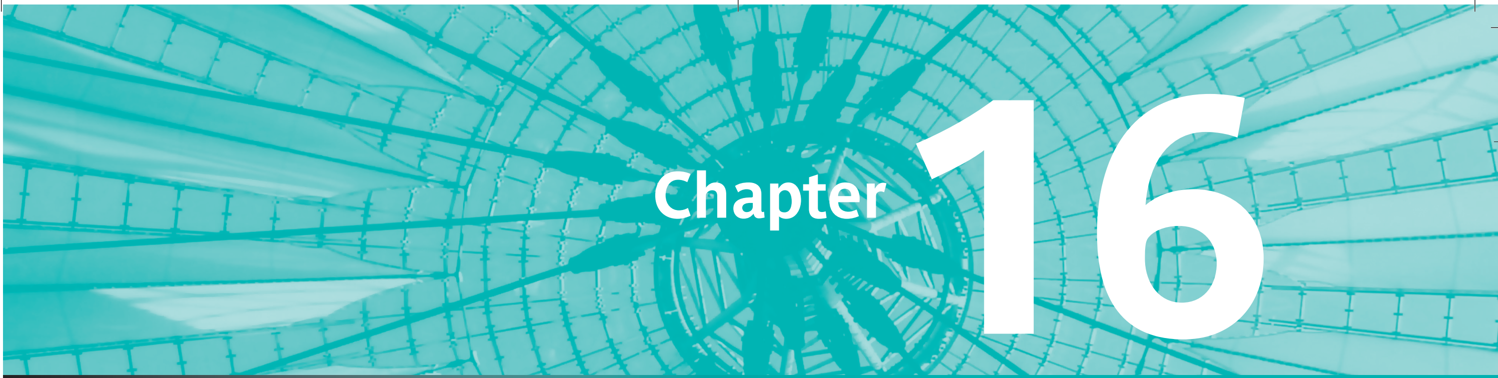

Chapter 16

Game Theory

When the One Great Scorer comes to write
against your name –
He marks – not that you won or lost – but
how you played the game.

Grantland Rice

LEARNING OBJECTIVES

After studying this chapter you will have:

- Defined an economic game including players, strategy, actions and payoffs
- Learned how decision tress are created
- Examined what a strategy is and the development of a dominant strategy
- Applied game theory to oligopolistic markets
- Defined credible threats and commitments
- Analysed games with incomplete information and worked through a prisoners' dilemma game
- Examined games which are repeated and limit pricing models

France and Italy had played through extra time and were still drawing 1–1 in the 2006 World Cup Final and once again, the World Cup was being decided by a penalty shoot-out. The score was 2-1 to Italy as David Trezuguet, the Juventus striker, stepped up to take the next kick for

France. Gianluigi Buffon, Italy's goalie, knew that once the ball left Trezuguet's foot, it would be too late to react. He had to make up his mind to throw his body one way or the other just as Trezuguet kicked the ball. Which way should Buffon go? It all depended on what he thought Trezuguet was going to do. Which way should Trezuguet kick the ball? It all depended on what he thought Buffon was going to do.

game A situation in which strategic behaviour is an important part of decision making.

non-cooperative game theory A set of tools for analysing decision making in situations where strategic behaviour is important.

The situation just described is a "game" in two senses of the word. One, it is a sport. Two, it is a strategic situation: each decision maker has to take into account what he or she thinks the other is going to do. Economists call any strategic situation – including, for example, oligopoly – a **game**. The notion of Nash equilibrium that we used in Chapter 15 to analyse oligopoly is part of a larger set of tools for analysing strategic behaviour – in economics, politics, card games, and other arenas of conflict – known as **non-cooperative game theory**. This theory is labelled "non-cooperative" because each decision maker acts solely in his or her own self-interest. Despite the label, the theory is relevant to the analysis of co-operation. Even "selfish" economic agents will co-operate if doing so is in their self-interest. For example, a self-enforcing agreement among firms to co-operate in restricting industry output is non-cooperative in the technical sense of the word – each firm adheres to the agreement solely because it is in the firm's self-interest to do so.

In this chapter, we develop a useful way to represent strategic situations graphically, and we use this representation to analyse oligopoly further. In particular, we use game theory to investigate the behaviour of oligopoly when there is a threat of entry. We will also see how game theory provides important insights into behaviour in a variety of other strategic situations.

16.1 Some Fundamentals of Game Theory

players The decision makers in a game.

strategy A player's plan of action in a game.

actions The particular things that are done according to a player's strategy for a game.

payoffs The rewards enjoyed by a player at the end of a game.

In any game there are decision makers; they are called **players**. In blackjack, the players are the dealer and the bettors. In oligopoly, the players are the firms in the industry. The choices that a player makes are called a **strategy**, and the particular things done according to a strategy are called **actions**. A blackjack player's strategy must indicate whether to stand or take another card when he or she already has 16 points. The action taken is then either "stick" or "twist" (take a card). An oligopolist's strategy may specify how it will respond to cheating by another firm. The action taken in a given period might then consist of setting a particular price. At the end of the game, players get **payoffs**, depending on what has happened. The bettors have winnings or losses; the oligopolists, profits. Of course, the game must be played according to some set of rules. In blackjack, the rules are explicitly spelt out. The rules in an oligopoly game are somewhat more difficult to discern, and we will have more to say about them below.

Game Trees: Decision Trees for Strategic Situations

We need a convenient way to represent the rules of the game, such as who moves when and what each player knows when it is his or her turn to move. If we simply were to list the rules, they might be very complicated, and it could be difficult to find an equilibrium. In Chapter 6 we saw how a decision tree could be used to break down a problem and simplify finding a solution. Here, we will develop a very similar tool known as a **game tree**. The main

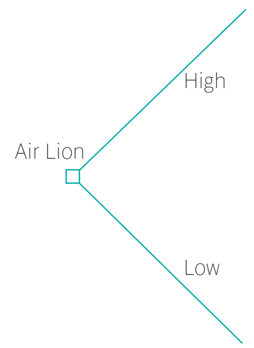

Figure 16.1 Air Lion's Decision

When Air Lion makes its output choice first and has only two possible output levels, "high" and "low", Air Lion's choice problem is represented by a single decision node with two branches.

difference between a game tree and a decision tree is that several different players make moves in a game tree, but only one player makes moves in a decision tree.

game tree
An extension of a decision tree that provides a graphical representation of a strategic situation.

To illustrate the use of a game tree, let's again consider the situation of Air Lion and Beta Airlines. In the previous chapter, we assumed that Air Lion and Beta made their decisions in any given period simultaneously. Here, however, we initially assume that Air Lion chooses its output level first, and then Beta responds. To keep things simple, assume that each firm has only two possible output levels: "high" and "low". Air Lion's decision is depicted in Figure 16.1. A little square called a *decision node* is used to represent a point at which a decision has to be made. Because there are two players, we have to provide a label that indicates whose decision node it is. If it were simply a decision tree for Air Lion, we would put payoffs at the end of each of the two branches. But the situation we are modelling is a game, and we need to take Beta's actions into account before we can calculate Air Lion's payoffs.

Beta has the same possible actions, "high" and "low", as Air Lion. But there is an important difference between the two firms. Before choosing its output level, Beta gets to see what choice Air Lion has made. As a result, Beta has two different decision nodes in Figure 16.2, one for the contingency that Air Lion chooses "high" and one for the contingency that Air Lion chooses "low".

The payoffs for each possible outcome (combination of actions) are represented at the end of each final branch. Since there are two players, we have to list two payoffs at the end of each branch, one for each player. The first number in each pair is Air Lion's payoff for that outcome; the second is Beta's. For example, if Air Lion's output is "high" and Beta's output is "low", then Air Lion earns a profit of €6,000 and Beta earns €1,000.

Although we have looked at the *actions* that the two firms can take, we have not yet looked at their *strategies*. A firm's strategy specifies the actions the firm will take in any situation that it might face in the game. In other words, a strategy specifies what action the firm will take at each of its decision nodes. Air Lion chooses its output level first and has only one decision node. For Air Lion, the strategy is simply: *produce "high" or produce "low"*. (We will use italics to distinguish strategies from actions.) Beta's strategy is more complicated than Air Lion's. As we just saw, Beta has two decision nodes in the game tree because Beta gets to see what choice Air Lion has made before making its own choice. Beta's strategy has to specify what the firm will do at each of its decision nodes. Rather than choosing a simple action ("high" or "low"), Beta chooses

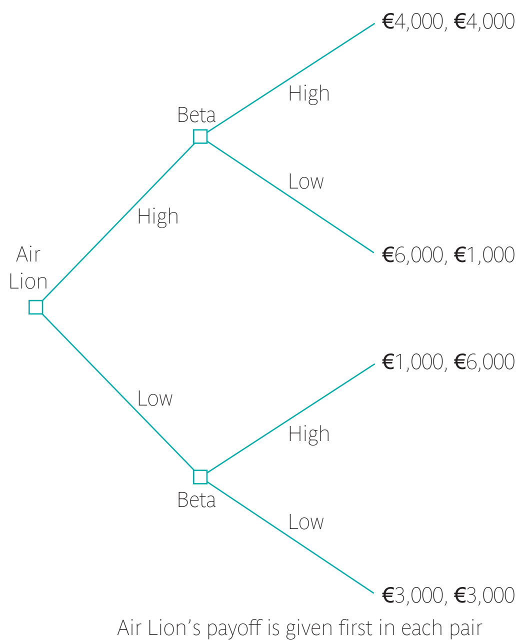

Figure 16.2 A Game in which Both Firms have Dominant Strategies

When Beta gets to see what choice Air Lion has made, Beta has two different decision nodes, one for the contingency that Air Lion chooses "high", and one for the contingency that Air Lion chooses "low". Since Air Lion's strategy (*produce "high"*) and Beta's strategy (*produce "high" no matter what Air Lion does*) are dominant strategies, there is a dominant strategy equilibrium in which the two firms play these strategies.

decision rule
A strategy that specifies what action will be taken conditional on what happens earlier in the game.

a **decision rule** that specifies what action it will take conditional on what Air Lion has done; that is, it specifies which action is to be taken at each of Beta's two decision nodes. One possible strategy for Beta is the following: *if Air Lion produces "high", then I will produce "low", and if Air Lion produces "low", then I will produce "high"* (PC 16.1).

Dominant Strategy Equilibrium

Now that we have a way to represent the rules of the game and its outcome, we can find the equilibrium. What would we expect Air Lion and Beta to do in the game illustrated by Figure 16.2? Consider Beta's strategy first. Suppose that Air Lion has chosen "low". Then Beta would earn €6,000 by choosing "high" and €3,000 by choosing "low". Thus, conditional on Air Lion's choosing "low", Beta's payoff is maximized by choosing "high". If Air Lion had chosen "high", then Beta would choose "high" for a payoff of €4,000 rather than "low" for a payoff of €1,000. Hence, no matter what Air Lion's strategy is, one best strategic response for Beta is to *produce "high" no matter what Air Lion does*. A strategy that works at least as well as any other one, no matter what the other player does, is known as a **dominant strategy**. There is no reason for players to use anything other than their dominant strategy, *if* they have one (this is a big "if," because in many situations there is no dominant strategy). Hence, in equilibrium, we would expect Beta to choose the strategy *produce "high" no matter what Air Lion does*.

dominant strategy
A strategy that works at least as well as any other one, no matter what any other player does.

Progress Check 16.1

There are three other decision rules that Beta could choose if it wanted to do so. Identify these three strategies.

What about Air Lion's strategy? If Air Lion chooses "low", its payoff is €3,000 if Beta responds by producing "low", or €1,000 if Beta responds by producing "high". Similarly, if Air Lion chooses "high", then its payoff is either €6,000 or €4,000, depending on how Beta responds. Notice that Air Lion does better by choosing "high", no matter how Beta responds. Hence, *produce "high"* is a dominant strategy for Air Lion to follow, and this is what we would expect Air Lion to do in equilibrium.

In this situation, each firm has a dominant strategy and would have to be crazy to play anything else. We conclude that *when each player has a dominant strategy, the only reasonable equilibrium outcome is for each player to use its dominant strategy*. The set of dominant strategies and the resulting outcome are known as a **dominant strategy equilibrium**. In the game played by our two airlines, the pair of strategies *produce "high"* for Air Lion and *produce "high" no matter what Air Lion does* for Beta constitute a dominant strategy equilibrium.

dominant strategy equilibrium An outcome in a game in which each player follows a dominant strategy.

At this point, you may be wondering how this notion of equilibrium corresponds to that used in Chapter 15. There, we required an equilibrium to satisfy two conditions: the Nash condition and the credibility condition. Recall that the Nash condition requires that no firm be able to gain by *unilaterally* changing what it is doing – each firm must play a best response to what the other firm is doing. In the language of game theory, each player's equilibrium strategy must be a best response to the equilibrium strategy chosen by the other player. Since a dominant strategy is a best response to anything, a dominant strategy equilibrium clearly satisfies the Nash condition.

What about the credibility condition? It requires that each time a firm is called upon to make a move (that is, at each of the firm's decision nodes), it is in the firm's self-interest to carry out the action called for by its strategy. This property is clearly satisfied by Air Lion's strategy: since there is only one part to Air Lion's strategy, the Nash condition alone guarantees that Air Lion would want to carry out that strategy. The credibility condition is more complicated for Beta, because its strategy has two parts that must be checked. First, the credibility condition is satisfied by Beta's decision to produce "high" in response to a high output level by Air Lion. Since this is the action the firm actually takes in equilibrium, it too is taken care of by the Nash condition. The real issue is whether Beta's *threatened* response of high output to low output by Air Lion is credible. Looking at the game tree in Figure 16.2, we see that it is; if Air Lion chose "low", then Beta would earn €6,000 by producing "high" and only €3,000 by producing "low". It is no coincidence that Beta's strategy is credible. Since Beta is playing a dominant strategy, we know that no matter what Air Lion does, Beta's equilibrium strategy does at least as well as any other. It would be in Beta's self-interest to carry out any part of its strategy if called upon to do so. In summary, the dominant strategy equilibrium that we found satisfies both the Nash condition and the credibility condition.

Perfect Equilibrium

The oligopoly game just examined in Figure 16.2 works out simply because there is a dominant strategy equilibrium. Unfortunately, in most games there is no dominant strategy equilibrium.

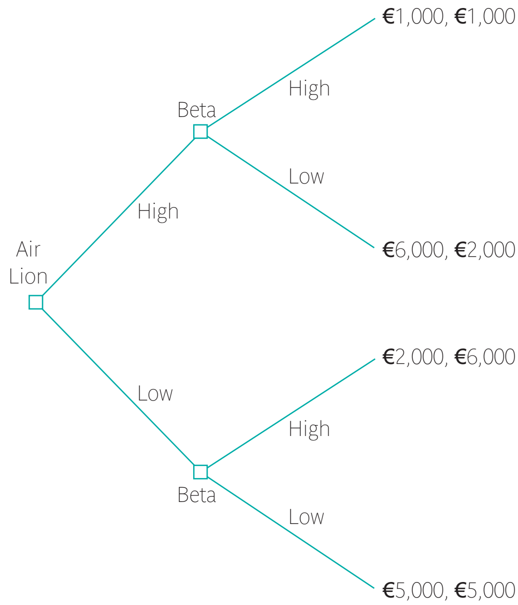

Figure 16.3 A Game in which Only Beta has a Dominant Strategy

Two outcomes can arise when the firms use strategies satisfying the Nash condition. In one, Air Lion plays *produce "high"* and Beta plays *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*. In the other, Air Lion chooses *produce "low"* and Beta chooses *produce "high" no matter what Air Lion does*. The first pair of strategies also satisfies the credibility condition. The second does not – Beta's threat to produce "high" in response to "high" is not credible. Hence, only the first pair of strategies constitutes a perfect equilibrium for this game.

Figure 16.3 illustrates one such game. In this new game, *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*, is a dominant strategy for Beta. Air Lion, however, does not have a dominant strategy. Suppose Beta's strategy is *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*. Given Beta's strategy, Air Lion would earn €6,000 by producing "high" and €2,000 by producing "low". Air Lion's best response is to *produce "high"*. But now suppose that Beta's strategy is to *produce "high" no matter what Air Lion does*. In this case, Air Lion would earn €1,000 by producing "high" and €2,000 by producing "low". Air Lion's best response is *produce "low"*. Air Lion's best response depends on the strategy chosen by Beta.

The problem faced by Air Lion's managers is that they have to make their decision first. What should they expect Beta to do? We just argued that the most reasonable expectation is that a firm will play a dominant strategy if it has one. Although Air Lion does not have a dominant strategy, Beta does: *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*. Game theory is based on the assumption that each player believes that the other players are rational. Hence, Air Lion should expect Beta to play its dominant strategy because Beta would have to be irrational to play anything else. Since Air Lion expects Beta to play the strategy *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*, Air Lion chooses the strategy *produce "high"*.

Let's step back for a moment and consider how we have found the equilibrium for this game. We noted that Air Lion Airlines has to form a prediction about what Beta is going to do. We concluded that the only rational thing to expect is that Beta will respond to Air Lion's move by taking the action that maximizes Beta's profit, *taking as given what Air Lion has done*. In terms of the game tree, the procedure is the following one. Find the last decision that a player makes before receiving the payoffs. For each of these decision nodes, find the action that maximizes that decision maker's profit. Construct a strategy by saying that the firm chooses its profit-maximizing action at each of its decision nodes. Now, use the resulting strategy to calculate what the player moving earlier should do. Because we start at the end of the tree rather than at the beginning, this process is known as *backward induction*. You may recall that this is the same procedure used back in Chapter 6 (page xxx) to solve decision problems that entailed sequential choices.

If you think about this procedure, you will recognize that it implicitly forces Air Lion to ignore any incredible threats or promises by Beta; when it is Beta's turn to move, the airline is going to do what is in its self-interest at that time. Indeed, this example shows why we need the credibility condition in addition to the Nash condition. If all we did was apply the Nash condition, there would be two candidates for the equilibrium outcome. One outcome is where Air Lion's output is high and Beta's is low. This outcome arises when Air Lion's strategy is *produce "high"* and Beta's strategy is *if Air Lion produces "high", I will produce "low", and if Air Lion produces "low", I will produce "high"*.[1] The other outcome is where Air Lion's output is low and Beta's is high. This outcome arises when Air Lion's strategy is *produce "low"*, and Beta's strategy is *produce "high" no matter what Air Lion does*.

The story behind the second candidate for equilibrium is unconvincing. Air Lion chooses a low level of output to keep Beta from later hurting *both* firms by playing "high" in response to "high". It does not seem sensible for Air Lion to take Beta's threat seriously. Recall that our notion of credibility is the following: a threat is credible only if it would be in the firm's self-interest to carry the threat out. Interpreted within the context of a game, this condition requires that, whenever it is a particular player's turn to move, the action called for by that player's equilibrium strategy must be one that is in the player's self-interest to take *at the time that the move is made*. Beta has threatened to choose "high" if Air Lion chooses "high". But suppose that Air Lion called Beta's bluff by choosing "high". What should Air Lion predict that Beta would do in response? If Air Lion has chosen "high", the threat has not worked. *At the time that Beta makes its move*, its self-interest dictates that it choose "low" rather than the threatened "high". Hence, Beta's threat to produce "high" is not credible. Knowing this, Air Lion should expect Beta to choose "low" in response to Air Lion's choice of "high". Air Lion should therefore choose "high" because this will lead to Beta's choosing "low" and Air Lion's earning €6,000 rather than €2,000. As in our discussion of co-operation and punishment in Chapter 15, incredible threats should be ignored, and we can reject this equilibrium as unreasonable.

We conclude that the equilibrium outcome is for Air Lion to produce the high output level and earn a profit of €6,000, while Beta produces a low output level and earns €2,000. As we have seen, *produce "high"* is Air Lion's best response to Beta's dominant strategy. And, by definition, Beta's dominant strategy is a best response to Air Lion's strategy. We have shown that this pair of strategies satisfies both the Nash condition and the credibility condition. An equilibrium outcome that satisfies these two conditions is known as a **perfect equilibrium**.[2]

perfect equilibrium A set of strategies that satisfies both the Nash condition and the credibility condition.

16.2 Applying Game Theory: Oligopoly with Entry

In many oligopolistic markets the incumbent firms face the threat of entry. When it started making plain-paper copiers, Xerox was alone in the industry. Today, Canon, Mita, Sharp and others are active participants in this market. In the 1960s a handful of domestic car producers dominated the Italian car market. Since then there has been dramatic entry by a variety of foreign firms. In this section, we apply the tools of game theory to analyse oligopoly when entry is possible.

We analyse the case of a single incumbent producer who faces the threat of entry by a single potential entrant. Suppose that the incumbent firm, Liege Pharmaceutical, is the sole producer of a drug on which the patent has just expired. General Generic is considering whether to enter this market now that anyone is free to manufacture and sell the drug. It might seem like we are back to looking at a simple monopolized industry since there initially is only one firm in the market; but the situation is, in reality, much closer to oligopoly without entry than it is to monopoly without entry. The similarity between this case and oligopoly without entry arises because both are instances in which strategic behaviour is important. A monopolist who does not face the threat of entry is not engaged in a strategic situation. A single incumbent facing a potential entrant is.

In deciding whether to come into the market, General Generic has to form beliefs about the post-entry equilibrium. As a manager at General, you would have to ask: "What will happen after we come into the market?" If the expected post-entry equilibrium would yield positive economic profit, then General Generic should enter the market. If not, the firm should stay out. Now, if the market under consideration were a perfectly competitive one, you could figure out General's potential profit simply by taking the market price of output as given. You could do this because, as an entrant into a perfectly competitive industry, General Generic would be so small relative to the rest of the market that the entry would make essentially no difference to the market equilibrium. The equilibrium price before entry would thus be a good predictor of the price after entry. Things are not so simple for a firm such as General Generic that is considering entry into an industry with only one or a few incumbents. The entrant must form more sophisticated expectations of how the incumbent or incumbents will react to the entrant's coming into the market.

For its part, an incumbent would like to scare off the entrant by threatening an unpleasant welcome to the industry. The incumbent could threaten to shoot the kneecaps of the entrant's managers or to blow up their offices. While firms have, at times, been accused of using such tactics, the typical threats used by incumbents are considerably less harsh. For example, Liege Pharmaceutical might threaten to produce a large amount of the drug and drive down the market price. As usual, we must ask whether this threat is credible. Will the entrant be scared off or will General Generic call the incumbent's bluff by coming in? If the threat is found to be incredible, is there some way for the incumbent to make its threat credible? The tools of game theory can help us answer these questions.

Figure 16.4 illustrates the game tree representing the entrant's decision whether to come in and the incumbent's output response. As the tree shows, General Generic moves first, and Liege Pharmaceutical makes its output decision after seeing whether General Generic has decided to enter. The strategy for the potential entrant is either to "*enter*" or "*stay out*". Since it makes its decision second, the incumbent's strategy specifies what it will do contingent on the action taken by the potential entrant. The following is an example of a strategy available to the management of Liege Pharmaceutical: *if General Generic chooses "enter", then we will choose "high output", and if General Generic chooses "stay out", then we will choose "low output"*.

There are two pairs of strategies that satisfy the Nash condition for this game. In one, General Generic plays the strategy "*keep out*" and Liege Pharmaceutical plays *produce "high output" no matter what General Generic does*. To verify that this is a Nash equilibrium, we need to check that each firm is choosing a best response to what the other firm is doing. Could either firm increase its profit by changing its strategy while the other firm's strategy remains the same? In the

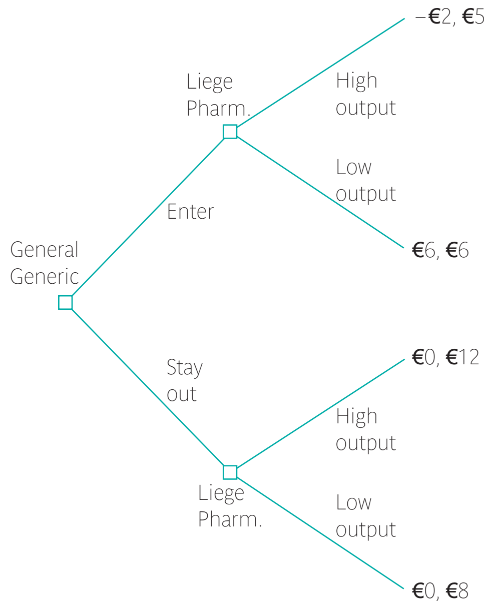

Figure 16.4 An Entry Game

Two outcomes in this game satisfy the Nash condition. In one, General Generic plays the strategy "*stay out*" and Liege Pharmaceutical plays *produce "high output" no matter what General Generic does*. In the other, General Generic plays "*enter*" and Liege Pharmaceutical plays *if General Generic chooses "enter", then we will choose "low output", and if General Generic chooses "stay out", then we will choose "high output"*. Only the second outcome satisfies our credibility condition. The threat to produce "*high output*" in response to "*enter*" in the first Nash outcome is not credible.

outcome being considered, General Generic earns a payoff of €0, whereas Liege Pharmaceutical earns €12 million. If General Generic were to enter, then its profit would be –€2 million since Liege's strategy entails producing a high output level in response to entry. Hence, General Generic has no incentive to change its strategy. What about Liege Pharmaceutical? Given that the potential entrant is staying out of the market, the incumbent's profit is maximized by producing "high output" rather than "low output" (€12 million is more than €8 million). Liege Pharmaceutical has no incentive to change its strategy. Since each firm is choosing a best response to the strategy of the other, this pair of strategies satisfies the Nash condition.

In the other outcome, General Generic plays "*enter*" and Liege Pharmaceutical plays the strategy *if General Generic chooses "enter", we will choose "low output", and if General chooses "stay out", we will choose "high output"*.[3] When the two firms play these strategies, General Generic and Liege Pharmaceutical each earn a profit of €6 million (PC 16.2).

Progress Check 16.2

Verify that this second pair of strategies also satisfies the Nash condition that each strategy be a best response to the other.

Chips with Everything?

Markets are fiercely defended by incumbent firms, especially when they have market power. Keeping out prospective new entrants into their markets is just as important as dealing with the competitors already in the market. However, there are times when this becomes a difficult distinction to maintain as market boundaries change. How does this affect behaviour? Using the analysis in the chapter so far, it would effectively make the values of the choices facing potential new entrants change. Accordingly, the outcome of the decision-making process could alter as a consequence of changes in the market brought about by a host of different factors.

To examine this, consider the market for microprocessors.[a] The market was essentially dominated by two firms for many years, namely Intel and AMD, both of which are US based. They both competed for the market to supply microprocessors for the PC, although Intel was the much larger of the two. They competed to supply PC and server chips, but there are also other major chip, markets – for mobile phones and also for high-graphics. Here other firms operate – Nvidia and ATI dominate the graphics market and mobile chips are dominated by ARM Holdings. Such market distinction reflects the differing needs for the chips produced.

As it stood, the most likely entrant into one of the sub-markets was a firm from one of the other sub-markets. However, the variety of needs of the different markets made this relatively remote as an option. In our analysis, the decision branch to enter would not have been very attractive financially.

Two things changed, however: first, the fear among the firms in their respective sub-markets that stagnation or market maturity would lead to diminishing profits in the medium term; second, as mobile devices have become more like PCs, thus requiring high-quality graphics, the specifications for the chips have become less sharply defined. This implies greater transferability of production processes and, hence, products to the different markets.

How have the firms responded? Intel has started to produce a new range of chips and AMD has bought out ATI and developed Puma, a new range of chips for laptops. Nvidia has seen this new move on its territory and it too has started to diversify its product range. All of the microchip firms are now starting to move into the markets of their major users, such as scientific instruments and mobile phones, thus changing the market dynamic there too. What this shows is that, while we can represent a choice using the decision-tree approach, we have to recognize that this is a static representation of what is a dynamic and ongoing process, so we can only catch a glimpse or snapshot of the decision-making process using this analysis.

[a] This draws on the article "Battlechips", *The Economist*, 7 June 2008.

Clearly, Liege Pharmaceutical prefers the outcome under which General Generic stays out of the market, whereas General Generic prefers the outcome under which it comes into the market. Is there some way to choose between the two outcomes? There is. Only the second one, in which entry occurs, satisfies our credibility condition. In the first outcome, the potential entrant chooses to stay out because the incumbent has threatened to choose the high level of output in response to entry. But suppose that General Generic tested this threat by entering. Once General has actually entered the market, it is not in Liege's self-interest to carry out the threat – Liege can earn €6 million by producing the low output level but only €5 million by producing the high output level. Therefore, the threat to produce a high level of output if entry occurs is not credible.

Knowing that this threat will not be carried out, General will expect to earn €6 million by entering and will choose to do so. The only perfect equilibrium (that is, the only strategy pair that satisfies both our Nash condition and our credibility condition) is for General Generic to play "*enter*" and Liege Pharmaceutical to play *if General Generic chooses "enter", we will choose "low output", and if General Generic chooses "stay out", we will choose "high output"*. In the unique perfect equilibrium, General Generic enters the market and Liege Pharmaceutical produces the low output level.

Note the irony in this situation. When an incumbent has the ability to collude with the entrant in setting post-entry output levels, that ability may make the incumbent worse off. The reason is that the potential entrant will take the possibility (and profitability) of collusion into account when deciding whether to come into the market. Knowing that the post-entry equilibrium will be a collusive one makes entry attractive. Even under the fully collusive duopoly outcome, however, the incumbent earns less than it would if it had retained its monopoly position.

Credible Threats and Commitment

In the example above, the incumbent would like to scare off the entrant by threatening to produce "high output" in response to a choice of "enter". This threat is not credible, however, and we would expect the entrant to come into the market. This is an example of a phenomenon that was present in our earlier discussions of cartels and oligopolistic behaviour: in many instances, one firm would like to threaten another, but incredible threats will be ignored. In this section, we will see how a firm can take actions to make otherwise incredible threats credible. Taking such actions is known as engaging in commitment. **Commitment** is the process whereby a firm irreversibly alters its payoffs in advance so that it will be in the firm's self-interest to carry out the threatened (or promised) action when the time comes. In this example, the incumbent would like to commit itself to responding to entry by producing "high output".

commitment The process whereby a player irreversibly alters its payoffs in advance so that it will be in the player's self-interest to carry out a threatened (or promised) action when the time comes.

What sorts of actions could Liege Pharmaceutical take to make this threat credible? One possibility is for the incumbent to incur sunk expenditures to construct a large plant with low marginal costs, so that producing the high level of output is the profit-maximizing response to entry. Suppose with the large plant, the payoffs look like those in Figure 16.5. Once this plant has been constructed, the incumbent's threat to produce "high output" in response to "enter" is credible. If entry occurs, it is in Liege Pharmaceutical's self-interest to choose "high output", which yields a payoff of €4 million, instead of choosing "low output", which results in a profit of only €3 million. Thus, in the unique perfect equilibrium, the potential entrant decides to "stay out" and the incumbent chooses "high output". The equilibrium strategies are "*stay out*" for General Generic and *produce "high output" no matter what General Generic does* for Liege Pharmaceutical.

We have seen what would happen if the incumbent had a small plant (Figure 16.4) and what would happen if it had a large plant (Figure 16.5). Now suppose that, before entry can occur, the incumbent gets to choose which size plant to construct. Figure 16.6 illustrates the resulting game tree. The figure is formed by joining the trees from the two figures for the small and large plants. How should Liege Pharmaceutical expect General Generic to react to Liege's plant choice? And when General Generic chooses whether to enter the market, what reaction should it expect from Liege Pharmaceutical? We can answer these questions by extending our earlier procedure for finding a perfect equilibrium – we work backwards through the tree.

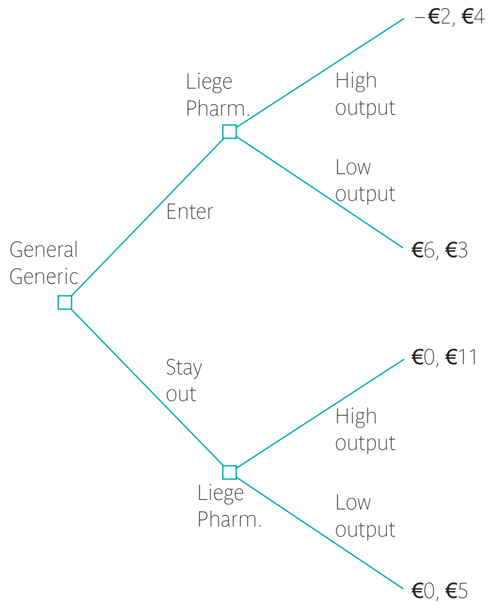

Figure 16.5 An Entry Game in which the Incumbent has a Large Plant

Once the large plant has been constructed, the incumbent's threat to produce "high output" in response to "enter" is credible – it yields Liege Pharmaceutical €4 million, whereas "low output" yields €3 million. The equilibrium strategies are *"stay out"* for General Generic and produce *"high output" no matter what General Generic does* for Liege Pharmaceutical.

We start by finding the last decisions in the game. Fortunately, we have already done most of the work. From our analysis of the two games in which the plant size was fixed, we know that General Generic will "enter" and Liege Pharmaceutical will produce "low output" if the incumbent constructs the small plant. We also know that General Generic will "stay out" and Liege Pharmaceutical will produce "high output" if the incumbent constructs the large plant. Because it earns €11 million from choosing "large plant" and only €6 million from choosing "small plant", Liege Pharmaceutical chooses "large plant". In equilibrium, Liege Pharmaceutical chooses "large plant", General Generic chooses "stay out", and Liege produces "high output".

Note that, *given that entry does not occur*, the incumbent would rather have the smaller plant. The €12 million payoff at the end of the branch (small plant, stay out, high output) in Figure 16.6 is larger than the €11 million payoff at the end of the branch (large plant, stay out, high output). The incumbent is not being irrational, however – the plant had to be built to keep the potential entrant out. The incumbent does better with the large plant and no entry than it would do with the small plant and entry (€11 million is greater than €6 million). The extra cost associated with the large plant can be thought of as an investment in entry deterrence.

Building a large plant is not the only way for the incumbent to deter entry by committing itself to an aggressive response. The incumbent might instead invest in cost-reducing research and development (R&D), and in that way make producing the high output level a rational response to entry. Or the firm might sign contracts with existing customers that legally bind the

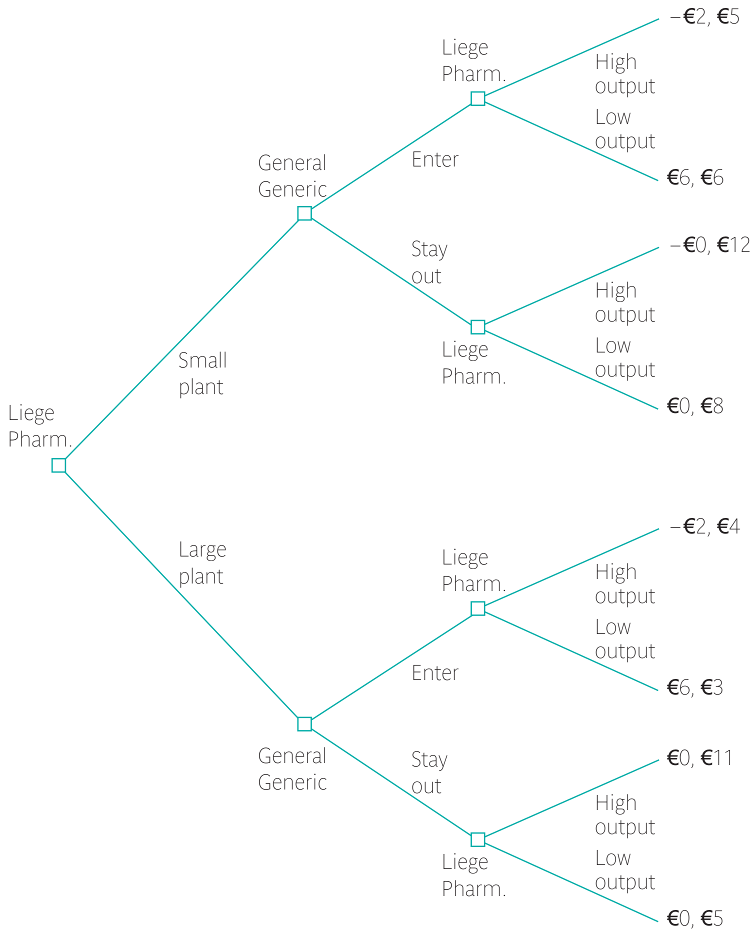

General Generic's payoff is given first in each pair (all figures in millions)

Figure 16.6 commitment in an Entry Game

When the incumbent can choose plant size, Liege Pharmaceutical can use a large plant as a form of commitment to deter entry. In equilibrium, Liege Pharmaceutical constructs the "large plant", General Generic chooses "stay out", and Liege produces "high output".

firm to match any future offer made by an entrant, thus committing the incumbent to fight any attempt by an entrant to take customers away.

More on Strategic Investment in Oligopoly

We have discussed the use of a strategic investment as a form of commitment in the context of a single incumbent facing the threat of entry. Strategic investment effects are present in markets where there are multiple incumbents as well, even if no additional entry is anticipated. When a firm engages in cost-reducing R&D, the firm will enjoy lower production costs in the future. A fall in production costs will increase profit for several reasons. First, the firm will have lower costs for any given level of production. Thus, at its old output level, the firm's profits would rise.

Second, the firm may adjust its output level to reflect its new cost structure. When marginal costs fall as a result of R&D, the firm will expand its output until it is at the point where the marginal revenue curve crosses the new, lower marginal cost curve. Because, over the range of expansion, marginal revenue is greater than marginal cost, profit increases as the firm expands. These first two effects arise whether the firm is a perfect competitor, a monopolist or an oligopolist.

When the firm is an oligopolist in a strategic market situation, the reduction in marginal cost has a third effect. The firm whose marginal costs have fallen has an increased incentive to produce output. The other firms in the industry must take this fact into account when they choose their output levels. Consequently, the other firms may cut back their output levels in the face of a more aggressive (lower-cost) rival.[4] The reduction in its rivals' output drives up the price that the firm conducting the R&D receives for its output. Thus, the firm counts this "strategic effect" as one of the benefits of R&D. In both an oligopolistic market and a market with a single firm facing the threat of entry, strategic investment can commit a firm to behave more aggressively, which then makes its rivals (the other oligopolists or the potential entrant) retreat (PC 16.3).

Playing Cards to Win

While strategic investment is often manifest in research and development that ultimately aims at reducing marginal costs, there are occasions where it actually pays the firm to increase its costs if it intends to increase customer numbers and loyalty to the firm. How can this be the case? An example from the UK food retailing sector perhaps illustrates this most clearly.

In the early 1990s the UK food retail sector exhibited elements of our textbook definition of oligopoly with a few large firms dominating the market and with growing strength of the brand name of the companies. The market leader was J Sainsbury with a share of about 16 per cent, followed by Tesco with, a share of about 12 per cent. This rival relationship between the two had developed in the 1980s and over time while both had increased their individual share of the market, the relative gap never seemed to close despite Tesco offering discounts, sales and so forth.

However, a step change occurred in 1995 that altered all this. Tesco undertook a major decision to introduce a loyalty card called the "Tesco Clubcard" which gave customers a point for every £1 (€1.25) spent in the store plus extra points on promotional items. The aim of this scheme was two-fold. First, it gave an added incentive for customers to go to Tesco because the points could be redeemed at a later stage for a range of goods and services. By this means Tesco hoped to increase its customer base. Second, as people joined the scheme, data on their shopping patterns could be used by Tesco to target promotions and other ideas which could be mailed to customers' home addresses. Such a scheme was not cheap to establish or run and there was a risk that Tesco could lose money from the clubcard if customers did not respond.

In terms of our analysis, this was a one-stage game with Sainsbury facing a choice – should it copy its rival straightaway and spend money on its own card (but with a risk of losing if the scheme didn't work), or should it ignore its rival and hope that it would gain if customers, seeing the Tesco scheme as an expensive gimmick, instead moved to Sainsbury who could cut prices as an added inducement to them?

Crucially, Sainsbury decided that the scheme would not be attractive and decided against implementing a rival version[a]. In fact, the scheme was a success and with it Tesco started to grow more rapidly than Sainsbury and indeed overtook the long-time market leader in late 1995. Despite Sainsbury finally beginning its own scheme in late 1996, Tesco has not moved from being the market leader since that time. While Sainsbury now has a market share that is still about 16 per cent of the market, Tesco has seen its share grow to over 30 per cent of the market. While this cannot be entirely attributed to the clubcard, the initial growth it created enabled further developments to take place using this momentum – for example, the diversification out of food into non-grocery items that were again targeted at "typical" Tesco customers on the basis of their shopping habits found from the clubcard data.

[a] Based on Andrew Seth and Geoffrey Randall (1999), *The Grocers: the rise and rise of the supermarket chains*, London: Kogan Page.

16.3 Games of Imperfect and Incomplete Information

Thus far we have examined game trees in which the firms move one after the other, and each firm can observe all of the earlier actions (if any) taken by its rival. Real life need not be so neat. When it is a firm's turn to choose a price or output level, it may not be able to tell what its rival has done or is doing. One reason that a firm may not know about its rival's move is that the two firms make their choices simultaneously. A second reason is that even though the other firm has already chosen its action, the firm choosing second is unable to observe the first firm's decision before having to make its own choice. A game in which some player must make a move but is unable to observe the earlier or simultaneous move of some other player is said to be a **game of imperfect information**.

game of imperfect information
A game in which some player must make a move but is unable to observe the earlier or simultaneous move of some other player.

Progress Check 16.3

Do strategic investment effects arise under perfect competition? Under monopolistic competition? How about monopoly with blocked entry?

The games we considered so far also assumed that each player knew everything there was to know about the other one. In our duopoly game, Air Lion predicted what Beta Airlines would do by thinking about which actions would be in Beta's self-interest. But it could be the case that Air Lion has only an imprecise idea of what Beta's payoffs look like. For instance, Air Lion might be unsure of Beta's cost level. Whenever one or more of the players is unsure about some part of the tree (such as the other player's payoffs), the situation is said to be a **game of incomplete information**.

game of incomplete information
A game in which some player is unsure about some of the underlying characteristics of the game, such as another player's payoffs.

You can think about the difference between these two new types of games as follows. In a game of imperfect information, a player is unsure about an earlier move made by someone else – the player, when it is his or her turn to move, is not sure exactly where he or she is in the tree. In contrast, in a game of incomplete information, the player is not sure exactly what the tree looks like. In this section, we will see how to extend game theory to deal with games of imperfect information and with games of incomplete information. In doing so, we will greatly increase the set of real-world situations into which game theory can provide valuable insights.

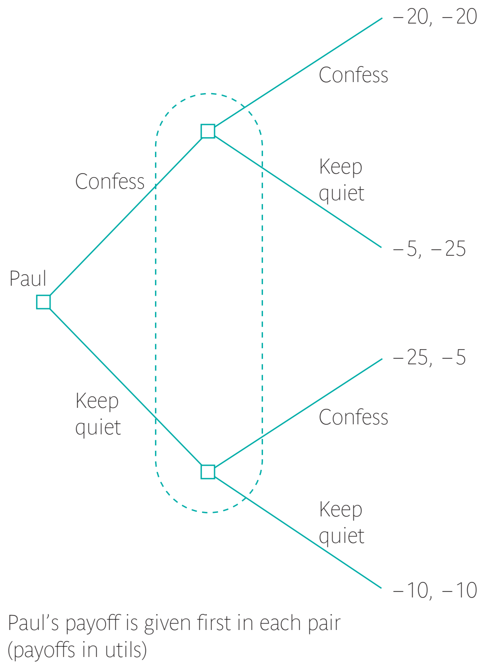

Figure 16.7 The Prisoners' Dilemma

When Simon and Paul make their decisions simultaneously, the dashed oval around Paul's two decision nodes represents the fact that Paul cannot distinguish between these two points at the time he makes his decision. For each robber, "*confess*" is a dominant strategy. The unique equilibrium outcome is for both to confess, even though both would be better off if they both kept quiet.

The Prisoners' Dilemma: a Game of Imperfect Information

Consider this. Two bank robbers, Simon and Paul, carry out a number of robberies, including one on a big bank as well as several on smaller banks. Unfortunately for them, they are caught by the police and put in separate interrogation rooms. The public prosecutor has enough evidence to convict both Simon and Paul for one of the small bank robberies, but she wants to convict them of the more serious and high-profile charge, that of the big bank robbery. For this she needs additional evidence. The public prosecutor goes to Paul and offers him a deal: a reduced prison term in return for testifying against Simon (which will increase Simon's prison term). *Simultaneously*, the public prosecutor's assistant goes to Simon and offers the same deal if he will turn in Paul. Each inside trader must choose between "confess" and "keep quiet". Figure 16.7 illustrates the tree for this game of imperfect information. The dashed oval around the two decision nodes for Paul is used to represent the fact that Paul cannot distinguish between these two points at the time he makes his decision. In other words, he cannot see whether Simon has chosen to "confess" or "keep quiet". Since he cannot see what Simon has done, Paul cannot make his choice contingent on Simon's. Thus, Paul must choose between the strategies, "*keep quiet*" and "*confess*".

Once we have a way to represent Paul's lack of information about Simon's action, finding the equilibrium is straightforward. Notice that the payoffs are measured in utils, so a longer prison term means a lower utility level, *ceteris paribus*. For each player (robber), "*confess*" is a dominant strategy. The unique equilibrium outcome is for both to confess, even though both bank robbers are better off when they keep quiet than when they confess (PC 16.4).

Progress Check 16.4

Suppose that instead of sending her assistant, the public prosecutor goes to see Simon herself, and *afterwards* she offers a plea bargain to Paul. Moreover, suppose that the public prosecutor refuses to tell Paul what Simon has done (Paul would not believe the public prosecutor anyway). How would the game tree for Paul and Simon change?

Based on this sort of story, such a situation – in which the two players each have a dominant strategy, but playing these strategies leads to an outcome in which both sides are worse off than if they collectively chose other strategies – has become known as the **prisoners' dilemma**, even if the players are not literally prisoners. The prisoners' dilemma structure applies to many important situations. Let's go back to our two airlines. Suppose that Air Lion and Beta Airlines have to make their output choices simultaneously. Figure 16.8 illustrates the game tree for this game of *imperfect* information. Again, the dashed oval around the two decision nodes for Beta is used to represent Beta's inability to distinguish between these two points at the time that it makes its decision: that is, when choosing its output level, Beta does not know if Air Lion has chosen "high" or "low". Therefore, Beta *cannot* adopt a strategy such as *produce "low" if Air Lion produces "high" and produce "high" if Air Lion produces "low"*, in which Beta's choice of action is contingent on what Air Lion has chosen. Rather, Beta's strategy is simply either *produce "high"* or *produce "low"*.

prisoners' dilemma
A strategic situation in which the two players each have a dominant strategy, but playing this pair of strategies leads to an outcome in which both sides are worse off than they would be if they co-operated by playing alternative strategies.

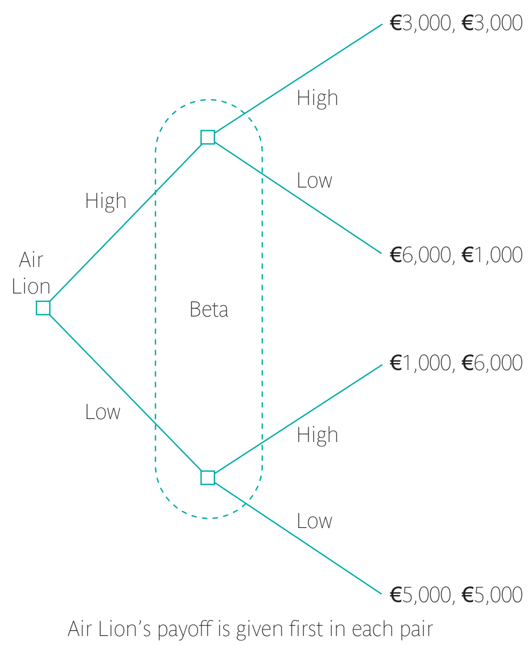

Figure 16.8 The Duopolists' Dilemma

When Air Lion and Beta make their output choices simultaneously, the dashed oval around Beta's two decision nodes represents the fact that Beta cannot distinguish between these two points at the time that it makes its decision. Since *produce "high"* is the dominant strategy for each firm, the unique equilibrium in this game is for Air Lion and Beta each to choose *produce "high"*.

For each firm, the better strategy, no matter what the other firm does, is to choose the high level of output. In other words, *produce "high"* is a dominant strategy for each firm. In the resulting dominant strategy equilibrium, each firm chooses "high" and earns a profit of €3,000. Notice how the structure of the payoffs in this game gives rise to a tension between the gains from co-operation and the incentives to compete. In equilibrium, each firm produces "high" even though each firm would be better off if both produced "low" – each firm would earn a profit of €5,000 instead of €3,000.

If the firms could sign a binding agreement enforced by a third party, we would expect them to choose the outcome under which each firm chose "low". But the firms cannot rely on the courts to enforce their agreement. Instead, they have to rely on self-enforcing agreements. The problem is that an agreement to produce "low" is not self-enforcing. To see why, suppose the two airlines agreed to have each produce "low". If Beta expected Air Lion to choose "low", Beta would have an incentive to cheat by producing "high" – Beta would earn €6,000 rather than €5,000. But Air Lion would never stick to the agreement to produce "low" in the first place; producing "high" is more profitable. The only self-enforcing agreement leads to both firms' producing "high". This game and the associated equilibrium constitute yet another demonstration that co-operation among self-interested parties may be difficult to achieve; the incentives to cheat may prevent the firms from enjoying the potential gains from co-operation. As you can see, the oligopolists' problem has the same structure as the inside traders' problem. Consequently, oligopolists often are described as facing a prisoners' dilemma.

Mixed Strategies

When we left them, Trezuguet and Buffon were trying to figure out which way Trezuguet was going to kick the ball in the World Cup Final. Because the goalie cannot wait to see which way the kick is going, this is a game of imperfect information. Figure 16.9 illustrates a game tree for this situation. The numbers are in utils and are chosen solely to capture the fact that Trezuguet wants to score and Buffon wants to stop him.

Now, let's look for an equilibrium. Suppose Trezuguet kicks to the left. Then the goalie should go left. But if the goalie goes left, then Trezuguet should kick to the right. So, it cannot be an equilibrium for Trezuguet to go left. But the same logic says that Trezuguet cannot go right. Trezuguet always wants to go the opposite direction of the goalie, but the goalie always wants to go the same direction as Trezuguet. They cannot both be satisfied at once. We have just shown that when both players' strategies are simply *always go "right"* or *always go "left"*, there is no equilibrium.

Always go "right" and *always go "left"* are each an example of what is known as a **pure strategy**. When you play a pure strategy, you have a definite action that you will take each time it is your turn to move. Another possibility is to randomize among the actions you will take at a particular decision node. For instance, Trezuguet can pursue a strategy of going right 30 per cent of the time and left 70 per cent of the time. When a player in a game randomizes across actions like this, he is said to be pursuing a **mixed strategy**.

> **pure strategy** A strategy that specifies a specific action at each decision point.

> **mixed strategy** A strategy that allows for randomization among actions at some or all decision points.

Although the penalty shoot-out has no equilibrium in pure strategies, there is an equilibrium in which the players choose mixed strategies. Let's find it. Suppose Trezuguet kicks right 70 per cent of the time. Then the goalie should go left all of the time, because that gives him the greatest chance of stopping the ball. But then

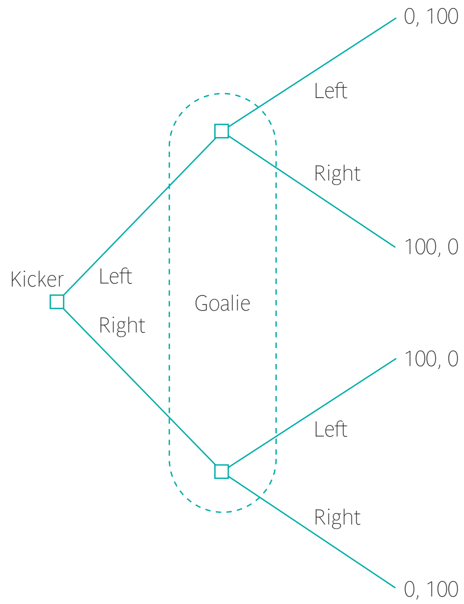

Kicker's payoff is given first in each pair (payoffs in utils)

Figure 16.9 Equilibrium in a Penalty Shoot-Out Entails Mixed Strategies

The goalie always wants to go the same way as the kicker, but the kicker always wants to go the opposite direction from the goalie. Hence, there is no equilibrium in pure strategies. There is, however, an equilibrium in mixed strategies: each player randomly goes "left" half of the time and "right" the other half. Neither player can make himself better off by switching his strategy, given the strategy being played by his opponent.

Trezuguet should always kick right. So kicking right 70 per cent of the time cannot be part of an equilibrium. Indeed, you should convince yourself that any percentage other than 50/50 is subject to the same problem: the goalie should always go the way the ball is kicked more often, but then the kicker should aim the other way.

Now suppose that Trezuguet kicks right half the time and left the other half. It doesn't matter which way the goalie goes, he will be correct half of the time no matter what. Now, if the goalie tends to go one way more than the other, the kicker will want to go the other direction all of the time. But then we are back to the problem above. Thus, the only candidate equilibrium strategy for the goalie is *go "right" half of the time and go "left" half of the time*. When this is the goalie's strategy, it doesn't matter which way Trezuguet chooses to kick.

What we have just shown is this. If one player randomizes 50/50 between "right" and "left", then it is a best response for the other player to randomize 50/50 between "right" and "left". In other words, each player's choosing to randomize equally between right and left is a Nash equilibrium (PC 16.5).

So what did Trezuguet do? He mishit the ball and kicked it over the goal, giving Italy the vital advantage in the shoot-out which they took and went on to win their fourth World Cup. Game theory is good stuff, but you still have to execute.

Mixed strategies are important in many other sports. For instance, in the choice of when to serve wide or down the middle is an important part of the tennis strategy. Whether to take the lead in a 1,500 metre race or to tuck in behind another runner is a strategic choice for each competitor. And in business rivalry it often is advantageous to keep your rivals guessing about what you are planning to do.

Progress Check 16.5

There actually are three choices: right, left or middle. Draw a tree for this game when the kicker and the goalie each have three choices. Describe the sets of possible strategies for each player. Explain why the equilibrium must be $\frac{1}{3}$ each way.[5]

These examples notwithstanding, mixed strategies strike many people as odd. If the other players all *believe* you are randomizing according to the theory, there is no need for you actually to go to the trouble. From your perspective, any one action you are supposed to randomize over is as good as any other (otherwise you would not be willing to randomize among them). So, why would anyone ever bother tossing a coin, rolling a die, or in some other way generate the randomness that game theory calls for?

In the light of this troubling question, how should we think about mixed strategies? It probably makes the most sense to think of them not as actual random strategies, but as a representation of other players' *beliefs*.[6] The model does not require that you literally toss a coin when choosing whether to serve wide or down the middle in tennis. Rather, you try to avoid falling into a predictable pattern, with the net result that the other player's *beliefs* are the same as they would be if you truly were randomizing. While mixed strategies may seem a bit strange to you, they capture the intuitive notion that it can be to one's strategic advantage to keep rival players guessing about what your future moves will be.

A bargaining game of incomplete information

Now, let's consider a game of incomplete information. Suppose there is a single seller bargaining with a single buyer. The seller, Liam, can produce one custom cuckoo clock for €1,000. The way the bargaining works is that Liam makes a take-it-or-leave-it offer to the buyer. The buyer then "accepts" or "rejects" the offer. If Liam knew the buyer's willingness to pay for the clock, he would set his price at that level, as long as it was at least €1,000. But what happens when the seller is unsure about the buyer's willingness to pay? In particular, he knows the buyer values the item at either €1,500 or €2,000, but he is unsure which. Thus, he does not know whether to set the price at €1,499 or €1,999.[7]

At first glance, it is not obvious that a single game tree can be drawn to represent this situation. There is one set of payoffs (and associated game tree) for a low-value buyer and another for a high-value buyer. Liam would not know which tree to use. Fortunately, we can model this situation as a single tree if we simply combine two earlier tricks. Figure 16.10 presents a game tree for this situation. As in decision trees, we can use a move by Nature to represent a player's uncertainty about some or all of the parameters of the game. Here, the seller's uncertainty about the buyer's valuation of the good is captured by having Nature make an unobservable move to pick the buyer's valuation. Of course, Liam cannot see what move Nature has made. Viewing Nature as a special player, we have a game of imperfect information and we draw a dashed oval around Liam's two decision nodes to represent the fact he cannot distinguish between them at the time he makes his decision.

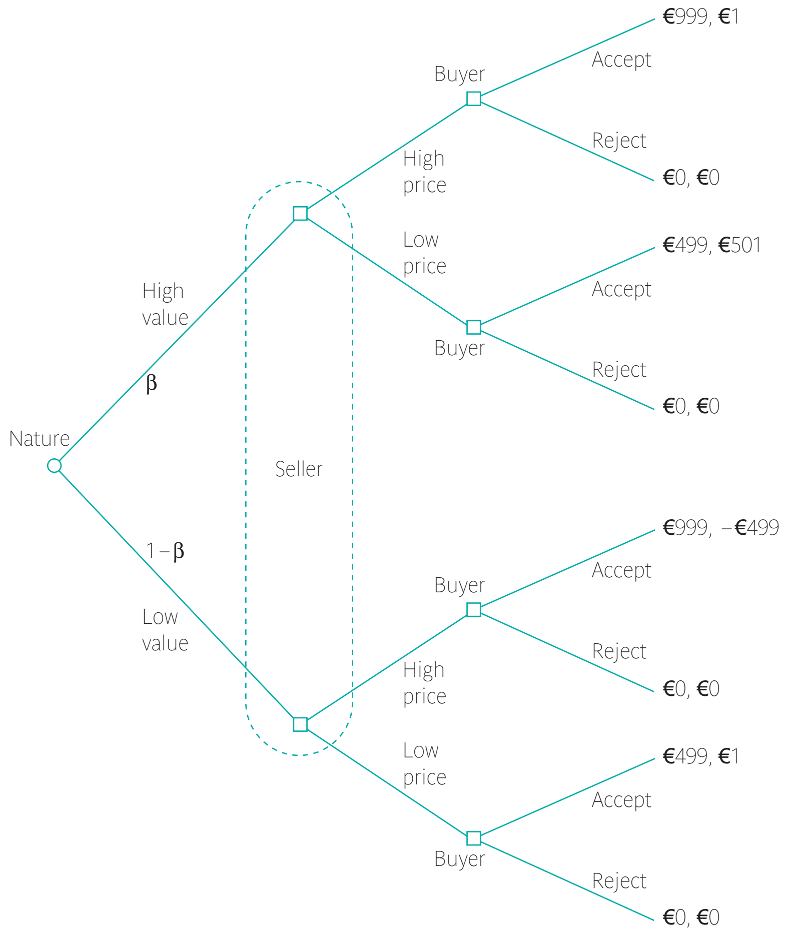

Figure 16.10 Bargaining under Incomplete information

The seller's uncertainty about the buyer's valuation of the good is captured by having Nature make an unobservable move to pick the buyer's valuation of the good. Nature chooses "high value" with probability β. Since the seller cannot see the move that Nature has made, we draw a dashed oval around his two decision nodes to represent the fact that he cannot distinguish between them. The expected profits from setting the high price are $\beta \times$ (€1,999 − €1,000) + $(1 - \beta) \times$ €0 = $\beta \times$ €999. Setting the low price yields profits of €499. Hence, the seller will set the high price if $\beta \geq 499/999$, and will set the low price otherwise.

To decide what price to charge, Liam has to form a prediction about how the buyer will react to his offer. When he sets the higher price, Liam's prediction will depend on whether he believes the buyer places a low or high value on the clock. If the buyer places the low value on the item, then the buyer will "reject" an offer of €1,999. But if the buyer places a high value on the clock, she will "accept". We reach this conclusion by working backward through the tree. Faced with a price of €1,999, a high-value buyer will act in her self-interest by choosing to "accept". And if the buyer places a low value on the good, she will respond to a price of €1,999 by choosing "reject". Similar reasoning establishes that both types of buyer will "accept" an offer of €1,499.

The seller has to decide whether to set a low price and make a sure sale or set a high price and make a sale only if the buyer turns out to place a high value on the clock. Setting the low price yields profits of €499 (= €1,499 − €1,000). To calculate the profitability of setting a high price, the seller has to form beliefs about Nature's move (that is, about the relative likelihood of the low and high valuations). Let β denote Liam's belief about the probability that the buyer has a high value for the clock, and let $1 - \beta$ denote his belief about the probability that the buyer places a low value on it. The expected profits from the high price are

$$\beta \times (€1{,}999 - €1{,}000) + (1 - \beta) \times €0 = \beta \times €999 \quad \textbf{(16.1)}$$

Hence, Liam will set the high price if $\beta \times €999$ is greater than €499, and will charge the low price otherwise. In other words, Liam sets the high price only if he is sufficiently optimistic that the buyer places a high value on the clock. Specifically, he sets the high price only if he believes $\beta \geq 499/999$.

Limit Pricing: A Game of Incomplete Information

Let's consider another game of incomplete information. Once again, suppose that a firm is considering entry into what is at present a monopolized industry. But now suppose that the potential entrant is unsure of whether the incumbent's marginal costs are low or high. The potential entrant is interested in the incumbent's costs because they affect the incumbent's payoffs and thus influence the incumbent's optimal reaction to entry.

Figure 16.11 presents a game tree for this situation. The entrant's uncertainty about the incumbent is captured by having Nature make an unobservable move to pick the incumbent's marginal cost function. Because the potential entrant cannot see what move Nature has made, we have a game of imperfect information and we draw a dashed oval around the potential entrant's two decision nodes to represent the fact that the firm cannot distinguish between them at the time it makes its decision.

You may have noticed that the structure of the tree in Figure 16.11 looks much like that of the tree in Figure 16.6, where the incumbent chose its plant size. The key difference is that here Nature, not the incumbent, is making the initial choice. Unlike the incumbent, Nature does not optimize its strategy. Rather, Nature's strategy simply represents the players' beliefs about the likely state of the world.

To decide whether to come into the market, the entrant has to form a prediction about the post-entry equilibrium. This prediction will depend on whether the entrant believes the incumbent has high or low marginal costs. If the incumbent has low costs and the entrant comes in, then the incumbent will choose "high output" and the entrant will suffer losses of −€3 million. Again, we reach this conclusion by working backwards through the tree – faced with entry, a low-cost incumbent will act in its self-interest by producing "high output". Similarly, if the incumbent has high costs, then it will respond to entry by choosing "low output", and the entrant will earn €4 million.

To assess the desirability of its entering, the entrant has to form beliefs about the relative likelihood of the high and low cost levels. Let p denote the entrant's belief about the probability that the incumbent has high costs, and let $1 - p$ denote the entrant's belief about the probability that the incumbent has low costs. The expected profits from entry are

$$p \times (€4 \text{ million}) + (1 - p) \times (-€3 \text{ million}) = €(7p - 3) \text{ million} \quad \textbf{(16.2)}$$

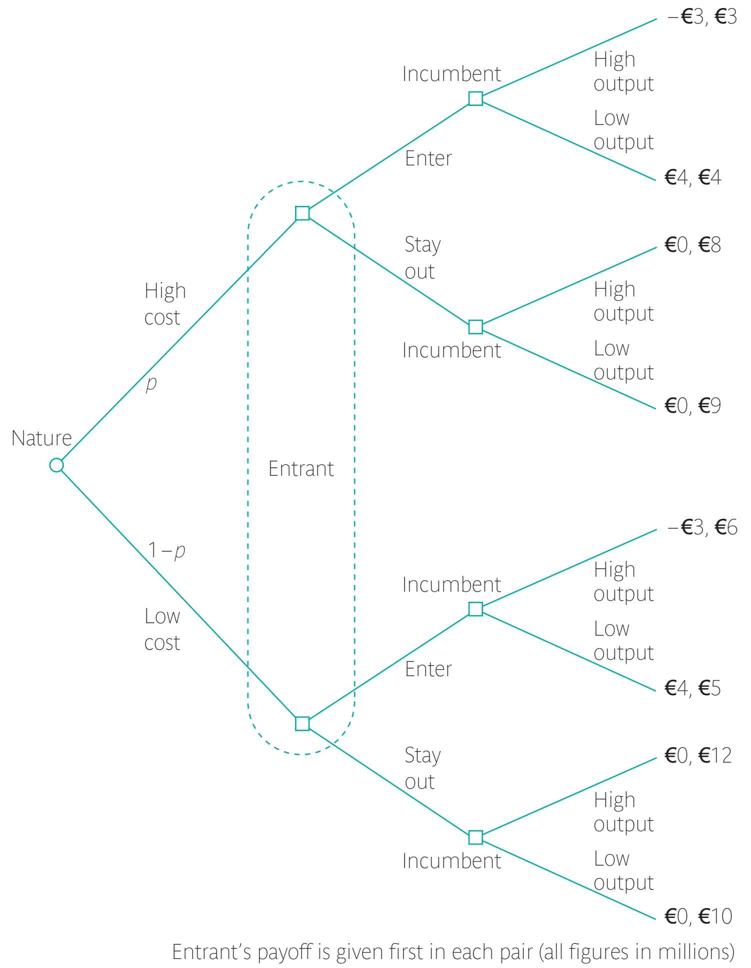

Figure 16.11 An Entry Game of Incomplete Information

The entrant's uncertainty about the incumbent's payoffs is captured by having Nature make an unobservable move to pick the incumbent's marginal cost level. Nature chooses "high cost" with probability P–. Since the potential entrant cannot see what move Nature has made, we draw a dashed oval around the potential entrant's two decision nodes. The expected profits from entry are $p \times$ (€4 million) + $(1 - p) \times$ (–€3 million) = €($7p - 3$) million. Staying out of the market yields €0. Hence, the potential entrant comes into the market if $p > \frac{3}{7}$, and stays out otherwise.

The expected profits from staying out of the market are €0. Hence, the potential entrant will come into the market if $p \geq \frac{3}{7}$ and will stay out otherwise. Intuitively, the potential entrant will come in if it is sufficiently optimistic that the incumbent has high costs and will be a weak rival.

The potential entrant would like to figure out what the incumbent's true costs are. If the potential entrant can observe some action by the incumbent prior to making the entry decision, the potential entrant may be able to draw inferences about the incumbent's underlying costs. Figure 16.12 illustrates an extended version of the game in Figure 16.11. Here, the incumbent chooses output in two periods. The key feature is that, before deciding whether to enter the market, the potential entrant gets to see the incumbent's initial output choice.

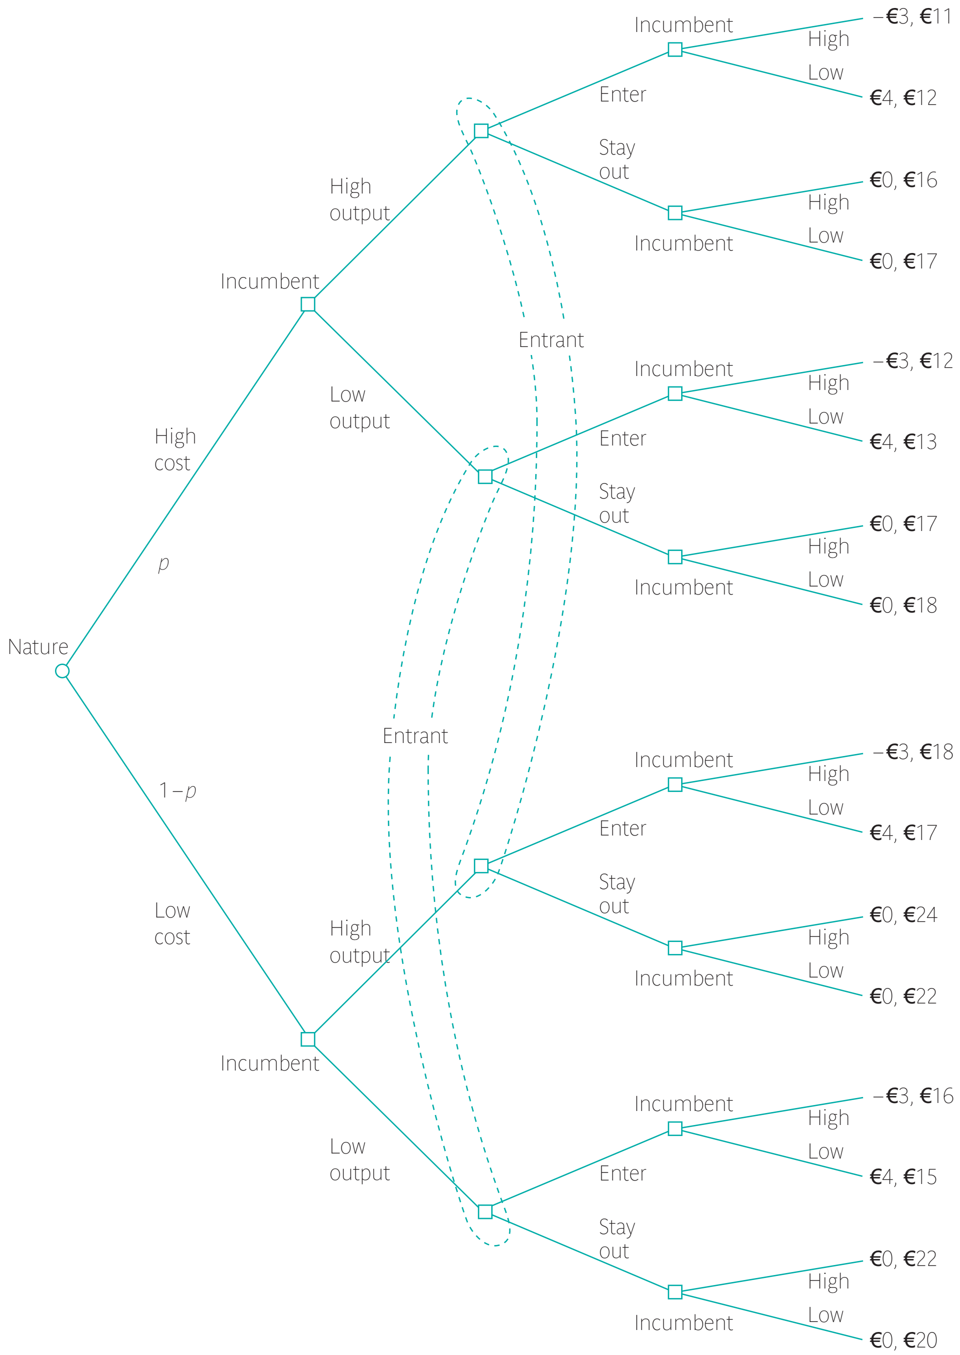

Entrant's payoff is given first in each pair

Figure 16.12 A Limit Pricing Game

The equilibrium strategies are as follows. The incumbent's equilibrium strategy is *produce "high output" in the first period; in the second period, produce "high output" if costs are low, and "low output" if costs are high*. The entrant's equilibrium strategy is *"enter" if and only if the incumbent produces "low output" in the first period*. Hence, the entrant stays out of the market.

Suppose that the entrant ignored the incumbent's first-period output choice. Further, suppose that $p = \frac{2}{7}$. In this case, the potential entrant would stay out. Knowing this fact, a low-cost incumbent would maximize its profit by choosing "high output" in each period, whereas a high-cost incumbent would choose "low output" in each period. The potential entrant would be foolish to ignore the incumbent's first-period output choice. By looking at the incumbent's initial choice, the potential entrant could infer the incumbent's cost level. If the potential entrant believed that the incumbent was following the strategy just outlined, then the potential entrant should adopt the strategy *"enter" if and only if the incumbent chooses "low output" in the first period*.

We are not done yet, however. The incumbent should account for the fact that the entrant is making such inferences. Suppose that both the high-cost and low-cost incumbent choose "high output" in the first period. Then the entrant would be able to draw no inference from the first-period output level. Given the low value of p, the potential entrant would stay out of the market. Thus, in making its first-period output choice, a high-cost incumbent compares the payoff at the end of the branch (low output, enter, low output) with the payoff at the end of the branch (high output, stay out, low output). Since €17 million is more than €13 million, it is more profitable to choose "high output" in the first period than to choose "low output".

The net result of all this is that the equilibrium strategies are as follows. The incumbent's equilibrium strategy is *produce "high output" in the first period; in the second period (whether or not entry occurs) produce "high output" if costs are low, and "low output" if costs are high*. The entrant's equilibrium strategy is *"enter" if and only if the incumbent produces "low output" in the first period*.

Notice how the high-cost incumbent distorts its behaviour to conceal its true cost level from the potential entrant. You might think that the incumbent is trying to trick the potential entrant into thinking that the incumbent is a low-cost firm. But a rational entrant anticipates that a high-cost incumbent will do this. The value to the incumbent is not that the entrant is fooled into thinking that the incumbent has low costs. Rather, the advantage is that the entrant is prevented from obtaining the information about costs that he or she would like and, given the lack of information, chooses to stay out of the market. The practice of setting a high output level or a low price to deter entry is known as **limit pricing**.

> **limit pricing**
> The practice of setting a high output level, or a low price, to deter entry.

This example is just one of many situations in which a potential entrant looks at the incumbent's actions to make inferences about some underlying characteristic of the incumbent or its market. In the example, the entrant looks at the incumbent's price and output in order to infer the incumbent's cost level. In other cases, a potential entrant might want to forecast market growth. Since the incumbent has experience in the industry, its prediction of market growth might be particularly valuable to the entrant. One way for the potential entrant to obtain this information may be to look at what investment in new capacity the incumbent is making.[8]

Trying to Prevent a Net Gain[a]

In markets which are growing and where there are significant profits to be made, it is not surprising that new firms wish to enter and share in those profits. Equally unsurprising, those firms already in such markets would like to keep the profits for themselves and, the larger the firm is in relation to the overall market, the easier they find it to do this, but it must be done legally.

However, if the advantage of economies of scale cannot be found then, as we have seen in this chapter, one way to keep rivals out is to maintain prices at levels that make it unprofitable for new firms to enter. This policy would of course reduce the incumbent's own profits if the price charged was below the profit-maximizing level. It is a short-term policy to gain a long-term advantage by keeping potential competitors out of the market.

Take the case of Wanadoo and its provision of internet services in France. Wanadoo is a subsidiary of France Telecom and had a dominant position in the French internet market even before the market for high-speed provision began its rapid growth and expansion. Wanadoo wished to launch its new service, extense, supplying high-speed internet services, but to do so would cost it significant sums. It was keen therefore to ensure that it could earn revenues large enough to make a positive return on its investment. How could this be done, knowing that other firms would be attracted into the industry by large profits?

One option was simply to be better and cheaper at supplying than competitors but . . . it could not guarantee this even if it knew lots of its competitors' cost and output information, which is highly unlikely. An alternative was to use its current market position to keep competitors from entering the more lucrative market. Wanadoo decided to pursue a policy of selling its ADSL service – the basic internet service – at prices below average costs from the end of 1999 to October 2002. In effect, this is similar to the limit pricing/predatory pricing we have seen earlier in the chapter. Why did it do this?

This is a key question as, during the period in question, the firm made significant losses. It did so in order to maintain its leading position for the new high-speed internet market. As the European Commission said in its report on the investigation of this predatory pricing policy, between January 2001 and September 2002, Wanadoo's market share rose from 47 to 72 per cent while the market grew five-fold. At the same time, the extent of financing needed by other firms to compete grew hugely and consequently no other firm had more than a 10 per cent share of the market. The investigation only come about when the prices in the wholesale market for internet services were dropped by France Telecom and the discrepancies between costs and prices became apparent. The Commission found Wanadoo guilty and fined them €10.35 million as a result.

What this case shows is that, as with cartel behaviour, even though an activity is illegal, there can be a strong incentive for a firm to resort to it as the profits that can be generated are significant. Equally, though, it shows that the expected gain can swiftly be wiped out by fines and regulatory intervention, so as to make the original decision appear unwise.

[a] Based on "High Speed Internet: the Commission Imposes a Fine on Wanadoo for Abuse of a Dominant Position", Communiqué de presse, European Commission, IP/03/1025, 16 July 2003.

16.4 Repeated Games

While decisions such as whether to enter a market involve one-time actions, there are many situations in which players find themselves repeatedly making the same decisions. For example, in the previous chapter we discussed informally the interaction of oligopolists who chose new output levels or prices each day. In this section we look in detail at one model of repeated price setting to examine carefully the costs and benefits of cheating. This model also demonstrates the dramatic way in which repeated oligopolistic interaction can affect the equilibrium outcome relative to situations in which firms make once-and-for-all choices. Let's reconsider the attempt of Air Lion and Beta Airlines to reach a self-enforcing agreement to hold up the price of a ticket. Now, however, let's allow for the firms' getting to choose new prices daily; specifically, at the start of each day the two firms simultaneously choose their prices for that day. The per day demand curve is $D(p)$. Each day's choice of prices constitutes a Bertrand game within the overall game. The overall game is made up of the repeated play of this component game, also known as a *stage game*.

Given this set-up, what equilibrium do we expect to emerge? One possibility is that each day the firms choose the same price that they would have done if they were each making only a single, once-and-for-all choice: that is, each day both airlines set their prices equal to the common value of marginal cost, c. To see that this is an equilibrium outcome, note that if each firm expects the other to set price equal to the common value of marginal cost, then it might as well do the same. Thus, there is no incentive to cheat on the agreement to set the price equal to marginal cost each day, and this agreement is self-enforcing (i.e. it is a Nash equilibrium). Unfortunately, from the firms' point of view, the agreement is not worth much – the firms earn zero economic profits.

Under the agreement just outlined, the firms do not take advantage of their repeated dealings with one another. One type of implicit agreement that does make use of the repeated interaction is known as a *grim-trigger* strategy. Suppose that the firms agree that each firm will charge price p_s daily as long as no one has cheated in the past (that is, p_s has always been charged in the past). If anyone does cheat, then the firms (including the cheater) "agree" to punish the cheater by setting all future prices equal to marginal cost. This type of strategy gets its name from the fact that the detection of cheating "triggers" an infinitely long punishment (a "grim" prospect indeed).

Let's analyse whether a grim-trigger strategy is self-enforcing and credible. To determine whether an agreement is self-enforcing, we need to know whether the benefits from cheating are less than the costs. As we saw in Chapter 15, the costs and benefits of cheating depend on the per day profits associated with sticking to the agreement, π^s, getting away with cheating, π^c, and getting punished, π^p.

If the firms have agreed to charge a price of p_s and neither firm cheats, then the two firms split the market sales of $D(p_s)$. Hence, the per day profit from sticking to the agreement is

$$\pi^s = \tfrac{1}{2}D(p_s) \times (p_s - c) \tag{16.3}$$

Now consider what a cheater earns each day that it escapes detection. Suppose that Beta abides by the agreement and sets its price equal to p_s. As we saw when analysing Bertrand duopoly, by just barely undercutting Beta's price, Air Lion can earn a per-day profit of approximately[9]

$$\pi^c = D(p_s) \times (p_s - c) \tag{16.4}$$

We will simplify the calculations below by making use of the fact that this value of π^c is exactly twice the value of π^s given by Equation 16.3.

Under the threatened punishment, a cheater earns no profit once it has been caught: $\pi^p = 0$. Before comparing the costs and benefits of cheating, we need to make sure that the threatened punishment is credible. Is the agreement to punish a cheater this way itself a self-enforcing agreement? Our earlier discussion showed that it is. Given that one firm prices at marginal cost in every period, it is in the other firm's interest also to price at marginal cost. Hence, if each firm expects the other to respond to cheating by setting price equal to marginal cost, it is in the firm's self-interest to do so as well. In short, the punishment is credible.

We are now ready to calculate the costs and benefits of cheating. Since the benefits accrue over time, we need to discount them. Let i denote the per day interest rate. A key determinant of these costs and benefits is how long it takes to detect cheating. Let's begin by analysing what happens when a firm can get away with cheating for only one day before being caught. The benefit of cheating is the present value of the extra profit that Air Lion earns during the time that it gets away with it. Since Air Lion goes undetected for just one day, the benefit is $\pi^c - \pi^s = \pi^s$ (recall π^c is twice π^s). The cost of cheating is the present value of the profits that are forgone when the firm is punished. Since the firm could have earned π^s each day but instead earns 0, the cost of the punishment in terms of forgone profit is

$$\frac{\pi^s}{(1+i)} + \frac{\pi^s}{(1+i)^2} + \frac{\pi^s}{(1+i)^3} + \ldots \qquad \textbf{(16.5)}$$

Applying our standard discounting formula for a perpetuity (see Chapter 5, page xxx), the cost of cheating is equal to π^s/i.

Comparing the benefit of cheating, π^s, with the cost, π^s/i, we see that unless Air Lion has an interest rate of at least 100 per cent per day, it will not cheat! When $i < 100$ per cent, these grim-trigger strategies constitute a perfect equilibrium that supports pricing at p_s. Remarkably, this result is *independent of the particular value of* p_s (as long as p_s is greater than c so firms do not suffer losses). This fact tells us that, if the firms can sustain a price 2 cents greater than marginal cost with a self-enforcing agreement, then they can sustain the full cartel price as well. What a difference repeat play can make! Bertrand firms, who make once-and-for-all price choices, set price equal to marginal cost and earn no profit. Firms that set prices repeatedly, however, can end up charging the monopoly price and splitting the monopoly profits. The difference arises entirely because Bertrand firms cannot punish one another for cheating, but firms that interact repeatedly can.

This example suggests collusion should be easy. But the example misses some important features of actual markets that we discussed in the previous chapter. One of the most important is that a firm may be able to escape detection from cheating for more than one day before being punished. To see the effects of such a punishment lag, suppose that a firm could get away with cheating for two days before being punished. In this case, a cheater gains $\pi^c - \pi^s$ for two days before being punished. Since $\pi^c - \pi^s = \pi^s$, getting away with cheating for a second day is worth $\pi^s/(1+i)$ in present-value terms. Hence, the present value of the total benefit of cheating rises to

$$\pi^s + \frac{\pi^s}{(1+i)} = \pi^5 \times \frac{(2+i)}{(1+i)^3} \qquad \textbf{(16.6)}$$

What about the cost of cheating? Once punished, the per-day decline in the firm's profit is again $\pi^s - \pi^p = \pi^s$. Now, however, the punishment does not begin for two days, so its present value is

$$\frac{\pi^s}{(1+i)^2} + \frac{\pi^s}{(1+i)^3} + \frac{\pi^s}{(1+i)^4} + \ldots \qquad \textbf{(16.7)}$$

Comparing Expression 16.7 with Expression 16.5, we see that the cost of cheating when there is a two-period lag is equal to the cost when there is a one-period lag divided by $(1 + i)$ to account for the fact that the punishment starts one period later. Dividing π^s/i by $(1 + i)$, the cost of cheating is now $\pi^s/(i(1 + i))$, which is lower than when detection takes only one day.

Since the benefits of cheating have risen, and the costs of cheating have fallen, it is harder to deter cheating. The comparison of the benefits of cheating, $\pi^s(2 + i)/(1 + i)$, with the cost, $\pi^s/[i(1 + i)]$, depends on whether $(2 + i)$ is greater than $1/i$. Now, the critical value of the per-day interest rate i is $\sqrt{2} - 1$, which is approximately 0.41. When a firm can get away with cheating for two days, collusion will be successful if and only if the per-day interest rate is less than 41 per cent. This still is an extremely high interest rate, but it is much lower than the earlier critical value of 100 per cent per day. If the lag between cheating and punishment were longer, then the incentive to cheat would be larger still, and collusion would fail at even lower interest rates.

However, grim-trigger strategies are not the only strategies available to firms. For instance, the punishment might vary over time in some complicated way. It is worth considering whether the firms could increase their profits by adopting more complicated strategies. As we have already discussed, the harsher the penalty, the greater the degree of collusion that can be supported. A first question, therefore, is can the firms find a more severe credible punishment? In our example, the answer is no. Under the proposed strategies, cheating is met with marginal cost pricing for ever afterwards, and the cheater's profit following detection is zero. No stronger punishment is available to the firms because any attempt to drive the cheater's profit below zero would simply lead to the cheater's shutting down. Moreover, no one would believe that the firm carrying out the punishment would set its price below marginal cost anyway. For the market setting that we have examined here, grim-trigger strategies provide the harshest possible credible punishment, and thus support the greatest degree of collusion possible.

More complicated strategies may yield harsher penalties in other circumstances, however, as when firms set their quantities repeatedly or when products are differentiated. Moreover, in a world where mistakes are possible, firms do not always want to impose the harshest possible punishment. It often is difficult for firms to verify whether their rivals are abiding by an agreement. Since agreements to collude often are tacit, misunderstandings can occur. Further, it may take the firms in an industry a while to figure out their optimal collusive arrangement. There will be no benefits from colluding if simple mistakes or the learning process itself trigger punishments that last for ever. To err may be human, but to forgive can be profitable.

Finitely Repeated Games

In the repeated pricing game we just examined, the firms set new prices each day for ever. Equivalently, the players know that the game will end some day, but they don't know in advance when that will be. Because there is no set end point, such situations are sometimes called *infinitely repeated games*. It turns out that the situation is much different when both players know that the game will end on a certain date.

Recall that to find a perfect equilibrium, we worked back from the end of the game tree through a process known as *backward induction*. We can use the logic of backward induction to derive a powerful result for situations with repeated price setting for a *fixed* number of periods, an example of a so-called *finitely repeated game*. Suppose that Air Lion and Beta both know that Beta is going to shut down exactly two years from now. How should they behave in the meantime? To answer this question, we start by thinking through what will happen on the last day Beta is in business. At that point, the firms will be facing a standard, one-time Bertrand pricing game. We know from Chapter 15 that each firm will price at marginal cost and neither firm will earn economic profits. What about the day before that? No matter what a firm does on that day, it knows that the next day will have pricing at marginal cost. So, there is no scope for threats or promises or any sort of punishment for failing to hold price above cost on the penultimate day. Because the future price is essentially set, a firm should choose its price on the next-to-last day as if there were no other day. But then we get the standard Bertrand outcome for this day, too.

This same logic applies to the third day from the end. Indeed, as we keep working backwards we find that the outcome in every period is equivalent to the one-time Bertrand outcome. Not only that, but this logic applies equally well to any other repeated game for which there is only one Nash equilibrium for the stage game played once in isolation. We have just shown that *when there is a unique Nash equilibrium of the stage game, the unique perfect equilibrium of the finitely repeated game is simply the one-shot equilibrium repeated in every period.*[10]

Chapter Summary

Nonco-operative game theory provides a set of tools for analysing oligopoly as well as strategic behaviour in many other areas of economics and politics.

- Game trees provide a convenient way to represent strategic situations.
- The notion of a perfect equilibrium captures the two features of equilibrium that we required for oligopoly: the Nash condition and the credibility condition.
- Perfect equilibrium embodies the Nash condition by imposing the requirement that each player choose an optimal strategy, given the strategies chosen by the other players in the game.
- Perfect equilibrium embodies the credibility condition by requiring that each player would find it to be in her self-interest to carry out any part of her chosen strategy if called upon to do so.
- A game of imperfect information is a situation in which some player must make a move but is unable to observe an earlier or simultaneous move of another player.
- A game of incomplete information is a situation in which a player is not sure about some characteristic of the structure of the game being played.
- The tools of non-cooperative game theory help us understand the process of entry into oligopolistic markets. Incredible threats by an incumbent firm will not serve as entry deterrents. An incumbent firm may be able to make investments that make its threats credible. A high-capacity plant, for example, may commit an incumbent to responding aggressively to entry. Thus, building such a plant may be an investment in entry deterrence. An incumbent firm may also invest in entry deterrence by setting a low price to make potential entrants worry that the incumbent may be a low-cost firm.
- When a player randomizes over his or her choice of action to keep rivals guessing, the player is said to follow a mixed strategy.
- When players play the same game within a game again and again, the overall game is called a repeated game. In comparison with the Bertrand model, the infinitely repeated pricing game illustrates the tremendous difference that the ability for players to respond to one another can make.

Discussion Questions

16.1 There is a theorem proving that there is a "solution" to chess; that is, there is a perfect equilibrium in which one of the players can guarantee itself at least a tie. While one can prove that such an equilibrium exists, no one knows what it is. Hence, international grand masters are able to make a great deal of money playing chess. What does this tell us about applying game theory to real-world situations?

16.2 It is often suggested that competing firms are engaged in a prisoners' dilemma when it comes to choosing their advertising expenditures.

(*a*) Draw a game tree to illustrate a duopoly market in which the suppliers must simultaneously choose their advertising levels. Put in payoffs that give this game a prisoners' dilemma structure. Does this pattern of payoffs seem to fit the real-world situation?

EU law forbids cigarette advertising on television, but this has not always been the case. Some observers have claimed that the ban may actually have raised cigarette industry profits by limiting costly advertising campaigns that largely cancel one another out.

(*b*) Is this story consistent with the model of industry behaviour that you developed in part *a*?

16.3 Until EU telecommunication markets were liberalized, regulation typically prevented additional firms from entering national markets. However, once liberalization occurred, firms could enter where they liked. Many management consultants concluded that, compared with each staying out of the other's market, telecommunication companies would end up losing money if each went into the other's market. Yet, these consultants also concluded that each firm would find entering the other's market too attractive to resist.

(*a*) Use the insights of the prisoners' dilemma to explain these apparently contradictory predictions.

(*b*) What difference does it make, if any, that each firm can monitor the other's entry decisions and thus make its own decision contingent on that of the other firm in its local market?

16.4 Consider again the entry game examined in Section 16.2. Contrary to the situation depicted in Figure 16.6, suppose that General Generic makes its entry decision (that is, constructs its own plant) *before* Liege Pharmaceutical makes its plant investment decision.

(*a*) Draw a game tree for this new situation.

(*b*) Describe the equilibrium. Will General Generic choose to enter? Which size plant will Liege Pharmaceutical choose to construct?

16.5 Go back to the situation depicted in Figure 16.6, but make one change. Suppose that prior to Liege Pharmaceutical's making its plant choice, General Generic can announce whether it intends to enter this market. The announcement is nonbinding and is costless to make.

(*a*) Draw a game tree for this new situation.

(*b*) Describe the equilibrium. Will General Generic announce that it intends to enter? Which size plant will Liege Pharmaceutical choose to construct? Will General Generic enter? In answering this question, be sure to state carefully what constitutes each firm's equilibrium *strategy*.

16.6 Auctions are used to sell a wide variety of goods, ranging from paintings to licences for offering wireless telephone service. Under a so-called *Dutch auction* (in the Netherlands, they use it to sell fresh flowers at wholesale to florists), there is an auctioneer who announces a very high price and then calls out successively lower prices. The first bidder to accept the auctioneer's price is declared the winner and receives the item for the price the auctioneer had last called out. Under a sealed-bid auction, each buyer submits a secret bid stating how much he or she is willing to pay. The bids are then opened and the item is sold to the highest bidder at the price submitted. Show that in terms of the strategies available to the players, what each player knows about the others, and thus what the equilibrium outcome will be, a Dutch auction is equivalent to a sealed-bid auction.

16.7 Consider the following market. There is a single incumbent and a single potential entrant. Each firm has constant marginal costs of c per unit, and the product is undifferentiated. To enter the market, the new firm would have to incur a one-time cost of €1,000. Resolve the following apparent paradox. If the post-entry game is a Bertrand duopoly, then neither firm will make any profits, whereas under a Cournot duopoly they would. The incumbent, however, would prefer to be in a situation where the post-entry interaction is Bertrand rather than Cournot.

16.8 The United States stations thousands of troops in South Korea. Ostensibly, these troops are there to help the South Korean military repulse any attempted invasion by North Korea. There are, however, far too few troops to do the job. Some say the real role of these troops is to serve as a commitment by the United States to come to South Korea's aid in the event of an invasion. Draw a game tree that captures the notion of the American troops' serving as a form of commitment.

16.9 Consider the game in Figure 16.12 one more time. Now suppose that the entrant believes there is a $\frac{6}{7}$ chance that the incumbent has high costs (that is, $p = \frac{6}{7}$). Show that, in equilibrium, the incumbent will not try to confuse the potential entrant by producing "high output" in the first period when it has high costs.

16.10 Consider the model of repeated price setting in Section 16.4. Suppose that a cheater could get away with three periods of cheating before being found out. Assume that even if the firm cheated for only one period, it would be discovered two periods later.

(*a*) Is there a zero-profit equilibrium in this market?

(*b*) Can the full cartel outcome be supported by a self-enforcing agreement if $i = 0.2$? What if $i = 2.0$? What if $i = 20$? Find a critical value for the interest rate that determines whether or not the cartel outcome can be supported.

16.11 In Discussion Question 15.11, we looked at the rivalry between Jo's and Sophie's golf schools. Recall that the daily demand for golf lessons at Jo's and Sophie's were $D_J(P_J, P_S) = 100 - 2p_J + p_S$ and $D_S(p_J, p_S) = 100 - 2p_S + p_J$, respectively, where p_J is Jo's price and p_S is Sophie's price. In the previous chapter, we found the Bertrand equilibrium when the schools set price once and for all.

Now suppose that Jo and Sophie choose prices repeatedly, that is, every day they set new prices based on what happened the day before.

(*a*) Is there an equilibrium in which the schools simply charge the prices found in Chapter 15, where each school chooses a single price once and for all?

(*b*) Would you expect the schools to achieve a more profitable equilibrium than the one described in part (*a*)? How does your answer depend on whether the firms are unsure when they will go out of business or whether they know that there is a definite date on which one of them is going to go out of business?

Notes

1 This outcome also arises when Air Lion's strategy is *produce "high"* and Beta's strategy is *produce "low" no matter what*. Since these strategies give rise to exactly the same outcome as the strategies in the text, they will not be discussed further.

2 A perfect equilibrium is also sometimes called a *subgame perfect equilibrium*. Notice that any dominant strategy equilibrium is a perfect equilibrium, but (as this example shows) a perfect equilibrium need not be a dominant strategy equilibrium.

3 The strategies "*enter*" and *choose "low" no matter what the entrant does*, also satisfy the Nash condition. These strategies give rise to exactly the same outcome (the potential entrant comes in and the incumbent chooses the low level of output) as do the strategies in the text and so will not be discussed further.

4 This effect arises in Equation 15.8, for example.

5 The choice of three actions led to the following bizarre scene: as one of the players in the shoot-out shot, the goalie dived to the right – just in time to get out of the way and let the ball go into the goal exactly where the goalie had been standing.

6 You may recall from the previous chapter that we interpreted Cournot reaction functions in terms of reacting to beliefs, rather than the actual quantity chosen by the other firm.

7 By pricing just below the buyer's value, the seller is giving the buyer an incentive to accept the offer. You should convince yourself that it never makes sense to set a price that is greater than €1,500 but less than €1,999.

8 In the next chapter we will consider many more situations in which one economic decision maker tries to figure out what another one knows by looking at his or her actions.

9 Equation 16.4 is an approximation because Air Lion would have to undercut p_s slightly. It is sufficiently close that we will not worry about the difference.

10 This result does not hold when there are two or more equilibria in the stage game itself. In such situations firms can, in effect, make threats and promises about which stage-game equilibrium they will choose in the last period and then work backwards from there.