What you will learn in this **Module:**

- How our understanding of oligopoly can be enhanced by using game theory
- The concept of the prisoners' dilemma
- How repeated interactions among oligopolists can result in collusion in the absence of any formal agreement

Module 65 Game Theory

Games Oligopolists Play

In our duopoly example and in real life, each oligopolistic firm realizes both that its profit depends on what its competitor does and that its competitor's profit depends on what it does. That is, the two firms are in a situation of interdependence, whereby each firm's decision significantly affects the profit of the other firm (or firms, in the case of more than two).

In effect, the two firms are playing a "game" in which the profit of each player depends not only on its own actions but on those of the other player (or players). In order to understand more fully how oligopolists behave, economists, along with mathematicians, developed the area of study of such games, known as **game theory.** It has many applications, not just to economics but also to military strategy, politics, and other social sciences.

Let's see how game theory helps us understand oligopoly.

The Prisoners' Dilemma

Game theory deals with any situation in which the reward to any one player—the **payoff**—depends not only on his or her own actions but also on those of other players in the game. In the case of oligopolistic firms, the payoff is simply the firm's profit.

When there are only two players, as in a lysine duopoly, the interdependence between the players can be represented with a **payoff matrix** like that shown in Figure 65.1. Each row corresponds to an action by one player; each column corresponds to an action by the other. For simplicity, let's assume that each firm can pick only one of two alternatives: produce 30 million pounds of lysine or produce 40 million pounds.

The matrix contains four boxes, each divided by a diagonal line. Each box shows the payoff to the two firms that results from a pair of choices; the number below the diagonal shows Firm 1's profits, the number above the diagonal shows Firm 2's profits.

These payoffs show what we concluded from our earlier analysis: the combined profit of the two firms is maximized if they each produce 30 million pounds. Either firm can, however, increase its own profits by producing 40 million pounds if the other produces only 30 million pounds. But if both produce the larger quantity, both will have lower profits than if they had both held their output down.

The study of behavior in situations of interdependence is known as **game theory.**

The reward received by a player in a game, such as the profit earned by an oligopolist, is that player's **payoff.**

A **payoff matrix** shows how the payoff to each of the participants in a two-player game depends on the actions of both. Such a matrix helps us analyze situations of interdependence.

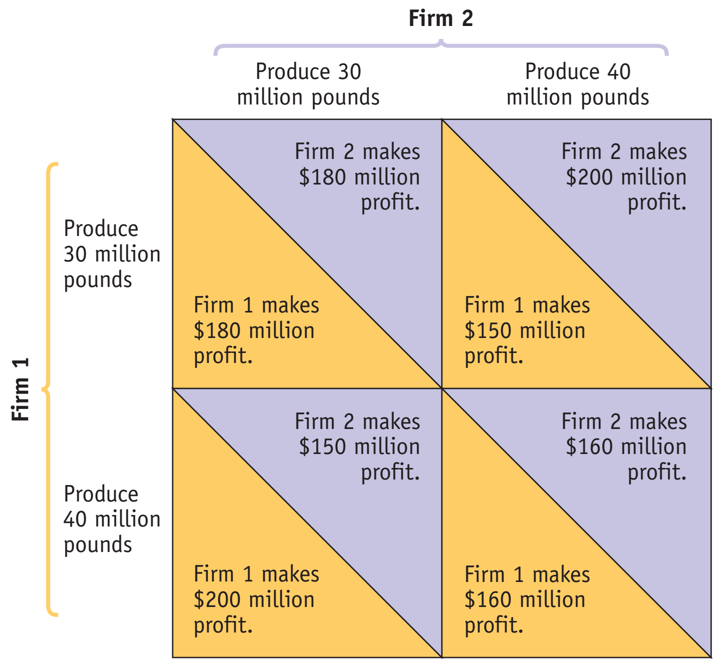

		Firm 2	
		Produce 30 million pounds	Produce 40 million pounds
Firm 1	Produce 30 million pounds	Firm 2 makes $180 million profit. Firm 1 makes $180 million profit.	Firm 2 makes $200 million profit. Firm 1 makes $150 million profit.
	Produce 40 million pounds	Firm 2 makes $150 million profit. Firm 1 makes $200 million profit.	Firm 2 makes $160 million profit. Firm 1 makes $160 million profit.

figure 65.1

A Payoff Matrix

Two firms must decide how much lysine to produce. The profits of the two firms are *interdependent:* each firm's profit depends not only on its own decision but also on the other's decision. Each row represents an action by Firm 1, each column one by Firm 2. Both firms will be better off if they both choose the lower output; but it is in each firm's individual interest to choose the higher output.

The particular situation shown here is a version of a famous—and seemingly paradoxical—case of interdependence that appears in many contexts. Known as the **prisoners' dilemma,** it is a type of game in which the payoff matrix implies the following:

- Each player has an incentive, regardless of what the other player does, to cheat—to take an action that benefits it at the other's expense.
- When both players cheat, both are worse off than they would have been if neither had cheated.

The prisoners' dilemma is a game based on two premises: (1) Each player has an incentive to choose an action that benefits itself at the other player's expense; and (2) When both players act in this way, both are worse off than if they had acted cooperatively.

The original illustration of the prisoners' dilemma occurred in a fictional story about two accomplices in crime—let's call them Thelma and Louise—who have been caught by the police. The police have enough evidence to put them behind bars for 5 years. They also know that the pair have committed a more serious crime, one that carries a 20-year sentence; unfortunately, they don't have enough evidence to convict the women on that charge. To do so, they would need each of the prisoners to implicate the other in the second crime.

So the police put the miscreants in separate cells and say the following to each: "Here's the deal: if neither of you confesses, you know that we'll send you to jail for 5 years. If you confess and implicate your partner, and she doesn't do the same, we reduce your sentence from 5 years to 2. But if your partner confesses and you don't, you'll get the maximum 20 years. And if both of you confess, we'll give you both 15 years."

Figure 65.2 on the next page shows the payoffs that face the prisoners, depending on the decision of each to remain silent or to confess. (Usually the payoff matrix reflects the players' payoffs, and higher payoffs are better than lower payoffs. This case is an exception: a higher number of years in prison is bad, not good!) Let's assume that the prisoners have no way to communicate and that they have not sworn an oath not to harm each other or anything of that sort. So each acts in her own self-interest. What will they do?

The Kobal Collection

The critically acclaimed 1991 movie *Thelma and Louise* was innovative in depicting two female characters running from the law.

figure 65.2

The Prisoners' Dilemma

Each of two prisoners, held in separate cells, is offered a deal by the police—a light sentence if she confesses and implicates her accomplice but her accomplice does not do the same, a heavy sentence if she does not confess but her accomplice does, and so on. It is in the joint interest of both prisoners not to confess; it is in each one's individual interest to confess.

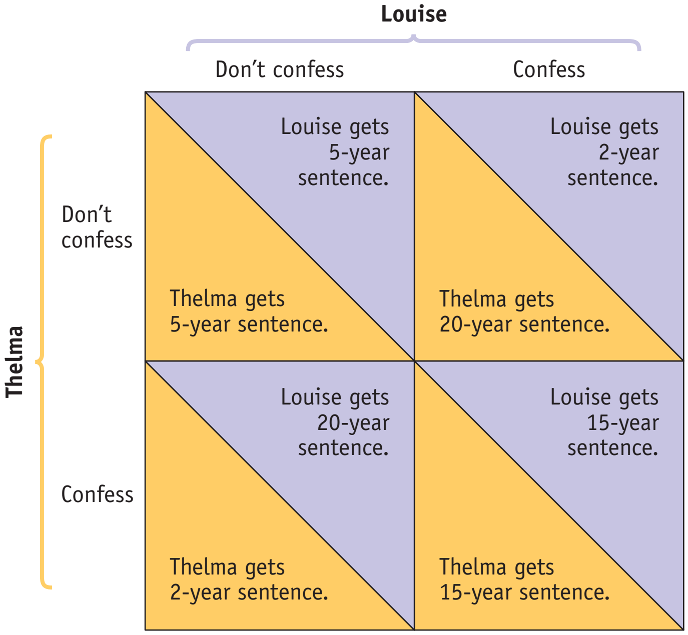

	Louise	
Thelma	Don't confess	Confess
Don't confess	Louise gets 5-year sentence. Thelma gets 5-year sentence.	Louise gets 2-year sentence. Thelma gets 20-year sentence.
Confess	Louise gets 20-year sentence. Thelma gets 2-year sentence.	Louise gets 15-year sentence. Thelma gets 15-year sentence.

An action is a **dominant strategy** when it is a player's best action regardless of the action taken by the other player.

A **Nash equilibrium,** also known as a **noncooperative equilibrium,** is the result when each player in a game chooses the action that maximizes his or her payoff, given the actions of other players.

The answer is clear: both will confess. Look at it first from Thelma's point of view: she is better off confessing, regardless of what Louise does. If Louise doesn't confess, Thelma's confession reduces her own sentence from 5 years to 2. If Louise *does* confess, Thelma's confession reduces her sentence from 20 to 15 years. Either way, it's clearly in Thelma's interest to confess. And because she faces the same incentives, it's clearly in Louise's interest to confess, too. To confess in this situation is a type of action that economists call a *dominant strategy*. An action is a **dominant strategy** when it is the player's best action regardless of the action taken by the other player. It's important to note that not all games have a dominant strategy—it depends on the structure of payoffs in the game. But in the case of Thelma and Louise, it is clearly in the interest of the police to structure the payoffs so that confessing is a dominant strategy for each person. As long as the two prisoners have no way to make an enforceable agreement that neither will confess (something they can't do if they can't communicate, and the police certainly won't allow them to do so because the police want to compel each one to confess), the dominant strategy exists as the best alternative.

So if each prisoner acts rationally in her own interest, both will confess. Yet if neither of them had confessed, both would have received a much lighter sentence! In a prisoners' dilemma, each player has a clear incentive to act in a way that hurts the other player—but when both make that choice, it leaves both of them worse off.

When Thelma and Louise both confess, they reach an *equilibrium* of the game. We have used the concept of equilibrium many times in this book; it is an outcome in which no individual or firm has any incentive to change his or her action. In game theory, this kind of equilibrium, in which each player takes the action that is best for her, given the actions taken by other players, is known as a **Nash equilibrium,** after the mathematician and Nobel Laureate John Nash. (Nash's life was chronicled in the best-selling biography *A Beautiful Mind,* which was made into a movie.) Because the players in a Nash equilibrium do not take into account the effect of their actions on others, this is also known as a **noncooperative equilibrium.**

In the prisoners' dilemma, the Nash equilibrium happens to be an equilibrium of two dominant strategies—a *dominant strategy equilibrium*—but Nash equilibria can exist

Associated Press/PLINIO LEPRI

Mathematician and Nobel Laureate John Forbes Nash proposed one of the key ideas in game theory.

when there is no dominant strategy at all. For example, suppose that after serving time in jail, Thelma and Louise are disheartened by the mutual distrust that led them to confess, and each wants nothing more than to avoid seeing the other. On a Saturday night, they might each have to choose between going to the nightclub and going to the movie theater. Neither has a dominant strategy because the best strategy for each depends on what the other is doing. However, Thelma going to the nightclub and Louise going to the movie theater is a Nash equilibrium because each player takes the action that is best given the action of the other. Thelma going to the movie theater and Louise going to the nightclub is also a Nash equilibrium, because again, neither wants to change her behavior given what the other is doing.

Now look back at Figure 65.1: the two firms face a prisoners' dilemma just like Thelma and Louise did after the crimes. Each firm is better off producing the higher output, regardless of what the other firm does. Yet if both produce 40 million pounds, both are worse off than if they had followed their agreement and produced only 30 million pounds. In both cases, then, the pursuit of individual self-interest—the effort to maximize profits or to minimize jail time—has the perverse effect of hurting both players.

Prisoners' dilemmas appear in many situations. The upcoming FYI describes an example from the days of the Cold War. Clearly, the players in any prisoners' dilemma would be better off if they had some way of enforcing cooperative behavior: if Thelma and Louise had both sworn to a code of silence, or if the two firms had signed an enforceable agreement not to produce more than 30 million pounds of lysine.

But we know that in the United States an agreement setting the output levels of two oligopolists isn't just unenforceable, it's illegal. So it seems that a noncooperative equilibrium is the only possible outcome. Or is it?

Overcoming the Prisoners' Dilemma: Repeated Interaction and Tacit Collusion

Thelma and Louise are playing what is known as a *one-shot* game—they play the game with each other only once. They get to choose once and for all whether to confess or deny, and that's it. However, most of the games that oligopolists play aren't one-shot games; instead, the players expect to play the game repeatedly with the same rivals. An oligopolist usually expects to be in business for many years, and knows that a decision today about whether to cheat is likely to affect the decisions of other firms in the future. So a smart oligopolist doesn't just decide what to do based on the effect on profit in the short run. Instead, it engages in **strategic behavior,** taking into account the effects of its action on the future actions of other players. And under some conditions oligopolists that behave strategically can manage to behave as if they had a formal agreement to collude.

Suppose that our two firms expect to be in the lysine business for many years and therefore expect to play the game of cheat versus collude shown in Figure 65.1 many times. Would they really betray each other time and again?

Probably not. Suppose that each firm considers two strategies. In one strategy it always cheats, producing 40 million pounds of lysine each year, regardless of what the other firm does. In the other strategy, it starts with good behavior, producing only 30 million pounds in the first year, and watches to see what its rival does. If the other firm also keeps its production down, each firm will stay cooperative, producing 30 million pounds again for the next year. But if one firm produces 40 million pounds, the other firm will take the gloves off and also produce 40 million pounds next year. This latter strategy—start by behaving cooperatively, but thereafter do whatever the other player did in the previous period—is generally known as **tit for tat.**

Playing "tit for tat" is a form of strategic behavior because it is intended to influence the future actions of other players. The "tit for tat" strategy offers a reward to

A firm engages in **strategic behavior** when it attempts to influence the future behavior of other firms.

A strategy of **tit for tat** involves playing cooperatively at first, then doing whatever the other player did in the previous period.

the other player for cooperative behavior—if you behave cooperatively, so will I. It also provides a punishment for cheating—if you cheat, don't expect me to be nice in the future.

The payoff to each firm of each of these strategies would depend on which strategy the other chooses. Consider the four possibilities, shown in Figure 65.3:

1. If one firm plays "tit for tat" and so does the other, both firms will make a profit of $180 million each year.
2. If one firm plays "always cheat" but the other plays "tit for tat," one makes a profit of $200 million the first year but only $160 million per year thereafter.
3. If one firm plays "tit for tat" but the other plays "always cheat," one makes a profit of only $150 million in the first year but $160 million per year thereafter.
4. If one firm plays "always cheat" and the other does the same, both firms will make a profit of $160 million each year.

figure 65.3

How Repeated Interaction Can Support Collusion

A strategy of "tit for tat" involves playing cooperatively at first, then following the other player's move. This rewards good behavior and punishes bad behavior. If the other player cheats, playing "tit for tat" will lead to only a short-term loss in comparison to playing "always cheat." But if the other player plays "tit for tat," also playing "tit for tat" leads to a long-term gain. So a firm that expects other firms to play "tit for tat" may well choose to do the same, leading to successful tacit collusion.

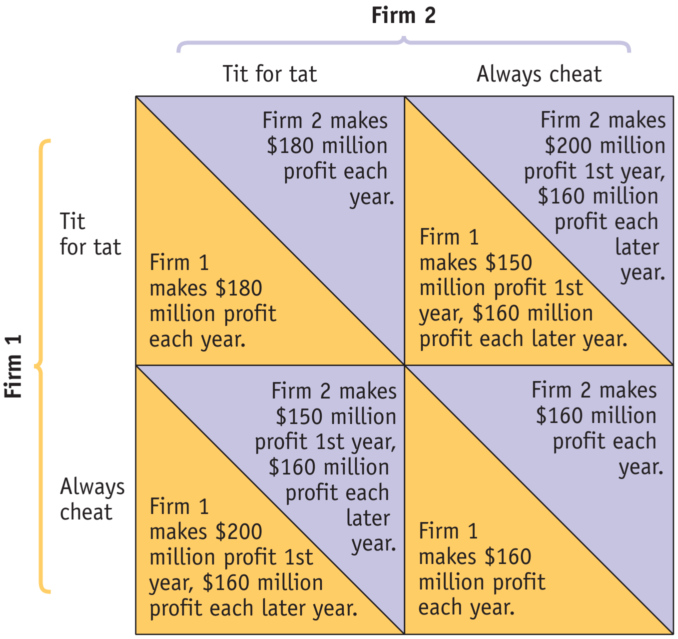

		Firm 2	
		Tit for tat	Always cheat
Firm 1	Tit for tat	Firm 2 makes $180 million profit each year. Firm 1 makes $180 million profit each year.	Firm 2 makes $200 million profit 1st year, $160 million profit each later year. Firm 1 makes $150 million profit 1st year, $160 million profit each later year.
	Always cheat	Firm 2 makes $150 million profit 1st year, $160 million profit each later year. Firm 1 makes $200 million profit 1st year, $160 million profit each later year.	Firm 2 makes $160 million profit each year. Firm 1 makes $160 million profit each year.

Which strategy is better? In the first year, one firm does better playing "always cheat," whatever its rival's strategy: it assures itself that it will get either $200 million or $160 million. (Which of the two payoffs it actually receives depends on whether the other plays "tit for tat" or "always cheat.") This is better than what it would get in the first year if it played "tit for tat": either $180 million or $150 million. But by the second year, a strategy of "always cheat" gains the firm only $160 million per year for the second and all subsequent years, regardless of the other firm's actions. Over time, the total amount gained by playing "always cheat" is less than the amount gained by playing "tit for tat": for the second and all subsequent years, it would never get any less than $160 million and would get as much as $180 million if the other firm played "tit for tat" as well. Which strategy, "always cheat" or "tit for tat," is more profitable depends on two things: how many years each firm expects to play the game and what strategy its rival follows.

If the firm expects the lysine business to end in the near future, it is in effect playing a one-shot game. So it might as well cheat and grab what it can. Even if the firm expects to remain in the lysine business for many years (therefore to find itself repeatedly playing this game) and, for some reason, expects the other firm will always cheat, it should also always cheat. That is, the firm should follow the old rule, "Do unto others before they do unto you."

But if the firm expects to be in the business for a long time and thinks the other firm is likely to play "tit for tat," it will make more profits over the long run by playing "tit for tat," too. It could have made some extra short-term profit by cheating at the beginning, but this would provoke the other firm into cheating, too, and would, in the end, mean less profit.

The lesson of this story is that when oligopolists expect to compete with each other over an extended period of time, each individual firm will often conclude that it is in its own best interest to be helpful to the other firms in the industry. So it will restrict its output in a way that raises the profit of the other firms, expecting them to return the favor. Despite the fact that firms have no way of making an enforceable agreement to limit output and raise prices (and are in legal jeopardy if they even discuss prices), they manage to act "as if" they had such an agreement. When this type of unspoken agreement comes about, we say that the firms are engaging in **tacit collusion.**

When firms limit production and raise prices in a way that raises each other's profits, even though they have not made any formal agreement, they are engaged in **tacit collusion.**

fyi

Prisoners of the Arms Race

Between World War II and the late 1980s, the United States and the Soviet Union were locked in a seemingly endless struggle that never broke out into open war. During this Cold War, both countries spent huge sums on arms, sums that were a significant drain on the U.S. economy and eventually proved a crippling burden for the Soviet Union, whose underlying economic base was much weaker. Yet neither country was ever able to achieve a decisive military advantage.

As many people pointed out, both nations would have been better off if they had both spent less on arms. Yet the arms race continued for 40 years.

Why? As political scientists were quick to notice, one way to explain the arms race was to suppose that the two countries were locked in a classic prisoners' dilemma. Each government would have liked to achieve decisive military superiority, and each feared military inferiority. But both would have preferred a stalemate with low military spending to one with high spending. However, each government rationally chose to engage in high spending. If its rival did not spend heavily, this would lead to military superiority; not spending heavily would lead to inferiority if the other government continued its arms buildup. So the countries were trapped.

The answer to this trap could have been an agreement not to spend as much; indeed, the two sides tried repeatedly to negotiate limits on some kinds of weapons. But these agreements weren't very effective. In the end the issue was resolved as heavy military spending hastened the collapse of the Soviet Union in 1991.

Unfortunately, the logic of an arms race has not disappeared. A nuclear arms race has developed between Pakistan and India, neighboring countries with a history of mutual antagonism. In 1998 the two countries confirmed the unrelenting logic of the prisoners' dilemma: both publicly tested their nuclear weapons in a tit-for-tat sequence, each seeking to prove to the other that it could inflict just as much damage as its rival.

TASS/Soufoto