10

Thinking Strategically: Competition among the Few

A US newspaper at the end of the 1990s[1] carried a story originating in that centre of wheeling, dealing and couch casting, Hollywood. At a Christmas Eve dinner party in 1997, actor Robert De Niro asked singer Tony Bennett if he would be willing to sing 'Got the World on a String' in the final scene of a film that De Niro would be acting in and which was to be produced by Warner Brothers. He was referring to the project that became the 1999 hit comedy *Analyze This*, in which the troubled head of a crime family, played by De Niro, seeks the counsel of a psychotherapist, played by Billy Crystal. In the script, both the mob boss and his therapist are big fans of Bennett's music. Bennett said he would be interested, and that was that ... for a year.

Then his son and agent, Danny Bennett, received a phone call from Warner Brothers to discuss terms. They had in mind a fee of $15,000 for Bennett Sr for singing the song in the final scene. That's not bad pay for an hour's work, so it was a very reasonable offer, and one any singer (or his agent) would be expected to accept in a semi-quaver. Unfortunately, the Warner negotiator let slip that the film was already in the can except for the final scene and the song, and the script clearly led up to this particular song and singer at the ending. Bennett Jr managed to get Warner Brothers up to $200,000. They were over the proverbial barrel. Had they put the offer to Bennett a year earlier, before shooting began, they would have been €185,000 better off. As they say: in life, timing is everything!

The payoff to many actions depends not only on the actions themselves but also on when they are taken and how they relate to actions taken by others. In Chapters 6–9, economic decision makers confronted an environment that was essentially fixed. This chapter will focus on cases in which people must consider the effect of their behavior on others. For example, an imperfectly competitive firm will in many circumstances want to weigh the likely responses of rivals when deciding whether to cut prices or to increase the advertising budget. *Interdependences* of this sort are the rule rather than the exception in economic and social life. To make sense of the world we live in, then, we must take these interdependences into account.

An analytical method for handling this type of problem is what is known as game theory. Its origins may be found in a book published in 1944 written by John von Neumann and

[1] *The New York Times*, 2 May 1999.

Oskar Morgenstern, *The Theory of Games and Economic Behavior*.[2] This book started from the premise that much economic behaviour can be analysed as a choice of a strategy in situations where people's interests do not coincide, so that *conflict between decision makers* is inevitable. Von Neumann and Morgenstern developed a mathematical methodology for analysing this type of situation. Then, in the early 1950s, John Nash,[3] a mathematician at Princeton, produced a couple of path-breaking papers dealing with the concept of an *equilibrium in a game* – meaning, loosely, an outcome that is stable and predictable given the motives of, and constraints facing, the players. Modern game theory has been built on these foundations, and Nash was subsequently awarded the Nobel Prize in Economics.

The theory of games

In chess, tennis, or any other game, the payoff to a given move depends on *what your opponent does in response*. In choosing your move, therefore, you must anticipate your opponent's responses, how you might respond and what further moves your own response might elicit.

Consider the following problem that shows how this idea applies in economics. You have decided to open a small supermarket in your home town neighbourhood, where there is already an established store belonging to a major national chain. The reason is that you have good information that there is 'cash on the table' in the form of profits to be appropriated if you can attract a sufficient number of customers from the incumbent. You do so because you estimate that you can offer a better value-for-money service. You could do this by undercutting the prices of the incumbent. You could do this simply by offering a different choice of goods. You could do this by offering a service that that will attract a sufficient number of higher-income shoppers to you who are not satisfied with the one-size-fits-all offerings of the established chain. The market share you hope to achieve makes the proposal profitable. But will you achieve it? That depends on how the incumbent firm reacts.

It could decide that the loss of market share is such that it must respond by lowering prices (increasing value for money). It could decide that there is room in the market for both of you, and it might be happy (if you go up-market) to leave the upper end of the market (with all the problems of dealing with better-off and more demanding purchasers) to you. It could decide that if you succeed in this venture it is probable that others will imitate you in other local markets, threatening the financial viability of its operations as a whole, and so launch a price war designed to force you out, on the basis that the sight of a corpse hanging from a gibbet deters imitation. Whether or not you enter the market, and the strategy you adopt in the market, will reflect your opinion as to what the other side will do in response to your decision. It's not at all clear that it makes sense to enter just because you see a profitable opportunity in the form of 'cash on the table'. In order to analyse and to predict outcomes in such situations, in which the payoffs to different actors depend on the actions their opponents undertake, economists and other behavioural scientists have devised the mathematical theory of games.

basic elements of a game the players, the strategies available to each player and the payoffs each player receives for each possible combination of strategies

The three elements of a game

Any game has three **basic elements:** the *players*, the list of possible actions (or *strategies*) each player can choose from and the *payoffs* the players receive for each combination of strategies. How these elements combine to form the basis of a theory of behaviour will become clear in the context of Examples 10.1–10.3.

[2] Von Neumann and Morgenstern (1944).

[3] Nash's sad life (he was diagnosed paranoid schizophrenic) was the subject of the movie *A Beautiful Mind*. For an economist, unfortunately, the movie is marred by its failure to make clear the significance, simple elegance and enormous analytical implications of his exposition of what is now known as the concept of a Nash equilibrium (see p. 276).

Example 10.1 Should Lufthansa spend more money on advertising?

Suppose that Lufthansa and Alitalia are the only air carriers that serve the Frankfurt–Milan route. Each currently earns an economic profit of €6,000 per flight on this route. If Lufthansa increases its advertising spending in this market by €1,000 per flight and Alitalia spends no more on advertising than it does now, Lufthansa's profit will rise to €8,000 per flight and Alitalia's will fall to €2,000. If both spend €1,000 more on advertising, each will earn an economic profit of €5,500 per flight. These payoffs are symmetric, so if Lufthansa stands still while Alitalia increases its spending by €1,000, Lufthansa's economic profit will fall to €2,000 per flight and Alitalia's will rise to €8,000. If each must decide independently whether to increase spending on advertising, what should Lufthansa do?

Think of this situation as a game. What are its three elements? The players are the two airlines, each of which must choose one of two strategies: to raise spending by €1,000 or to leave it the same. The payoffs are the economic profits that correspond to the four possible scenarios resulting from their choices. One way to summarise the relevant information about this game is to display the players, strategies, and payoffs in the form of a simple table called a **payoff matrix** (see Table 10.1).

payoff matrix a table that describes the payoffs in a game for each possible combination of strategies

Lufthansa's Choice	Alitalia's Choice: Raise ad spending	Alitalia's Choice: Leave ad spending the same
Raise ad spending	Alitalia gets €5,500; Lufthansa gets €5,500	Alitalia gets €2,000; Lufthansa gets €8,000
Leave ad spending the same	Alitalia gets €8,000; Lufthansa gets €2,000	Alitalia gets €6,000; Lufthansa gets €6,000

Table 10.1 The Payoff Matrix for an Advertising Game

Confronted with the payoff matrix in Table 10.1, what should Lufthansa do? The essence of strategic thinking is to begin by looking at the situation from the other party's point of view. Suppose Alitalia assumes that Lufthansa will raise its spending on advertising (the top row in Table 10.1). In that case, Alitalia's best bet would be to follow suit (the left column in Table 10.1). Why is the left column Alitalia's best response when Lufthansa chooses the top row? Alitalia's economic profits, given in the upper left cell of Table 10.1, will be €5,500 as compared with only €2,000 if it keeps spending level (see the upper right cell).

Alternatively, suppose Alitalia assumes that Lufthansa will keep its ad spending level (that is, Lufthansa will choose the bottom row in Table 10.1). In that case, Alitalia would still do better to increase spending, because it would earn €8,000 (the lower left cell) as compared with only €6,000 if it keeps spending level (the lower right cell). In this particular game, no matter which strategy Lufthansa chooses, Alitalia will earn a higher economic profit by increasing its spending on advertising. And since this game is perfectly symmetric, a similar conclusion holds for Lufthansa: no matter which strategy Alitalia chooses, Lufthansa will do better by increasing its spending on ads.

When one player has a strategy that yields a higher payoff no matter which choice the other player makes, that player is said to have a **dominant strategy**. Not all games involve dominant strategies, but both players in this game have one, and that is to increase spending on ads. For both players, to leave ad spending the same is a **dominated strategy** – one that leads to a lower payoff than an alternative choice, regardless of the other player's choice.

dominant strategy one that yields a higher payoff no matter what the other players in a game choose

dominated strategy any other strategy available to a player who has a dominant strategy

Notice, however, that when each player chooses the dominant strategy, the resulting payoffs are smaller than if each had left spending unchanged. When Lufthansa and Alitalia increase their spending on ads, each earns only €5,500 in economic profits as compared with the €6,000 each would have earned without the increase. (We'll say more below about this apparent paradox.)

Nash equilibrium

A game is said to be in equilibrium if each player's strategy is the best he or she can choose, given the other players' chosen strategies. This definition of equilibrium is sometimes called a **Nash equilibrium**, after the Nobel Laureate John Nash. When a game is in equilibrium, no player has any incentive to deviate from his or her current strategy.

Nash equilibrium any combination of strategies in which each player's strategy is his or her best choice, given the other players' strategies

If each player in a game has a dominant strategy, as in Example 10.1, equilibrium occurs when each player follows that strategy. But even in games in which not every player has a dominant strategy, we can often identify an equilibrium outcome. Consider, for instance, the following variation on the advertising game in Example 10.1.

Example 10.2 Should Alitalia or Lufthansa spend more money on advertising?

Once again, suppose that Lufthansa and Alitalia are the only carriers serving the Frankfurt–Milan route, and that the payoffs are as in Table 10.2. Has Lufthansa a dominant strategy? Has Alitalia? If each firm does the best it can, given the incentives facing the other, what will be the outcome of this game?

Lufthansa's Choice \ Alitalia's Choice	Raise ad spending	Leave ad spending the same
Raise ad spending	Alitalia gets €4,000; Lufthansa gets €3,000	Alitalia gets €3,000; Lufthansa gets €8,000
Leave ad spending the same	Alitalia gets €5,000; Lufthansa gets €4,000	Alitalia gets €2,000; Lufthansa gets €5,000

Table 10.2 Equilibrium when One Player Lacks a Dominant Strategy

In this particular game, no matter what Lufthansa does, Alitalia will do better to raise its ad spending, so raising the advertising budget is a dominant strategy for Alitalia. Lufthansa, however, does not have a dominant strategy. If Alitalia raises its spending, Lufthansa will do better to stand still; if Alitalia stands still, however, Lufthansa will do better to spend more. But even though Lufthansa hasn't a dominant strategy, we can still predict what is likely to happen in this game. After all, Lufthansa's managers know what the payoff matrix is, so they can predict that Alitalia will spend more on ads (since that is Alitalia's dominant strategy). Thus the best strategy for Lufthansa, given the prediction that Alitalia will spend more on ads, is to keep its own spending level. If both players do the best they can, taking account of the incentives each faces, this game will end in the lower left cell of the payoff matrix in Table 10.2: Alitalia will raise its spending on ads and Lufthansa will not. (Note that when both players are positioned in the lower left cell, neither has any incentive to change its strategy.)

Note that the choices corresponding to the lower left cell in Table 10.2 satisfy the definition of a Nash equilibrium. If Lufthansa found itself in that cell, its alternative would be to raise its ad spending, a move that would reduce its payoff from €4,000 to €3,000. So Lufthansa has no incentive to abandon the lower left cell. Similarly, if Alitalia found itself in the lower left cell of Table 10.2, its alternative would be to leave ad spending the same, a move that would reduce its payoff from €5,000 to €2,000. So Alitalia also has no incentive to abandon the lower left cell. And that means that the lower left cell of Table 10.2 is a Nash equilibrium – a combination of strategies for which each player's choice is the best available option, given the choice made by the other player.

Exercise 10.1 What should Lufthansa and Alitalia do if their payoff matrix is modified as follows?

Alitalia's Choice

Lufthansa's Choice	Raise ad spending	Leave ad spending the same
Raise ad spending	Alitalia gets €2,000 Lufthansa gets €3,000	Alitalia gets €3,000 Lufthansa gets €4,000
Leave ad spending the same	Alitalia gets €3,000 Lufthansa gets €2,000	Alitalia gets €4,000 Lufthansa gets €3,000

RECAP The theory of games

The three elements of any game are the players, the list of strategies from which they can choose and the payoffs to each combination of strategies. Players in some games have a *dominant strategy*, one that yields a higher payoff regardless of the strategies.

Equilibrium in a game occurs when each player's strategy choice yields the highest payoff available, given the strategies by other players. Such a combination of strategies is called a *Nash equilibrium*.

The prisoner's dilemma

prisoner's dilemma a game in which each player has a dominant strategy, and when each plays it, the resulting payoffs are smaller than if each had played a dominated strategy

The game in Example 10.1 belongs to an important class of games called the **prisoner's dilemma**. In the prisoner's dilemma, when each player chooses his dominant strategy, the result is unattractive to the group of players as a whole.

The original prisoner's dilemma

Example 10.3 recounts the original scenario from which the prisoner's dilemma drew its name.

Example 10.3 Should the prisoners confess?

Two prisoners, Horace and Jasper, are being held in separate cells for a serious crime that they did in fact commit. The prosecutor, however, has only enough hard evidence to convict them of a minor offence, for which the penalty is 1 year in jail. Each prisoner is told that if one confesses while the other remains silent, the confessor will be released without prosecution, and the other will spend 20 years in prison. If both confess, they will get an intermediate sentence of 5 years. (These payoffs are summarised in Table 10.3.) The two prisoners are not allowed to communicate with one another. Do they have a dominant strategy? If so, what is it?

		Jasper's Choice	
		Confess	Remain silent
Horace's Choice	Confess	Japser gets 5 years; Horace gets 5 years	Japser gets 20 years; Horace gets 0 years
	Remain silent	Japser gets 0 years; Horace gets 20 years	Japser gets 1 years; Horace gets 1 years

Table 10.3 The Payoff Matrix for a Prisoner's Dilemma

In this game, the dominant strategy for each prisoner is to confess. No matter what Jasper does, Horace will get a lighter sentence by speaking out. If Jasper confesses, Horace will get 5 years (upper left cell in Table 10.3) instead of 20 (lower left cell). If Jasper remains silent, Horace will go free (upper right cell) instead of spending a year in jail (lower right cell). Because the payoffs are perfectly symmetrical: Jasper will also do better to confess, no matter what Horace does. The difficulty is that when each follows his dominant strategy and confesses, both will do worse than if each had said nothing. When both confess, they each get 5 years (upper left cell) instead of the 1 year they would have received by remaining silent (lower right cell). Hence the name of this game, the prisoner's dilemma (or, indeed, prisoners' dilemma).

Exercise 10.2 GM and Chrysler must both decide whether to invest in a new process. Games 1 and 2 show how their profits depend on the decisions they might make. Which of these games is a prisoner's dilemma?

Game 1

GM's Choice \ Chrysler's Choice	Don't invest	Invest
Don't invest	10 for Chrysler; 10 for GM	12 for Chrysler; 4 for GM
Invest	4 for Chrysler; 12 for GM	5 for Chrysler; 5 for GM

Game 2

GM's Choice \ Chrysler's Choice	Don't invest	Invest
Don't invest	12 for Chrysler; 4 for GM	5 for Chrysler; 5 for GM
Invest	10 for Chrysler; 10 for GM	4 for Chrysler; 12 for GM

The prisoner's dilemma is one of the most powerful metaphors in all of human behavioural science. Countless social and economic interactions have payoff structures analogous to the one confronted by the two prisoners. Some of those interactions occur between only two players, as in the examples just discussed; many others involve larger groups. But regardless of the number of players involved, the common thread is one of conflict between the narrow self-interest of individuals and the broader interests of larger communities.

Prisoners' dilemmas confronting imperfectly competitive firms

A **cartel** is any coalition of firms or producers that conspires to restrict production for the purpose of earning an economic profit by increasing price and widening the gap between price and cost. The world's best-known (and for some reason widely tolerated) cartel is OPEC, the Organisation of Petroleum Exporting Countries, formed in 1960 to control oil production. Its members account for a very large share of the world's oil production. After the surge in oil prices in 2004 from \$20 to \$55 per barrel, due for the most to precautionary stockpiling because of the situation in Iraq, as oil prices fell back toward \$40, and looked as if they would continue to decline, OPEC announced that it intended to cut production in order to maintain prices at or around \$40. This was not treated very seriously by most industry observers because over the previous 30 years it had been shown that after initial increases in prices sparked by OPEC production cuts, prices always fell back as OPEC members broke ranks and non-members took advantage of OPEC cuts to raise production. OPEC is not only the best-known cartel, but the best-known example of the problem besetting all cartels: how to get them to work. History shows that cartels are notoriously unstable. As we shall see in Economic naturalist 10.1, the problem confronting oligopolists who are trying to form a cartel is a classic illustration of the prisoner's dilemma.

cartel a coalition of firms that agrees to restrict output for the purpose of earning an economic profit

Economic naturalist 10.1 Why are cartel agreements notoriously unstable?

Consider a market for bottled water served by only two firms, Aquapure and Mountain Spring. Each firm can draw water free of charge from a mineral spring located on its own land. Customers supply their own bottles. Rather than compete with one another, the two firms decide to collude by selling water at the price a profit-maximising pure monopolist would charge. Under their agreement (which constitutes a cartel), each firm would produce and sell half the quantity of water demanded by the market at the monopoly price (see Fig. 10.1). The agreement is not legally enforceable, however, which means that each firm has the option of charging less than the agreed price. If one firm sells water for less than the other firm, it will capture the entire quantity demanded by the market at the lower price.

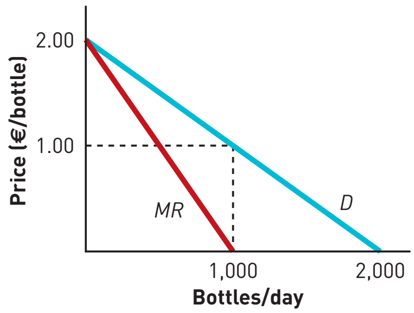

Figure 10.1 The Market Demand for Mineral Water. Faced with the demand curve shown, a monopolist with zero marginal cost would produce 1,000 bottles per day (the quantity at which marginal revenue equals zero) and sell them at a price of €1.00 per bottle.

Why is this agreement likely to collapse?

Since the marginal cost of mineral water is zero, the profit-maximising quantity for a monopolist with the demand curve shown in Fig. 10.1 is 1,000 bottles per day, the quantity for which marginal revenue equals marginal cost. At that quantity, the monopoly price is €1 per bottle. If the firms abide by their agreement, each will sell half the market total, or 500 bottles per day at a price of €1 per bottle, for an economic profit of €500 per day.

But suppose Aquapure reduced its price to €0.90 per bottle. By underselling Mineral Spring, it would capture the entire quantity demanded by the market which, as shown in Fig. 10.2 is 1,100 bottles per day. Aquapure's economic profit would rise from €500 per day to (€0.90 per bottle) (1,100 bottles per day) = €990 per day, almost twice as much as before. In the process, Mountain Spring's economic profit would fall from €500 per day to zero. Rather than see its economic profit disappear, Mountain Spring would match Aquapure's price cut, recapturing its original 50 per cent share of the market. But when each firm charges €0.90 per bottle and sells 550 bottles per day, each earns an economic profit of (€0.90 per bottle) (550 bottles per day) = €495 per day, or €5 per day less than before.

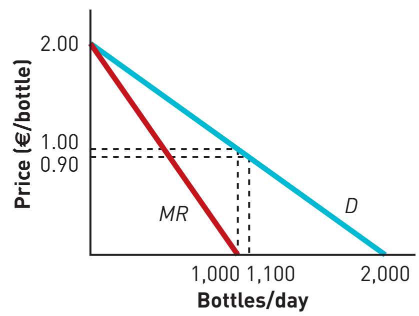

Figure 10.2 The Temptation to Violate a Cartel Agreement. By cutting its price from €1 per bottle to €0.90 per bottle, Aquapure can sell the entire market quantity demanded at that price, 1,100 bottle per day, rather than half the monopoly quantity of 1,000 bottles per day.

Suppose we view the cartel agreement as an economic game in which the two available strategies are to sell for €1 per bottle or to sell for €0.90 per bottle. The payoffs are the economic profits that result from these strategies. Table 10.4 shows the payoff matrix for this game. Each firm's dominant strategy is to sell at the lower price, yet in following that strategy each earns a lower profit than if each had sold at the higher price.

Aquapure's Choice	Mountain Spring's Choice: Charge €1	Mountain Spring's Choice: Charge €0.90
Charge €1	€500/day for Mountain Spring; €500/day for Aquapure	€990/day for Mountain Spring; €0/day for Aquapure
Charge €0.90	€0/day for Mountain Spring; €990/day for Aquapure	€495/day for Mountain Spring; €495/day for Aquapure

Table 10.4 The Payoff Matrix for a Cartel Agreement

The game does not end with both firms charging €0.90 per bottle. Each firm knows that if it cuts the price a little further, it can recapture the entire market, and in the process earn a substantially higher economic profit. At every step the rival firm will match any price cut, until the price falls all the way to the marginal cost – in this example, zero.

Cartel agreements confront participants with the economic incentives inherent in the prisoner's dilemma, which explains why such agreements have historically been so unstable. Usually a cartel involves not just two firms, but several, an arrangement that can make retaliation against price-cutters extremely difficult. In many cases, discovering which parties have broken the agreement is difficult. For example, OPEC, has no practical way to prevent member countries from secretly pumping oil offshore in the dead of night.

Economic naturalist 10.2 How did governments in Europe and North America unwittingly solve the television advertising dilemma confronting cigarette producers?

In 1970, the US Congress enacted a law making cigarette advertising on television illegal after 1 January 1971. Since the 1960s a similar restriction has been progressively introduced by all Western European countries and is now part of the Union's '*acquis communautaire*' (the body of rules and practices that have evolved and are accepted as binding on existing and new members). In the United States, given the evidence of the steadily declining proportion of Americans who smoke, this law seems to have achieved its stated purpose of protecting citizens against a proven health hazard. In Europe, this has been to some extent offset by an increase in the number of women who smoke, but is widely credited as having contributed to the decline in the total number of smokers and the total amount of smoking. But these laws also had an unintended effect, which was to increase the economic profit of cigarette makers, at least in the short run. In the year before the law's passage, US tobacco manufacturers spent more than $300 million on advertising (well over $1 billion in 2004 prices), about $60 million more than they spent during the year after the law was enacted. Much of the saving in advertising expenditures in 1971 was reflected in higher cigarette profits at the year-end. But if eliminating television advertising made companies more profitable, why didn't the manufacturers eliminate the ads on their own?

When an imperfectly competitive firm advertises its product, its demand curve shifts rightward, for two reasons. First, people who have never used that type of product learn about it, and some buy it. Second, people who consume a different brand of the product may switch brands. The first effect boosts sales industrywide; the second merely redistributes existing sales. Whatever the long-run effects on total sales, available evidence suggests that, in the short run, tobacco advertising in the main affects market shares rather than the size of the market. Its primary short-run effect is *brand switching*.

Thus the decision of whether to advertise confronts the individual firm with a prisoner's dilemma. Table 10.5 shows the payoffs facing a pair of US cigarette producers trying to decide whether to advertise. If both firms advertise on TV (the upper left cell), each earns a profit of only $10 million per year as compared with a profit of $20 million per year for each if neither advertises (the lower right cell). Clearly, both will benefit if neither advertises.

Yet note the powerful incentive that confronts each firm. RJR sees that if Philip Morris doesn't advertise, RJR can earn higher profits by advertising ($35 million per year) than by not advertising ($20 million per year). RJR also sees that if Philip Morris does advertise, RJR will again earn more by advertising ($10 million per year) than by not advertising ($5 million per year). Thus RJR's dominant strategy is to advertise. And because the payoffs are symmetric, Philip Morris' dominant strat-

RJR's Choice	Philip Morris' Choice: Advertise on TV	Philip Morris' Choice: Don't advertise on TV
Advertise on TV	$10 million/year for Philip Morris; $10 million/year for RJR	$5 million/year for Philip Morris; $35 million/year for RJR
Don't advertise on TV	$35 million/year for Philip Morris; $5 million/year for RJR	$20 million/year for Philip Morris; $20 million/year for RJR

Table 10.5 Cigarette Advertising as a Prisoner's Dilemma

egy is also to advertise. So when each firm behaves rationally from its own point of view, the two together do worse than if they had both shown restraint. The congressional ad ban forced cigarette manufacturers to do what they could not have accomplished on their own.

Prisoners' dilemmas in everyday life

As Economic naturalists 10.3 and 10.4 make clear, the prisoner's dilemma helps us to make sense of human behaviour not only in the world of business but in other domains of life as well.

Economic naturalist 10.3 Why do people often stand at rock concerts, even though they can see just as well when everyone sits?

A few years ago, an economic naturalist went with friends to hear Diana Ross sing. They bought good seats, some 20 rows from the stage. But before Ross had finished her first song, several people in front of them rose to their feet, presumably to get a better view. In doing so, they blocked the line of sight for others behind them, forcing those people to stand to see better. Before long, the entire crowd was standing. Then a few people in the front rows climbed on top of their seats, blocking the views of those behind them and forcing them to stand on their seats too. The seats had fold-up bottoms, so from time to time someone who stood too close to the pivot point would tumble as the seat popped into its vertical position. All things considered, the outcome was far less satisfactory than if everyone had remained seated. Why this pattern of self-defeating behaviour?

To understand what happened at the concert, note that standing is self-defeating only when viewed from the group's perspective. From the individual's perspective, however, standing passes the cost–benefit test. No matter what others do, an individual gets a better view by standing than by sitting. Suppose for the sake of discussion that you and other members of the audience would be willing to pay €2 to avoid standing and €3 to get a better view (or avoid having a worse one). In this multi-person prisoner's dilemma, you are one player and the rest of the audience is the other. The two strategies are to stand or to sit. Suppose everyone is seated to begin with. The payoffs you and others face will depend on the combination of strategies that you and others choose, as shown in Table 10.6.

Your Choice \ Other People's Choice	Stand	Sit
Stand	–€2 for others –€2 for you	–€3 for others –€1 for you
Sit	–€1 for others –€3 for you	€0 for others €0 for you

Table 10.6 Standing versus Sitting at a Concert as a Prisoner's Dilemma

The payoff of €0 in the lower right cell of the payoff matrix in Table 10.6 reflects the fact that when everyone remains seated, everyone is just as well off as before. Your payoff of –€3 in the lower left cell reflects the fact that if you sit while others stand, you will have a worse view. Your payoff of –€2 in the upper left cell reflects the fact that when you and others stand, you must endure the €2 cost of standing, even though you don't get a better view. Finally, your €1 payoff in the upper right cell represents the difference between your €3 benefit and your €2 cost of standing when you stand while others sit.

These payoffs mean that your dominant strategy is to stand. If others stand, you will get –€2 by standing, which is better than the –€3 you will get by sitting. If others sit, you will get €1 by standing, which is better than the €0 you will get by sitting. Since this game is symmetric, the dominant strategy for others is also to stand. Yet when everyone stands, everyone gets a payoff of –€2, which is €2 worse than if everyone had remained seated. As in all prisoner's dilemmas, the choice that is more attractive from the perspective of the individual turns out to be less attractive from the perspective of the group.

And here is something else to think about: Why does the problem of standing at critical moments not occur at the opera? After all, there is no reason to believe that most opera lovers have read much economics.

Economic naturalist 10.4 Why do people shout at parties?

Whenever large numbers of people gather for conversation in a closed space, the ambient noise level rises sharply. After attending such gatherings, people often complain of sore throats and hoarse voices. If everyone spoke at a normal volume at parties, the overall noise level would be lower, and people would hear just as well. So why do people shout?

Again, the problem involves the difference between individual incentives and group incentives. Suppose everyone starts by speaking at a normal level. But because of the crowded conditions, conversation partners have difficulty hearing one another, even when no one is shouting. The natural solution, from the point of view of the individual, is to simply raise one's voice a bit. But that is also the natural solution for everyone else. When everyone speaks more loudly, the ambient noise level rises, so no one hears any better than before.

No matter what others do, the individual will do better by speaking more loudly. Doing so is a dominant strategy for everyone, in fact. Yet when everyone follows the dominant strategy, the result is worse (no one can hear well) than if everyone had continued to speak normally. While shouting is wasteful, individuals acting alone have no better option. If anyone were to speak softly while others shout, that person wouldn't be heard. No one wants to go home with raw vocal cords, but people apparently prefer that cost to the alternative of not being heard at all.

Here's another version of the last problem: Do people shout at summer garden parties at Buckingham Palace?

Exercise 10.3

It was reported at the end of 2004 that (as many have suspected) the use of mobile phones on planes does not pose a safety hazard. Consequently legal restrictions on their use would in all probability be lifted. How likely is it that all airlines will completely lift the restriction even if it is no longer legally binding on them to impose it?

Tit-for-tat and the repeated prisoner's dilemma

When all players co-operate in a prisoner's dilemma, each gets a higher payoff than when all defect. So people who confront a prisoner's dilemma will be on the lookout for ways to create incentives for mutual co-operation. What they need is some way to *penalise* players who defect. When players interact with one another only once, this turns out to be difficult to achieve. But when they expect to interact repeatedly, new possibilities emerge.

A **repeated prisoner's dilemma** is a standard prisoner's dilemma that confronts the same players not just once but many times. Experimental research on repeated prisoner's dilemmas in the 1960s identified a simple strategy that proves remarkably effective at limiting defection. The strategy is called **tit-for-tat**, and here is how it works. The first time you interact with someone, you co-operate. In each subsequent interaction you simply do what that person did in the previous interaction. Thus, if your partner defected on your first interaction, you would then defect on your next interaction with her. If she then co-operates, your move next time will be to co-operate as well.

repeated prisoner's dilemma a standard prisoner's dilemma that confronts the same players repeatedly

tit-for-tat a strategy for the repeated prisoner's dilemma in which players co-operate on the first move, then mimic their partner's last move on each successive move

On the basis of elaborate computer simulations, University of Michigan political scientist Robert Axelrod showed that tit-for-tat was a remarkably effective strategy, even when pitted against a host of ingenious counter-strategies that had been designed for the explicit purpose of trying to exploit it. The success of tit-for-tat requires a reasonably stable set of players, each of whom can remember what other players have done in previous interactions. It also requires that players have a significant stake in what happens in the future, for it is the fear of *retaliation* that deters people from defecting.

Since rival firms in the same industry interact with one another repeatedly, it might seem that the tit-for-tat strategy would ensure widespread collusion to raise prices. And yet, as noted earlier, cartel agreements are notoriously unsuccessful. One difficulty is that tit-for-tat's effectiveness depends on there being only two players in the game. In competitive and monopolistically competitive industries there are generally many firms, and even in oligopolies there are often several. When there are more than two firms, and one defects now, how do the co-operators selectively punish the defector later? By cutting price? That will penalise everyone, not just the defector. Even if there are only two firms in an industry, these firms realise that other firms may enter their industry. So the would-be cartel members have to worry not only about each other, but also about the entire list of firms that might decide to compete with them. Each firm may see this as an impossible task and decide to defect now, hoping to reap at least some economic profit in the short run. What seems clear, in any event, is that the practical problems involved in implementing tit-for-tat have made it difficult to hold cartel agreements together for long.

RECAP The prisoner's dilemma

The *prisoner's dilemma* is a game in which each player has a dominant strategy, and in which the payoff to each player when each chooses that strategy is smaller than if each had chosen a dominated strategy. Incentives analogous to those found in the prisoner's dilemmas help to explain a broad range of behaviour in business and everyday life – among them, excessive spending on advertising, cartel instability, standing at concerts and shouting at parties. Co-operation in repeated prisoner's dilemmas can often be sustained by the *tit-for-tat* strategy, in which players co-operate on the first move and mimic their partner's previous move thereafter.

Games in which timing matters

In the games discussed so far, players were assumed to choose their strategies simultaneously, and which player moved first didn't particularly matter. For example, in the prisoner's dilemma, players would follow their dominant strategies even if they knew in advance what strategies their opponents had chosen. But in other situations, such as the negotiations between Warner Brothers and Tony Bennett described at the beginning of this chapter, timing is of the essence.

The ultimatum bargaining game

The following example illustrates another game in which timing plays a crucial role.

Example 10.4 Should Michael accept Tom's offer?

Tom and Michael are subjects in an experiment. The experimenter begins by giving €100 to Tom, who must then propose how to divide the money between himself and Michael. Tom can propose any division he chooses, provided the proposed amounts are whole euros and he offers Michael at least €1. Suppose Tom proposes €X for himself and €$(100 - X)$ for Michael, where X is a whole number no larger than 99. Michael must then say whether he accepts the proposal. If he does, each will get the proposed amount. But if Michael rejects the proposal, each player will get zero, and the €100 will revert to the experimenter. If Tom and Michael know they will play this game only once, and each wants to make as much money for himself as possible, what should Tom propose?

A payoff matrix is not a useful way to summarise the information in this game, because it says nothing about the timing of each player's move. For games in which timing matters, a **decision tree**, or **game tree**, is more useful. This type of diagram describes the possible moves in the sequence in which they may occur and lists the final payoffs for each possible combination of moves.

decision tree (or game tree) a diagram that describes the possible moves in a game in sequence and lists the payoffs that correspond to each possible combination of moves

The decision tree for the game in Example 10.4 is shown in Fig. 10.3. At *A*, Tom begins the game by making his proposal. At *B*, Michael responds to Tom's proposal. If he accepts (the top branch of the tree), Tom will get €*X* and Michael will get €(100 – *X*). If he refuses (the bottom branch of the tree), both will get nothing.

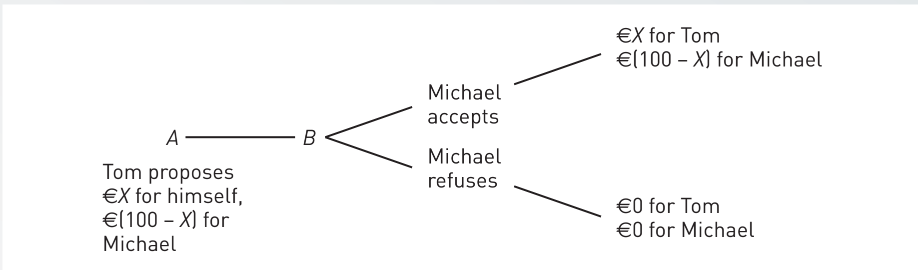

Figure 10.3 Decision Tree for Example 10.4. This decision tree shows the possible moves and payoffs for the game in Example 10.4 in the sequence in which they may occur.

In thinking strategically about this game, the key for Tom is to put himself in Michael's shoes and imagine how he might react to various proposals. Because he knows that Michael's goal is to make as much money as possible, he knows that Michael will accept his offer no matter how small, because the alternative is to reject it and get nothing. For instance, suppose that Tom proposes €99 for himself and only €1 for Michael (see Fig. 10.4). At *B*, Michael's best option is to accept the offer. This is a Nash equilibrium, because neither player has any incentive to deviate from the strategy he chose.

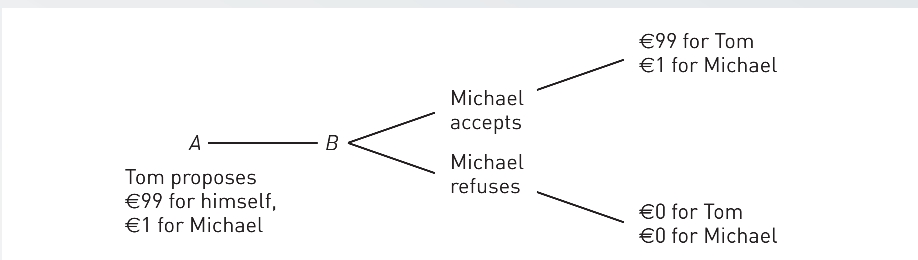

Figure 10.4 Tom's Best Strategy in an Ultimatum Bargaining Game. Because Tom can predict that Michael will accept any positive offer, Tom's income-maximising strategy at *A* is to offer Michael the smallest positive amount possible, €1.

ultimatum bargaining game one in which the first player has the power to confront the second player with a take-it-or-leave-it offer

This type of game has been called the **ultimatum bargaining game**, because of the power of the first player to confront the second player with a take-it-or-leave-it offer. Michael could refuse a one-sided offer from Tom, but doing so would make him worse off than if he accepted it.

Example 10.5 illustrates the importance of the *timing of moves* in determining the outcome of the ultimatum bargaining game.

Example 10.5 What should Michael's acceptance threshold be?

Suppose we change the rules of the ultimatum bargaining game slightly so that Michael has the right to specify *in advance* the smallest offer he will accept. Once Michael announces this number, he is bound by it. If Tom's task is again to propose a division of the €100, what amount should Michael specify?

This seemingly minor change in the rules completely alters the game. Once Michael announces that €Y is the smallest offer he will accept, his active role in the game is over. If Y is €60 and Tom proposes €X for himself and €$(100 - X)$ for Michael, his offer will be rejected automatically if X exceeds 40. The decision tree for this game is shown in Fig. 10.5.

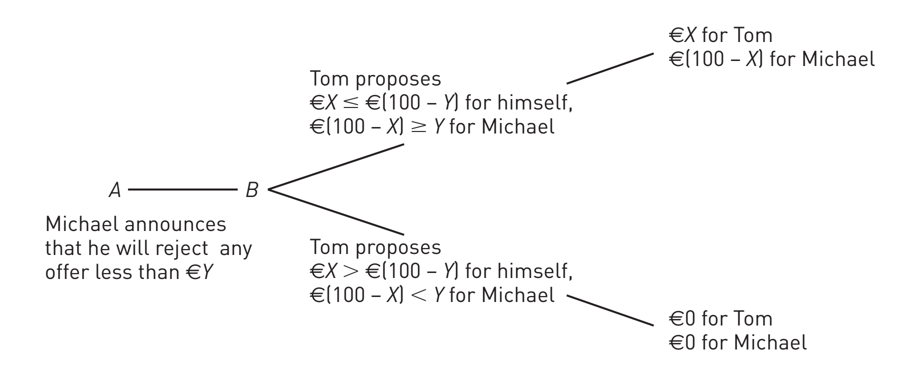

Figure 10.5 The Ultimatum Bargaining Game with an Acceptance Threshold. If Michael can commit himself to a minimum acceptable offer threshold at *A*, he will fare dramaticallly better than in the standard ultimatum bargaining game.

When Michael announces that €Y is the smallest offer he will accept, the best Tom can do is to propose €$(100 - Y)$ for himself and €Y for Michael. If he proposes any amount less than €Y for Michael, both will get nothing at all. Since this reasoning holds for any value of Y less than 100, Michael's best bet is to announce an acceptance threshold of €99 – the largest whole number that is less than €100. The equilibrium outcome of the game will then be €99 for Michael and only €1 for Tom, exactly the opposite of the outcome when Tom had the first move.

Credible threats and promises

credible threat a threat to take an action that is in the threatener's interest to carry out

Why couldn't Michael have threatened to refuse a one-sided offer in the original version of the game? While nothing prevented him from doing so, such a threat would not have been credible. In the language of game theory, a **credible threat** is one that is in the threatener's interest to carry out when the time comes to act. The problem in the original version of the game is that Michael would have no reason to carry out his threat to reject a one-sided offer in the event that he actually received one. Once Tom announced such an offer, refusing it would not pass the cost–benefit test.

The concept of a credible threat figured prominently in the negotiations between Warner Brothers managers and Tony Bennett over the matter of Bennett's fee for performing in *Analyze This*. Once most of the film had been shot, managers knew they couldn't threaten credibly to refuse Bennett's salary demand, because at that point adapting the film to another singer would have been prohibitively costly. In contrast, a similar threat made before production of the movie had begun would have been credible.

Example 10.6 is another case in which one person suffers as a result of the inability to make a credible threat.

Example 10.6 Is it safe to steal Veronica's briefcase?

When Veronica travels out of town on business, she usually brings along an expensive briefcase. A stranger takes a liking to her briefcase and assumes that because Veronica is an economist, she must be a self-interested, rational person.[4] If the cost to Veronica of pressing charges in the event that her briefcase is stolen exceeds the value of the briefcase, can the stranger safely steal it?

Provided that the thief's assumptions about Veronica are correct, he can get away with his crime. To press charges once her briefcase has been stolen, Veronica must call the police and will probably miss her flight home. Months later, she will have to return to testify at the thief's trial, and she may have to endure hostile cross-examination by the thief's lawyer. Since these costs clearly exceed the value of the briefcase, a rational, self-interested person would simply write it off. But if Veronica could somehow have made a credible threat to press charges in the event that her briefcase was stolen, she could have deterred the thief. The problem is that the thief knows the cost of retaliation will exceed the benefit, so the threat is not credible.

credible promise a promise that is in the interests of the promissor to keep when the time comes to act.

Just as in some games credible threats are impossible to make, in others **credible promises** are impossible.

Example 10.7 Should the business owner open a remote office?

The owner of a thriving business wants to start up an office in a distant city. If she hires someone to manage the new office, she can afford to pay a weekly salary of €1,000 – a premium of €500 over what the manager would otherwise be able to earn – and still earn a weekly economic profit of €1,000 for herself. The owner's concern is that she will not be able to monitor the manager's behaviour. The owner knows that by managing the remote office dishonestly, the manager can boost his take-home pay to €1,500 while causing the owner an economic loss of €500 per week. If the owner believes that all managers are selfish income maximisers, will she open the new office?

The decision tree for the remote office game is shown in Fig. 10.6. At *A*, the managerial candidate promises to manage honestly, which brings the owner to *B*, where she must decide whether to open the new office. If she opens it, they reach *C*, where the manager must decide whether to manage honestly. If the manager's only goal is to make as much money as he can, he will manage dishonestly (bottom branch at *C*), since that way he will earn €500 more than by managing honestly (top branch at *C*).

[4] Experiments with the ultimatum game have uncovered something that may not surprise you. When classroom experiments are played using sociology students, literature students and similar groups as test populations, the offers that are made are usually much closer to a 50/50 split than when they are carried out using economics students. There, the general trend is toward a 90/10 to 70/30 split. Figure that out! Does economics make you 'rational', or do more 'rational' people take economics?

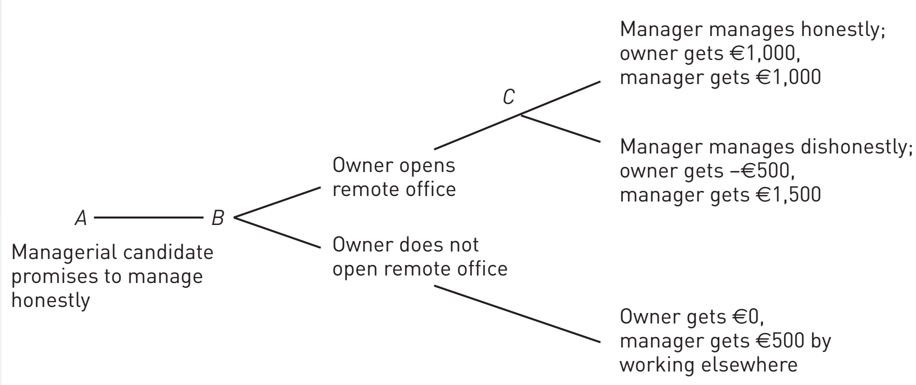

Figure 10.6 Decision Tree for the Remote Office Game. The best outcome is for the manager to open the office at *B* and for the manager to manage the office honestly at *C*. But if the manager is purely self-interested and the owner knows it, this path will not be an equilibrium outcome.

So if the owner opens the new office, she will end up with an economic loss of €500. If she had not opened the office (bottom branch at *B*), she would have realised an economic profit of zero. Since zero is better than –€500, the owner will choose not to open the remote office. In the end, the opportunity cost of the manager's inability to make a credible promise is €1,500: the manager's forgone €500 salary premium and the owner's forgone €1,000 return.

Exercise 10.4

Smith and Jones are playing a game in which Smith has the first move at *A* in the following decision tree. Once Smith has chosen either the top or bottom branch at *A* Jones, who can see what Smith has chosen, must choose the top or bottom branch at *B* or *C*. If the payoffs at the end of each branch are as shown, what is the equilibrium outcome of this game? If, before Smith chose, Jones could make a credible commitment to choose either the top or bottom branch when his turn came, what would he do?

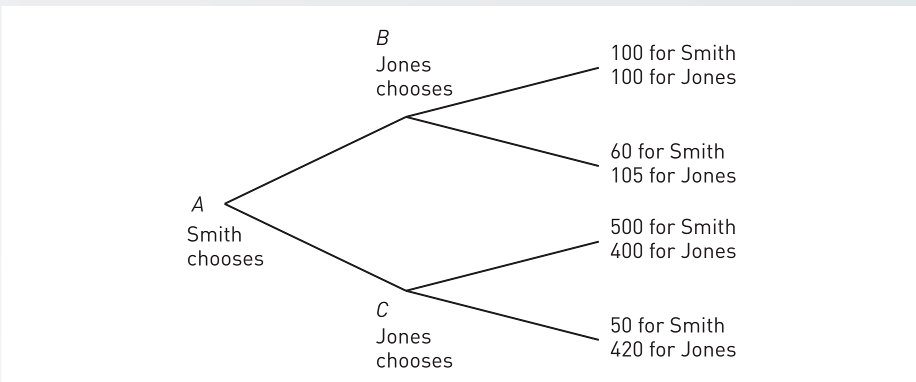

Commitment problems

Games like that in Exercise 10.3, as well as the prisoner's dilemma, the cartel game, the ultimatum bargaining game and the remote office game, confront players with a **commitment problem**, a situation in which they have difficulty achieving the desired outcome because they cannot make credible threats or promises. If both players in the prisoner's dilemma (Example 10.3) could make a binding promise to remain silent, both would be assured of a shorter sentence. Hence the logic of the underworld code of *omertà*, under which the family of anyone who provides evidence against a fellow mob member is killed. A similar logic explains the adoption of military arms control agreements, in which opponents sign an enforceable pledge to curtail weapons spending.

commitment problem a situation in which people cannot achieve their goals because of an inability to make credible threats or promises

The commitment problem in a game can be solved if the potential beneficiary can find some way of committing himself to a course of action in the future. For example, suppose firm *A* wants to discourage firm *B* from price cutting, and knows that to do so involves acting in a way that makes *B* confident that *A* will not overtly or tacitly engage in price cutting itself. It could sell firm *B* a 'put' option, whereby firm *B* could oblige firm *A* to buy specified quantities of its output at some critical price below today's price. A tacit promise not to cut prices would be made credible by this **commitment device**.

commitment device a way of changing incentives so as to make otherwise empty threats or promises credible

Business owners seem well aware of commitment problems in the workplace and have adopted a variety of commitment devices to solve them. Consider, for example, the problem confronting the owner of a restaurant. She wants her table staff to provide good service so that customers will enjoy their meals and come back in the future. And since good service is valuable to her, she would be willing to pay waiters extra for it. For their part, waiters would be willing to provide good service in return for the extra pay. The problem is that the owner cannot always monitor whether the waiters do provide good service. Her concern is that, having been paid extra for it, the waiters may slack off when she isn't looking. Unless the owner can find some way to solve this problem, she will not pay extra, the waiters will not provide good service, and she, they and the diners will suffer. A better outcome for all concerned would be for the waiters to find some way to commit themselves to good service.

Restaurateurs in many countries have tried to solve this commitment problem by encouraging diners to leave tips at the end of their meals. The attraction of this solution is that the diner is *always* in a good position to monitor service quality. The diner should be happy to reward good service with a generous tip, since doing so will help to ensure good service in the future. And the waiter has a strong incentive to provide good service, because he knows that the size of his tip may depend on it.

The various commitment devices just discussed – the underworld code of *omertà*, military arms control agreements, the tip for the waiter – all work because they change the material incentives facing the decision makers. But as Example 10.8 illustrates, sometimes this simple calculus of incentives is not a complete explanation.

Example 10.8 Will Federico leave a tip when dining on the road?

Federico has just finished a €30 dinner at Ristorante Stendhal,[5] just off the Milan–Ancona autostrada near Parma, some 300 km from home. The meal was superb, and the waiter provided good service. If Federico cares only about himself, will he leave a tip?

---

[5] It exists, in a little village on the banks of the Po. One of us has been there a couple of times. It's well worth the detour ... and you can have a great meal and have change out of €30 each (2004 prices). Watch out for the advertising sign beside the autostrada near Parma, or look it up on the web. And, yes, a tip was left!

Once the waiter has provided good service, there is no way for him to take it back if the diner fails to leave a tip. In restaurants patronised by local diners, failure to tip is not a problem, because the waiter can simply provide poor service the next time a non-tipper comes in. And no one wants to appear mean in front of people who might care. But the waiter lacks that leverage with out-of-town diners eating alone. Having already received good service, Federico must choose between paying €30 or €35 for his meal. If he is an essentially selfish person, the former choice may be a compelling one. But if you know that the waiter depends for much of his living on tips you are likely to tip anyway, even if not overgenerously, for the same reason as most people do not engage in shoplifting even when they know they would get away with it: our moral sense overrides our instinct for self-advancement.

RECAP Games in which timing matters

The outcomes in many games depend on the *timing* of each player's move. For such games, the payoffs are best summarised by a *decision tree* rather than a payoff matrix.

The inability to make credible threats and promises often prevents people from achieving desired outcomes in many games. Games with this property are said to confront players with *commitment problems*. Such problems can sometimes be solved by employing *commitment devices* – ways of changing incentives to facilitate making credible threats or promises.

The strategic role of preferences

In all the games we have discussed so far, players were assumed to care only about obtaining the best possible outcome for themselves. Thus each player's goal was to get the highest monetary payoff, the shortest jail sentence, the best chance of survival and so on. The irony, in most of these games, is that players do not attain the best outcomes. Better outcomes can sometimes be achieved by altering the material incentives selfish players face, but not always.

If altering the relevant material incentives is not possible, commitment problems can sometimes be solved by altering people's psychological incentives. As Example 10.9 illustrates, in a society in which people are strongly conditioned to develop moral sentiments – feelings of guilt when they harm others, feelings of sympathy for their trading partners, feelings of outrage when they are treated unjustly – commitment problems arise less often than in more narrowly self-interested societies.

Example 10.9 In a moral society, will the business owner open a remote office?

Consider again the owner of the thriving business who is trying to decide whether to open an office in a distant city (Example 10.7). Suppose the society in which she lives is one in which all citizens have been strongly conditioned to behave honestly. Will she open the remote office?

Suppose, for instance, that the managerial candidate would suffer guilt pangs if he embezzled money from the owner. Most people would be reluctant to assign a monetary value to guilty feelings. But for the sake of discussion, let us suppose that those feelings are so unpleasant that the manager would be willing to pay at least €10,000 to avoid them. On this assumption, the manager's payoff if he manages dishonestly will not be €1,500 but €1,500 – €10,000 = –€8,500. The new decision tree is shown in Fig. 10.7.

In this case, the best choice for the owner at *B* will be to open the remote office, because she knows that at *C* the manager's best choice will be to manage honestly. The irony, of course, is that the honest manager in this example ends up richer than the selfish manager in Example 10.7, who earned only a normal salary.

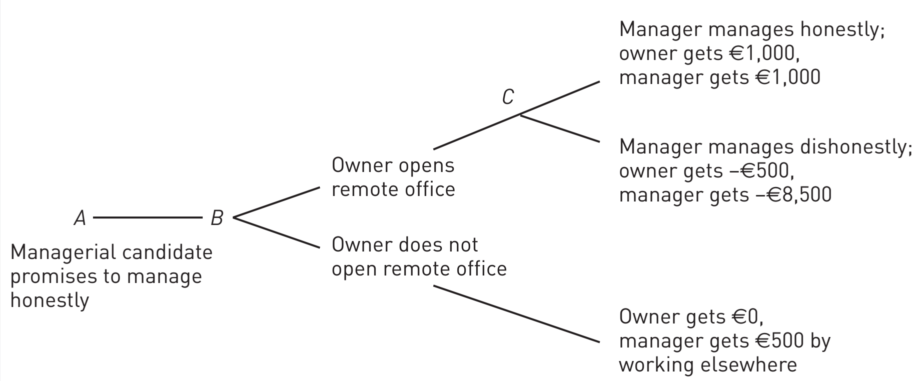

Figure 10.7 The Remote Office Game with an Honest Manager. If the owner can identify a managerial candidate who would choose to manage honestly at *C*, she will hire that candidate at *B* and open the remote office.

Are people fundamentally selfish?

As Examples 10.7 and 10.8 suggest, the assumption that people are 'self-interested' in the narrow sense of the term does not always capture the full range of motives that govern choice in strategic settings. Federico's case is well documented as quite normal. Researchers have found that tipping rates in restaurants patronised mostly by out-of-town diners are essentially the same as in restaurants patronised mostly by local diners.

Reflect also on how you would behave in some of the other games we have discussed. In the ultimatum bargaining game, what would you do if your partner proposed €99 for himself and only €1 for you? Would you reject the offer? If so, you are not alone. Two findings of extensive laboratory studies of the ultimatum bargaining game challenge the assumption that most players are narrowly self-interested. First, the most common proposal by the first player in this game is not a 99/1 split, but a 50/50 split. And second, on the few occasions when the first player does propose a highly one-sided split, the second player almost always rejects it. Subjects who reject the offer often mention the satisfaction they experienced at having penalised the first player for an 'unfair' offer.

Indeed, there are many exceptions to the outcomes predicted on the basis of the assumption that people are self-interested in the most narrow sense of the term. People who have been treated unjustly often seek 'revenge' even at ruinous cost to themselves. Every day people walk away from profitable transactions whose terms they believe to be 'unfair'.

American decision makers, immersed in the doctrine of strategic deterrence and mutually assured destruction (MAD), found it hard to understand the British decision to spend vast sums, lose lives and risk the core of the Royal Navy's surface fleet to recover the desolate Falkland Islands in 1982, even though they had little empire left against which to deter future aggression. Presumably the Argentine *junta* was of a similar mind. After all, as the Argentine writer Jorge Luis Borges observed, the Falkland War made about as much sense to North and South Americans as two bald men fighting over a comb.[6] It looked like a case of other values taking precedence over narrow self-interest. Possibly true: Mrs Thatcher was no ordinary Prime Minister, but rejoiced in the nickname of the 'Iron Lady'. At the time, Spain was putting pressure on Britain over Gibraltar, and Britain was facing difficult negotiations with China over the future administration of Hong Kong after the 99-year lease of most of the

[6] Quoted in Barnstone (1993).

colony's territory ran out in 1997, which made a return of Hong Kong to China unavoidable. And the UK government was facing internal opposition from the unionised coal miners who were threatening general strikes if the industry was rationalised. In these circumstances, does it seem so economically irrational to demonstrate that you will not be trampled on?

Preferences as solutions to commitment problems

Economists tend to view preferences as ends in themselves. Taking them as given, they calculate what actions will best serve those preferences. This approach to the study of behaviour is widely used by other social scientists and by game theorists, military strategists, philosophers and others. In its standard form, it assumes purely self-interested preferences for present and future consumption goods of various sorts, leisure pursuits and so on. Concerns about fairness, guilt, honour, sympathy and the like typically play no role.

Preferences clearly affect the choices people make in strategic interactions. Sympathy for one's trading partner can make a businessperson trustworthy even when material incentives favour cheating. A sense of justice can prompt a person to incur the costs of retaliation, even when incurring those costs will not undo the original injury. It can also induce people to reject one-sided offers, even when their wealth would be increased by accepting them.

Note, however, that although preferences can clearly shape behaviour in these ways, that alone does not solve commitment problems. The solution to such problems requires not only that a person *have* certain preferences, but also that others have some way of *discerning* them. Unless the business owner can identify the trustworthy employee, that employee cannot land a job whose pay is predicated on trust. Unless the predator can identify a potential victim whose character will motivate retaliation, that person is likely to become a victim. And unless a person's potential trading partners can identify him as someone predisposed to reject one-sided offers, he will not be able to deter such offers.

From among those with whom we might engage in ventures requiring trust, can we identify reliable partners? If people could make *perfectly* accurate character judgements, they could always steer clear of dishonest persons. That people continue to be victimised, at least occasionally, by dishonest persons suggests that perfectly reliable character judgements are either impossible to make or prohibitively expensive.

Vigilance in the choice of trading partners is an essential element in solving (or avoiding) commitment problems, for if there is an advantage in being honest and being perceived as such, there is an even greater advantage in only *appearing* to be honest. After all, a liar who appears trustworthy will have better opportunities than one who glances about furtively, sweats profusely and has difficulty making eye contact. Indeed, the liar will have the same opportunities as an honest person but will get higher payoffs because the liar will exploit them to the full.

In the end, the question of whether people can make reasonably accurate character judgements is an empirical one. Experimental studies have shown that even on the basis of brief encounters involving strangers, subjects are adept at predicting who will co-operate and who will defect in prisoners' dilemma games. For example, in one experiment in which only 26 per cent of subjects defected, the accuracy rate of predicted defections was more than 56 per cent. One might expect that predictions regarding those we know well would be even more accurate.

Do you know someone who would return an envelope containing €1,000 in cash to you if you lost it at a crowded concert? If so, then you accept the claim that personal character can help people to solve commitment problems. As long as honest individuals can identify at least some others who are honest and can interact selectively with them, honest individuals can prosper in a competitive environment.

RECAP The strategic role of preferences

Most applications of the theory of games assume that players are 'self-interested' in the narrow sense of the term. In practice, however, many choices, such as leaving tips in out-of-town restaurants, appear inconsistent with this assumption.

The fact that people seem driven by a more complex range of motives makes behaviour more difficult to predict but also creates new ways of solving commitment problems. *Psychological incentives* can often serve as commitment devices when changing players' material incentives is impractical. For example, people who are able to identify honest trading partners and interact selectively with them are able to solve commitment problems that arise from lack of trust.

Games, timing, beliefs and behaviour: oligopolistic markets

This section of the chapter contains some advanced analysis that can be skipped if you wish. It expands what has already been covered and introduces you to how economics approaches the analysis of firm behaviour in small-number markets, an analysis that underlies contemporary competition policy.

The concepts developed up to this point in this chapter are helpful in explaining, and even predicting, the behaviour of firms in markets where a small number of large firms determine product characteristics, prices and output levels. The menu of possible outcomes in such markets is long and varied, reflecting permutations and combinations of firm strategies, firm beliefs, modes of competition, and numbers and sizes of firms. In this section, we shall look at some basic models that are of use in indicating how differences in structure of markets, firms' beliefs as to other firms' responses and the manner in which firms compete can affect prices and outputs in those markets.

- **Structure** By 'market structure' we mean the degree of *similarity or difference between* firms (players, in game theory terms) and the *number of firms*. For example, in most parts of Britain or France the everyday grocery trade is shared by a small number of similar-sized large firms (small convenience stores abound, especially in Britain, but they largely sell into another market). On the other hand, retail sales of clothing is divided between a small number of large chain stores (Marks and Spencer, C&A, Debenhams, Galeries Lafayette …) and a very large number of much smaller outlets.
- **Beliefs** In the prisoner's dilemma, the outcome reflected the *knowledge and beliefs of each of the players about the other*. The outcome depends on the idea that each expects the other to behave independently in a one-off situation to maximise his own utility subject to no external influence or concern about the future. This may (or may not) be plausible in terms of suspects in police custody, but a wider set of beliefs, etc. is plausible in the case of firms interacting in markets.
- **Competition** Firms also differ as to how they *compete*, usually reflecting the products they are engaged in producing. For example, car producers can plausibly be modelled as deciding on a volume of output of a particular type of car, for which they tool up (commit themselves to produce, hoping to be able to sell them). Quantity is the decision competition variable: the firm lets market demand determine how much it realises for a given volume of production, or average revenue, meaning unit price. If firms compete by setting quantities, competition between firms is called 'Cournot competition' (see p. 298 below). A life assurance producer is best thought of as developing a financial product range, pricing it and waiting to see how many units of the product it can sell. Price is the decision variable, with market demand determining the amount sold. If firms compete by setting prices, competition is called 'Bertrand competition' (again, see p. 298)

Two basic tools of analysis

The residual demand curve

To understand this concept, consider the following problem. A large firm faces a number of smaller firms serving a market. The large firm has the advantage of lower costs of production, but has to settle on a profit-maximising price and output level in the knowledge that the price adopted will determine how much the smaller firms offer for sale.

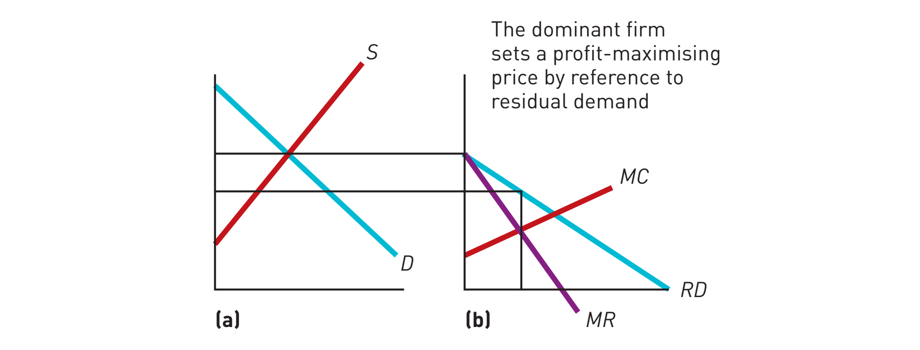

Figure 10.8 Residual Demand.

In panel (a) of Fig. 10.8, we show the supply curve for the combined output of the small firms, and the market demand curve. From this, we can see how much of total demand at any price will be met by the output of the small firms. Thus at P_1 they will produce only Q_1, leaving the balance of demand at that price to be supplied (if it wants to do so) by the large firm. At P_2, the small firms would produce sufficient output to satisfy market demand, and at that price there would be no room for the large firm to produce any output, given what the small firms would produce.

In panel (b) of Fig. 10.8 we draw the demand curve facing the large curve based on panel (a). It shows market demand at any price after the output of the smaller firms is subtracted. Thus, at P_2 there is no demand for output from the large firm, and at P_1 the demand for its output is the difference between total demand in panel (a) minus the output of the smaller firms at that price. Based on this **'residual' demand curve** we draw a marginal revenue curve; and the profit-maximising output and price for the large ('dominant') firm is given by the intersection of its *MC* curve with this *MR* curve.

residual demand curve shows the price and output combinations available to a firm given other firms' decisions

The concept of the residual demand curve can be derived in any market. Thus, in a two-firm market (a 'duopoly') we can derive for each firm the residual demand curve reflecting the market space left to it for different prices facing the other firm. This would describe part of the strategic problem facing, for example, Boeing and Airbus in the markets for very large commercial passenger and freight planes, for which they are the only producers. At the smaller end of the market (B737, Airbus 318–321) they both face competition from other suppliers (e.g. Embraer in Brazil, De Havilland Canada, SAAB in Sweden and a few other regional jet producers, as well as emerging producers in some Asian countries).

Exercise 10.5

Use residual demand curve analysis to illustrate Boeing's decision not to compete head-to-head with Airbus when the latter had committed itself to producing the A380, a super-jumbo with a capacity 50–70 per cent greater than Boeing's B747 jumbo.

The reaction function

In the simple game theory examples used earlier we considered whether a dominant strategy existed for either or both players. In those cases, there were two possible actions by each player, and payoffs derived from these produced a single-equilibrium outcome. It described the best choice of action contingent on what the other player might do. However, where pricing or output decisions are concerned there is a theoretical infinity of choices a firm can make, with related payoffs. Hence we have a range of *best values* for one firm depending on the choices of the other. This leads us to develop the concept of a 'reaction function' or a 'reaction curve'. This will show the preferred action of a given firm as a function of the action of the other firm(s).

reaction function shows the preferred response of a firm in terms of a decision variable as a response to a value of that variable chosen by the other firm(s)

A widely used variant of this concept describes how output levels are determined in markets in which firms choose profit-maximising outputs. We develop it here for a duopoly in which firms compete by setting quantities which they offer to the market.

Using a residual demand curve for Firm *I* (Fig. 10.9) we can determine its profit-maximising output for any output by Firm *II*. This is indicated on the horizontal axis in Fig. 10.9. Firm *II*'s output is indicated on the vertical axis. Thus, if firm *II* produced the perfectly competitive output (Q_{pc}), with price equal to average cost, and no profits, it would not pay Firm *I* to produce any output. If Firm *II* chose to produce nothing, leaving the entire market to firm *I* the best output for firm *I* is the monopoly output. With some not too restrictive assumptions about demand curves and cost curves, the reaction curve for firm *I* will be a straight line as drawn in Fig. 10.9. It runs from the monopoly output on Firm *I*'s axis to the competitive output on Firm *II*'s axis. It shows the profit-maximising output for Firm *I* for any output by Firm *II*. The further out a firm is on its reaction curve the lower are its profits. For either firm, the closer the output it produces is to the monopoly output on its own axis the higher are its profits.

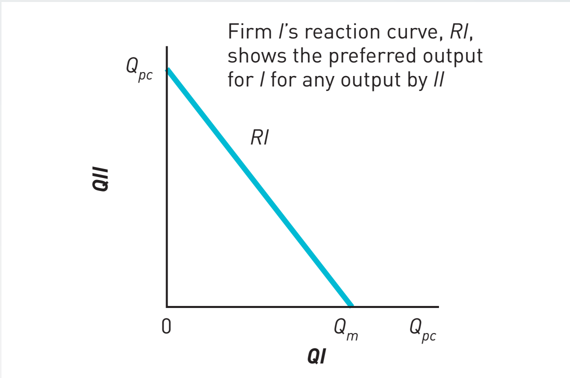

Figure 10.9 **Firm *I*'s Reaction Curve, *RI*.**

In Fig. 10.10 we show the reaction curves for both firms. For ease of analysis we assume that the firms have identical costs and the products are highly substitutable (so that they face similar demand conditions). Notice also the broken lines joining Q_{pc} on each axis and Q_m on each axis. These represent the competitive output and monopoly output divided between the two firms in proportion to the shares indicated by any point on the line. They also, therefore, represent the competitive level of profits (zero) and the monopoly level similarly divided between them. The firm outputs given by the intersec-

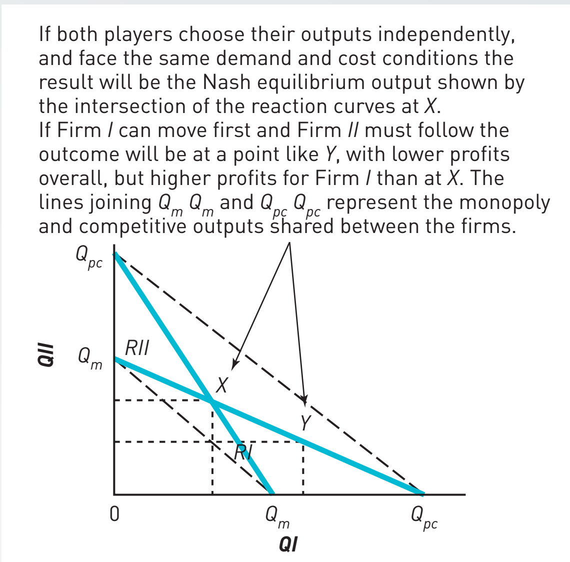

Figure 10.10 **If Both Players Choose Their Outputs Independently.**

tion of the two reaction curves involves a level of profits below the monopoly level and above the competitive level shared equally,

What will be the outcome? That depends on firms' beliefs and how the firms can choose how and when to move.

Figure 10.10 illustrates the choices we looked at in the prisoner's dilemma game (Example 10.3). If they can co-operate (collude rather than compete) they can share the monopoly profits. If they have to choose independently what to do, and each believes that the other faces the same problem and the same payoffs, and they can be thought of as deciding what to do simultaneously (as in the prisoner's dilemma), neither can do better than select the output indicated by the intersection of the two reaction curves. This is a dominant strategy for both, and the outcome is a Nash equilibrium outcome.

Suppose, however, that Firm *I* moves first, and believes (correctly) that Firm *II* will follow its move, rather than move independently (a leader and follower game). Firm *I* knows that Firm *II* will choose an output on the basis of its reaction curve. That curve becomes the set of outputs for the two firms among which Firm *I* has to choose. The Nash equilibrium is a feasible choice for Firm *I*, but an output level and share that is preferable for Firm *I* exists. It is at a point like *Y* in Fig. 10.10. Total output is higher, and price and industry profits are lower than at the Nash equilibrium. This shown by the fact that a line drawn through *Y* parallel to the line joining the competitive output on both axes (the total output shared between the two firms as indicated by point *X*) lies nearer to the competitive output line than the Nash equilibrium output line, meaning an output level that is higher and industry profits that are lower than at the Nash equilibrium output. There exists some point like *X* such that Firm *I*'s output share is sufficiently large to give Firm *I* a higher level of profit despite lower prices than at the Nash equilibrium, meaning of course that Firm *II* does much worse.

This illustrates the fact that firms' beliefs and the timing of moves affect the end result.

Competition as in this case when firms choose *output levels* is described as **Cournot competition**, after the French nineteenth-century economist, A. A. Cournot, who considered the problem of duopoly decision making on output levels as long ago as 1838, and arrived at the conclusion that the solution to the problem was what would today be described as the Nash equilibrium solution.

Cournot competition firms choose an output and accept the market price

Bertrand competition firms choose a price and accept that quantity sold depends on demand at that price

Suppose, however, that firms compete on *price*. This possibility was analysed by another French mathematician and economist, Joseph Bertrand, who argued that the Cournot solution did not generally hold and derived an alternative result where firms decided on prices rather than output levels – this is **Bertrand competition.**

We saw that under prisoner's dilemma circumstances in the Cournot game the outcome was a level of output between the competitive and the monopoly levels. Now suppose we think about a firm setting a price when its product is identical to that of the other firm, and they each have the same costs. Each firm knows that for any price it chooses which is higher than marginal cost the best response of the other firm is to set a price that is a little lower. If it sets the same price it shares industry profits equally, since it will sell the same amount as the competitor. But a slightly lower price will give it the whole market and all the profits.

You can try this as a problem, but a little thought will result in the following conclusion. Each firm will choose a price at which profits are zero. Price will be set equal to marginal cost, and with constant costs this means zero profits.

This leads to what has been called the 'Bertrand Paradox'. As long as there are at least two similar firms in the market the dominant strategy is for each to set price equal to cost. Increasing the number beyond two, or reducing the number as long as there are at least two results in a zero price cost margin. Industry structure has no apparent effect on price.

9707 CH10.QXD 1/5/06 2:10 PM Page 299

This result points to different conclusions depending on whether firms can best be treated as competing on quantity or price.

But if we look around, we can see that there are cases where firms can plausibly be treated as competing in terms of setting *prices* (Bertrand as opposed to Cournot competitors), and profits are not zero. We can also see examples of higher prices and profit margins in Bertrand markets when there are fewer competitors. Petrol retailing in local markets is an obvious example of both these market characteristics.

How do we square this with the Bertrand model just described?

There are two bases on which to suggest that the Bertrand Paradox outcome is not an equilibrium one. The first is that the products of the firms may not be *perfectly substitutable*. If that is so the products will not have to be sold at the same price (different brands of toothpaste are sold on the same supermarket shelves at different prices). When this is the case, it is possible to look at a firm's pricing decision in the same way as we looked at the output decision in the basic Cournot model. We construct a reaction curve for each firm, showing its preferred price given any price set by the competitor (Fig. 10.11).

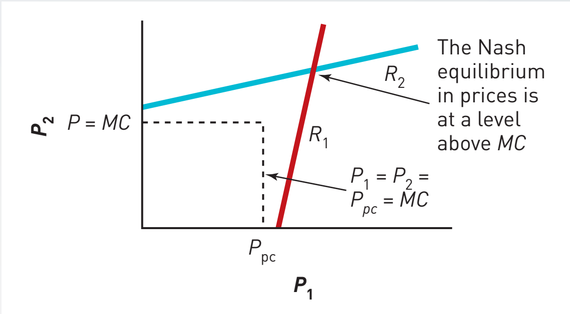

Figure 10.11 The Nash Equilibrium in Prices.

In this case, R_1 shows Firm *I*'s preferred price for its own product given (as a function of) Firm *II*'s price. R_2 is Firm *II*'s reaction curve. Neither will set a price below marginal cost, and each will set a price equal to marginal cost only if the other prices below marginal cost. Assuming constant unit cost and price elasticity of demand symmetry the reaction curves will (conveniently) be straight lines, and a Nash equilibrium in prices is where they intersect. Once again, this price (and the joint profits of the firms) will be lower than if they colluded to set the monopoly price and shared the market. However, the point is that the Nash equilibrium does not mean that prices are set equal to marginal cost. Positive profits can be earned in equilibrium without collusion. Furthermore, if there are more firms/products on offer, meaning that each firm/product faces closer substitutes, a firm's reaction curve becomes steeper, and the Nash equilibrium approaches the perfectly competitive price ($P = MC$). In the absence of co-operative behaviour, the more petrol stations there are in a town the lower we would expect prices to be.

There is a second, and even more important, reason for not expecting prices to fall to marginal cost under Bertrand competition, even if the products of each firm are very similar to those of others. That is the fact that the analysis so far treats the pricing decision as if it were taken just once, and never changed. In fact, of course, firms have to make decisions on price frequently through time. As supply costs move, and/or as demand conditions change, the 'correct' price for a firm to set changes too. We must assume that firms have memories. We must also assume that firms know how to signal how they will respond to price cutting, and how to indicate by their pricing behaviour whether or not they intend competing aggressively or adopting a 'live and let live' approach. Put another way, pricing decisions involve choosing not just a set of prices for your products, but deciding on a *pricing strategy*, or adopting a rule on setting prices as circumstances change.

The consequences of this approach to pricing include a high degree of 'price stickiness' in small-number markets where players can observe accurately what their competitors are doing as far as prices are concerned. This is sometimes illustrated in terms of what is called the 'kinked oligopoly demand curve' (Fig. 10.12). Given the price being charged, a firm will be reluctant to change it because it fears that an increase in price will leave it exposed to a col-

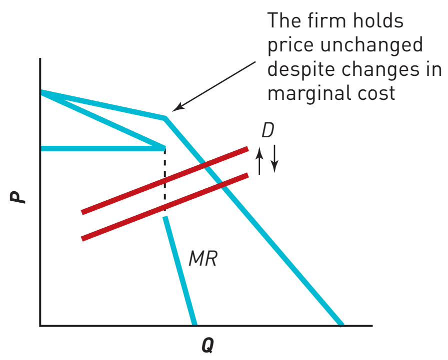

Figure 10.12 Kinked Demand Curve.

lapse in market share, while if it lowers its price the others will follow. Hence, if costs rise or fall, the firm holds its price.

There are several limitations of this approach. It does not explain how the observed price emerged. Nor does it suggest that firms will respond to market conditions by changing prices, something that we observe as a widespread phenomenon even in markets with a small number of similar-sized firms (i.e. there is no dominant firm price leader as in Fig. 10.8).

Instead, it makes more sense to analyse the pricing behaviour of firms in these markets in terms of repeated games, pricing rules, signalling and tacit collusion. This approach is also useful because it helps us understand when and why this *tacitly co-ordinated behaviour* breaks down, and we see price wars breaking out.

Repeated games

Think of a cartel, and ask why a member of it might decide to break ranks. One way to do this is to analyse the cartel as a *repeated game*, in which in each period any player has to choose whether or not to stick by the cartel rules, restrict output and sell at a price at which the firm's marginal revenue exceeds its marginal cost. The single-period dominant strategy (and the consequent Nash equilibrium) suggests breaking the rules, because by breaking them while everyone else obeys the rules you stand to clean up. However, if all the players know the game is going to be repeated indefinitely into the future, the calculation changes. Even if you clean up this week, by undermining the cartel you know that profits in future weeks will be lower. As a result, it is reasonable to believe that you might compare the present value of two profit streams into the future: profits if you undercut the other players this week, but everyone sells at a much lower price into the future, and profits if you – and, you hope, the others – stick by the rules.

Of course, you might reason that since you are a small player, and provided you don't lower prices by too much, the others will prefer to ignore your behaviour – or, better still, may not even be aware of it. Suppose, however, that either by experience, or by some explicit and credible threat, it was clear to you that cheating would not be tolerated, and selective price cuts by large and close competitors were the expected response to any deviation from the rule. How would that affect your choice of strategy? Suppose trade magazines published frequent, reliable and detailed information on who was charging what prices for which goods. Would that encourage you to shave your price?

On the other hand, if it was difficult to observe the quality-adjusted price because the product as sold differs among customers, does that help price fixing? If you had very poor information as to what the market was going to look like in the future, so that future profits are highly uncertain, what effect would that have? Or suppose you were in debt to a financial institution that was charging a high rate of interest on your loan, but the loan could be paid off by a quick killing in the market. High interest rates act as a disincentive to stick by cartel rules.

In practice, of course, we are not usually dealing with an actual cartel. Under competition law in Europe, the United States and other OECD countries, such arrangements are illegal. Participation not only results in being liable to substantial financial penalties, but in some circumstances can result in a prison sentence for those responsible. However, and especially in

tacit collusion firms behaving in a manner that resembles what might emerge from a collusive agreement because they recognise their interdependence

small-number situations, firms can decide to operate a pricing policy that is based on a recognition that how they behave affects others' behaviour and vice versa. Recognition of *interdependence of decisions* leads to a situation in which the observed behaviour resembles what might flow from an explicit agreement not to compete aggressively on price. For this reason this behaviour is frequently described as **tacit collusion**.

Economic naturalist 10.5 Tacit collusion or a cartel? The great cement scandal

There had long been rumours that all was not well with competition among cement producers in Europe when these were confirmed by the findings of the EU Commission's Competition Directorate in 1994.[7] Dismissing arguments about parallel pricing and tacit collusion, it found that Europe's largest cement producers had been operating a cartel for several years – fixing prices, sharing the EU market and administering production quotas. The industry was one in which all the conditions for tacit collusion were present. The product was homogeneous across producers, technology and costs were similar among producers, pricing was transparent, in any country producers were large and few in number and faced similar costs, and entry was difficult. However, from whistle-blower and dawn raid evidence, the Commission was satisfied that parallel pricing and output decisions were not simply the result of tacit collusion but the consequence of a fully orchestrated cartel covering most of the major producers. The Commission identified a starting date of January 1983, when, at a meeting, senior representatives of the European cement industry agreed on a rule of 'non-transhipment to home markets'. This was designed to prohibit any export of cement within Europe which might 'destabilise' neighbouring markets. A fine of about €250 million was imposed on the firms involved.

Tacit collusion in oligopolistic and imperfectly competitive industries is frequently supported by *pricing* and *signalling* strategies designed to discourage competition and reinforce co-operative behaviour. One common device, which to an uninformed observer looks like aggressive competition on price, is a commitment to match or beat any price by guaranteeing to be the lowest-price seller. This is very widespread as a promotional device but, on closer examination may in fact operate to support existing prices rather than ensure that buyers face the lowest possible prices. When a large white goods retailer makes such a promise he is signalling credibly to competitors small and large that if they lower prices, he will lower his.

Economic naturalist 10.6 Why was Statoil forced to abandon a rebate scheme that would have helped petrol retailers cut their prices?

In 2002 the Competition Authority in Dublin threatened legal action against Statoil as a result of an arrangement it had entered into with its retailers in a country town in Ireland.[8] Statoil's scheme involved (a) fixing a maximum price at which petrol

[7] The text of the decision and the facts and reasoning behind it are published as 94/815/EC, 30/11/1994, available on the EU website or in the EU *Official Journal*, OJ L33/1 1995.

[8] The Competition Authority has described the confrontation and its outcome in Decision E/03/002, which is available on its website: www.TCA. ie.

could be sold, and (b) offering a rebate to its dealers to enable them to cut their prices if other petrol outlets cut their prices below a marker price. While the first might have been part of a campaign to persuade people that Statoil petrol was a good buy (people notoriously appear not to pay much attention to actual prices per litre), the second element meant that other dealers knew that there was no point reducing price below the marker price since Statoil dealers could costlessly follow them down. This would clearly dampen price competition. Statoil withdrew the scheme.

Conditions facilitating tacit collusion do not always exist, even where there are few firms and they recognise the fact that their behaviour affects that of the others. First of all, firms may face different cost structures. This can result in significant differences between them in terms of what pricing structure and behaviour maximise profits when demand conditions change. For example, the lower a firm's variable costs relative to its fixed costs, the more likely it is to prefer to cut price rather than to cut output if demand weakens.

Pricing may not be transparent, with each unit being sold at a price negotiated with the purchaser on a confidential basis. Tacit collusion on pricing becomes meaningless. This factor on its own is probably the main reason why we see cut-throat competition between Boeing and Airbus.

Differences in product characteristics may mean that firms face different price elasticities of demand. This means they differ in terms of how they view the consequences of a change in prices.

As a result, we do not observe high prices supported by tacit collusion as inevitably following an observation of a small number of large producers in a market. And even when it does occur, it is frequently undermined by repeated episodes of price wars.

When do price wars break out, and how do they end?

In mid-2004 the Chinese Ministry of Information announced that it was stepping up its regulatory supervision of pricing in the telecoms sector in China in order to avoid price wars. Until earlier in the year mobile telephony charges had been set by the Ministry, which was also the regulator. In order to encourage the development of the sector, these price controls were relaxed. As elsewhere in the world, mobile telephony in China was dominated by a small number of players, all large. As soon as price controls were relaxed, a price war broke out, as firms slashed prices to increase market share in what was seen as an expanding market. The government, however, wanted the companies to maintain high profit margins and use them to finance qualitative upgrading of the mobile network. When the government intervened to restrict price cutting, of course, the industry greeted this warmly.[9] Regulation produced higher profits.

One possibility here is that under a tight regulatory regime firms had not been able to learn how others would react to price changes, nor been able to estimate demand elasticities correctly. This would interpret the price war as flowing from poor information as to demand conditions and the degree of interdependence of pricing. Another possibility is that, even with good information, tacit collusion and weak price competition was unlikely under deregulation because of the extraordinary degree of non-transparency of pricing: mobile phone pricing packages are numerous, complex and difficult to compare.

The impact of poor information is a useful starting point. Suppose a firm sees its sales falling, and has poor information as to market conditions. Three possible reasons will arise. The first is that market demand has weakened, and it is sharing the pain with other firms.

[9] The story can be found in the English-language *China Daily*, at its website www2.chinadaily.cn/English/doc/2004-07/03/content_345216.htm.

The second is that for some reason it has lost competitiveness. The third is that supply has increased as a competitor has increased production. Unfortunately, the firm is not able with any degree of certainty to establish what exactly is going on.

If it had been content not to compete aggressively on price up to this point, it now faces a problem. Suppose its best guess is that there has been a fall in demand. It must choose between (a) tacit collusion on the basis that this is a temporary blip in demand, everyone is in the same boat and price cutting will achieve nothing as it will be imitated; (b) price cutting anyway because the demand reduction may be permanent so that even if everyone is similarly affected there are fewer seats in the boat than potential passengers, and the last to cut prices exits the market. Suppose, however, it suspects that it has lost competitiveness. The options are not to do anything and accept a lower volume of sales (and possibly further falls) or reduce price to make its products more attractive. Finally, it might suspect that the problem is not its own offerings, or demand in the market, but a competitor increasing production. If it does nothing, it accepts a lower market share. If it wishes to protect its market share and signal its intention to do so to anyone who wants to raise theirs, it will respond by reducing price.

What will happen? In all three scenarios, cutting price looks a better bet, even if a case can be made for not cutting price. Betting against a price cut under uncertainty as to what is going on looks unwise. Hence, it is easy to understand why price wars can be triggered by falling demand or by an overestimate of market demand.

Notice, too, that the scenarios point to a further reason for price wars: market disciplining and market 'restructuring'. Responding to price cuts by one supermarket operator by a 'race to the bottom' may not seem to make much sense until you reflect that the expectation that this will happen in future will be reinforced by an aggressive response. It can also be a mechanism to secure a shake-up in an industry leading to a consolidation among players.

Economic naturalist 10.7 How the structure of the US tobacco industry was changed by a price war in the 1990s

Philip Morris cut the price of the world's best-selling cigarette, Marlboro, in the spring of 1993 by almost 20 per cent. Price competition from generic and mid-range brands over a period of years had reduced significantly the share of the premium cigarette producers, who at first accommodated the expansion of the cheaper producers and even increased prices in what was a declining market overall. The price cut by Philip Morris hit the mid-range producers worst, and the main firm producing in this range, American Tobacco, exited the market and sold its business to British American Tobacco. The generic producers appear to have taken the hint, and by 1995 Philip Morris had regained its lost market share and both premium and generic cigarette prices were rising together, if slowly.[10] In effect the Marlboro price managed simultaneously to exacerbate and end an ongoing price war, while restoring Philip Morris' position as market leader and consolidating the market.

And finally, of course, an episode of price cutting can arise as a consequence of changes on the supply side. New entrants to a market can disturb a collusive equilibrium. Even with tacit collusion, a fall in input costs can result in the collusive equilibrium level of prices falling. Alternatively, a significant fall in one firm's costs that is not enjoyed by other firms can lead to the lower-cost firm seeking to increase share on the back of lower costs.

[10] A full account of the episode, on which Economic naturalist 10.6 is based, is available in an article by John Kay, *Financial Times*, 5 July 1996.

Do oligopolistic firms always behave in a fashion that looks rational from a game theory perspective? Example 10.10 suggests otherwise.

Example 10.10 The Inverness bus war of 1988–91

After bus deregulation in 1986, many municipal bus companies in Britain found themselves facing competition from private sector bus companies. In May 1988, an entrant, ITL, started to compete with the incumbent public sector operator, HSO, in Inverness in the Scottish Highlands, by running minibuses on the routes operated by HSO. HSO responded by increasing its frequency (the volume of service) and, eventually, by replacing some of its conventional buses by minibuses. ITL in August increased its frequencies, and HSO responded in kind. Both firms made substantial losses, and by mid-1989 ITL was effectively bankrupt and had to sell out to a national bus operator. The latter continued to operate in Inverness, and in 1991 HSO withdrew from the town. Britain's Monopolies and Mergers Commission (MMC) investigated the matter and concluded that HSO had engaged in predatory pricing to drive out the entrant. Predatory pricing is selling at below cost to eliminate competition so as to be able to raise prices later and maintain monopoly profits.

In 1993, an academic paper appeared that undermined the MMC's conclusions, and indicated that the unfortunate experience of ITL and HSO was due to a series of miscalculations.[11] The economist authors estimated the Cournot reaction curves for the two firms in the market, and showed that (a) if ITL had entered at a lower level of frequency it could have done so profitably, and (b) that the level at which it entered made it impossible for the incumbent or the entrant to cover costs. Worse, when the losses mounted, and HSO had cut its costs by introducing minibuses, both sides decided to try to establish a leadership position and force the other to cut back capacity. In effect, in terms of Fig. 10.10, both tried to get to a point like *Y* on the other's reaction curve. If you add the two firm's production at points like *Y* you will see that total output exceeds the competitive output so that both must make losses. ITL buckled first, but faced with the might of a major national concern operating ITL's services, HSO had to exit the market a couple of years later.

Summary

- Economists use the mathematical theory of games to analyse situations in which the payoffs of one's actions depend on the actions *taken by others*. Games have three basic elements: the *players*; the list of *possible actions*, or *strategies*, from which each player can choose; and the *payoffs* the players receive for those strategies. The *payoff matrix* is the most useful way to summarise this information in games in which the timing of the players' moves is not decisive. In games in which the timing of moves does matter, a *decision tree* summarises the information in a much more useful format.
- A *dominant strategy* is one that yields a higher payoff regardless of the strategy chosen by the other player. In some games, such as the prisoner's dilemma, each player has a dominant strategy. The equilibrium occurs in such games when each player chooses his or her dominant strategy. In other games, not all players have a dominant strategy.
- Although the equilibrium outcome of any game is any combination of choices in which each player does the best he can, given the choices made by others, the result is often unattractive from the perspective of players as group. The prisoner's dilemma has this feature. The *incentive structure* of this game

[11] Dodgson, Katsoulakos and Newton (1993).

helps explain such disparate social dilemmas as excessive advertising, military arms races and failure to reap the potential benefits of interactions requiring trust.

- Individuals can often resolve these dilemmas if they can make *binding commitments* to behave in certain ways. Some commitments, such as those involved in military arms control agreements, are achieved by altering the material incentives confronting the players. Other commitments can be achieved by relying on psychological incentives to counteract material payoffs. Moral sentiments such as guilt, sympathy and a sense of justice often foster better outcomes than can be achieved by narrowly self-interested players. For this type of commitment to work, the relevant moral sentiments must be discernible by one's potential trading partners.
- Building on the idea of payoffs and games enables economists to construct models that indicate the importance of such things as beliefs and modes of competition in understanding how *small-number markets operate*.
- These models yield interesting and plausible conclusions that explain some features of *market behaviour* that are not explained by the simple models of perfect competition or monopolistic competition.

Key terms

basic elements of a game (274)
Bertrand competition (298)
cartel (279)
Cournot competition (298)
commitment device (291)
commitment problem (291)
credible promise (289)
credible threat (288)
decision tree (or game tree) (287)
dominant strategy (276)
dominated strategy (276)
Nash equilibrium (276)
payoff matrix (275)
prisoner's dilemma (278)
reaction function (297)
repeated prisoner's dilemma (285)
residual demand curve (296)
tacit collusion (301)
tit-for-tat (285)
ultimatum bargaining game (288)

Review questions

1. Explain why a military arms race is an example of a prisoner's dilemma.
2. Why did Warner Brothers make a mistake by waiting until the filming of *Analyze This* was almost finished before negotiating with Tony Bennett to perform in the final scene?
3. Suppose General Motors is trying to hire a small firm to manufacture the door handles for Opel saloon cars. The task requires an investment in expensive capital equipment that cannot be used for any other purpose. Why might the president of the small firm refuse to undertake this venture without a long-term contract fixing the price of the door handles?
4. Would you be irrational to refuse a one-sided offer in an ultimatum bargaining game if you knew that you would be playing that game many times with the same partner?
5. Can you show where (approximately) prices would settle if firm 1 were a leader and firm 2 a follower in the price-setting game (the equivalent to point *Y* in Fig. 10.10)?
6. Describe the commitment problem that narrowly self-interested diners and waiters confront at restaurants located on interstate highways. Given that in such restaurants tipping does seem to assure reasonably good service, do you think people are always selfish in the narrowest sense?

Problems

Problems marked with an asterisk (*) are more difficult.

1. In studying for his economics final, Sam is concerned about only two things: his grade and the amount of time he spends studying. A good grade will give him a benefit of 20; an average grade, a benefit of 5; and a poor grade, a benefit of 0. By studying a lot, Sam will incur a cost of 10; by studying a little, a cost of 6. Moreover, if Sam studies a lot and all other students study a little, he will get a good grade and they will get poor ones. But if they study a lot and he studies a little, they will get good grades and he will get a poor one. Finally, if he and all other students study the same amount of time, everyone will get average grades. Other students share Sam's preferences regarding grades and study time.

- **a.** Model this situation as a two-person prisoner's dilemma in which the strategies are to study a little and to study a lot, and the players are Sam and all other students. Include the payoffs in the matrix.
- **b.** What is the equilibrium outcome in this game? From the students' perspective, is it the best outcome?

2. Consider the following 'dating game', which has two players, *A* and *B*, and two strategies, to buy a movie ticket or a baseball ticket. The payoffs, given in points, are as shown in the matrix below. Note that the highest payoffs occur when both *A* and *B* attend the same event.

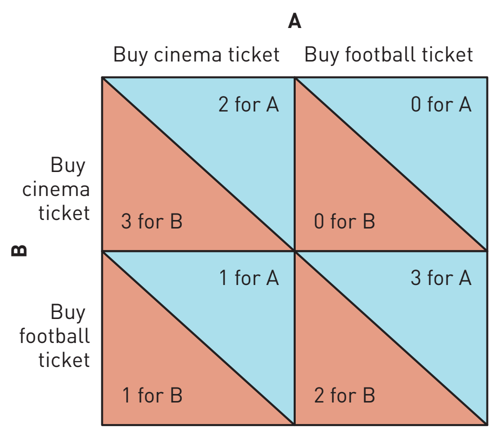

Assume that players *A* and *B* buy their tickets separately and simultaneously. Each must decide what to do knowing the available choices and payoffs but not what the other has actually chosen. Each player believes the other to be rational and self-interested.

- **a.** Does either player have a dominant strategy?
- **b.** How many potential equilibria are there? (**Hint:** To see whether a given combination of strategies is an equilibrium, ask whether either player could get a higher payoff by changing his or her strategy.)
- **c.** Is this game a prisoner's dilemma? Explain.
- **d.** Suppose player *A* gets to buy her ticket first. Player *B* does not observe *A*'s choice but knows that *A* chose first. Player *A* knows that player *B* knows she chose first. What is the equilibrium outcome?
- **e.** Suppose the situation is similar to part (d), except that player *B* chooses first. What is the equilibrium outcome?

3. Blackadder and Baldrick are rational, self-interested criminals imprisoned in separate cells in a dark medieval dungeon. They face the prisoner's dilemma displayed in the matrix below.

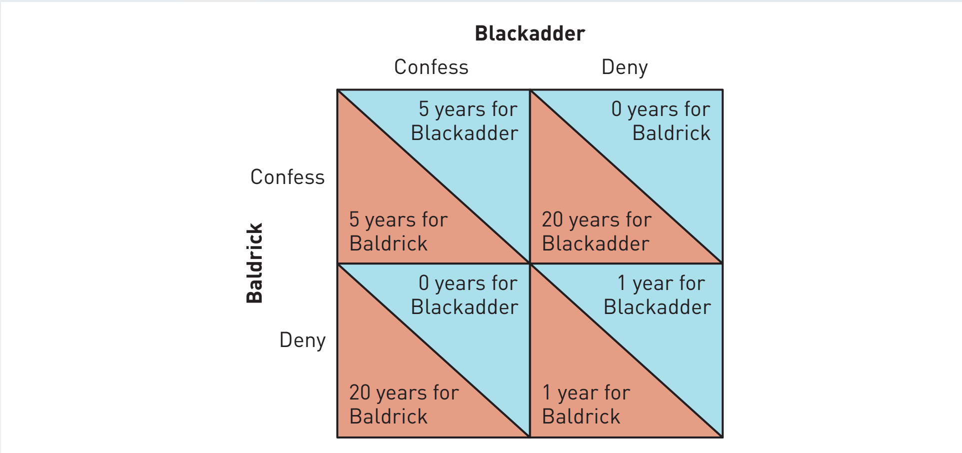

Assume that Blackadder is willing to pay 1,000 ducats for each year by which he can reduce his sentence below 20 years. A corrupt jailer tells Blackadder that before he decides whether to confess or deny the crime, he can tell him Baldrick's decision. How much is this information worth to Blackadder?

4. The owner of a thriving business wants to open a new office in a distant city. If he can hire someone who will manage the new office honestly, he can afford to pay that person a weekly salary of €2,000 (€1,000 more than the manager would be able to earn elsewhere) and still earn an economic profit of €800. The owner's concern is that he will not be able to monitor the manager's behaviour and that the manager will therefore be in a position to embezzle money from the business. The owner knows that if the remote office is managed dishonestly, the manager can earn €3,100 while causing the owner an economic loss of €600 per week.

a. If the owner believes that all managers are narrowly self-interested income-maximisers, will he open the new office?

b. Suppose the owner knows that a managerial candidate is a devoutly religious person who condemns dishonest behaviour and who would be willing to pay up to €15,000 to avoid the guilt she would feel if she were dishonest. Will the owner open the remote office?

5. Imagine yourself sitting in your car in a university car park that is currently full, waiting for someone to pull out so that you can park your car. Somebody pulls out, but at the same moment a driver who has just arrived overtakes you in an obvious attempt to park in the vacated spot before you can. Suppose this driver would be willing to pay up to €10 to park in that spot and up to €30 to avoid getting into an argument with you. (That is, the benefit of parking is €10, and the cost of an argument is €30.) At the same time the other driver guesses, accurately, that you too would be willing to pay up to €30 to avoid a confrontation and up to €10 to park in the vacant spot.

a. Model this situation as a two-stage decision tree in which the other driver's bid to take the space is the opening move and your strategies are (1) to protest and (2) not to protest. If you protest (initiate an argument), the rules of the game specify that the other driver has to let you take the space. Show the payoffs at the end of each branch of the tree.

b. What is the equilibrium outcome?

c. What would be the advantage of being able to communicate credibly to the other driver that your failure to protest would be a significant psychological cost to you?

6. Newfoundland's fishing industry has declined sharply due to overfishing, even though fishing companies were supposedly bound by a quota agreement. If all fishing companies had abided by the agreement, yields could have been maintained at high levels.

a. Model this situation as a prisoner's dilemma in which the players are Company *A* and Company *B* and the strategies are to keep the quota and break the quota. Include appropriate payoffs in the matrix. Explain why overfishing is inevitable in the absence of effective enforcement of the quota agreement.

b. Provide another environmental example of a prisoner's dilemma.

c. In many potential prisoners' dilemmas, a way out for a would-be co-operator is to make reliable character judgements about the trustworthiness of potential partners. Explain why this solution is not available in many situations involving degradation of the environment.

7. Consider the following game, called 'matching pennies', which you are playing with a friend. Each of you has a penny hidden in your hand, facing either heads up or tails up (you know which way the one in your hand is facing). On the count of 'three' you simultaneously show your pennies to each other. If the face-up side of your coin matches the face-up side of your friend's coin, you get to keep the two pennies. If the faces do not match, your friend gets to keep the pennies.

a. Who are the players in this game? What are each player's strategies? Construct a payoff matrix for the game.

b. Is there a dominant strategy? If so, what?

c. Is there an equilibrium? If so, what?

8. Consider the following game. Harry has four 20 pence pieces. He can offer Sally from one to four of them. If she accepts his offer, she keeps the coins Harry offered her and Harry keeps the others. If Sally declines Harry's offer, they both get nothing. They play the game only once, and each cares only about the amount of money he or she ends up with.

a. Who are the players? What are each player's strategies? Construct a decision tree for this ultimatum bargaining game.

b. Given their goal, what is the optimal choice for each player?

9. Two aeroplane manufacturers are considering the production of a new product, a 550-passenger jet. Both are deciding whether to enter the market and produce the new plane. The payoff matrix is as shown (payoff values are in million euros).

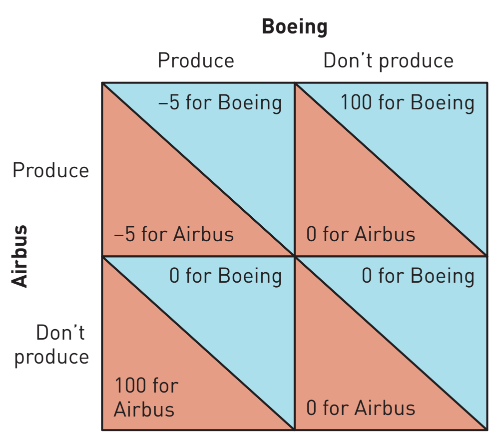

The implication of these payoffs is that the market demand is large enough to support only one manufacturer. If both firms enter, both will sustain a loss.

a. Identify two possible equilibrium outcomes in this game.

b. Consider the effect of a subsidy. Suppose the European Union decides to subsidise the European producer, Airbus, with a subsidy of €25 million if it enters the market. Revise the payoff matrix to account for this subsidy. What is the new equilibrium outcome?

c. Compare the two outcomes (pre- and post-subsidy). What qualitative effect does the subsidy have?

10.* Jill and Jack both have two pails that can be used to carry water down from a hill. Each makes only one trip down the hill, and each pail of water can be sold for €5. Carrying the pails of water down requires considerable effort. Both the children would be willing to pay €2 each to avoid carrying one bucket down the hill and an additional €3 to avoid carrying a second bucket down the hill.

a. Given market prices, how many pails of water will each child fetch from the top of the hill?

b. Jill and Jack's parents are worried that the two children don't co-operate enough with one another. Suppose they make Jill and Jack share their revenues from selling the water equally. Given that both are self-interested, construct the payoff matrix for the decisions Jill and Jack face regarding the number of pails of water each should carry. What is the equilibrium outcome?

Answers to in-chapter exercises

10.1 No matter what Alitalia does, Lufthansa will do better to raise ad spending. No matter what Lufthansa does, Alitalia will do better to leave ad spending the same. So each player will play its dominant strategy: Lufthansa will raise its ad spending, and Alitalia will leave its ad spending the same, as in the table below.

Lufthansa's Choice	Alitalia's Choice: Raise ad spending	Alitalia's Choice: Leave ad spending the same
Raise ad spending	Alitalia gets €2,000; Lufthansa gets €3,000	Alitalia gets €3,000; Lufthansa gets €4,000
Leave ad spending the same	Alitalia gets €3,000; Lufthansa gets €2,000	Alitalia gets €4,000; Lufthansa gets €3,000

10.2 In game 1, no matter what Chrysler does, General Motors (GM) will do better to invest, and no matter what GM does, Chrysler will do better to invest (see the table below). Each has a dominant strategy, but in following it, each does worse than if it had not invested. So game 1 is a prisoner's dilemma. In game 2, no matter what Chrysler does, GM again will do better to invest; but no matter what GM does, Chrysler will do better *not* to invest. Each has a dominant strategy, and in following it, each gets a payoff of 10 – that is, 5 more than if each had played its dominated strategy. So game 2 is not a prisoner's dilemma.

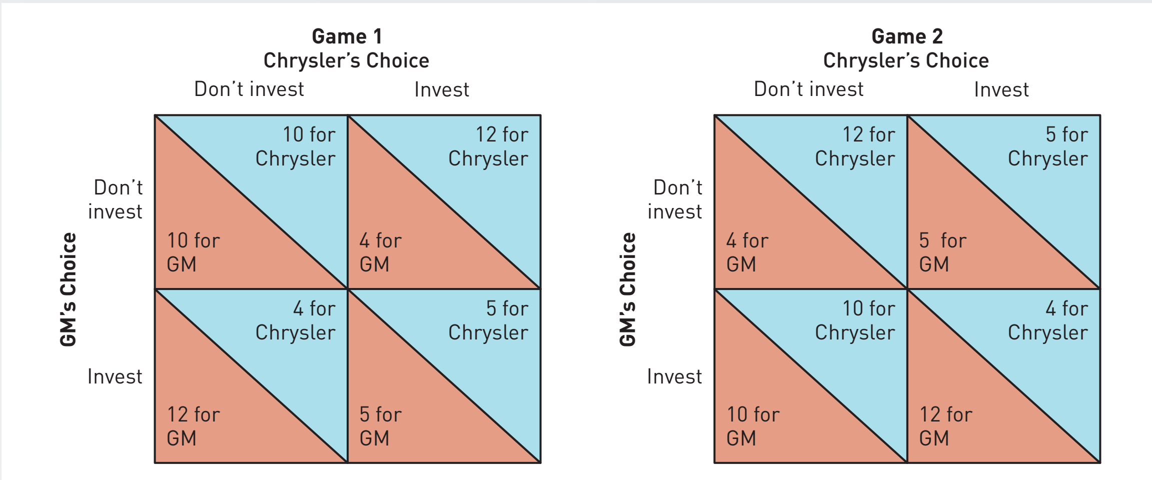

10.3 People may well collectively prefer the absence of mobile phones while personally wishing to be able to use them. This another prisoner's dilemma problem: everyone would be better of if no one used a phone, but each person's dominant strategy is to use it. In these circumstances an airline could well conclude that its services would have a competitive advantage if it prohibited mobile phone usage. It could even seek to introduce a 'hurdle' to price discriminate on the basis of permission to use mobile phones in a separate cabin. It would be surprising if the absence of a legal prohibition resulted in a free for all. We shall see.

10.4 Smith assumes that Jones will choose the branch that maximises his payoff, which is the bottom branch at either *B* or *C*. So Jones will choose the bottom branch when his turn comes, no matter what Smith chooses. Since Smith will do better (60) on the bottom branch at *B* than on the bottom branch at *C* (50), Smith will choose the top branch at *A*. So the equilibrium in this game is for Smith to choose the top branch at *A* and Jones to choose the bottom branch at *B*. Smith gets 60, and Jones gets 105. If Jones could make a credible commitment to choose the top branch no matter what, both would do better. Smith would choose the bottom branch at *A* and Jones would choose the top branch at *C*, giving Smith 500 and Jones 400 (see the decision tree below).

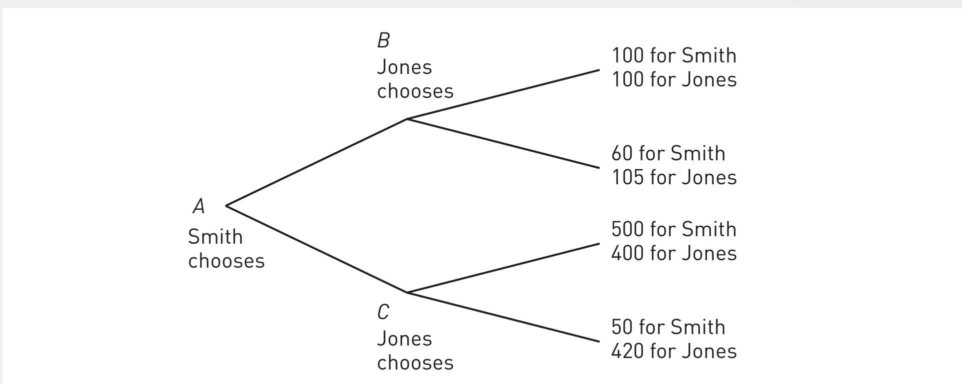

9707 CH10.QXD 1/5/06 2:10 PM Page 311

10.5 Given its expected demand for super-jumbos over the life of the production run, Airbus thinks that it makes a profit by producing and selling 400 A380s and reckons that the maximum possible volume of profitable sales is 500 units. It would maximise its profits if it could sell 450. Suppose it now commits itself to sell 450, and expects to get €300 million per plane. The position facing Boeing is that it can sell planes only if it accepts a maximum price (for one plane) of €300 million, and to sell more (given that Airbus is committed to producing 450) can do so only if it is prepared to accept prices indicated by the residual demand curve in the diagram below. However, the residual demand curve lies below the long-run average cost curve, so it can never be profitable for Boeing to commit itself (incur sunk costs) to producing its planes once Airbus has done so provided it does not believe that Airbus has underestimated demand (in fact, Boeing believed that Airbus was being wildly optimistic).

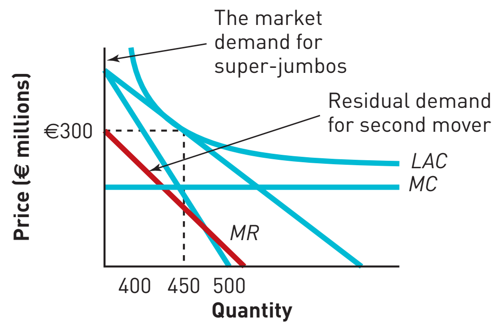

References

Barnstone, W. (1993) *With Borges on an Ordinary Evening in Buenos Aires* (Champaign, IL: University of Illinois Press).

Dodgson, J. S., Y. Katsoulakos and C. R. Newton (1993) 'An Application of the Economic Modeling Approach to the Investigation of Population', *Journal of Transport Economics and Policy*, pp. 153–70.

Von Neumann, J. and O. Morgenstern (1944) *The Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press).