American League pitchers are less likely to face direct retaliation (since they do not have to appear at the plate themselves) if they hit or come close to hitting an opposing batsman with a pitch. The result is as one would predict. Since the introduction of the DH rule, American League pitchers have been hitting 10 to 15 percent more batters with their pitches than National League pitchers (holding all other factors constant).

[13]This application is based on "Why Four-Wheel Drive Isn't Always Safer in Snow," *Wall Street Journal*, February 22, 1995, pp. B1 and B6; and Brian L. Goff, William F. Shughart II, and Robert D. Tollison, "Moral Hazard and the Effects of the Designated Hitter Rule Revisited," *Economic Inquiry*, 36 No. 4 (October 1998), pp. 688–692.

14.6 LIMITED PRICE INFORMATION

Consumers may lack information about product prices as well as product qualities. In the competitive model, where consumers are fully informed, if one firm raises its price above the competitive level (the price charged by the other firms) it will lose all its sales because its customers know that the other firms charge a lower price and so will purchase from them. When consumers have perfect information about prices, all firms have to charge the same price and thus each firm faces a horizontal demand curve.

By contrast, if consumers do not know the prices at stores other than the one where they are currently shopping, retailers can raise prices without losing all customers. Consumers' lack of information about competitors' prices means that each store confronts a

downward-sloping demand curve, which, in turn, gives it some market power. Exactly how such a market functions depends on a number of factors, including the extent of consumer ignorance and the cost of acquiring price information. In general, however, we expect that the firms in markets where many or all consumers are uninformed will be charging different prices. There is likely to be a range of prices or **price dispersion** for the same product.

PRICE DISPERSION
a range of prices for the same product, usually as a result of customers' lacking price information

SEARCH COSTS
the costs that customers incur in acquiring information

Now consider consumer behavior in a market where there is price dispersion. Consumers wish to purchase from the firm offering the product for the lowest price, but they don't know which firm it is. They can find out, but there is a cost of acquiring the information. **Search costs** are the costs that consumers incur in acquiring information; they include such things as time (making telephone calls, buying newspapers, reading the ads) and transportation between stores. There is also a benefit from acquiring price information, of course, because consumers can buy the product for a lower price. However, consumers are unlikely to search until they are fully informed about all the prices being charged by various stores. The reason for this is that the expected marginal benefit from additional search declines the longer the search goes on.

A simple example illustrates this point. Suppose 20 stores sell a compact disc you want to purchase. You call one store and find its price is $12. Should you call a second store? If all stores charge different prices, the probability that the second store will charge a lower price is about one-half. So you call a second store and find its price is $13; you lost the coin flip. Now the probability that a third telephone call will generate a price lower than either of these is only about one-third. Even after calling 19 stores, there is a slight probability that the lowest-price store is the one you haven't called, but you would be unlikely to take the time to call all 20 stores just to guarantee you find the lowest price. In general, because the expected marginal benefit of additional search declines with the amount of search, you will not keep placing calls until you become fully informed. You may stop after three or four telephone calls and purchase from the lowest-priced source you have located, although that will often not be the lowest price available in the market.

For a given consumer, the amount of search undertaken will have no effect on the actual price dispersion in the market. But if many or all consumers increase search intensity, becoming better informed about price, the price dispersion will be reduced. As consumers become better informed, high-priced firms lose business relative to low-priced firms, and high-priced firms are forced to reduce their prices. Taken to the limit, if consumers become fully informed, only one price can prevail.

How does this help us understand the amount of price dispersion for a given product? The theory predicts that the dispersion falls when consumers search more (that is, become better informed). They will search more when the benefit from search is higher than the cost. A little thought will convince you that the benefit will be higher the greater the product's price. Finding a store that sells a compact disc for 1 percent less than another may not be worth an extra phone call, but saving 1 percent on the price of a car is likely to be worth it. Thus, it is not surprising that empirical studies have found less relative price dispersion (price variation compared with the average price) for higher-priced products. Prices are less widely dispersed for cars of the same make and model, for example, than for washing machines.[14]

Finally, we should not ignore the possibility that when consumers are not fully informed and there is price dispersion in a market, it may be in the interest of the low-price firms to inform consumers about price. This is one reason why firms advertise, which leads us to the next section.

[14]George J. Stigler, "The Economics of Information," *Journal of Political Economy*, 69 No. 2 (June 1961), pp. 213–225.

APPLICATION 14.9 THE INTERNET AND THE PRICE OF LIFE INSURANCE

Economists Jeffrey Brown and Austan Goolsbee analyzed the effect of the Internet on the price of life insurance.[15] By examining individuals' life insurance policies and controlling for other factors, the researchers found that increases in Internet use during the 1990s reduced the average price of life insurance policies by 8 to 15 percent. This decrease appears to reflect the fact that the Internet makes the life insurance market more competitive by enabling price comparisons online and thereby reducing search costs. The results of the study also indicate that price dispersion first increased with the initial, limited introduction of Internet search sites and then diminished as more and more prospective insurance customers acquired access to the Internet.

[15]Jeffrey R. Brown and Austan Goolsbee, "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, 110 No. 3 (2002), pp. 481–507.

14.7 ADVERTISING

Firms advertise in an attempt to increase the demand for their products. On that point economists agree. But there is some disagreement over whether advertising expenditures serve the useful function of providing information to consumers or the baleful function of wasting resources and distorting consumption choices.

Economists have a long tradition of skepticism regarding the benefits of advertising. In this view, competitive firms have no need to advertise because they can sell as much as they want at the market price (although this point is disputed), so the very existence of advertising implies that firms have some monopoly power. But advertising itself may not only be a symptom of market power—it can also help firms achieve and maintain market power. Firms may use Madison Avenue's tools of persuasion to convince consumers that their products are different from and better than those of competitors. This is sometimes referred to as **artificial product differentiation** and, if successful, increases the demand for the product and also makes demand more inelastic, conferring additional market power on the firm. For example, the chemical composition of Bayer aspirin is nearly indistinguishable from that of generic aspirin, but many consumers think Bayer is better and pay a substantially higher price for that brand.

ARTIFICIAL PRODUCT DIFFERENTIATION the use of advertising to differentiate products that are essentially the same

It is also charged that advertising can operate as a barrier to entry. If a new firm attempts to enter a profitable industry, it may find that advertising by the established firms has created a captive audience of consumers who are reluctant to try a new brand. It may be necessary for an entrant to wage a massive advertising campaign to get consumers to give its product a try, and the prospect of that cost could be an effective deterrent. On the other hand, advertising can also be a means of breaking into an entrenched market, giving an entrant a way to increase sales quickly so that economies of scale can be realized.

In its most extreme form, criticism of advertising holds that it manipulates consumers and leads them to choose products they don't want or need. In this view, consumers' tastes are not formed independently, but are actually created by advertisers. Thus, instead of consumers' underlying wants being the factor behind the position of demand curves (and thereby the pattern of production), producers are held to play the central role of determining consumers' wants through their promotional activities.

Although few economists hold this extreme view on advertising today, there continues to be a belief that advertising may enhance market power, deter entry, and lead to more

concentrated industries. Numerous studies have tried to find a connection between advertising and industrial concentration or profitability. By and large, their results have been inconclusive: about as many find advertising to be unassociated with industrial concentration or profitability as the reverse. This evidence has dispelled the worst fears of the critics of advertising, but there remains the possibility that advertising has harmful effects that are simply not large enough to be empirically identified.

Advertising as Information

A view that advertising is benign in its effects has more recently emerged as an outgrowth of research on the economics of information. Advertising is held to be a low-cost way of providing information to consumers about the availability of products and their prices and qualities. It may make markets more competitive and even lead to lower prices for consumers.

To see why, imagine an industry where there is no advertising. As we explained in the previous section, firms will then face downward-sloping demand curves because it is costly for consumers to find out about alternatives. In Figure 14.1, D_1 is a typical firm's demand curve when no firm in the industry advertises. Now if this firm *alone* advertises, its demand curve shifts outward to D_2 and may become more inelastic, as drawn. This appears to support the negative view of advertising—that market power is enhanced as demands become more inelastic. However, we must recognize that other firms are also free to advertise, and the effects of all firms advertising may be quite different from the effects when only one does. So, suppose that our typical firm continues to advertise, but now its competitors also try to persuade consumers of the virtues of their products. This will cause this firm's demand curve to shift leftward from D_2 and become more elastic, as shown by D_3, the demand curve when all firms are advertising. It is more elastic than demand when no firms advertise (D_1) because consumers will be more aware of the alternatives available on the market, and an increase in this firm's price will cause more consumers to shift to other products than when the consumers are not aware of the alternatives. Thus, it is possible that advertising, when undertaken by many or all competing firms, actually confronts firms with more elastic demand curves and reduces their market power.

FIGURE 14.1

Advertising and the Firm's Demand Curve
When there is no advertising in an oligopoly, a firm's demand curve is D_1. If only this firm advertises, its demand curve shifts out to D_2. If all firms advertise, the firm's demand curve becomes D_3.

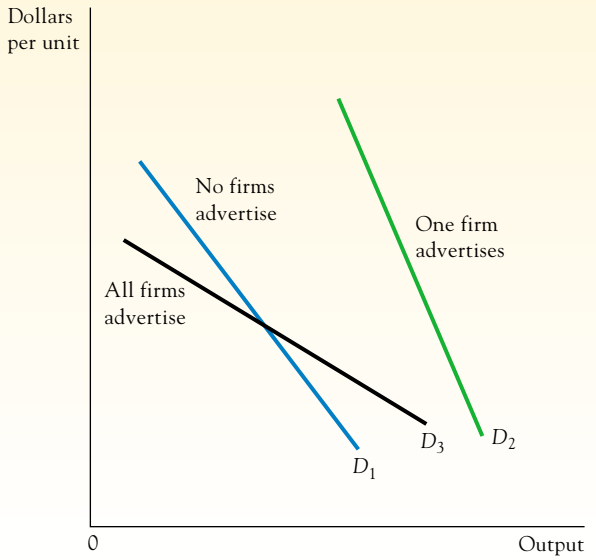

In the previous section, we explained how price dispersion arises in a market when consumers are uninformed. The existence of price dispersion gives consumers an incentive to incur search costs to obtain information. It also, however, gives firms an incentive to advertise. For example, if some firms charge a higher price to uninformed consumers and if firms charging lower prices could inform consumers of this fact, the low-price firms can enhance the demand for their products and potentially their profits. Moreover, if consumers switch from high-price to low-price firms, it will reduce the average price in the market. It could also lead the high-price firms to reduce their prices, which reduces the price dispersion and further lowers the average price. In this way, advertising can lead to lower prices.

It is also possible for advertising to solve the lemons problem discussed in Section 14.4. Recall that in this model, consumers cannot determine the quality of the product before purchase, and as a result low-quality products drive out high-quality ones. Through advertising, high-quality sellers can inform consumers that their products are of high quality. The problem arises in deciding whether such claims are true. Won't the low-quality sellers make the same claims? Possibly so, but surprisingly, high-quality sellers will normally have a greater incentive to advertise. Such is the importance of repeat purchases and word-of-mouth endorsements to a firm's future sales. A firm selling a lemon can at best convince a consumer to purchase one time, but it can be sure there will be no future sales to that customer, and it may lose future sales from other consumers that hear of this buyer's bad experience. By contrast, a firm selling a high-quality product can gain future sales in addition to the current sale, so it has a greater incentive to promote its product.

One interesting aspect of this analysis is that it suggests consumers may gain useful information from advertising even when that advertising does nothing more than assert a product is "great tasting" (that is, of high quality). Higher-quality (or higher-value) products will tend to be more heavily advertised, and if consumers are influenced by the amount of advertising they may be led to try the better products.

There is some evidence that advertising does work to the benefit of consumers, as this analysis suggests. For example, according to one study, the average price of eyeglasses in states permitting advertising is about 25 percent lower than the prices in states where such advertising is prohibited.[16] Eyeglass prices are also lower in states prohibiting advertisers from mentioning price, but permitting claims of high quality (prices are lowest in states where prices can be advertised). Thus, even when advertising does not mention prices, it can apparently contribute to lower prices for consumers. Other studies have also produced evidence that advertising about price lowers the average price consumers pay for products such as liquor, drugs, contact lenses, toys, and gasoline.

Even when advertising does not lead to lower prices, it may be advantageous for consumers. In the absence of advertising, consumers have to incur search costs to find out about products. The true prices they pay are then the sum of the money price and the search costs they bear; this is sometimes referred to as the **full price** of the product. One effect of advertising is that it is a substitute for the consumer's own search efforts. Thus, advertising can reduce consumers' search costs. Even if the money price is unchanged or rises somewhat, the full price consumers pay may fall as a result of advertising. For example, even if they do not lead to a decrease in the money price of products, advertisements on the Internet promise to significantly decrease the full price consumers pay for products by lowering consumers' search costs.

FULL PRICE
the sum of the money price and the search costs that consumers incur

The advertising-as-information view suggests that advertising is a low-cost way of conveying useful information to consumers about alternative products and their prices, and thus makes markets work more efficiently. Not all economists accept this positive view,

[16]Lee Benham, "The Effect of Advertising on the Price of Eyeglasses," *Journal of Law and Economics*, 15 No. 2 (October 1972), pp. 45–74.

and television advertising has been singled out for special criticism. (Significantly, however, only 22 percent of all advertising expenditures are devoted to television; newspapers account for a larger share.) This is not surprising, because television advertising is more intrusive than most other forms and is more difficult to target to interested consumers than, say, advertising in specialized magazines. Moreover, there is the lingering suspicion that some consumers (of course, not us) may be swayed by the emphasis on visual images and emotional appeals.

In sum, probably each view of advertising contains an element of truth. It may be that advertising works well in certain types of markets and for certain types of products, but has more deleterious effects in other cases.

APPLICATION 14.10 A NEWSPAPER STRIKE'S EFFECT ON FOOD PRICES

In the fall of 1979, a press operators' strike temporarily halted publication of New York City's three major dailies, the *Times*, *Post*, and *Daily News*. *Newsday*, in neighboring Long Island, continued publishing—and thus carrying local supermarkets' food advertisements. A study examining the prices of various food items in the New York City area during the first week of the strike found that prices increased by 3.4 percent less in supermarkets on Long Island than in other New York City boroughs experiencing reduced access to food advertisements.[17] Significantly, the difference in price changes was observed for supermarkets, which advertise in newspapers, but not for small fruit and vegetable stores, which do not. Thus, for supermarket food items in New York City, newspaper advertising appears to promote lower prices.

[17]Amihai Glazer, "Advertising, Information, and Prices—A Case Study," *Economic Inquiry*, 19 No. 3 (October 1981), pp. 661–671.

SUMMARY

- Game theory is a method of analyzing situations involving strategic interactions among decisionmakers, a setting characteristic of oligopolistic markets.
- A payoff matrix clarifies the nature of the problem confronting decisionmakers and helps us identify the equilibrium.
- Two types of equilibrium are dominant strategy and Nash.
- A particularly important type of game is the prisoner's dilemma, in which self-interest on the part of each player can lead to a result in which all players are worse off then they could be if different choices were made.
- The prisoner's dilemma game helps us see why firms have an incentive to cheat on a cartel agreement. When the game is repeated, however, the analysis becomes more complicated, and there is a greater possibility of collusion.
- Markets can work very differently when consumers or firms are not fully informed about prices and/or product characteristics. When consumers cannot determine the quality of a product before purchase, the lemons model suggests that low-quality products will predominate.
- Market forces limit the extent of the lemons problem.
- In the case of insurance, firms may have less information than do consumers. Adverse selection can then lead to only high-risk customers being insured.
- When consumers are not informed about the prices charged by all firms, it is possible for more than one price to prevail in the market. It is not necessary, however, for all customers to be fully informed for a single price to emerge. In general, the larger the proportion of informed consumers, the less the price dispersion in the market.
- Advertising is a particularly important way in which information is provided to consumers by firms. It can lead markets to operate more efficiently, but there is also the possibility that it can distort consumer choices and make it difficult for new firms to enter profitable industries.

REVIEW QUESTIONS AND PROBLEMS

Questions and problems marked with an asterisk have solutions given in Answers to Selected Problems at the back of the book (page 577).

14.1. What is a dominant-strategy equilibrium? What is a Nash equilibrium? Is it possible for a Nash equilibrium to exist where neither player has a dominant strategy?

***14.2.** Construct and explain the payoff matrix for two firms that operate in a competitive market. How does it differ from the situation illustrated in Table 14.1?

14.3. What is a prisoner's dilemma game? Why is it relevant in evaluating the likelihood of cheating in a cartel?

14.4. Tables 14.1 and 14.4 both involve two firms each choosing between low and high outputs, but only one of the tables illustrates the prisoner's dilemma. Explain why the nature of the market in which firms interact may sometimes produce a prisoner's dilemma and sometimes not.

14.5. Is a repeated- or single-period game more appropriate for the study of oligopolies? In which setting is collusion more likely to be a stable outcome? Explain your answer.

14.6. Construct a payoff matrix to examine the determination of outputs in the Cournot duopoly model. What type of equilibrium exists for this model? Does the game-theoretic approach make this model any more plausible?

14.7. Why do you think that game theory has become the preferred method of analyzing oligopolistic markets? What advantages does it have over simply assuming, say, Cournot behavior?

14.8. What is the basic assumption about information in the lemons model?

14.9. In the lemons model, there is only one price even though the products differ in quality. Why is that? What factors determine that price? How does the price affect the quantities traded of the different quality goods?

14.10. College instructors know more about the quality of their courses than prospective students. Does this mean that the lemons model is appropriate? How does this market differ from the one assumed in the pure lemons model?

14.11. What is adverse selection in insurance markets and how does it relate to the lemons model?

***14.12.** Consider insurance covering the costs of cancer when there is no way to determine how likely it is that any given individual will contract the disease. How will the price of the policy be determined? Now suppose that it is determined that smokers have 10 times the risk of nonsmokers. How will the price be affected if insurance companies cannot determine who smokes and who doesn't?

14.13. How does the moral hazard problem differ from the adverse selection problem in markets for medical insurance?

14.14. Why don't consumers become fully informed about the prices different firms charge? If consumers are not fully informed, why is a firm likely to possess some degree of market power?

14.15. Suppose that a college town has a large number of firms selling a homogeneous product—pizza—and that there are two types of consumers in the town. The town's permanent residents are fully informed about the prices charged by all firms and always shop at the firm or firms with the lowest price. On the other hand, the students attending college in the town (temporary residents) are completely uninformed; they do not know anything about prices and simply choose among firms on a random basis. Explain why, in such a setting, a single price may prevail in the market for pizza.

14.16. How does advertising affect the demand curve confronting a single firm? How does the outcome depend on whether other firms also advertise? If all firms in an industry advertise, how will this shift the industry demand curve for the product?

***14.17.** "Because advertising adds to firms' costs of production, it cannot lead to lower product prices." True or false? Explain.

14.18. According to the English poet Alfred Lord Tennyson, "Tis better to have loved and lost than never to have loved at all." Does Tennyson's observation imply that love is a dominant strategy? Explain why or why not.

14.19. In multidivision corporations where division heads are allocated an annual budget, explain why the "use-it-or-lose-it" phenomenon occurs and is a reflection of a prisoner's dilemma.

14.20. If there is asymmetric information between the owners of a baseball team for which a given player plays and other teams' owners, would you predict that players who opt to become free agents and end up getting traded to another team will spend more days on the disabled list, after being traded, than players who remain with their existing team? Explain why or why not.

14.21. In 1995, President Bill Clinton unveiled sweeping proposals to regulate cigarette advertising. In response, the largest cigarette manufacturer, Philip Morris, whose leading brand, Marlboro, accounted for 30 percent of all cigarette sales in the United States, lambasted the proposals and filed suit against the federal government. Notwithstanding its actions, why do you think some tobacco industry analysts and even Philip Morris executives believed that the proposed regulations would end up benefiting the world's biggest cigarette supplier?

14.22. Imagine that several players are each asked to pick a number from zero to 100. The prize associated with winning the game is awarded to the player coming closest to the number that is half the average of what all the other players select. Viewed from a game-theoretic setting, what number will be selected by the players in a Nash equilibrium? (Economist Hal Varian of Berkeley notes that real people playing such a game

do not typically make such a selection and that this calls into question the assumption that players are always fully rational in game-theoretic settings.)

14.23. As mentioned in Chapter 5 (Application 5.1), economists have proposed retaining the tax-exempt status of health care while lifting the requirement that employees spend all of the tax-free fringe benefits solely on health care. If retirees had a similar option to select such "medisave" accounts versus retain their traditional Medicare coverage, explain why average government payments to retirees opting to retain their traditional Medicare coverage could end up rising.

14.24. In 2002, Allergan, the company making wrinkle-fighter Botox, put together a $50 million advertising campaign to hook the masses on the anti-aging drug. Botox is derived from the neurotoxin that causes botulism, a deadly food poisoning. Injected under the skin, it relaxes facial muscles that cause lines. Botox was approved by the Food and Drug Administration in 1989 to treat crossed eyes and uncontrolled blinking. Word spread, however, in the late 1990s regarding its ability to also control wrinkles. A 15-minute doctor's office treatment costs around $500. Customers have to get repeated treatments since the effects wear off in roughly four months. In light of this, evaluate the statement made by a political activist stating that Botox advertising should be banned by the federal government since it promotes a product with "no socially redeeming value and [which] merely manipulates vulnerable consumers into buying a product with short-lived positive effects that they don't really need."

14.25. A considerable number of initial public offerings (IPOs) of stock in a company evidence a substantial run-up in price during early trading. Explain why asymmetric information between the investment banks organizing the offerings and prospective investors may be the reason for such a phenomenon.