who live nearby. These last two features may allow the corner shop to charge a few pence more for a jar of coffee than the supermarket in the main shopping area some distance away. But, if prices are higher by more than a few pence, even shoppers who live nearby will make the trip to the supermarket.

As with most definitions, the lines between these types of market structure are a little blurred. A major reason is the ambiguity about the relevant definition of the market. Is British Rail a monopoly in railways or an oligopolist in transport? Similarly, when a country trades in a competitive world market, even the sole domestic producer may have little influence on market price. We can never fully remove these ambiguities, but Table 9-1 shows some things to bear in mind as we proceed through this chapter. Notice that the table includes the ease with which new firms can enter the industry. This has a crucial bearing on the ability of existing firms to maintain high prices and supernormal profits in the long run. We have already seen its importance when contrasting the long-run behaviour of perfectly competitive industries and pure monopolies.

9-1 WHY MARKET STRUCTURES DIFFER

We have already drawn attention to the influence of government legislation on market structure. In the UK the nationalized industries, for example coal, rail, and electricity generation, are legal monopolies; they are the sole licensed producers. Patent laws may confer temporary monopoly on producers of a new process. Ownership of a raw material, as in the case of de Beers and diamonds, may also confer monopoly status on a single producer. Having noted these interesting special cases, we now develop a general theory of how the economic factors of demand and cost interact to determine the likely structure of a particular industry.

The motor car industry is not an oligopoly one day but perfectly competitive the next. It is in long-run influences that we must seek the causes of different market structures. Similarly, although a particular firm may have a temporary advantage in technical know-how or workforce skill, in the long run one firm can hire another's workers and learn its technical secrets. In the long run all firms or potential entrants to an industry essentially have access to the same cost curves.

Figure 9-1 shows the demand curve DD for the output of an industry. Suppose first that in the long run all firms and potential entrants face the average cost curve LAC_1. At the price P_1, free entry and exit ensures that each firm produces q_1. Given the demand curve DD, the industry output is Q_1 and the industry can support N_1 firms where $N_1 = Q_1/q_1$. If q_1, the minimum average cost output on LAC_1, lies sufficiently far to the left relative to DD, then N_1 will be a very large number of firms. It will be reasonable for each firm to act on the assumption that it has a trivial effect on industry supply and market price. We have discovered a perfectly competitive industry.

Now suppose that each firm has the cost curve

TABLE 9-1
MARKET STRUCTURE

		IMPERFECT COMPETITION		
CHARACTERISTIC	PERFECT COMPETITION	MONOPOLISTIC COMPETITION	OLIGOPOLY	MONOPOLY
Number of firms	Many	Many	Few	One
Ability to affect price	None	Limited	Some	Considerable
Entry barriers	None	None	Some	Complete
Example	Fruit stalls in Covent Garden	Corner grocer	Cars	de Beers

FIGURE 9-1 DEMAND, COSTS, AND MARKET STRUCTURE. DD is the industry demand curve. In a competitive industry, minimum efficient scale occurs at an output level q_1 when firms have average cost curves LAC_1. The industry can support a very large number of firms whose total output is Q_1 at the price P_1. When LAC_3 describes average costs, the industry will be a natural monopoly. When a single firm produces the entire industry output, no other firm can break into the market and make a profit. For intermediate positions such as LAC_2 the industry can support a few firms in the long run, and no single firm can profitably meet the entire demand. The industry will be an oligopoly.

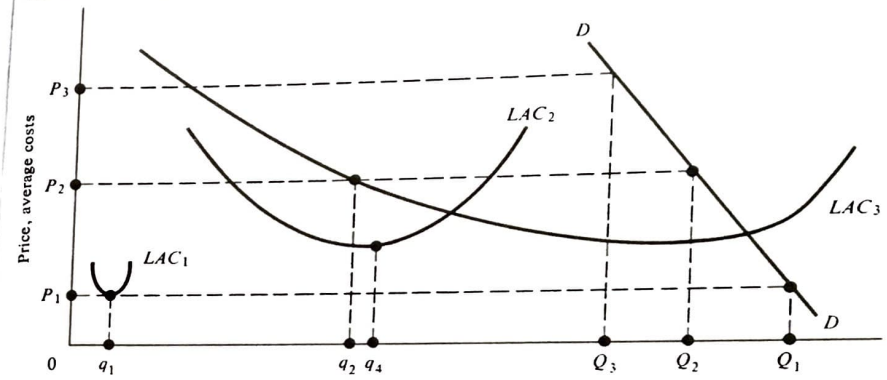

LAC_3. Economies of scale are very large relative to the market size. The lowest point on LAC_3 occurs at an output large relative to the demand curve DD. Suppose initially there are two producers each producing q_2. Market output Q_2 is twice as large. The market clears at P_2 and both firms break even. However, if one firm expands a little its average costs will fall. It will also bid the price down. With lower average costs, that firm will survive and the other firm will lose money. The firm that expands will gobble up the whole market, undercut its competitor, and eventually drive the other firm out of business.

We have discovered an industry that is a natural monopoly. Suppose that Q_3 is the output at which its marginal cost and marginal revenue coincide. The price is P_3 and the natural monopoly makes supernormal profits. Yet there is no room in the industry for other firms with access to the same LAC_3 curve. A new entrant needs a large output to get average costs down. Extra output on this scale would so depress the price that both firms would make losses. The potential entrant is powerless to break in. The natural monopolist can completely disregard the threat of entry.

Finally, we show the LAC_2 curve with more economies of scale than a competitive industry but fewer than a natural monopoly. This industry will support at least two firms enjoying economies of scale near the lowest point of their LAC_2 curves. It will be an oligopoly. Attempts to expand either firm's output beyond q_4 quickly encounter decreasing returns to scale and prevent it from expanding to drive its competitor out of business.

In Chapter 7 we introduced the notion of the minimum efficient scale.

> The *minimum efficient scale* is the output at which a firm's long-run average cost curve stops falling.

We now see that the crucial determinant of market

structure is the output at minimum efficient scale relative to the size of the total market as represented by the demand curve. Table 9-2 summarizes our discussion. It is the interaction of market size and the output at minimum efficient scale that matters. When the demand curve shifts to the left, an industry previously supporting many firms may have room for only a few. Similarly, an increase in fixed costs which increases the output at minimum efficient scale will reduce the number of producers. In the 1950s there were a large number of European aircraft manufacturers, and several even in the UK. Today, the research and development costs of a major commercial airliner are enormous. Apart from the co-operative European venture Airbus Industrie, which has been heavily subsidized by European governments, only the American giants Boeing, Lockheed, and McDonnell-Douglas survive.

Table 9-2 does not explicitly show monopolistic competition. In one sense such industries lie midway between oligopoly and perfect competition. But it is the fact that monopolistic competitors all supply slightly different products, such as the location in which you do your shopping, that makes them special.

Evidence on Market Structure

The larger the minimum efficient scale relative to the market size, the fewer will be the number of plants – and probably the number of firms – in the industry. What is the number of plants (*NP*) operating at minimum efficient scale that the current market size could allow? In Chapter 7 we discussed how economists have tried to estimate the minimum efficient scale for plants in different industries. By looking at the total quantity of consumption of a product we can estimate the market size. Hence we can construct estimates of *NP* for each industry.

How do we measure how many firms there are in an industry? Even industries that essentially have only a few very large firms may have some small firms on the fringe. Large transport costs may allow a few small local suppliers to survive in very isolated parts of the country. The number of firms in the industry tells us nothing about their size or importance. It might be a misleading indicator of the essential structure of the industry. For this reason, economists use the *N*-firm concentration ratio to measure the number of important firms in the industry.

> The *N-firm concentration ratio* is the market share of the largest *N* firms in the industry.

Thus the 3-firm concentration ratio tells us the percentage of the total market supplied by the largest three firms in the industry. If there are basically only three firms that matter, they will supply almost 100 per cent of the total market for the product. If the industry is perfectly competitive, the largest three firms will still account for only a tiny share of the total market for the product.

Table 9-3 looks at the evidence for selected industries in three European countries, the UK, France, and West Germany. *CR* is the 3-firm concentration ratio, the market share of the top three firms. *NP* is the number of plants at minimum efficient scale which the market size would allow. Nothing guarantees that all plants are being operated at minimum efficient scale. Nevertheless, if our theory of market structure is correct, industries with large economies of scale relative to the market size – a very low value of *NP* – should exhibit a large *CR*. Such industries should have only a few important firms. The top three firms should account for most of the market. Conversely, where *NP* is very high, economies of scale are relatively unimportant and the largest three firms should control a much smaller market share. *CR* should be much lower.

Table 9-3 confirms that our theory of market structure is compatible with the facts. In industries such as refrigerator and cigarette manufacture there is room for only very few plants operating at minimum efficient scale, and these industries exhibit higher degrees of concentration. The largest three firms control almost the whole market. Economies of scale are still substantial in industries such as brewing and petroleum refining and the top three firms control around half the market on average. In fabric manufacture about fifty plants can operate at minimum efficient scale and the top three firms control only about one-quarter of the market.

Industries such as shoe manufacture quickly encounter rising average cost curves; they have room for a large number of factories operating at minimum efficient scale, and consequently are much closer to competitive industries. The top three firms in shoe manufacturing control less than one-fifth of the market.

TABLE 9-2
DEMAND, COSTS, AND MARKET STRUCTURE

	MINIMUM EFFICIENT SCALE RELATIVE TO MARKET SIZE		
	TINY	INTERMEDIATE	LARGE
ndustrial structure	Perfect competition	Oligopoly	Natural monopoly

TABLE 9-3
CONCENTRATION AND SCALE ECONOMIES IN THREE EUROPEAN COUNTRIES, 1967–70*

	UK		FRANCE		W. GERMANY	
INDUSTRY	*CR*	*NP*	*CR*	*NP*	*CR*	*NP*
Refrigerators	65	1	100	2	72	3
Cigarettes	94	3	100	2	94	3
Petroleum refining	79	8	60	7	47	9
Brewing	47	11	63	5	17	16
Fabrics	28	57	23	57	16	52
Shoes	17	165	13	128	20	197

*CR = % market share of 3 largest firms; *NP* = market size divided by output of minimum efficient scale.
Source: F. M. Scherer *et al.*, *The Economics of Multiplant Operation*, Harvard University Press, 1975, and F. M. Scherer, *Industrial Market Structure and Economic Performance*, Rand McNally, 1980.

9-2 MONOPOLISTIC COMPETITION

The essence of oligopoly is interdependence. Large firms must guess what their large rivals are up to. Before turning to this exciting branch of economic analysis, however, we begin with a simpler case.

The theory of monopolistic competition[1] envisages a large number of quite small firms so that each firm can neglect the possibility that its own decisions provoke any adjustment in other firms' behaviour. We also assume free entry and exit from the industry in the long run. In these respects the framework resembles our earlier discussion of *perfect* competition. What distinguishes monopolistic competition is that each firm faces a *downward*-sloping demand curve.

Monopolistic competition describes an industry in which each firm can influence its market share to some extent by changing its price relative to its competitors. Its demand curve is not horizontal because different firms' products are only limited substitutes. We have given one example, the location of corner grocers. A lower price attracts some customers from another shop, but each shop will always have some local customers for whom

[1] This theory was independently invented in the early 1930s by E. H. Chamberlin in the United States and by Joan Robinson in Britain.

the convenience of a nearby shop is more important than a few pence on the price of a jar of coffee.

Economists sometimes say that monopolistically competitive industries exhibit *product differentiation*. For corner grocers this differentiation is based on location, but in other cases it is based on brand loyalty. The special features of a particular restaurant or hairdresser may allow that firm to charge a slightly different price from other producers in the industry without losing all its customers or completely taking over the entire market for the industry.

Although brand loyalty and produce differentiation may also be important in many other industries these need not be monopolistically competitive. Brand loyalty limits the substitution between Ford and Vauxhall in the car industry but, with so few producers, the key feature of the industry remains the oligopolistic interdependence of the decisions of different firms. Monopolistic competition requires not merely product differentiation, but also limited opportunities for economies of scale so that there are a great many producers who can largely neglect their interdependence with any particular rival. Hence many of the best examples of monopolistic competition are service industries where economies of scale are small.

The industry demand curve shows the total industry output which would be demanded at each price if every firm in the industry charged that price. The market share of each firm depends on the number of firms in the industry and on the price it charges. For a given number of firms, a shift in the industry demand curve will shift the demand curve for the output of each individual firm. For a given industry demand curve, an increase (decrease) in the number of firms in the industry will shift the demand curve of each firm to the left (right) as its market share falls (rises). But each firm faces a downward-sloping demand curve. For a given industry demand curve, number of firms, and price charged by all other firms, a particular firm can increase its market share to some extent by charging a lower price and inducing some consumers to switch to its particular product.

Figure 9.2 shows the supply decision of a firm. Given its own demand curve DD and marginal revenue curve MR the firm produces Q_0 at a price P_0 making short-run profits equal to $Q_0 \times (P_0 - AC_0)$. In the long run these profits attract new entrants, who dilute the market share of each firm in the industry, shifting their demand curves to the left. Entry stops when each firm's demand curve has shifted so far to the left that price equals average cost and firms are just breaking even. In Figure 9.2 this occurs when demand has shifted to DD' and the firm produces Q_1 at a price P_1 to reach the tangency equilibrium at F.

> In monopolistic competition the long-run *tangency equilibrium* occurs where each firm's demand curve is tangent to (just touches) its AC curve at the output level at which MC equals MR. Each firm is maximizing profits but just breaking even. There is no further entry or exit.

Notice two things about the firm's long-run equilibrium at F. First, the firm is *not* producing at minimum average cost. It has excess capacity. It could reduce average costs by further expansion. However, its marginal revenue would be so low this would not be profitable. Second, the firm retains some monopoly power because of the special feature of its particular brand or location. Price exceeds marginal cost.

This second observation helps explain why firms are usually eager for new customers prepared to buy additional output at the *existing* price. In Robert Bishop's phrase, it explains why 'we are a race of eager sellers and coy buyers'. It is purchasing agents who get Christmas presents from sales reps, not the other way round.[2] Remarkably enough, under perfect competition the firm does not care if another buyer shows up at the existing price. With price equal to marginal cost, the firm is already selling as much as it wants.

The theory of monopolistic competition yields interesting insights when there are many goods each of which is a close but not perfect substitute for the other. For example, it explains why Britain exports Jaguars and Rovers to Germany and Sweden but simultaneously imports Volvos and Mercedes. There are large economies of scale in making cars. In the absence of trade the domestic car market would have room for only a few varieties. Producing a large number of brands at low output would enormously raise average costs. International trade allows each country to specialize in a few types of car and produce a much larger output of that brand than the home market alone could support. By swapping these cars between countries, it is possible to give consumers a wider range from which to choose while allowing each individual producer to enjoy economies of scale and hold prices down.

[2] Quotation from Professor Bishop's unpublished magnum opus 'Microeconomic Theory', on which generations of MIT economics graduates were raised.

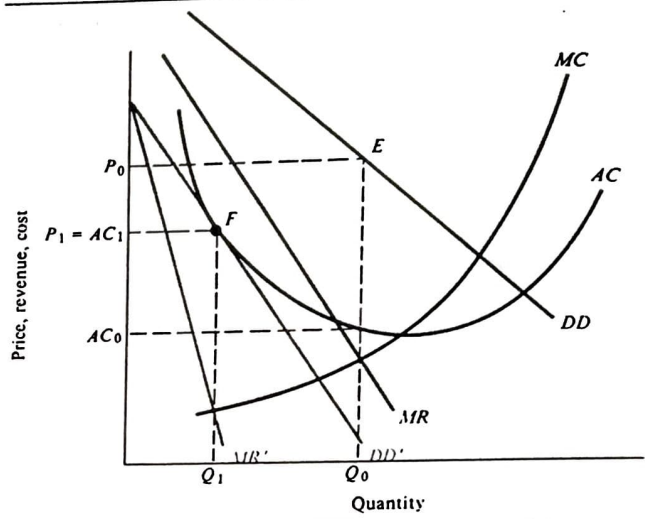

FIGURE 9-2 EQUILIBRIUM FOR A MONOPOLISTIC COMPETITOR. In the short run the monopolistic competitor faces the demand curve DD and sets MC equal to MR to produce Q_0 at a price P_0. Profits are $Q_0 \times (P_0 - AC_0)$. Profits attract new entrants and shift each firm's demand curve to the left. When the demand curve reaches DD' we reach the long-run tangency equilibrium at F. The firm sets MC equal MR' to produce Q_1 at which P_1 equals AC_1. Firms are breaking even and there is no further entry.

9-3 OLIGOPOLY AND INTERDEPENDENCE

Under perfect competition or monopolistic competition, there are so many firms in the industry that no single firm need worry about the effect of its own actions on rival firms. However, the very essence of an oligopolistic industry is the need for each firm to consider how its own actions will affect the decisions of its relatively few competitors.

Although in the last chapter we used a hypothetical example of a *monopoly* airline, in practice of course airlines are oligopolists. Even on the popular transatlantic routes, British Airways, Air France, and TWA have significant market shares, and the position of each of their demand curves depends critically on how their rivals behave and can be induced to behave. In contemplating a cut-price deal, each airline needs to consider whether or not other airlines will follow suit. Similarly, when new airlines try to break into the market by offering cheap fares – Laker in the 1970s and People's Express in the 1980s – these entrants' prospects depend on how existing airlines respond. Laker, for example, may have miscalculated how other airlines would react, failing to foresee the extent to which they would cut prices to drive it out of business.

What makes oligopoly so fascinating is that the optimal supply decision of a particular firm depends on its guess about how its rivals will react. Exciting recent developments in economics shed important insight into what constitutes a smart guess. First, however, we introduce the basic tension between competition and collusion which lies beneath all oligopolistic situations.

Collusion is an explicit or implicit agreement between existing firms to avoid competition with one another.

Initially, for simplicity, we neglect the possibility of entry and focus only on the behaviour of existing firms.

The Profits from Collusion

The existing firms will maximize their *joint* profits if they behave as if they were a multi-plant monopolist. A monopolist or sole decision-maker would organize the output from the industry to maximize total profits. Hence, if the few producers in an industry collude to behave as if they were a monopolist, their *total* profit will be maximized.

Figure 9-3 shows an industry where each firm, and the entire industry, has constant average and marginal costs at the level P_c. In the last chapter we saw that a competitive industry would produce Q_c at a price P_c but a multi-plant monopolist would maximize profits by producing Q_m at a price P_m. If the oligopolists collude jointly to produce Q_m we say they are acting as a *collusive monopolist*. Having thus decided industry output, there will then be some negotiation backstage to divide up output and profits between individual firms.

However, it is hard to stop individual firms cheating on the collective agreement. In Figure 9-3 joint profit is maximized when aggregate output is restricted to Q_m and the price forced up to P_m. Yet each firm can expand at marginal cost P_c. If one firm expands production by undercutting the agreed price P_m, its profits will rise since its marginal revenue will exceed its marginal cost. But this firm's gain is at the expense of its collusive partners. Industry output is now higher than Q_m, total profits are lower, and other firms must suffer.

Hence oligopolists are torn between the desire to collude, thus maximizing joint profits, and the desire to compete, in the hope of increasing market share and profits at the expense of rivals.

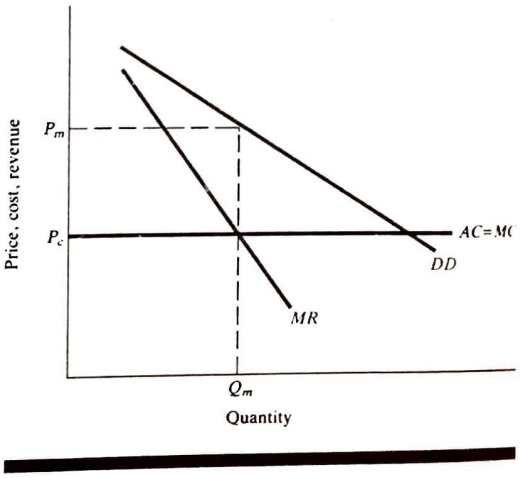

FIGURE 9-3 COLLUSION VERSUS COMPETITION. By colluding to restrict industry output to Q_m, joint profits are maximized and equal to those which a multi-plant monopolist would obtain. But each firm, with a marginal cost of P_c, has an incentive to cheat on the collusive agreement and expand its output.

Yet if all firms compete, joint profits will be low and no firm is likely to do very well. Therein lies the dilemma.

Cartels

Collusion or co-operation between firms is easiest when formal agreements are legally permitted. Such arrangements are called *cartels*. In the late nineteenth century cartels were common, and they agreed market shares and prices in many industries. Such practices are now outlawed in Europe, the United States, and many other countries. Although there are usually large penalties for being caught, informal agreements and secret deals in smoke-filled rooms are not unknown even today.

The most famous cartel is OPEC, the Organization of Petroleum Exporting Countries. Active since 1973, its members (of which the UK is not one) meet regularly to set price and output levels. Initially, OPEC was very successful in organizing quantity reductions to force up the price of oil. Real OPEC revenues rose 340 per cent between 1974 and 1980. Yet almost from the start, many economists have predicted that OPEC, like most cartels, would quickly collapse. Usually, the incentive to cheat is too strong to resist, and once somebody breaks ranks others tend to follow.

In practice, one reason OPEC was successful for so long was the willingness of Saudi Arabia, the largest oil producer, to restrict its output further when smaller members insisted on expansion. By 1986, however, Saudi Arabia was no longer prepared to play by these rules, and refused to prop up the price any longer. The oil price collapsed from just under $30 to $9 a barrel before recovering a little. Whether this signals the end of OPEC as a major force we shall discuss shortly.

The Kinked Oligopoly Demand Curve

Collusion is much harder if there are many firms in the industry, if the product is not standardized, and if demand and cost conditions are changing rapidly. In the absence of collusion, each firm's demand curve depends on how competitors react. Firms must guess how their rivals will behave. Before undertaking a serious analysis of how firms might make intelligent guesses, we introduce a simple model which highlights the key feature of this interdependence.

Suppose that each firm believes that its own price cut will be matched by all other firms in the industry but that an increase in its own price will induce no price response from competitors.[3] Figure 9-4 shows the demand curve *DD* that each firm would then face. The current price is P_0 and the firm is producing Q_0. Since competitors do not follow suit, a price increase will lead to a large loss of market share to other firms. The firm's demand curve is elastic above *A* at prices above the current price P_0. Conversely, a price cut is matched by other firms and market shares are unchanged. Sales increase only because the industry as a whole moves down the market demand curve as prices fall. The demand curve *DD* is much less elastic for price reductions from the initial price P_0.

The key feature of Figure 9-4 is that the marginal revenue curve *MR* is discontinuous at the output Q_0. Below Q_0 the elastic part of the demand curve is relevant, but at the output Q_0 the firm suddenly encounters the inelastic portion of its kinked demand curve and marginal revenue suddenly falls. Q_0 is the profit-maximizing output for the firm, given its belief about how competitors will respond.

The model has one important implication. Suppose the *MC* curve of a single firm shifts up or down by a small amount. Since the *MR* curve has a discontinuous vertical segment at the output Q_0, it will remain optimal to produce Q_0 and charge the price P_0. In contrast, a monopolist facing a continuously downward-sloping *MR* curve would adjust quantity and price when the *MC* curve shifted. The kinked demand curve model may

[3] This model was independently invented in 1939 by Paul Sweezy in the United States and R. L. Hall and C. J. Hitch in the UK.

FIGURE 9-4 THE KINKED DEMAND CURVE. An oligopolist believes rivals will match price cuts but not price rises. The oligopolist's demand curve is kinked at *A*. Price rises lead to a large loss of market share, but price cuts increase quantity only by increasing industry sales. Marginal revenue is discontinuous at Q_0. The oligopolist produces Q_0, the output at which *MC* crosses the *MR* schedule.

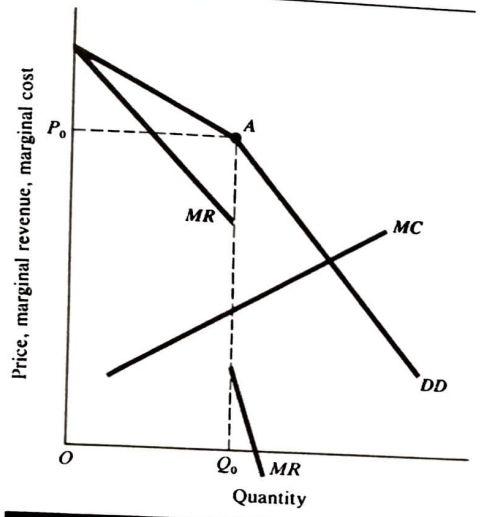

explain the empirical finding that firms do not always adjust prices when they face a change in costs.[4]

The model does not explain what determines the initial price P_0. One possible interpretation is that it is the collusive monopoly price. Each firm believes that an attempt to undercut its rivals will provoke them to co-operate among themselves and retaliate in full. However, its rivals will be happy for it to charge a higher price and see its market share destroyed. The model can be applied in other circumstances where there is less co-operation between firms but then we require an additional theory to explain what determines the initial price P_0.

One advantage of interpreting P_0 as the collusive monopoly price is that it contrasts the effect of a cost change for a single firm and a cost change for all firms. The latter will shift the marginal cost curve up for the industry as a whole and increase the collusive monopoly price. Each firm's kinked demand curve will shift upwards since the monopoly price P_0 has increased. Thus we can reconcile the stickiness of a single firm's prices with respect to changes in its own costs alone, and the speed with which the entire industry marks up prices when all firms' costs are increased by higher taxes (as in the cigarette industry) or inflationary wage settlements in the whole industry.

9-4 GAME THEORY AND INTERDEPENDENT DECISIONS

A good poker player sometimes bluffs. Sometimes you can clean up with a bad hand, provided your opponents misread it for a good hand. Similarly, by having bluffed in the past and been caught, you may persuade them to keep betting even when you have a terrific hand.

Like poker players, oligopolists have to try to second-guess their rivals' moves to determine their own best action. To study how such interdependent decisions are made, we use *game theory*.

> A *game* is a situation in which intelligent decisions are necessarily interdependent.

The *players* in the game try to maximize their own *payoffs*. In an oligopoly, the firms are the players and their payoffs are their profits in the long run. Each player must choose a strategy.

> A *strategy* is a game plan describing how the player will act or *move* in every conceivable situation.

Being a pickpocket is a strategy. Lifting a particular wallet is a move.

In game theory, as elsewhere in economics, we are interested in equilibrium. In most games, each player's best strategy depends on the strategies chosen by other players. It is silly to be a pickpocket in an area where the police have TV cameras. Equilibrium occurs when each player chooses the best strategy, *given* the strategies being followed by other players. This description of equilibrium, invented by John Nash, is called Nash equilibrium. Nobody in the game wants to change their strategy, since other people's strategies have already been figured into the calculation of each player's best strategy.

Sometimes, but not usually, a player's best strategy is independent of those chosen by others. If so, it is called a *dominant strategy*. To introduce the use of game theory in understanding oligopoly, we begin with an example in which each player has a dominant strategy.

Collude or Cheat? The Cartel Example Again

Figure 9-5 shows a game[5] which we can imagine is between the only two members of a cartel like OPEC. Each firm can select a high-output or low-output strategy. In each box of Figure 9-5 the blue number shows firm A's profits and the black number, firm B's profits for that output combination of the two firms.

When both have high output, industry output is high, the price is low, and each firm makes a small profit of 1. When each has low output, the outcome is more like the collusive monopoly of Figure 9-3. Prices are high and each firm does better, making a profit of 2. Each firm does best (a profit of 3) when it alone has high output; for then, the other firm's low output helps hold down industry output and keep up the price. In this situation we assume the low-output firm makes a profit of 0.

Now we can see how the game will unfold. Consider firm A's decision. If firm B has a high-output strategy, firm A does better also to have high output. In the two left-hand boxes of Figure 9-5, firm A gets a profit of 1 by choosing high but a profit of 0 by choosing low. Now suppose firm B chooses a low-output strategy. From the two right-hand boxes, Firm A still does better by choosing high, since this yields it a profit of 3 whereas low yields it a profit of only 2. Hence firm A has a dominant strategy. Whichever strategy firm B adopts, firm A does better to choose a high-output strategy.

Firm B also has a dominant strategy to choose high output. Use Figure 9-5 to check for yourself that A does better to go high whichever strategy B selects. Since both firms choose high, the equilibrium of the game is the top left-hand box. Each firm gets a profit of 1.

Yet both firms would do better, getting a profit of 2, if they colluded to form a cartel and both produced low output—the bottom right-hand box. But neither can afford to take the risk of going low. Suppose firm A goes low. Firm B, comparing the two boxes in the bottom row, will then go high, preferring a profit of 3 to a profit of 2. And firm A will get screwed, earning a profit of 0 in that event. Firm A can figure all this out in advance, which is why its dominant strategy is to go high.

This is a particularly clear illustration of the tension between collusion and competition which we discussed earlier. In this example, it appears

FIGURE 9-5 THE PRISONERS' DILEMMA GAME. The blue and black numbers in each box indicate profits to firms A and B, respectively. Whether B pursues high or low output, A makes more profit going high; so does B, whichever A adopts. In equilibrium both go high. Yet both would make greater profits if both went low!

		Firm B output: High		Firm B output: Low	
Firm A output: High		1	1	3	0
Firm A output: Low		0	3	2	2

[4] This argument was first explored by George Stigler, 'The Kinky Oligopoly Demand Curve and Rigid Prices', *Journal of Political Economy*, October 1947.

[5] The game is usually called the Prisoners' Dilemma, because it was first used to analyse the choices facing two people arrested and in different cells, each of whom could plead guilty or not guilty to the only crime that had been committed. Each prisoner would plead innocent if only he or she knew the other would plead guilty.