CHAPTER 28

GAME THEORY

The previous chapter on oligopoly theory presented the classical economic theory of strategic interaction among firms. But that is really just the tip of the iceberg. Economic agents can interact strategically in a variety of ways, and many of these have been studied by using the apparatus of **game theory**. Game theory is concerned with the general analysis of strategic interaction. It can be used to study parlor games, political negotiation, and economic behavior. In this chapter we will briefly explore this fascinating subject to give you a flavor of how it works and how it can be used to study economic behavior in oligopolistic markets.

28.1 The Payoff Matrix of a Game

Strategic interaction can involve many players and many strategies, but we'll limit ourselves to two-person games with a finite number of strategies. This will allow us to depict the game easily in a **payoff matrix**. It is simplest to examine this in the context of a specific example.

Suppose that two people are playing a simple game. Person A will write one of two words on a piece of paper, "top" or "bottom." Simultaneously, person B will independently write "left" or "right" on a piece of paper. After they do this, the papers will be examined and they will each get the payoff depicted in Table 28.1. If A says top and B says left, then we examine the top left-hand corner of the matrix. In this matrix the payoff to A is the first entry in the box, 1, and the payoff to B is the second entry, 2. Similarly, if A says bottom and B says right, then A will get a payoff of 1 and B will get a payoff of 0.

Person A has two strategies: he can choose top or he can choose bottom. These strategies could represent economic choices like "raise price" or "lower price." Or they could represent political choices like "declare war" or "don't declare war." The payoff matrix of a game simply depicts the payoffs to each player for each combination of strategies that are chosen.

What will be the outcome of this sort of game? The game depicted in Table 28.1 has a very simple solution. From the viewpoint of person A, it is always better for him to say bottom since his payoffs from that choice (2 or 1) are always greater than their corresponding entries in top (1 or 0). Similarly, it is always better for B to say left since 2 and 1 dominate 1 and 0. Thus we would expect that the equilibrium strategy is for A to play bottom and B to play left.

In this case, we have a **dominant strategy**. There is one optimal choice of strategy for each player no matter what the other player does. Whichever choice B makes, player A will get a higher payoff if he plays bottom, so it makes sense for A to play bottom. And whichever choice A makes, B will get a higher payoff if he plays left. Hence, these choices dominate the alternatives, and we have an equilibrium in dominant strategies.

A payoff matrix of a game.

		Player B	
		Left	Right
Player A	Top	1, 2	0, 1
	Bottom	2, 1	1, 0

If there is a dominant strategy for each player in some game, then we would predict that it would be the equilibrium outcome of the game. For a dominant strategy is a strategy that is best no matter what the other player does. In this example, we would expect an equilibrium outcome in which A plays bottom, receiving an equilibrium payoff of 2, and B plays left, receiving an equilibrium payoff of 1.

28.2 Nash Equilibrium

Dominant strategy equilibria are nice when they happen, but they don't happen all that often. For example, the game depicted in Table 28.2 doesn't have a dominant strategy equilibrium. Here when B chooses left the payoffs to A are 2 or 0. When B chooses right, the payoffs to A are 0 or 1. This means that when B chooses left, A would want to choose top; and when B chooses right, A would want to choose bottom. Thus A's optimal choice depends on what he thinks B will do.

A Nash equilibrium.

		Player B	
		Left	Right
Player A	Top	2, 1	0, 0
	Bottom	0, 0	1, 2

However, perhaps the dominant strategy equilibrium is too demanding. Rather than require that A's choice be optimal for *all* choices of B, we can just require that it be optimal for the *optimal* choices of B. For if B is a well-informed intelligent player, he will only want to choose optimal strategies. (Although, what is optimal for B will depend on A's choice as well!)

We will say that a pair of strategies is a **Nash equilibrium** if A's choice is optimal, given B's choice, *and* B's choice is optimal given A's choice.[1] Remember that neither person knows what the other person will do when he has to make his own choice of strategy. But each person may have

[1] John Nash is an American mathematician who formulated this fundamental concept of game theory in 1951. In 1994 he received the Nobel Prize in economics, along with two other game theory pioneers, John Harsanyi and Reinhard Selten. The 2002 film *A Beautiful Mind* is loosely based on John Nash's life; it won the Academy Award for best movie.

some expectation about what the other person's choice will be. A Nash equilibrium can be interpreted as a pair of expectations about each person's choice such that, when the other person's choice is revealed, neither individual wants to change his behavior.

In the case of Table 28.2, the strategy (top, left) is a Nash equilibrium. To prove this note that if A chooses top, then the best thing for B to do is to choose left, since the payoff to B from choosing left is 1 and from choosing right is 0. And if B chooses left, then the best thing for A to do is to choose top since then A will get a payoff of 2 rather than of 0.

Thus if A chooses top, the optimal choice for B is to choose left; and if B chooses left, then the optimal choice for A is top. So we have a Nash equilibrium: each person is making the optimal choice, *given* the other person's choice.

The Nash equilibrium is a generalization of the Cournot equilibrium described in the last chapter. There the choices were output levels, and each firm chose its output level taking the other firm's choice as being fixed. Each firm was supposed to do the best for itself, assuming that the other firm continued to produce the output level it had chosen—that is, it continued to play the strategy it had chosen. A Cournot equilibrium occurs when each firm is maximizing profits given the other firm's behavior; this is precisely the definition of a Nash equilibrium.

The Nash equilibrium notion has a certain logic. Unfortunately, it also has some problems. First, a game may have more than one Nash equilibrium. In fact, in Table 28.2 the choices (bottom, right) also comprise a Nash equilibrium. You can either verify this by the kind of argument used above, or just note that the structure of the game is symmetric: B's payoffs are the same in one outcome as A's payoffs are in the other, so that our proof that (top, left) is an equilibrium is also a proof that (bottom, right) is an equilibrium.

The second problem with the concept of a Nash equilibrium is that there are games that have no Nash equilibrium of the sort we have been describing at all. Consider, for example, the case depicted in Table 28.3. Here a Nash equilibrium of the sort we have been examining does not exist. If player A plays top, then player B wants to play left. But if player B plays left, then player A wants bottom. Similarly, if player A plays bottom, then player B will play right. But if player B plays right, then player A will play top.

28.3 Mixed Strategies

However, if we enlarge our definition of strategies, we can find a new sort of Nash equilibrium for this game. We have been thinking of each agent as choosing a strategy once and for all. That is, each agent is making one choice and sticking to it. This is called a **pure strategy**.

A game with no Nash equilibrium (in pure strategies).

		Player B	
		Left	Right
Player A	Top	0, 0	0, −1
	Bottom	1, 0	−1, 3

Another way to think about it is to allow the agents to *randomize* their strategies—to assign a probability to each choice and to play their choices according to those probabilities. For example, A might choose to play top 50 percent of the time and bottom 50 percent of the time, while B might choose to play left 50 percent of the time and right 50 percent of the time. This kind of strategy is called a **mixed strategy**.

If A and B follow the mixed strategies given above, of playing each of their choices half the time, then they will have a probability of 1/4 of ending up in each of the four cells in the payoff matrix. Thus the average payoff to A will be 0, and the average payoff to B will be 1/2.

A Nash equilibrium in mixed strategies refers to an equilibrium in which each agent chooses the optimal frequency with which to play his strategies given the frequency choices of the other agent.

It can be shown that for the sort of games we are analyzing in this chapter, there will always exist a Nash equilibrium in mixed strategies. Because a Nash equilibrium in mixed strategies always exists, and because the concept has a certain inherent plausibility, it is a very popular equilibrium notion in analyzing game behavior. In the example in Table 28.3 it can be shown that if player A plays top with probability 3/4 and bottom with probability 1/4, and player B plays left with probability 1/2 and right with probability 1/2, this will constitute a Nash equilibrium.

EXAMPLE: Rock Paper Scissors

But enough of this theory. Let's look at an example that really matters: the well-known pastime of "rock paper scissors." In this game, each player simultaneously chooses to display a fist (rock), a palm (paper), or his first two fingers (scissors). The rules: rock breaks scissors, scissors cuts paper, paper wraps rock.

Throughout history, countless hours have been spent in playing this game. There is even a professional society, the RPS Society, that promotes the game. It offer both a Web site and a movie documenting the 2003 championships in Toronto.

Of course, game theorists recognize that the equilibrium strategy in rock paper scissors is to randomly choose one of the three outcomes. But humans are not necessarily so good at choosing totally random outcomes. If you can predict your opponent's choices to some degree, you can have an edge in making your own choices.

According to the somewhat tongue-in-cheek account of Jennifer 8. Lee, psychology is paramount.[2] In her article she writes that "most people have a go-to throw, reflective of their character, when they are caught off guard. Paper, considered a refined, even passive, throw, is apparently favored by literary types and journalists."

What is the go-to throw of economists, I wonder? Perhaps it is scissors, since we like to cut to the essential forces at work in human behavior. Should you play rock against an economist, then? Perhaps, but I wouldn't rely on it ...

28.4 The Prisoner's Dilemma

Another problem with the Nash equilibrium of a game is that it does not necessarily lead to Pareto efficient outcomes. Consider, for example, the game depicted in Table 28.4. This game is known as the **prisoner's dilemma**. The original discussion of the game considered a situation where two prisoners who were partners in a crime were being questioned in separate rooms. Each prisoner had a choice of confessing to the crime, and thereby implicating the other, or denying that he had participated in the crime. If only one prisoner confessed, then he would go free, and the authorities would throw the book at the other prisoner, requiring him to spend 6 months in prison. If both prisoners denied being involved, then both would be held for 1 month on a technicality, and if both prisoners confessed they would both be held for 3 months. The payoff matrix for this game is given in Table 28.4. The entries in each cell in the matrix represent the utility that each of the agents assigns to the various prison terms, which for simplicity we take to be the negative of the length of their prison terms.

Put yourself in the position of player A. If player B decides to deny committing the crime, then you are certainly better off confessing, since then you'll get off free. Similarly, if player B confesses, then you'll be better off confessing, since then you get a sentence of 3 months rather than a sentence of 6 months. Thus *whatever* player B does, player A is better off confessing.

[2] Yes, "8" really is her middle name. "Rock, Paper, Scissors: High Drama in the Tournament Ring" was published in the *New York Times* on September 5, 2004.

The prisoner's dilemma.

		Player B: Confess	Player B: Deny
Player A	Confess	-3, -3	0, -6
	Deny	-6, 0	-1, -1

The same thing goes for player B—he is better off confessing as well. Thus the unique Nash equilibrium for this game is for both players to confess. In fact, both players confessing is not only a Nash equilibrium, it is a dominant strategy equilibrium, since each player has the same optimal choice independent of the other player.

But if they could both just hang tight, they would each be better off! If they both could be sure the other would hold out, and both could agree to hold out themselves, they would each get a payoff of -1, which would make each of them better off. The strategy (deny, deny) is Pareto efficient—there is no other strategy choice that makes both players better off—while the strategy (confess, confess) is Pareto inefficient.

The problem is that there is no way for the two prisoners to coordinate their actions. If each could trust the other, then they could both be made better off.

The prisoner's dilemma applies to a wide range of economic and political phenomena. Consider, for example, the problem of arms control. Interpret the strategy of "confess" as "deploy a new missile" and the strategy of "deny" as "don't deploy." Note that the payoffs are reasonable. If my opponent deploys his missile, I certainly want to deploy, even though the best strategy for both of us is to agree not to deploy. But if there is no way to make a binding agreement, we each end up deploying the missile and are both made worse off.

Another good example is the problem of cheating in a cartel. Now interpret confess as "produce more than your quota of output" and interpret deny as "stick to the original quota." If you think the other firm is going to stick to its quota, it will pay you to produce more than your own quota. And if you think that the other firm will overproduce, then you might as well, too!

The prisoner's dilemma has provoked a lot of controversy as to what is the "correct" way to play the game—or, more precisely, what is a reasonable way to play the game. The answer seems to depend on whether you are playing a one-shot game or whether the game is to be repeated an indefinite number of times.

If the game is going to be played just one time, the strategy of defecting—in this example, confessing—seems to be a reasonable one. After all, whatever the other fellow does, you are better off, and you have no way of influencing the other person's behavior.

28.5 Repeated Games

In the preceding section, the players met only once and played the prisoner's dilemma game a single time. However, the situation is different if the game is to be played repeatedly by the same players. In this case there are new strategic possibilities open to each player. If the other player chooses to defect on one round, then you can choose to defect on the next round. Thus your opponent can be "punished" for "bad" behavior. In a repeated game, each player has the opportunity to establish a reputation for cooperation, and thereby encourage the other player to do the same.

Whether this kind of strategy will be viable depends on whether the game is going to be played a *fixed* number of times or an *indefinite* number of times.

Let us consider the first case, where both players know that the game is going to be played 10 times, say. What will the outcome be? Suppose we consider round 10. This is the last time the game will be played, by assumption. In this case, it seems likely that each player will choose the dominant strategy equilibrium, and defect. After all, playing the game for the last time is just like playing it once, so we should expect the same outcome.

Now consider what will happen on round 9. We have just concluded that each player will defect on round 10. So why cooperate on round 9? If you cooperate, the other player might as well defect now and exploit your good nature. Each player can reason the same way, and thus each will defect.

Now consider round 8. If the other person is going to defect on round 9 ... and so it goes. If the game has a known, fixed number of rounds, then each player will defect on every round. If there is no way to enforce cooperation on the last round, there will be no way to enforce cooperation on the next to the last round, and so on.

Players cooperate because they hope that cooperation will induce further cooperation in the future. But this requires that there will always be the possibility of future play. Since there is no possibility of future play in the last round, no one will cooperate then. But then why should anyone cooperate on the next to the last round? Or the one before that? And so it goes—the cooperative solution "unravels" from the end in a prisoner's dilemma with a known, fixed number of plays.

But if the game is going to be repeated an indefinite number of times, then you *do* have a way of influencing your opponent's behavior: if he refuses to cooperate this time, you can refuse to cooperate next time. As long as both parties care enough about future payoffs, the threat of non-cooperation in the future may be sufficient to convince people to play the Pareto efficient strategy.

This has been demonstrated in a convincing way in a series of experiments run by Robert Axelrod.[3] He asked dozens of experts on game theory to submit their favorite strategies for the prisoner's dilemma and then ran a "tournament" on a computer to pit these strategies against each other. Every strategy was played against every other strategy on the computer, and the computer kept track of the total payoffs.

The winning strategy—the one with the highest overall payoff—turned out to be the simplest strategy. It is called "tit for tat" and goes like this. On the first round, you cooperate—play the "deny" strategy. On every round thereafter, if your opponent cooperated on the previous round, you cooperate. If your opponent defected on the previous round, you defect. In other words, do whatever the other player did in the last round.

The tit-for-tat strategy does very well because it offers an immediate punishment for defection. It is also a forgiving strategy: it punishes the other player only once for each defection. If he falls into line and starts to cooperate, then tit for tat will reward the other player with cooperation. It appears to be a remarkably good mechanism for achieving the efficient outcome in a prisoner's dilemma that will be played an indefinite number of times.

28.6 Enforcing a Cartel

In Chapter 27 we discussed the behavior of duopolists playing a price-setting game. We argued there that if each duopolist could choose his price, then the equilibrium outcome would be the competitive equilibrium. If each firm thought that the other firm would keep its price fixed, then each firm would find it profitable to undercut the other. The only place where this would not be true was if each firm were charging the lowest possible price, which in the case we examined was a price of zero, since the marginal costs were zero. In the terminology of this chapter, each firm charging a zero price is a Nash equilibrium in pricing strategies—what we called a Bertrand equilibrium in Chapter 27.

The payoff matrix for the duopoly game in pricing strategies has the same structure as the prisoner's dilemma. If each firm charges a high price, then they both get large profits. This is the situation where they are both cooperating to maintain the monopoly outcome. But if one firm is charging

3 Robert Axelrod is a political scientist from the University of Michigan. For an extended discussion, see his book *The Evolution of Cooperation* (New York: Basic Books, 1984).

a high price, then it will pay the other firm to cut its price a little, capture the other fellow's market, and thereby get even higher profits. But if both firms cut their prices, they both end up making lower profits. Whatever price the other fellow is charging, it will always pay you to shave your price a little bit. The Nash equilibrium occurs when each fellow is charging the lowest possible price.

However, if the game is repeated an indefinite number of times, there may be other possible outcomes. Suppose that you decide to play tit for tat. If the other fellow cuts his price this week, you will cut yours next week. If each player knows that the other player is playing tit for tat, then each player would be fearful of cutting his price and starting a price war. The threat implicit in tit for tat may allow the firms to maintain high prices.

Real-life cartels sometimes appear to employ tit-for-tat strategies. For example, the Joint Executive Committee was a famous cartel that set the price of railroad freight in the United States in the late 1800s. The formation of this cartel preceded antitrust regulation in the United States, and at the time was perfectly legal.[4]

The cartel determined what market share each railroad could have of the freight shipped. Each firm set its rates individually, and the JEC kept track of how much freight each firm shipped. However, there were several occasions during 1881, 1884, and 1885 where some members of the cartel thought that other member firms were cutting rates so as to increase their market share, despite their agreement. During these periods, there were often price wars. When one firm tried to cheat, all firms would cut their prices so as to "punish" the defectors. This kind of tit-for-tat strategy was apparently able to support the cartel arrangement for some time.

EXAMPLE: Tit for Tat in Airline Pricing

Airline pricing provides an interesting example of tit-for-tat behavior. Airlines often offer special promotional fares of one sort or another; many observers of the airline industry claim that these promotions can be used to signal competitors to refrain from cutting prices on key routes.

A senior director of marketing for a major U.S. airline described a case in which Northwest lowered fares on night flights from Minneapolis to various West Coast cities in an effort to fill empty seats. Continental Airlines interpreted this as an attempt to gain market share at its expense and responded by cutting *all* its Minneapolis fares to Northwest's night-fare

[4] For a detailed analysis, see Robert Porter, "A Study of Cartel Stability: the Joint Executive Committee, 1880–1886," *The Bell Journal of Economics*, 14, 2 (Autumn 1983), 301–25.

level. However, the Continental fare cuts were set to expire one or two days after they were introduced.

Northwest interpreted this as a signal from Continental that it was not serious about competing in this market, but simply wanted Northwest to retract its night-fare cuts. But Northwest decided to send a message of its own to Continental: it instituted a set of cheap fares to the West Coast for its flights departing from Houston, Continental's home base! Northwest thereby signaled that it felt its cuts were justified, while Continental's response was inappropriate.

All these fare cuts had very short expiration dates; this feature seems to indicate that they were meant more as messages to the competition than as bids for larger market share. As the analyst explained, fares that an airline doesn't want to offer "should almost always have an expiration date on them in the hopes that the competition will eventually wake up and match."

The implicit rules of competition in duopoly airline markets seem to be the following: if the other firm keeps its prices high, I will maintain my high prices; but if the other firm cuts its prices, I will play tit for tat and cut my prices in response. In other words, both firms "live by the Golden Rule": do unto others as you would have them do unto you. This threat of retaliation then serves to keep all prices high.[5]

28.7 Sequential Games

Up until now we have been thinking about games in which both players act simultaneously. But in many situations one player gets to move first, and the other player responds. An example of this is the Stackelberg model described in Chapter 27, where one player is a leader and the other player is a follower.

Let's describe a game like this. In the first round, player A gets to choose top or bottom. Player B gets to observe the first player's choice and then chooses left or right. The payoffs are illustrated in a game matrix in Table 28.5.

Note that when the game is presented in this form it has two Nash equilibria: (top, left) and (bottom, right). However, we'll show below that one of these equilibria isn't really reasonable. The payoff matrix hides the fact that one player gets to know what the other player has chosen before he makes his choice. In this case it is more useful to consider a diagram that illustrates the asymmetric nature of the game.

Figure 28.1 is a picture of the game in **extensive form**—a way to represent the game that shows the time pattern of the choices. First, player A

[5] Facts taken from A. Nomani, "Fare Warning: How Airlines Trade Price Plans," *Wall Street Journal,* October 9, 1990, B1.

The payoff matrix of a sequential game.

		Player B	
		Left	Right
Player A	Top	1, 9	1, 9
	Bottom	0, 0	2, 1

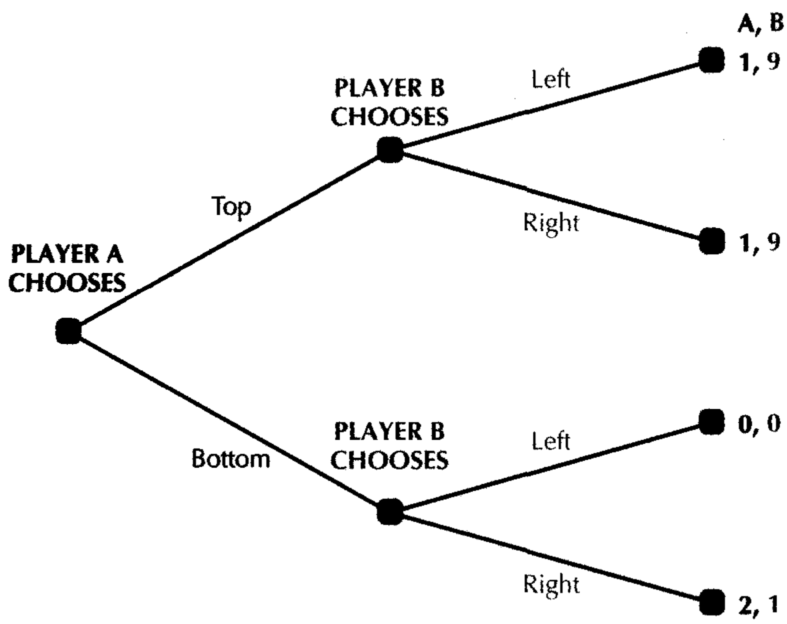

Extensive form of the game. This way of depicting a game indicates the order in which the players move.

has to choose top or bottom, and then player B has to choose left or right. But when B makes his choice, he will know what A has done.

The way to analyze this game is to go to the end and work backward. Suppose that player A has already made his choice and we are sitting in one branch of the game tree. If player A has chosen top, then it doesn't matter what player B does, and the payoff is (1,9). If player A has chosen bottom, then the sensible thing for player B to do is to choose right, and the payoff is (2,1).

Now think about player A's initial choice. If he chooses top, the outcome will be (1,9) and thus he will get a payoff of 1. But if he chooses bottom, he gets a payoff of 2. So the sensible thing for him to do is to choose bottom. Thus the equilibrium choices in the game will be (bottom, right), so that the payoff to player A will be 2 and to player B will be 1.

The strategies (top, left) are not a reasonable equilibrium in this sequential game. That is, they are not an equilibrium given the order in which the players actually get to make their choices. It is true that if player A chooses top, player B could choose left—but it would be silly for player A to ever choose top!

From player B's point of view this is rather unfortunate, since he ends up with a payoff of 1 rather than 9! What might he do about it?

Well, he can *threaten* to play left if player A plays bottom. If player A thought that player B would actually carry out this threat, he would be well advised to play top. For top gives him 1, while bottom—if player B carries out his threat—will only give him 0.

But is this threat credible? After all, once player A makes his choice, that's it. Player B can get either 0 or 1, and he might as well get 1. Unless player B can somehow convince player A that he will really carry out his threat—even when it hurts him to do so—he will just have to settle for the lower payoff.

Player B's problem is that once player A has made his choice, player A expects player B to do the rational thing. Player B would be better off if he could *commit* himself to play left if player A plays bottom.

One way for B to make such a commitment is to allow someone else to make his choices. For example, B might hire a lawyer and instruct him to play left if A plays bottom. If A is aware of these instructions, the situation is radically different from his point of view. If he knows about B's instructions to his lawyer, then he knows that if he plays bottom he will end up with a payoff of 0. So the sensible thing for him to do is to play top. In this case B has done better for himself by *limiting* his choices.

28.8 A Game of Entry Deterrence

In our examination of oligopoly we took the number of firms in the industry as fixed. But in many situations, entry is possible. Of course, it is in the interest of the firms in the industry to try to prevent such entry. Since they are already in the industry, they get to move first and thus have an advantage in choosing ways to keep their opponents out.

Suppose, for example, that we consider a monopolist who is facing a threat of entry by another firm. The entrant decides whether or not to come into the market, and then the incumbent decides whether or not to cut its price in response. If the entrant decides to stay out, it gets a payoff of 1 and the incumbent gets a payoff of 9.

If the entrant decides to come in, then its payoff depends on whether the incumbent fights—by competing vigorously—or not. If the incumbent fights, then we suppose that both players end up with 0. On the other hand, if the incumbent decides not to fight, we suppose that the entrant gets 2 and the incumbent gets 1.

Note that this is exactly the structure of the sequential game we studied earlier, and thus it has a structure identical to that depicted in Figure 28.1. The incumbent is player B, while the potential entrant is player A. The top strategy is to stay out, and the bottom strategy is to enter. The left strategy is to fight and the right strategy is not to fight. As we've seen in this game, the equilibrium outcome is for the potential entrant to enter and the incumbent *not* to fight.

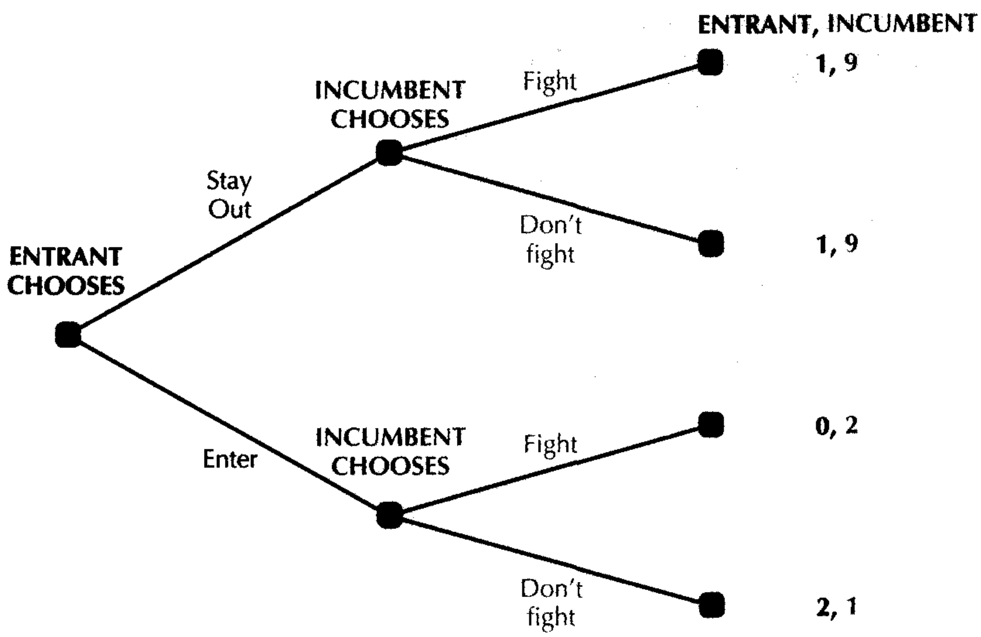

The new entry game. This figure depicts the entry game with the changed payoffs.

The incumbent's problem is that he cannot precommit himself to fighting if the other firm enters. If the other firm enters, the damage is done and the rational thing for the incumbent to do is to live and let live. Insofar as the potential entrant recognizes this, he will correctly view any threats to fight as empty.

But suppose that the incumbent can purchase some extra production capacity that will allow him to produce more output at his current marginal cost. Of course, if he remains a monopolist, he won't want to actually use this capacity since he is already producing the profit-maximizing monopoly output.

But, if the other firm enters, the incumbent will now be able to produce so much output that he may well be able to compete much more successfully against the new entrant. By investing in the extra capacity, he will lower his costs of fighting if the other firm tries to enter. Let us assume that if he purchases the extra capacity and if he chooses to fight, he will make a profit of 2. This changes the game tree to the form depicted in .

Now, because of the increased capacity, the threat of fighting is credible. If the potential entrant comes into the market, the incumbent will get a payoff of 2 if he fights and 1 if he doesn't; thus the incumbent will rationally choose to fight. The entrant will therefore get a payoff of 0 if he enters, and if he stays out he will get a payoff of 1. The sensible thing for the potential entrant to do is to stay out.

But this means that the incumbent will remain a monopolist and never have to use his extra capacity! Despite this, it is worthwhile for the monopolist to invest in the extra capacity in order to make credible the *threat* of fighting if a new firm tries to enter the market. By investing in "excess" capacity, the monopolist has signaled to the potential entrant that he will be able to successfully defend his market.

Summary

1. A game can be described by indicating the payoffs to each of the players for each configuration of strategic choices they make.

2. A dominant strategy equilibrium is a set of choices for which each player's choices are optimal *regardless* of what the other players choose.

3. A Nash equilibrium is a set of choices for which each player's choice is optimal, given the choices of the other players.

4. The prisoner's dilemma is a particular game in which the Pareto efficient outcome is strategically dominated by an inefficient outcome.

5. If a prisoner's dilemma is repeated an indefinite number of times, then it is possible that the Pareto efficient outcome may result from rational play.

6. In a sequential game, the time pattern of choices is important. In these games, it can often be advantageous to find a way to precommit to a particular line of play.

REVIEW QUESTIONS

1. Consider the tit-for-tat strategy in the repeated prisoner's dilemma. Suppose that one player makes a mistake and defects when he meant to cooperate. If both players continue to play tit for tat after that, what happens?

2. Are dominant strategy equilibria always Nash equilibria? Are Nash equilibria always dominant strategy equilibria?

3. Suppose your opponent is *not* playing her Nash equilibrium strategy. Should you play your Nash equilibrium strategy?

4. We know that the single-shot prisoner's dilemma game results in a dominant Nash equilibrium strategy that is Pareto inefficient. Suppose we allow the two prisoners to retaliate after their respective prison terms. Formally, what aspect of the game would this affect? Could a Pareto efficient outcome result?

5. What is the dominant Nash equilibrium strategy for the repeated prisoner's dilemma game when both players know that the game will end after one million repetitions? If you were going to run an experiment with human players for such a scenario, would you predict that players would use this strategy?

6. Suppose that player B rather than player A gets to move first in the sequential game described in this chapter. Draw the extensive form of the new game. What is the equilibrium for this game? Does player B prefer to move first or second?