9781405176323_4_009.qxd 10/19/07 8:10 PM Page 193

Part III
Oligopoly and Strategic Interaction

Part III starts our analysis of markets populated by more than one but still just a few firms, i.e., oligopolies. In such a setting, the actions of any one firm can change the market environment, e.g., the market price, not just for itself but for all firms. Hence, such actions will induce reactions that will in turn prompt further actions and so on. This interaction is of course recognized by each firm and plays a crucial role in determining each firm's strategic choice. In short, we now enter the world of strategic interaction for which the standard analytical tool is game theory. Accordingly, the next three chapters present formal models of oligopoly behavior each of which is rooted in game theory principles and, in particular, yields a market equilibrium consistent with the Nash (1951) concept.

Chapter 9 begins with a brief presentation of game theory and the basic Nash equilibrium solution. We then consider the earliest formal model of oligopoly, namely, the Cournot model. Although conceived a century before Nash's seminal work, the Cournot equilibrium outcome has all the features of the Nash solution and has become a workhorse in economic theory.

An important insight of game theory is that the outcome of any game is heavily dependent on the rules of the game. In the Cournot model, a key rule or assumption is that the firms compete in quantities or production levels. In contrast, the Bertrand model of Chapter 10 assumes that the firms compete in prices. Because price competition can be particularly fierce when firms compete in homogenous goods, the Bertrand assumption gives firms an important motivation for differentiating their products. Hotelling's (1929) spatial model is a useful approach to modeling product differentiation. Therefore, we return to that model here and use it to understand what happens when firms compete vigorously in prices in a product-differentiated industry. Of course, products may be differentiated vertically as well as horizontally. Which sort of differentiation is relevant depends on the nature of consumer preferences. We explore the implications of this point with an empirical study of gasoline prices in southern California in the 1990s.

Finally, in Chapter 11, we consider a different alteration of the Cournot analysis, namely, the order of play. Both the Cournot and Bertand models assume that firms move simultaneously—choosing either production levels or prices at the same time. In contrast, the Stackelberg model of Chapter 11 retains the Cournot assumption of quantity competition but now assumes that one firm plays first, i.e., chooses its production level before its rivals. This permits consideration of the benefits of incumbency and first mover advantages, more generally.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 194

The analysis in the next three chapters is central to all that follows in the rest of the text. Whether the topic is collusion, mergers, advertising or innovation, all modern industrial economics builds on the game theoretic models described in Chapters 9 through 11. Therefore, it is essential to understand this material before proceeding further.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 195

9

Static Games and Cournot Competition

One of the most successful companies in the history of business is Coca-Cola. Indeed, "Coca-Cola" is said to be the second most well-known phrase in the world, the first being "okay."[1] Yet despite its iconic status in American popular culture Coca-Cola is not a monopoly. Coca-Cola shares the carbonated soft drink market share with its archrival PepsiCo. An ongoing battle for market share has engaged these two companies for around a hundred years. The cola wars have been fought with a number of strategies, one of which is the frequent introduction of new soft drink products. Pepsi launched Pepsi Vanilla in the summer of 2003 in response to the year-earlier introduction of Vanilla Coke. In 2006, Coke initiated its biggest new brand campaign in 22 years for its new diet drink, Coke Zero. This followed Pepsi's revitalization of its Pepsi One brand made with Splenda sweetener instead of Aspartame.

In fighting these cola wars each company must identify and implement the strategy that it believes is best suited to gaining a competitive advantage in the soft drink industry. If Coca-Cola were a monopoly, it would not have to worry about the entry of Pepsi products. Life is simpler when you do not have to worry about how rivals will react to your decisions. The simpler life is a feature common to both monopoly and perfect competition. When either a monopoly or a competitive firm chooses how much output to produce, neither has to worry about how that decision affects others. In a pure monopoly there *are* no other firms. In a perfectly competitive market, there are other firms but no one firm needs to be concerned about the effect its output decision will have on the others. Each firm is so small that its output decision will cause not even a ripple in the industry.

The truth is, however, that Coke, Pepsi, and many other firms are neither monopolists nor perfect competitors. These firms, perhaps the majority of corporations, live in the middle ground of *oligopoly* where firms have visible rivals with whom strategic interaction is a fact of life. Each firm is aware that its actions affect others, and therefore, prompt *re*actions. Each firm must, therefore, take these interactions into account when making a decision about prices, or output, or other business actions. Decisions in such an interactive setting are called *strategic* decisions, and *game theory* is the branch of social science that formally analyzes and models strategic decisions. As a result, it is not surprising that game theory and the study

[1] "Coca-cola is okay" has been claimed to be understood in more places by more people than any other sentence, Tedlow (1996).

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 196

of oligopoly are closely intertwined. A central goal of this chapter is to introduce some basic game theoretic analysis and to show how it may be used to understand oligopoly markets.

Game theory itself is divided into two branches: *noncooperative* and *cooperative* game theory.[2] The essential difference between these two branches is that in noncooperative games, the unit of analysis is the individual decision-maker or player, e.g., the firm. By contrast, cooperative game theory takes the unit of analysis to be a group or a coalition of players, e.g., a group of firms. We will focus almost exclusively on noncooperative game theory. The individual player will be the firm. The *rules of the game* will define how competition between the different players, or firms, takes place. The noncooperative setting means that each player is concerned only with doing as well as possible for herself, subject to the rules of the game. The player is not interested in advancing a more general group interest. As we shall see though, such noncooperative behavior can sometimes look very much like cooperative behavior because cooperation sometimes turns out to maximize the well-being of each individual player as well.

Two basic assumptions underlie the application of noncooperative game theory to oligopoly. The first is that *firms are rational*. They pursue well-defined goals, principally profit maximization. The second basic assumption is that firms apply their rationality to the process of *reasoning strategically*. That is, in making its decisions, each firm uses all the knowledge it has to form expectations regarding how other firms will behave. The motivation behind these assumptions is that our ultimate goal is to understand and predict how real firms will act. We assume that firms are rational and reason strategically because we suspect that real firms do precisely this or will be forced to do so by market pressures. Hence, understanding what rational and strategic behavior implies ought to be useful for understanding and predicting real-world outcomes.

There is one caution that any introduction to the study of oligopoly must include. It is that, unlike the textbook competition and monopoly cases, there is no single, standard oligopoly model. Differences in the rules of the game, the information available to the various players, and the timing of each player's actions all conspire to yield a number of possible scenarios. Yet while there is not a single theory or model of oligopoly, common themes and insights from the various models of oligopoly do emerge. Understanding these broad concepts is our goal for the next three chapters. Moreover, we should add that the lack of one single oligopoly model is not entirely a disadvantage. Rather, it means that one has a rich assortment of models from which to choose for any particular investigation. One model will be appropriate for some settings, a different model for other settings. Because the real business world environment is quite diverse, it is useful to have a variety of analyses on which to draw. We will present three different oligopoly models. In this chapter we introduce the Cournot (1836) model of oligopoly, in the next chapter the Bertrand model, and then in Chapter 10 the Stackelberg model.

9.1 STRATEGIC INTERACTION: INTRODUCTION TO GAME THEORY

In game theory, each player's decision or plan of action is called a *strategy*. A list of strategies showing one particular strategy choice for each player is called a *strategy combination*.

[2] A good textbook that offers a more formal treatment of game theory and its applications to economics is Rasmusen (2007).

Any given strategy combination determines the *outcome* of game, which describes the payoffs or final net gains earned by each player. In the context of oligopoly theory, these payoffs are naturally interpreted as each firm's profit.

For a game to be interesting, at least one player must be able to choose from more than one strategy so that there will be more than one possible *strategy combination*, and more than one possible outcome to the game. Yet while there may be many possible outcomes, not all of these will be *equilibrium* outcomes. By equilibrium we mean a *strategy combination* such that no firm has an incentive to *change* the strategy it is currently using given that no other firm changes its current strategy. If this is the case, then the combination of strategies across firms will remain unaltered since no one is changing her behavior. The market or game will come to rest. Nobel Laureate John Nash developed this notion of an equilibrium strategy combination for a noncooperative game. In his honor, it is commonly referred to as the Nash equilibrium concept.[3]

In the oligopoly models studied in the next three chapters, a firm's strategy focuses on either its price choice or its output choice. Each firm chooses either the price it will set for its product or how much of that product to produce. A corresponding Nash equilibrium will, therefore, be either a set of prices, one for each firm, or, a set of production levels, again one for each firm, for which no firm wishes to change its price (quantity) decision given those of all the other firms.

We note parenthetically here that, unlike the monopoly case, the price strategy outcome differs from the quantity strategy outcome in oligopoly models. For a monopolist, the choice of price implies—via the market demand curve—a unique output. In other words, the monopolist will achieve the same market outcome whether he picks the profit-maximizing price or the profit-maximizing output.[4] Matters are different in an oligopoly setting. When firms interact strategically, the market outcome obtained when each firm chooses price will usually differ from the outcome obtained when each firm chooses the best output level. The fact that the outcome depends on whether the rules of the game specify a price strategy or a quantity one is just one of the reasons that the study of oligopoly does not yield a unique set of theoretical predictions.

Since interaction is the central fact of life for an oligopolist, rational strategic action requires that such interaction be recognized. For example, when one firm in an oligopoly market lowers its price, the effect will be noticed by its rivals as they lose customers to the price-cutter. If these firms then lower their price too, they may win back their original customers. Because prices have fallen throughout the industry, the quantity demanded at each firm may well increase. However, each firm will now be meeting that demand at a lower price that earns a lower mark-up. Our assumption that the oligopoly firm is a rational strategic actor means that firm will understand and anticipate this chain of events *and* that the firm will include this information in making the decision whether or not to lower her price in the first place.

Our opening story about carbonated beverages is an example of such interaction except that instead of a price decision, Coca-Cola and Pepsi were making product design choices.

3 Nash shared the 1994 prize with two other game theorists, R. Selten and J. Harsanyi. The award to the three game theorists served as widely publicized recognition of the importance game theory has achieved as a way of thinking in economic analysis.

4 Competitive firms have no option as to which choice variable—price or quantity—to select. Competitive firms by definition cannot make a price choice. They are price-takers and can only choose the quantity of output they sell.

In doing so, each forms some idea as to how its rival will react. It would be *irrational* for Coke to anticipate no reaction from Pepsi, when, in fact, Coke understands that not reacting is not in Pepsi's interest. Similarly if Coke lowers the price of its soft drinks it doesn't make sense for Coke to hope that Pepsi will continue to charge a high price if Coke knows that Pepsi would do better to match its price reduction.

How can an oligopolist anticipate what the response of its rivals will be to any specific action? The best way to make such a prediction is to have information regarding the structure of the market and the strategy choices available to other firms. In a symmetric situation, where all firms are identical, such information is readily available. Any one firm can proceed by asking itself, "What would I do if I were the other player?" Sometimes, even when firms are not symmetric, they will still have enough experience or business "savvy" or other information to be fairly confident regarding their rivals' behavior. As we shall see later, precisely what information firms have about each other is a crucial element determining the final outcome of the game.

Another crucial element in determining the outcome of the game is the time-dimension of the strategic interaction. In a two-firm oligopoly or *duopoly*, like Coca-Cola and Pepsi, we can imagine that one firm, say Coca-Cola, makes its choice—introduce Vanilla Coke first. Then in the next period, the other firm, Pepsi, follows with its choice. In that case, the strategic interaction is *sequential*. Each firm moves in order and each, when its turn comes, must think strategically about how the course of action it is about to choose will affect the future action of the other firm and how those *re*actions will then feed back on its own future choices. Chess and Checkers are each a classic example of a two-person, sequential game. Sequential games are often called dynamic games.

Alternatively, both players might make their choices *simultaneously*, thereby acting without knowledge as to what the other player has actually done.[5] Yet even though the other player's choice is unknown, knowledge of the strategy choices available to the other player permits a player to think rationally and strategically about what other players will choose. The childhood game, "Rock—Scissors–Paper" is an example of a simultaneous two-person game. Such simultaneous games are often called static games.

Whether the game is sequential or simultaneous, the requirement that the strategic firm rationally predicts the choices of its rivals is the same. Once it has done this, the firm may then choose what action is in its own best interest. In other words, being rational means that the firm's choice of strategy is the optimal (profit-maximizing) choice against the anticipated actions of its rivals. When each firm does this, and when each has, as a result of rational strategizing, correctly predicted the choice of the others; we will obtain a Nash equilibrium. In this chapter we will focus on solving for Nash equilibria in simultaneous or static games.

9.2 DOMINANT AND DOMINATED STRATEGIES

Sometimes Nash equilibria are rather easy to determine. This is because some of a firm's possible strategies may be *dominated*. For example, suppose that we have two firms, A and B, in a market and that one of A's strategies is such that it is *never* a profit-maximizing strategy regardless of the choice made by B. That is, there is always an alternative strategy for

[5] The important aspect of simultaneous games is not that the firms involved actually make their decisions at the same time. Rather, it is that no firm can *observe* any other firm's choice before making its own. This lack of information makes the actions of each firm effectively simultaneous.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 199

firm A that yields higher profits than does the strategy in question. Then we say that the strategy in question is dominated: rationally speaking that it will never be chosen. Player A would never choose a dominated strategy since to do so would be to guarantee that A's profit was not maximized. No matter what B does, the dominated strategy does worse for A than one of A's other strategies. In turn, this means that in determining the game's equilibrium, we do not have to worry about any strategy combinations that include the dominated strategy. Since these will never occur, they cannot possibly be part of the equilibrium outcome.

Dominated strategies can be eliminated one by one. Once the dominated strategies for one firm have been eliminated, we can turn to the other firms to see if any of their strategies are dominated given the strategies still remaining for the first firm that we examined. We can proceed firm-by-firm eliminating all dominated strategies until only non-dominated ones remain available to each player. Often but not always, this iterative procedure of eliminating dominated strategies leaves one or more players with only one strategy choice remaining.[6] It is then a simple matter to determine the game's outcome since, for such firms, their course of action is clear.

As an example, consider the case of two airlines, Delta and American, each offering a daily flight from Boston to Budapest. We assume that each firm has already set a price for the flight but that the departure time is still undecided. Departure time is the strategy choice in this game. We also assume that the two firms choose departure times simultaneously. Neither can observe the departure time selected by the other before it makes its own departure time selection. Managers for each airline do realize, however, that at the very time American's managers are meeting to make their choice, Delta's managers are too. The two firms are engaged in a strategic game of simultaneous moves.

In part, the choice of departure time will depend upon consumer preferences. Suppose that market research has shown that 70 percent of the potential clientele for the flight would prefer to leave Boston in the evening and arrive in Budapest the next morning. The remaining 30 percent prefer a morning Boston departure and arrival in Budapest late in the evening of the same day. Both firms know this distribution of consumer preferences. Both also know that, if the two airlines choose the same flight time, they split the market. Profits at each carrier are directly proportional to the number of passengers carried so that each wishes to maximize its share of the market.

If they are rational and strategic, Delta's managers will reason as follows: If American flies in the morning, then we at Delta can either fly at night and serve 70 percent of the market or, like American, depart in the morning in which case we (Delta) will serve 15 percent of the market (half of the 30 percent served by the two carriers in total). On the other hand, if American chooses an evening flight time, then we at Delta may choose either a night departure as well, and serve 35 percent (half of 70 percent) of the market or, instead, offer a morning flight and fly 30 percent of the market.

A little reflection will make clear that Delta does better by scheduling an evening flight *no matter which departure time American chooses*. In other words, choosing a *morning* departure time is a dominated strategy. If Delta is interested in maximizing profits, it will never select the morning flight option. But of course, American's managers will reason similarly. They will recognize that flying at night is their best choice regardless of Delta's selection. So, it seems clear that the only equilibrium outcome for this game is to have both airlines choose an evening departure time.

6 If the process continues until only one strategy remains for each player then we have found an iterated dominance equilibrium.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 200

Table 9.1 Strategy combinations and firm payoffs in the flight departure game

		American	
		Morning	Evening
Delta	Morning	(15, 15)	(30, 70)
	Evening	(70, 30)	(35, 35)

Table 9.1 illustrates the logic just described and the reasons as to why the outcome in which both Delta and American both choose the evening flight must be the equilibrium. The table shows four entries, each consisting of a pair of values. These entries describe the payoffs or market shares associated with the four feasible strategy combinations of the game. American's strategy choices are shown as the columns, while Delta's choices are shown as the rows. The pair of values at each row-column intersection gives the payoffs to each carrier if that particular strategy combination occurs. The first (left-hand) value of each pair is the payoff—the percent of the total potential passenger market—that goes to Delta. The second (right-hand) value is the payoff to American.

Now we put ourselves in the shoes of Delta's managers and ask first what Delta should do if American chooses a morning flight. The answer is obvious. If Delta also chooses a morning flight then Delta's market share will be 15 percent whereas if Delta chooses an evening flight its market share will be 70 percent. The evening flight is clearly the better choice. Now consider Delta's response should American choose an evening flight. If Delta opts for a morning departure its market share is 30 percent whereas if it goes for an evening departure its market share is 35 percent. Once again, the evening departure is the better choice. In other words, no matter what American does, Delta will never choose to depart in the morning. Whatever the equilibrium outcome is, it must involve Delta choosing an evening flight.

If we now place ourselves in American's shoes, a similar result obtains. Again we start by considering American's best response should Delta choose a morning flight. The answer is that American should choose an evening flight to gain 70 percent of the market as compared to the 15 percent that a morning departure would generate. Similarly, should Delta choose an evening departure, American should do likewise since this will give it 35 percent of the market as against the 30 percent that a morning departure would give. As in Delta's case, we discover that flying in the morning is a dominated strategy for American since it never does as well as flying in the evening no matter what Delta does. Hence, just like Delta, American will always choose the evening departure time.

The outcome of the game is now fully determined. Both carriers will choose an evening departure and share equally the 70 percent of the potential Boston-to-Budapest flyers who prefer that time. That this is a Nash equilibrium is easy to see by virtue of the dominated strategy argument. Clearly, neither carrier has an incentive to change its choice from evening to morning since neither carrier would ever choose a morning flight time in any case.

Solving the flight departure game was easy because each carrier had only two strategies and for each player one of the strategies—the morning flight—was dominated. To put it another way, we might refer to the evening departure strategy as *dominant*. A dominant strategy is one that outperforms all of a firm's other strategies *no matter what its rivals do*. That is, it

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 201

leads to higher profits (or sales, or growth, or whatever the objective is) than any other strategy the firm might pursue regardless of the strategies selected by the firm's rivals. This does not imply that a dominant strategy will lead a firm to earn higher profits than its competitors. It only means that the firm will do the best it possibly can if it chooses such a strategy. Whether its payoff is as good as, or better than the payoffs obtained by its rivals depends on the structure of the game.

Except when the number of strategy choices is two, a firm may have some *dominated* strategies, or choices which are never good ones because better ones are available—but not have any *dominant* strategy, or a choice that always yields better results than all others. Sometimes, a firm will have neither a dominant nor a dominated strategy. But for a firm that has a dominant strategy, the choice is clear. Use it! Such a firm really does not have to think very much about what other firms do.

Let's rework the departure time game so that at least one firm has no dominated strategies (and so, since the number of strategies is just two, has no dominant strategy). To do this, we will now suppose that because of a frequent flyer program, some of the potential Boston-to-Budapest flyers prefer Delta even if the two carriers fly at the same time. Specifically, assume now that departing at the same time does not yield an even split of customers between the two carriers. Instead, whenever the two carriers schedule identical departure times, Delta gets 60 percent of the passengers and American gets only 40 percent. Table 9.2 depicts the new payoffs for each strategy combination.

As can be seen from the table, a morning flight is still a dominated strategy for Delta. It always carries more passengers by choosing an evening flight than it would by choosing a morning flight, regardless of what American does. However, American's strategy choices are no longer so clear. If Delta chooses a morning flight, American should fly at night. But if Delta chooses an evening departure time, American does better by flying in the morning.

It may appear that American cannot easily determine its own best course without knowing Delta's choice. Yet this is not the case. There is a self-evident way for American to make a selection even without waiting to see what Delta does. This is because each carrier knows the payoff structure shown in Table 9.2. Accordingly American can readily determine that its rival, Delta, is never going to select a morning flight. Since a morning flight is a dominated strategy for Delta, there is no question of this strategy ever being that carrier's choice. Knowing that Delta will never choose the morning departure, it is then an easy matter for American to select a morning departure as its best response since it knows that Delta will choose an evening departure. The equilibrium outcome for this modified departure time game is therefore just as clear as that for the earlier version. In this case, the equilibrium involves Delta choosing an evening flight and American opting to fly in the morning. Again, it is easily verified that this equilibrium satisfies the Nash criteria.

Table 9.2 Strategy combinations and firm payoffs in the modified flight departure game

		American	
		Morning	Evening
Delta	Morning	(18, 12)	(30, 70)
	Evening	(70, 30)	(42, 28)

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 202

In solving both the previous games, we made extensive use of the ability to rule out dominated strategies and, when possible, to focus on dominant ones.[7] We showed that the outcomes obtained by this process were Nash equilibrium outcomes. However in many games no dominated or dominant strategies can be found. In such cases, the Nash equilibrium concept becomes more than just a criterion to check our analysis. It becomes part of the solution procedure itself. This is because rational, strategic firms will use the Nash concept to determine the reactions of their rivals to their own strategic choice. In the modified departure time game just described, for instance, Delta can work out that if it selects an evening departure, then its rival American will choose a morning flight. Delta can infer that the strategy combination of both carriers flying at night can never be an equilibrium—in the Nash sense—because if that outcome occurred, American would have a clear incentive to change its choice.

9.3 NASH EQUILIBRIUM AS A SOLUTION CONCEPT

In order to understand how to use the Nash equilibrium concept to solve a game, let's change the Boston-to-Budapest game one more time. This time we will change the decision variable from the choice of flight time to one regarding the ticket price. We assume now that consumers are indifferent about the time of departure and, instead, care only about the price they pay for the flight. Specifically, we will suppose that there are 60 consumers with a reservation price of \$500 for the flight, and another 120 with the lower reservation price of \$220. If the two carriers set a common price, they share equally all those customers willing to pay that fare. On the cost side, we will suppose that the unit cost of serving a single passenger for either airline is \$200 whether the flight leaves in the morning or the evening. Let us also assume that each airline is flying a plane with a 200-seat capacity.

Although many price strategies are available to each firm, let's limit ourselves to just two. One is to set a high price of \$500. Another is to set a low price of \$220. So, as before, each firm has two strategies and there are four possible strategy combinations. Each such strategy combination will have associated with it a set of profits. If both Delta and American set the high price of \$500, then each airline will serve half of the 60 passengers willing to pay that fare, or 30 passengers. Because each such passenger involves a cost of \$200, each airline will earn profits of (\$500 − \$200) × 30 = \$9,000. On the other hand, if each sets a price of \$220, they will each share equally in a market of 180 customers and therefore carry 90 passengers apiece. Because of the smaller price cost margin profits to each firm are only (\$220 − \$200) × 90 = \$1,800.

What happens if one airline sets a high price and the other a low one? If say, Delta sets a fare of \$500 and American sets a fare of \$220, Delta will carry *no* passengers. All 180 consumers willing to pay the fare of \$220 or higher will choose American. Delta's profits will be zero. American's profits will be given by (\$220 − \$200) × 180 = \$3,600. Obviously, just the reverse will occur if instead American sets the high price and Delta sets the low one.

[7] Some care needs to be taken in ruling out dominated strategies. While one can eliminate *strictly* dominated strategies as a rational choice, *weakly* dominated strategies cannot be so ruled out. A strategy is *weakly* dominated if there exists some other strategy, which is possibly better but never worse, yielding a higher payoff in some strategy combinations and never yielding a lower payoff. The Nash equilibrium may be affected by the order of exclusion of *weakly* dominated strategies. See Mas-Colell et al. (1995), pp. 238–41.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 203

Table 9.3 Payoff matrix for the airfare game

		American	
		P_H = \$500	P_L = \$220
Delta	P_H = \$500	(\$9,000, \$9,000)	(\$0, \$3,600)
	P_L = \$220	(\$3,600, \$0)	(\$1,800, \$1,800)

The payoff matrix for the new airfare game is shown in Table 9.3. As before, the entries in each row-column intersection show the profit to each firm associated with that strategy combination, with Delta's profit listed as the first entry in each case.

The first thing to notice is that there is no dominant or dominated strategy for either firm. If American selects a high price, Delta should also select a high price. But if American selects a low price, Delta's best bet is to match this price reduction. So we cannot rely on eliminating dominated strategies to identify the outcome of the game. What can we do? Again, we can place ourselves in the shoes of each company's managers. We'll start with Delta. The managers of Delta will look at the payoff matrix of Table 9.3 and reason, as we just did, that their best bet is to choose the same fare as American does. The issue then becomes one of predicting what American will do. Let us suppose that Delta's managers expect American to set a low fare. Then their best choice is to also set the low fare of \$220. But when would this expectation make sense? It will only do so if Delta also believes that American's management team is likewise persuaded that it, i.e., Delta, is going to set a low fare. Delta can go a small step further and reason that if this is in fact the expectation of American, then it may as well go ahead and set the low fare because that is exactly what American is going to do. In other words, Delta's expectation that American will set a low fare because American, in turn, expects Delta to do so will in fact induce Delta to set a low fare. The low fare strategy is Delta's *best response* to its prediction of American's strategy, and that predicted strategy is also the *best response* to Delta's *best response* to that predicted strategy.

In the language of game theory, the strategy combination (*low fare, low fare*) is a Nash equilibrium. If each firm chooses the low fare strategy then neither firm will have any incentive to change its behavior given that the other firm does not change. Each will be pursuing its best course of action given what the other is doing. However, in the airfare game of Table 9.3 there are *two* such equilibria. Following precisely the same reasoning as above, we can work out that the strategy combination (*high fare, high fare*) is also a Nash equilibrium. This game does not have a unique Nash equilibrium.

As we shall see, the existence of more than one Nash equilibrium for a game is not uncommon. But the fact that a unique Nash equilibrium does not always exist does not diminish the usefulness of the concept. To begin with, even if focusing on Nash equilibria does not completely solve the game it certainly narrows the list of potential outcomes. In the airfare game just described, the requirement that the solution be a Nash equilibrium has permitted us to eliminate two, i.e., half of the possible strategy combinations from consideration. Moreover, there are often good, largely intuitive means for determining *which* Nash equilibrium is most likely. The book by Nobel laureate, Thomas Schelling, *The Strategy of Conflict*, offers much guidance in this respect.

Consider the airfare game once more. We want to know which Nash equilibrium (*low fare, low fare*) or (*high fare, high fare*) is more likely to be the outcome. As Schelling observed, taking account of other factors such as the past experience and learning of each firm's managers may be helpful. If the managers of both sides are "old pros" who have dealt with each other for many years, they may be able to avoid the "price war" outcome and coordinate to achieve the more profitable (*high fare, high fare*) outcome. But if the management of either or both sides is new and inexperienced, it will be harder to determine which Nash equilibrium will occur.[8]

We should note that the foregoing analysis of Nash equilibria is relevant to *pure* strategy equilibria. In game theory, a strategy choice is *pure* if a player picks it with certainty, e.g., always calls "heads" in a coin toss. Such pure strategies should be distinguished from *mixed* strategies in which the player uses a probabilistically weighted mixture of two or more strategies, e.g., calling "heads" half the time and "tails" the other half. In some games, mixed strategies or randomizing among strategies makes the most sense. However, we focus primarily on market games in which the only sensible Nash equilibria are those involving pure strategies.

9.1 **Practice Problem**

Firm 1 and firm 2 are movie producers. Each has the option of producing a blockbuster romance or a blockbuster suspense film. The payoff matrix displaying the payoffs for each of the four possible strategy combinations (in thousands) is shown below, with firm 1's payoff listed first. Each firm must make its choice without knowing the choice of its rival.

		Firm 2	
		Romance	Suspense
Firm 1	Romance	(\$900, \$900)	(\$400, \$1,000)
	Suspense	(\$1,000 \$400)	(\$750, \$750)

9.4 STATIC MODELS OF OLIGOPOLY: THE COURNOT MODEL

All the games of the previous section are single period or static. Delta and American, for example, are assumed to choose either their departure times or their airfares simultaneously and without regard to the possibility that, at some later date, they might play the game again. This is a feature of earlier work on modeling oligopoly markets. Firms in these models "meet only once" and the market clears once-and-for-all. There is no sequential movement over time and no repetition of the interaction. These may be limitations. Yet the analysis is still

[8] Alternatively, we might think about the "regret" either player would feel if she plays the wrong strategy. If, for example, Delta chooses P_H expecting American will, too only to discover that American actually chooses P_L, it will earn zero. But if it chooses P_L expecting American also to set a low price and then discovers that American chooses P_H, Delta will earn a profit of \$3,600. In other words, Delta will have much less regret when it assumes the Nash equilibrium will be (P_L, P_L) than when it assumes it will be (P_H, P_H). The same is of course true for American. This thinking suggests that the low price Nash equilibrium will prevail.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 205

capable of generating important insights. Moreover, studying such static models is a good preparation for later examining dynamic models.

The most well known static oligopoly models are the Cournot and Bertrand models, each named after its respective author who did their work in the late nineteenth century. Interestingly enough, these models incorporate modern game theoretic elements. The solution proposed by each author implies the concept of a Nash equilibrium, even though the two models were developed well before the formal development of game theory. In the Cournot model the choice or strategic variable which firms choose when they compete is the quantity of output, whereas in the Bertrand model, the strategic variable chosen is price. We now turn to a presentation of the Cournot model, leaving the Bertrand analysis to Chapter 10.

The work of Augustin Cournot, a French mathematician in the mid-nineteenth century, is now understood as a cornerstone of modern industrial organization theory despite the fact that it went largely unrecognized for about one hundred years after its publication in 1836. The Cournot duopoly model anticipates Nash's concept of an equilibrium, and so not surprisingly, Cournot's work is regarded as a classic in game theoretic analysis.

The story that Cournot told to motivate his analysis went as follows. Assume a single firm wishes to enter a market currently supplied by a monopoly. The entrant is able to offer a product that is identical in all respects to that of the incumbent monopolist and to produce it at the same unit cost. Entry is attractive because under the assumption of constant and identical costs, we know that the monopolist is producing where price is greater than marginal cost, which means that the price also exceeds the marginal cost of the would-be entrant. Hence, the entrant firm will see that it can profitably sell some amount in this market. However, the new entrant will, Cournot reasoned, choose an output level that maximizes its profit, *after taking account of the output being sold by the monopolist*.

Of course, if entry occurred and the new firm produced its chosen output, the monopolist would react. Before entry, the monopolist chose a profit-maximizing output assuming no other rivals. Now, the former monopolist will have to re-optimize and choose a new level. In so doing, the monopolist will (as did the new entrant previously) choose an output level that maximizes profits *given the output sold by the new rival firm*.

This process of each firm choosing an output conditional on the other's output choice is to be repeated—at least as a mental exercise. For every output choice by the incumbent, firm 1, the entrant, firm 2, is shown to have a unique, profit-maximizing response and vise-versa. Cournot called the graph representations of these responses Reaction Curves. Each firm has its own Reaction Curve that can be graphed in the q_1q_2 quadrant. That Cournot anticipated Nash is evidenced by the fact that he described the equilibrium outcome of this process as that pair of output levels at which each firm's output choice is the profit-maximizing response to the other's quantity. Otherwise, Cournot reasoned, at least one firm would wish to change its production level. A further appealing aspect of Cournot's duopoly model is that the equilibrium price resulting from the output choices of the two firms is below that of the pure monopoly outcome. Yet it is also greater than that which would occur if there were not two firms but many firms and pure competition prevailed.

To present Cournot's analysis more formally we assume that the industry *inverse* demand curve[9] is linear, and can be described by:

$$P = A - BQ = A - B(q_1 + q_2) \tag{9.1}$$

[9] By an inverse demand curve we mean a demand curve in which price is expressed as a function of quantity rather than quantity being expressed as a function of price.

where Q is the sum of each firm's production, i.e., the total amount sold on the market, q_1 is the amount of output chosen by firm 1, the incumbent firm, and q_2 is the amount of output chosen by firm 2, the new competitor. As noted earlier, we shall also assume that each firm faces the same, constant marginal cost of production, c.

If we now consider firm 2, alone, and take firm 1's output, q_1, as given, the inverse demand curve, facing firm 2 is:

$$P = A - Bq_1 - Bq_2 \tag{9.2}$$

which is formally identical to (9.1). However, from firm 2's perspective, the first two terms on the right-hand side are not part of its decision-making, and can be taken as given. In other words, those two terms together form the intercept of firm 2's perceived demand curve so that firm 2 understands that the only impact *its* output choice has on price is given by the last term of the equation, namely, $-Bq_2$. Note, however, that any change in the anticipated output choice of firm would be communicated to firm 2 by means of a shift in firm 2's perceived demand curve. Figure 9.1 illustrates this point.

As we can see from Figure 9.1, a different choice of output by firm 1 will imply a different demand curve for firm 2 and, correspondingly, a different profit-maximizing output for firm 2. Thus, for each choice of q_1 there will be a different optimal level of q_2. We can solve for this relationship algebraically, as follows. Associated with each demand curve illustrated in Figure 9.1 there is a marginal revenue curve that is twice as steeply sloped: this was discussed in Chapter 2, and is adapted to the present model in the inset. That is, firm 2's marginal revenue curve is also a function of q_1 given by:

$$MR_2 = (A - Bq_1) - 2Bq_2 \tag{9.3}$$

Marginal cost for each firm is constant at c. Setting marginal revenue MR_2 equal to marginal cost c, as required for profit-maximization, and solving for q_2^* yields firm 2's Reaction Curve. So we have $MR_2 = c$, which implies that $A - Bq_1 - 2Bq_2^* = c$ or $2Bq_2^* = A - c - Bq_1$. Further simplification then gives the Reaction Function for firm 2:

Derivation Checkpoint

Review of Marginal Revenue and Demand

Assume that the inverse demand curve facing firm 2 is:

$$P = A - Bq_1 - Bq_2$$

Then total revenue is:

$$TR_2 = (A - Bq_1 - Bq_2)q_2 = Aq_2 - Bq_1q_2 - Bq_2^2.$$

Marginal revenue is the differential of total revenue with respect to output, so that:

$$MR_2 = \partial TR_2/\partial q_2 = A - Bq_1 - 2Bq_2$$

This has the same price intercept as the inverse demand function but twice the slope.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 207

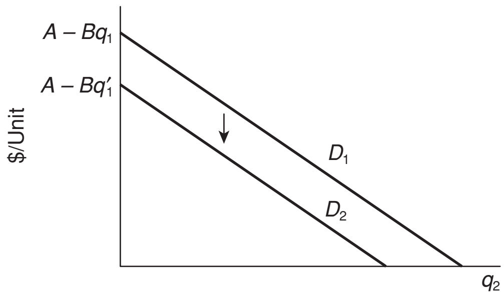

Figure 9.1 Firm 2's demand curve in the Cournot Duopoly game depends on firm 1's output
An increase in q_1 to q_1' shifts D_2, the demand curve facing firm 2, downwards.

$$q_2^* = \frac{(A - c)}{2B} - \frac{q_1}{2} \tag{9.4}$$

Equation 9.4 describes firm 2's best output choice, q_2^*, for every choice of q_1. Note that the relationship is a negative one. Every increase in firm 1's output lowers firm 2's demand and marginal revenue curves and, with a constant marginal cost, also lowers firm 2's profit-maximizing output.

Of course, matters work both ways. We may symmetrically re-work the industry demand curve to show that firm 1's individual demand depends similarly on firm 2's choice of output, so that as q_2 changes, so does the profit-maximizing choice of q_1. Then, we may analogously derive firm 1's Reaction Curve giving its best choice of q_1 for each alternative possible value of q_2. By symmetry with firm 2, this is given by:

$$q_1^* = \frac{(A - c)}{2B} - \frac{q_2}{2} \tag{9.5}$$

As was the case for firm 2, firm 1's profit maximizing output level q_1^* falls as q_2 increases.[10] The Reaction Curve for each firm is shown in Figure 9.2 in which the strategic variables for each firm, outputs, are on the axes.

Consider first, the Reaction Curve of firm 1, the initial monopolist. This curve says that if firm 2 produces nothing, then firm 1 should optimally produce quantity $\frac{(A - c)}{2B}$, which is, in fact, the pure monopoly level, at which we assumed firm 1 to be producing in the first place. Now consider the Reaction Curve for firm 2. That curve shows that if firm 1 were producing at the assumed level of $\frac{(A - c)}{2B}$, then firm 2's best bet is to produce at level $\frac{(A - c)}{4B}$, that is, firm 2 should enter the market. However, if firm 2 does choose that level then firm

[10] We could alternatively solve for q_2^* by writing firm 2's profit function, Π^2, as revenue less cost, or: $\Pi^2(q_1, q_2) = (A - Bq_1 - Bq_2)q_2 - cq_2 = (A - Bq_1 - c)q_2 - Bq_2^2$. When we differentiate this expression with respect to q_2 and set the result equal to 0, (first order condition for maximization), and then solve for q_2^* we get the same result as equation (9.4). A similar procedure may be used to obtain q_1^*.

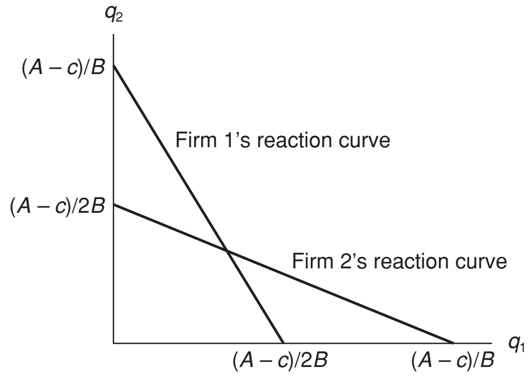

Figure 9.2 Best-response (reaction) curves for the Cournot Duopoly model

1 will no longer do best by producing the monopoly level. Instead, firm 1 will maximize profits by selecting quantity $q_1 = \dfrac{3(A - c)}{8B}$.

As Cournot understood, none of the output or strategy combinations just described corresponds to an equilibrium outcome. In each case, the *Reaction* of one firm is based upon a choice of output for the other firm that is not, itself, that other firm's *Reaction*. For the outcome to be an equilibrium, it must be the case that each firm is responding optimally to the (optimal) choice of its rival. We want each firm to choose a Reaction based upon a prediction about what the other firm will produce and, in equilibrium, we want each firm's prediction to be correct. Put more simply, equilibrium requires that *both* firms be on their respective Reaction Curves. This happens at only one point in Figure 9.2, namely, the intersection of the two Reaction Curves.

To see how this works, recall the Reaction Function for firm 2: $q_2^* = \dfrac{(A - c)}{2B} - \dfrac{q_1}{2}$. We know and firm 2 knows that in an equilibrium, firm 1 must also be on its Reaction Function, or that $q_1^* = \dfrac{(A - c)}{2B} - \dfrac{q_2}{2}$. Substituting this into firm 2's Reaction Function allows firm 2 (and also us) to solve for: $q_2^* = \dfrac{A - c}{2B} - \dfrac{1}{2}\left(\dfrac{A - c}{2B} - \dfrac{q_2^*}{2}\right) = \dfrac{A - c}{4B} - \dfrac{q_2^*}{4}$ so that $\dfrac{3q_2^*}{4} = \dfrac{A - c}{4B}$. In turn, this implies: $q_2^* = \dfrac{(A - c)}{3B}$. Symmetry implies that $q_1^* = \dfrac{(A - c)}{3B}$ as well. We leave it as an end-of-chapter exercise for the reader to verify that this equilibrium also satisfies the Nash criterion.

Total output is for the market is $Q^* = \dfrac{2(A - c)}{3B}$. Substituting this into the demand function gives the equilibrium price: $P = A - BQ = \dfrac{A + 2c}{3}$. Profit for each firm is total revenue less total cost, which can be solved as $\pi_i = \dfrac{(A - c)^2}{9B}$.

As Figure 9.2 makes clear, the Cournot duopoly model just presented has a unique Nash equilibrium. Hence, in terms of our earlier discussion regarding the strategy of solving games, we can solve the Cournot duopoly game simply by focusing on its Nash equilibrium. Since there is only one Nash equilibrium this must be the outcome of the game. The importance of this insight is difficult to overstate.

To see the power of the Nash concept, let us briefly reflect on the initial Cournot setup. We had two firms, each choosing quantity as its strategic variable. If, as also postulated, each knows the industry demand curve and the fact that each has an identical constant marginal cost, how should each firm act? Our discussion of Reaction Curves borrowed from Cournot suggests a kind of trial-by-learning process by which the two firms act and react until the equilibrium is achieved. But the power of the Nash equilibrium is that it makes such an iterative procedure played out in real time unnecessary. Recall the basic game theory assumptions that firms are rational and strategic. In choosing its own production level, firm 1 *must anticipate* that firm 2 will do whatever maximizes firm 2's profits. An expectation, for instance, by firm 1 that firm 2 will produce 0 and that therefore firm 1 should choose the monopoly output would ***not*** be rational because the reaction curve tells us that 0 is not firm 2's best response to that situation. Hence firm 1 would never predict 0 as firm 2's output choice. Similarly, firm 1 also ought never to predict $q_2 = \frac{(A - c)}{4B}$. Here, such a prediction would lead to an inconsistency because it would imply firm 1 choosing a profit-maximizing output $q_1^* = \frac{3(A - c)}{8B}$, for which the predicted value of $q_2 = \frac{(A - c)}{4B}$ is again not optimal. In short, there is only one prediction for q_2 that firm 1 can possibly make if it is to act rationally. This prediction is that $q_2 = \frac{(A - c)}{3B}$ the value of q_2 in the Nash equilibrium. This is the only prediction, which, if made, will actually induce the behavior consistent with that expectation being fulfilled. If firm 1 expects q_2 to be equal to $\frac{(A - c)}{3B}$, then firm 1 will *optimally* choose that output level, too. In turn, this output choice by firm 1 is such that firm 2 should indeed produce at the level of $\frac{(A - c)}{3B}$ if it wishes to maximize its profits.

To put it another way, what we are saying is that rational and strategic firms can work through the Cournot model as a pure thought experiment, without any time-consuming real world trials and errors. When they do, such firms will quickly realize that the only sensible prediction is that each will produce the unique Nash equilibrium output value, $q_i^* = \frac{(A - c)}{3B}$. It is only when each firm makes and acts upon that particular expectation that each firm will find that its prediction comes true.

Many economists, including us, prefer to use the term "*best response function*" instead of "Reaction Curve." The point is to emphasize that the correct interpretation of the Cournot model is one of *simultaneous* and *not sequential* output choice. The Cournot equilibrium is one in which each seller's predictions are consistent both with profit-maximization and with the actual market outcome.[11]

[11] Friedman (1977) includes a brief discussion of these issues, particularly valuable to those interested in the history of economic thought. He notes that Cournot's fate was not quite one of total obscurity owing to his friendship with the father of the French economist Walras. The English economist Marshall apparently was also well aware of and influenced by Cournot's work.

As a numeric example, consider two firms, Untel and Cyrox, who supply the market for computer chips for toaster ovens. Untel's chips are perfect substitutes for Cyrox's chips and vice versa. Market demand for chips is estimated to be $P = 120 - 20Q$, where Q is the total quantity (in millions) of chips bought. Both firms have a constant marginal cost equal to 20 per unit of output. Untel and Cyrox independently choose what quantity of output to produce. The price then adjusts to clear the market of the total quantity of chips produced. What quantity of output will Untel produce? What quantity of output will Cyrox produce? What will be the price of computer chips and how much profit will each firm make?

Let's put ourselves on the management team at Untel to see the problem from its perspective. The demand curve that Untel faces can be written as $P = 120 - 20q_c - 20q_u$, where q_c is the output of Cyrox and q_u is the output of Untel. Untel's marginal revenue curve is $MR_u = 120 - 20q_c - 40q_u$. To maximize profit Untel chooses a quantity of output q_u^* such that its marginal revenue is equal to marginal cost. That is, $120 - 20q_c - 40q_u^* = 20$. This condition for profit maximization implies that:

$$q_u^* = \frac{120 - 20}{40} - \frac{20}{40}q_c \ \ or\ \ q_u^* = \frac{5}{2} - \frac{1}{2}q_c. \tag{9.6}$$

This is Untel's Reaction Function for any given level of output by Cyrox. In other words, Untel knows that its profit-maximizing choice of output depends on what its rival, Cyrox, chooses to produce. Untel wants to predict what Cryox is going to do, and then respond to it in a way that maximizes Untel's profit. Of course Untel knows that Cyrox is also a profit maximizer, and so Untel anticipates that Cyrox will want to produce q_c^* to satisfy the condition for profit maximization at Cyrox. By precisely the same argument that we have just gone through, Untel knows that Cyrox's Reaction Function is $q_c^* = \frac{120 - 20}{40} - \frac{20}{40}q_u \ \ or\ \ q_c^* = \frac{5}{2} - \frac{1}{2}q_u$. Untel can recognize that Cyrox's choice of output depends on Untel's. Untel also knows that Cyrox knows that Untel is a profit-maximizer, and that Cyrox will anticipate that Untel will choose a profit-maximizing level of output q_u^*. Therefore, Untel predicts that Cyrox will choose $q_c^* = \frac{5}{2} - \frac{1}{2}q_u^*$. Substituting this prediction into Untel's Reaction Curve, equation (9.6), leads Untel to produce $q_u^* = \frac{5}{2} - \frac{1}{2}\,q_c^* = \frac{5}{2} - \frac{1}{2}\left(\frac{5}{2} - \frac{1}{2}q_u^*\right) \Rightarrow q_u^* = \frac{5}{3}$.

Now let's put ourselves on the management team at Cyrox and repeat the exercise. Because the two firms are *identical* there is no reason why Cyrox would do anything different from Untel, and so we can quickly jump to the conclusion that Cyrox will also produce $q_c^* = \frac{5}{3}$. Note that when Untel produces $\frac{5}{3}$, Cyrox's best response is to produce $q_c^* = \frac{5}{3}$, and similarly when Cyrox produces $\frac{5}{3}$, Untel's best response is to produce $q_u^* = \frac{5}{3}$.

Aggregate market output is $Q^* = \frac{10}{3}$, and so the price that clears the market is $P^* = 120 - 20\left(\frac{10}{3}\right) = \$53.33$. For each firm the margin of price over unit cost is \$33.33 so that each firm makes a profit of \$55.55.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 211

9.2 Practice Problem

Assume that there are two identical firms serving a market in which the inverse demand function is given by $P = 100 - 2Q$. The marginal costs of each firm are \$10 per unit. Calculate the Cournot equilibrium outputs for each firm, the product price and the profits of each firm.

Cournot's model is insightful in its treatment of the interaction among firms and remarkably modern in its approach. Yet these are not its only strengths. Cournot's analysis has the further advantage that the results also blend well with economic intuition. In the simple Cournot duopoly model described above each firm produces its Nash equilibrium output of $\frac{(A - c)}{3B}$, implying that total industry output is $\frac{2(A - c)}{3B}$. This is clearly greater than the monopoly output for the industry, which would be $Q^M = \frac{(A - c)}{2B}$. Yet it is also less than the perfectly competitive output, $Q^C = \frac{(A - c)}{B}$, where price equals marginal cost. Accordingly, the market-clearing price in Cournot's model $P = \frac{(A + 2c)}{3}$ is less than the monopoly price $P^M = \frac{(A + c)}{2}$ but it is higher than the competitive price, c, which is equal to marginal cost. That is, Cournot's duopoly model has the intuitively plausible result that the interaction of two firms yields more industry output at a lower price than would occur under a monopoly, but not as much as the output produced under perfect competition.

9.5 VARIATIONS ON THE COURNOT THEME: MANY FIRMS AND DIFFERENT COSTS

Cournot's model can be enriched in several ways. One attractive feature of the model is its prediction that the addition of a second firm moves the industry outcome away from the monopoly result and toward that which obtains under perfect competition. A natural question then arises. Would introducing a third firm bring the industry still closer to the competitive ideal? What about a fourth? Or a fifth? Is the Cournot analysis consistent with the notion that when there are many firms the price converges to marginal cost?

To explore the Cournot model's implications when we vary the number of competing firms, let us work with the general case of N firms. These firms are, as before, assumed to be identical. Each produces the same homogenous good and each has the same, constant marginal cost c. Industry demand is again given by $P = A - BQ$ where Q is aggregate output. However, now we have that $Q = q_1 + q_2 + \ldots + q_N = \sum_{i=1}^{N} q_i$ so that $P = A - B\sum_{i=1}^{N} q_i$, where q_i is the output of the ith firm. In turn, this means that we can write the demand curve facing just a single firm, say firm 1, as: $P = (A - Bq_2 - Bq_3 - \ldots - Bq_N) - Bq_1$. The parenthetical expression reflects the fact that for firm 1, this term is beyond its control and merely appears as the intercept in firm 1's demand curve. It is conventional to use the notation Q_{-1} as a shorthand method of denoting the sum of all industry output *except* that of firm 1's. Using this notation, we can write firm 1's demand curve even more simply as: $P = A - BQ_{-1} - Bq_1$.

Clearly, firm 1's profits depend on both Q_{-1}, over which it has no control, and its own production level, q_1, which it is free to choose. Given its constant unit cost of c, firm 1's profits Π^1 can be written as: $\Pi^1(Q_{-1}, q_1) = (A - BQ_{-1} - Bq_1)q_1 - cq_1$.

Profit maximization requires that firm 1 chooses its output level where marginal revenue equals marginal cost. Since marginal revenue is given by a curve with the same intercept but twice as steeply sloped as firm 1's demand curve, the condition for profit maximization at firm 1 is:

$$(A - BQ_{-1}) - 2Bq_1^* = c \tag{9.7}$$

Solving this equation for q_1^* gives us the Reaction Curve or what we will now call the *best response* function for firm 1 of:

$$q_1^* = \frac{(A - c)}{2B} - \frac{Q_{-1}}{2} \tag{9.8}$$

Since all firms are identical, we can extend this same logic to develop the best response function for any firm. Using the same shorthand notation, we can use Q_{-i} to mean the total industry production *excluding that of firm i*. This means that the demand function for firm i, taking the output of all other firms as given, is:

$$P = (A - BQ_{-i}) - Bq_i$$

The associated marginal revenue function of firm i is

$$MR_i = (A - BQ_{-i}) - 2Bq_i.$$

Equating marginal revenue with marginal cost gives the *best response function* for firm i:

$$q_i^* = \frac{(A - c)}{2B} - \frac{Q_{-i}}{2} \tag{9.9}$$

In a Nash equilibrium, each firm i chooses a best response, q_i^* that reflects a correct prediction of the outputs that the other $N - 1$ firms will choose. Denote by Q_{-i}^* the sum of all the outputs excluding q_i^* when each element in that sum is *each firm's best output response decision*. Then an algebraic representation of the Nash equilibrium is:

$$q_i^* = \frac{(A - c)}{2B} - \frac{Q_{-i}^*}{2}; \quad \text{for } i = 1, 2, \ldots N \tag{9.10}$$

Recall however that the N firms are identical. They each produce the *same* good at the *same* unit marginal cost, c. From this it follows that, in equilibrium, each will produce the *same* output, i.e., $q_1^* = q_2^* = \ldots = q_N^*$, or just q^* for short. So, noting that $Q_{-i}^* = (N - 1)q^*$, we can rewrite equation (9.10) as:

$$q^* = \frac{(A - c)}{2B} - \frac{(N - 1)q^*}{2} \tag{9.11}$$

from which it follows that the equilibrium output for each firm, what we refer to as the *Cournot–Nash* equilibrium output, is:

$$q^* = \frac{(A - c)}{(N + 1)B} \tag{9.12}$$

There are N firms each producing q^* as given by equation (9.12). From this we may derive both the Cournot-Nash equilibrium industry output, $Q^* = Nq^*$, and the Cournot–Nash equilibrium industry price, $P^* = A - BQ^*$, as:

$$Q^* = \frac{N(A - c)}{(N + 1)B};\ P^* = \frac{A}{(N + 1)} + \frac{N}{(N + 1)}c. \tag{9.13}$$

Examine the two equations in (9.13) carefully. When $N = 1$, industry output is $\frac{(A - c)}{2B}$ and the corresponding price is $\frac{(A + c)}{2}$. But this is just the monopoly outcome, as of course it should be. When N increases to two we obtain the duopoly output and price levels derived in our earlier analysis. What happens when the number of firms rises above two? In particular, what happens when N gets very large?

Consider first the Cournot–Nash equilibrium price, P^*. As N gets larger and larger, the term $\frac{A}{(N + 1)}$ gets closer and closer to zero and, in the limit, vanishes. Similarly, as N increases the term $\frac{N}{(N + 1)}$ becomes arbitrarily close to 1. Thus, equation (9.13) says that when the number of industry firms gets very large, the industry equilibrium price, P^*, converges to marginal cost, c. But this is just the perfectly competitive result! Confirmation of this result is further obtained by noting that total industry output (the first part of equation 9.13) is similarly close to the competitive output of $\frac{(A - c)}{B}$ when N is large.

Consider the following numerical example, if the inverse demand curve is: $P = 100 - 2Q$, so that $A = 100$, and $B = 2$; and if the unit cost $c = 4$, then the monopoly output Q^M and price P^M are: $Q^M = 24$ and $P^M = 52$. Moving from a monopoly to a duopoly raises the equilibrium output, $Q^D = 32$, and lowers the price to $P^D = 36$. If the number of firms increases to 99 then the price falls to $P^{99} = \frac{100}{100} + \frac{99}{100} \times 4 = \$4.96$. As we increase the number of firms selling in the market the Cournot equilibrium market output continues to rise and the price continues to fall until, with many firms, we approximate the competitive equilibrium with $Q = 48$ and $P = 4$.

In short, the Cournot model implies that as the number of *identical* firms in the market grows, the industry equilibrium gets closer and closer to that prevailing under perfect competition. Of course, this result seems quite natural since, as N increases, each Cournot firm becomes smaller relative to the market. It is an appealing feature of Cournot's analysis that it predicts a plausible relationship between market structure and market performance. Market outcomes improve as market concentration falls and the competitive standard is approached.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 214

What if the firms competing in the market are not identical? Specifically, what if each firm has a different marginal cost? We first handle this question for the case of two firms. Assume that the marginal costs of firm 1 are c_1 and of firm 2 are c_2. We use the same approach as before with the duopoly model, starting with the demand function for firm 1, which we can write as:

$$P = (A - Bq_2) - Bq_1$$

The associated marginal revenue function is

$$MR_1 = (A - Bq_2) - 2Bq_1.$$

As before, firm 1 maximizes profit by equating marginal revenue with marginal cost. So setting $MR_1 = c_1$ and solving for q_1 gives the best response function for firm 1 as:

$$q_1^* = \frac{(A - c_1)}{2B} - \frac{q_2}{2} \tag{9.14a}$$

By an exactly symmetric argument, the best response function for firm 2 is:

$$q_2^* = \frac{(A - c_2)}{2B} - \frac{q_1}{2} \tag{9.14b}$$

Notice that the only difference from our initial analysis of the Cournot model is that now each firm's best response function reflects its own specific marginal cost.

An important feature of these best response functions that is obscured when the firms are identical is that the *position* of each firm's best response function is affected by its marginal cost. For example, if the marginal cost of firm 2 increases from say, c_2 to c_2', its best response curve will shift inwards.

Figure 9.3 illustrates this point. It shows the best response function for each firm assuming initially that each firm has identical costs as in Figure 9.2. It then shows what happens when firm 2's unit cost rises. As equation (9.14b) makes clear, this cost increase *lowers* firm

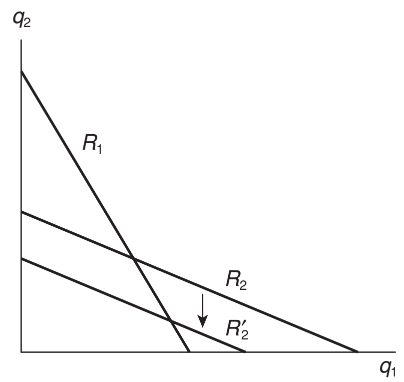

Figure 9.3 The Cournot Duopoly model with different costs across firms
A rise in firm 2's unit cost shifts the firm 2 best response function downward from R_2 to R_2'. In the new equilibrium, firm 1 produces more and firm 2 produces less than previously.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 215

2's best output response for any given level of q_1. That is, it shifts firm 2's best response curve inward. This change in firm 2's best response function affects the equilibrium outputs that the two firms will choose. As you can see from the diagram, an increase in firm 2's marginal cost leads to a new equilibrium in which firm 1 produces more than it did in the initial equilibrium and firm 2 produces less. This makes intuitive sense. We should expect that low-cost firms will generally produce more than high-cost firms. The changes are not offsetting, however. Firm 2's output falls by more than firm 1's production rises so that the new equilibrium is characterized by less output in total than was the original equilibrium. (Can you say why?)

The Cournot–Nash equilibrium can be obtained as before by substituting the expression for q_2^* into firm 1's best response to solve for q_1^*. Then we may use this value to solve for q_2^*. In other words, we have:

$$q_1^* = \frac{(A - c_1)}{2B} - \frac{1}{2}\left(\frac{(A - c_2)}{2B} - \frac{q_1^*}{2}\right)$$

which can be solved for q_1 to give the equilibrium:

$$q_1^* = \frac{(A + c_2 - 2c_1)}{3B} \tag{9.15a}$$

By an exactly symmetric argument, the equilibrium output for firm 2 is:

$$q_2^* = \frac{(A + c_1 - 2c_2)}{3B} \tag{9.15b}$$

It is easy to check that the relative outputs of these two firms are determined by the relative magnitudes of their marginal costs. The firm with the lower marginal costs will have the higher output.

Let's return to our Untel and Cyrox example of the two firms who produce computer chips for toaster ovens but now change this story a bit. While we still assume that Untel's chips are perfect substitutes for Cyrox's chips and vice versa, we no longer assume that they have identical costs. Instead, we now assume that Untel is the low cost firm with a constant unit cost of 20, and Cyrox is the high cost producer with a constant unit cost of 40. Market demand for chips is still estimated to be $P = 120 - 20Q$, where Q is the total quantity (in millions) of chips bought. What now happens when Untel and Cyrox independently choose the quantity of output to produce? What quantity of output will Untel produce? What quantity of output will Cyrox produce?

Again we put ourselves on the management team at Untel to see the problem from Untel's perspective. The demand curve that Untel faces is still $P = 120 - 20q_c - 20q_u$, where q_c is the output of Cyrox and q_u is the output of Untel. Untel's marginal revenue is again $MR_u = 120 - 20q_c - 40q_u$. To maximize profit Untel should sell a quantity of output q_u^* such that at that quantity marginal revenue is equal to marginal cost. That is, $120 - 20q_c - 40q_u^* = 20$, and so the condition for profit maximization implies that:

$$q_u^* = \frac{5}{2} - \frac{1}{2}q_c. \tag{9.16}$$

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 216

Untel's profit-maximizing choice of output still depends on the output that the higher cost rival, Cyrox, chooses to produce. Equally importantly, comparison of equation (9.16) with (9.6) indicates that Untel's best response function is unaffected by the assumed increase in Cryox's marginal cost. What about Cryox? By the same argument, the demand curve that Cryox faces is $P = 120 - 20q_u - 20q_c$ and its marginal revenue curve is $MR_c = 120 - 20q_u - 40q_c$. Equating this with marginal cost of 40 and solving for q_c gives the best response function for Cryox of $q_c^* = \frac{120 - 40}{40} - \frac{20}{40}q_u$ *or* $q_c^* = 2 - \frac{1}{2}q_u$. As we expected, the best response function for Cryox is shifted downward by the assumed increase in its marginal cost.

Untel knows that higher cost Cyrox is also a profit maximizer and therefore anticipates that Cyrox will want to produce q_c^* that maximizes its profit. It is also the case, as it was before, that Untel knows that Cyrox knows that Untel is a profit-maximizer, and so knows that Cyrox will anticipate that Untel will choose a profit-maximizing level of output q_u^*. All of this implies that Untel predicts that Cyrox will choose $q_c^* = 2 - \frac{1}{2}q_u^*$. Substituting this new prediction into Untel's best response function leads Untel to produce $q_u^* = \frac{5}{2} - \frac{1}{2}q_c^* = \frac{5}{2} - \frac{1}{2}\left(2 - \frac{1}{2}q_u^*\right) \Rightarrow q_u^* = 2.$

Now we put ourselves on the management team at Cyrox and repeat the exercise. To cut to the chase, we know that Cyrox's best response is $q_c^* = 2 - \frac{1}{2}q_u$. Moreover we know that Cyrox will predict that Untel will produce a best response that is based on a prediction that Cyrox will also produce a best response. That is, Cryrox predicts that Untel will produce $q_u^* = \frac{5}{2} - \frac{1}{2}q_c^*$. Substituting this prediction into Cyrox's best response function leads to: $q_c^* = 2 - \frac{1}{2}q_u^* = 2 - \frac{1}{2}\left(\frac{5}{2} - \frac{1}{2}q_c^*\right) \Rightarrow q_c^* = 1$. Again note that when Untel produces 2, Cyrox's best response is to produce $q_c^* = 1$, and similarly when Cyrox produces 1, Untel's best response is to produce $q_u^* = 2$.

Although the foregoing analysis is limited to just two firms, it still yields important insights. One of these is that in the Cournot model, firms with higher costs have smaller market shares and smaller profits. This means that a Cournot firm benefits when its rival's costs go up, as in our example above. Moreover, when costs vary across firms, the equilibrium Cournot output Q^* is not only too low (i.e., less than the competitive level), it is also produced inefficiently. As we know from Chapter 4, efficient production among two or more firms would allocate output such that, in the final configuration, each firm's marginal cost is the same. This would be the outcome, for example, if the industry were comprised of a single, profit maximizing multi-plant monopolist. It will also obtain under perfect competition. However, as we have just seen, the Cournot-Nash equilibrium does not require that firms' marginal costs be equalized.[12] Hence, output allocation in a Cournot equilibrium with different costs between firms is not an efficient one.

[12] Our example assumed constant but different marginal costs across firms. The same insight could be easily obtained for the more general presentation in which the marginal cost of firm i, c_i, is a general function of its output, q_i, as in $c_i = c_i(q_i)$.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 217

9.3 **Practice Problem**

What is aggregate output, market price and the Untel's profit and Cyrox's profit for the above case in which Untel is the low cost producer and Cyrox the high cost one? Compare your answers to the ones you work out when the two firms are identical and have a constant unit cost of 20.

9.6 CONCENTRATION AND PROFITABILITY IN THE COURNOT MODEL

Let us now try to combine the case of many firms together with the assumption of *non*-identical costs. That is, let us analyze the Cournot model with N firms, each with its own (constant) marginal cost such that the marginal cost of firm i is c_i. We can use the first order condition for profit-maximization for each firm i, equation (9.7), and substitute c_i for c in this equation. This gives us the following:

$$A - BQ_{-i} - 2Bq_i^* - c_i = 0 \tag{9.17}$$

where Q_{-i} again is shorthand for the industry production accounted for by all firms other than the ith one.

In a Nash equilibrium, the equilibrium output q_i^* *for each firm i* must satisfy the first-order profit-maximizing condition. Hence, in the Nash equilibrium, the term Q_{-i} must be the sum of the *optimal* outputs q_j^* for each of the "not i" firms. Denote this equilibrium sum as Q_{-i}^*. Then we can re-write (9.17) as:

$$A - BQ_{-i}^* - 2Bq_i^* - c_i = 0 \tag{9.18}$$

By definition, the total equilibrium output, Q^*, equals the sum of Q_{-i}^* and q_i^*. Hence, (9.18) implies that

$$A - B(Q^* - q_i^*) - 2Bq_i^* - c_i = 0$$

Which can be reorganized to give:

$$A - BQ^* - c_i = Bq_i^* \tag{9.19}$$

We also know that the Nash equilibrium price, P^*, is obtained by substituting the Nash equilibrium output into the industry demand curve yielding, $P^* = A - BQ^*$. Substitution into equation (9.19) then yields:

$$P^* - c_i = Bq_i^* \tag{9.20}$$

Dividing both sides of equation (9.20) by P^*, and multiplying the *right-hand side* by $\frac{Q^*}{Q^*}$, we obtain:

$$\frac{P^* - c_i}{P^*} = \frac{BQ^*}{P^*} s_i^* \tag{9.21}$$

where $s_i^* = \dfrac{q_i^*}{Q^*}$ is the ith firm's market share in equilibrium.

Let us consider equation (9.21) step-by-step. The left-hand-side term is the difference between price and firm i's marginal cost as a proportion of market price. This is just the *Lerner Index of Monopoly Power* that we met in Chapter 3. The notion is that the greater firm i's market power, the greater it's ability to keep price above marginal cost.

The right-hand-side of (9.21) has two terms. The first is the slope of industry demand curve times the ratio of industry output to price. But the slope is just $B = dP/dQ$ so that we have $\dfrac{BQ^*}{P^*} = \dfrac{dP}{dQ} \cdot \dfrac{Q^*}{P^*}$. Recall the definition of the price elasticity of demand: $\eta = \dfrac{dQ}{dP} \cdot \dfrac{P}{Q}$. So the first term on the right-hand side of equation (9.21) is just the inverse of the price elasticity of demand. The second term is just the market share of the ith firm, i.e. its output relative to total industry output. Hence, equation (9.21) may be rewritten as:

$$\frac{P^* - c_i}{P^*} = \frac{s_i^*}{\eta} \tag{9.22}$$

where η is the price elasticity of industry demand.

Equation (9.22) is a further implication of the Cournot model, now extended to allow for many firms with differing costs. What it says is this. A firm that produces in an industry where demand is relatively inelastic and where it has a relatively large market share will

Reality Checkpoint

Cournot Theory and Public Policy: The 1982 Merger Guidelines

In our review of antitrust policy in Chapter 1, we noted the dramatic change in policy regarding the treatment of mergers that occurred in 1982. In that year, the Department of Justice issued a new version of its *Horizontal Merger Guidelines*. This version replaced the original guidelines issued in 1968. Like that first set of guidelines, the 1982 document specified the conditions under which the government would challenge horizontal mergers. Unlike their predecessor, however, the new guidelines were based explicitly on the Herfindahl Index. Specifically, they stated that a merger would not be challenged if the industry Herfindahl Index was less than 1,000. A merger would also not be challenged if the index was over 1,000 but less than 1,800 *and* if the merger did not raise the Herfindahl Index by over 100 points. If the Herfindahl Index exceeded 1,800 points, then any merger that raised the index by over 50 points would cause concern and likely be challenged.

We will discuss these guidelines and their more recent modifications again in Chapter 16. For now, the point to note is that the explicit use of the Herfindahl Index may be viewed as a bow to the Cournot model which, as shown in the text, directly connects that index to the price–cost margin measure of monopoly power.

Source: Department of Justice, *Horizontal Merger Guidelines* (1982, 1984).

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 219

also be a firm with a substantial degree of market power as measured by the Lerner Index or the firm's price-marginal cost distortion.

The relationship described in equation (9.22) tells us about market power at the level of the firm. In Chapter 3, we discussed the structure-conduct-performance (SCP) paradigm in industrial organization that linked market power, as measured by the Lerner Index, to the structure of the *industry*. The question that remains is whether we can extend the relationship in (9.22) at the firm level to the level of the entire industry.

To see how let us first multiply each side of equation (9.22) by the firm's market share, s_i^*. Then add together the modified equation for firm 1 with that for firm 2 and that for firm 3 and so on until we add together all N modified (9.22) equations. The left-hand side of this sum of N equations is:

$$\sum_{i=1}^{N} s_i^* \left(\frac{P^* - c_i}{P^*}\right) = \frac{\left(\sum_{i=1}^{N} s_i^* P^* - \sum_{i=1}^{N} s_i^* c_i\right)}{P^*} = \frac{P^* - \bar{c}}{P^*}$$

where $\bar{c}$ is the weighted average unit cost of production, the weights being the market shares of the firms in the industry. The right-hand side of the summed N equations is:

$$\frac{\sum_{i=1}^{N} (s_i^*)^2}{\eta} = \frac{H}{\eta}$$

where H is the Herfindahl Index that we defined as a measure of concentration in Chapter 4 (here expressed using fractional shares, e.g., a 10 percent share is recorded as $s_i = 0.10$. Therefore equation (9.22) aggregated at the level of the industry implies that:

$$\frac{(P^* - \bar{c})}{P^*} = \frac{H}{\eta} \tag{9.23}$$

Our generalized Cournot model thus gives theoretical support for the view that as concentration (here measured by the industry's Herfindahl Index) increases, prices also rise farther and farther above marginal cost.

A variant of the relationship in (9.23) was tested in Marion, Mueller, Cotterill, Geithmann, and Schmelzer (1979) for food products. They collected price data for a basket of 94 grocery products, and market share data for 36 firms operating in 32 U.S. Standard Metropolitan Statistical Areas, and found that price is significantly higher in markets with a higher Herfindahl Index. Likewise Marvel (1989) found that for 22 U.S. cities, concentration in the retail market for gasoline, as measured by the Herfindahl Index, had a significant impact on the average price of gasoline.

Summary

For industries populated by a relatively small number of firms, strategic interaction is a fact of life. Each firm is aware of the fact that its decisions have a significant impact on its rivals. Each firm will want to take account of the anticipated response of its rivals when determining its course of action. It is reasonable to believe that firms' anticipations or expectations are rational.

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 220

Game theory is the modern formal technique for studying rational strategic interaction. Each player in a game has a set of strategies to choose from. A strategy combination is a set of strategies—one for each player. Each such strategy combination implies a particular payoff or final outcome for each player. A Nash equilibrium is a strategy combination such that each player is maximizing her payoff *given* the strategies chosen by all other players. In a Nash equilibrium, no player has an incentive to change his behavior unilaterally.

In this chapter we presented the well-known Cournot model of competition. It is a static or single market period model of oligopoly. Although this model was developed prior to the formal development of game theory, the outcome proposed by Cournot captures basic game theoretic principles, specifically the Nash equilibrium solution.

The Cournot model makes clear the importance of firms recognizing and understanding their interdependence. The model also has the nice intuitive implication that the degree of departure from competitive pricing may be directly linked to the structure of the industry as measured by the Herfindahl index. However, as pointed out before in Chapter 4, market structure is endogenous. Strategies that generate above normal profits for existing firms will induce new firms to enter. At the same time, incumbent firms may be able to take actions that deter such entry. We need to extend our analysis in ways that allow us to examine these issues.

The Cournot model studied in this chapter has firms interacting only once. The reality, of course, is that firms are involved in strategic interactions repeatedly. In such a setting, issues such as learning, establishing a reputation, and credibility can become quite important. We turn to a consideration of how the nature of strategic interaction over time affects market structure in Chapters 11–13.

Problems

1. Harrison and Tyler are two students who met by chance the last day of exams before the end of the spring semester and the beginning of summer. Fortunately, they liked each other very much. Unfortunately, they forgot to exchange addresses. Fortunately, each remembers that they spoke of attending a campus party that night. Unfortunately, there are two such parties. One party is small. If each attends this party, they will certainly meet. The other party is huge. If each attends this one, there is a chance they will not meet because of the crowd. Of course, they will certainly not meet if they attend separate parties. Payoffs to each depending on the combined choice of parties are shown below, with Tyler's payoffs listed first.

		Harrison	
		Go to small party	Go to large party
Tyler	Go to small party	(1,000, 1,000)	(0, 0)
	Go to large party	(0, 0)	(500, 500)

   a. Identify the Nash equilibria for this problem.
   b. Identify the Pareto optimal outcome for this "two party" system.

2. Suppose that the small party of Problem 1 is hosted by the "Outcasts," twenty men and women students trying to organize alternatives to the existing campus party establishment. All 20 Outcasts will attend the party. But many other students—not unlike Harrison and Tyler—only go to a party to which others (no one in particular, just people in general) are expected to come. As a result, total attendance A at the small party depends on just how many people X everyone *expects* to show up. Let the relationship between A and X be given by: $A = 20 + 0.6X$.
   a. Explain this equation. Why is the intercept 20? Why is the relation between A and X positive?
   b. If the equilibrium requires that party-goers' expectations be correct, what is the equilibrium attendance at the Outcasts' party?

3. A game known well to both academics and teenage boys is "Chicken." Two players each drive their car down the center of a road in opposite directions. Each chooses either *Stay*

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 221

or *Swerve*. Staying wins adolescent admiration and a big payoff *if* the other player chooses *Swerve*. Swerving loses face and has a low payoff when the other player stays. Bad as that is, it is still better than the payoff when both players choose *Stay* in which case they each are badly hurt. These outcomes are described below with player A's payoffs listed first.

		Player B	
		Stay	Swerve
Player A	Stay	(–6, –6)	(2, –2)
	Swerve	(–2, 2)	(1, 1)

a. Find the Nash equilibria in this game.
b. This is a good game to introduce mixed strategies. If player A adopts the strategy, *Stay* one-fifth of the time, and *Swerve* four-fifths of the time. Show that player B will then be indifferent between either strategy, *Stay* or *Swerve*.
c. If *both* players use this probability mix, what is the chance that they will both be badly hurt?

4. You are a manager of a small "widget"-producing firm. There are only 2 firms including yours that produce "widgets." Moreover your company and your competitor's are identical. You produce the same good and face the same costs of production described by the following total cost function: total cost = 1,500 + $8q$ where q is the output of an individual firm. The market-clearing price, at which you can sell your widgets to the public, depends on how many widgets both you and your rival choose to produce. A market research company has found that market demand for widgets can be described as: $P = 200 - 2Q$ where $Q = q_1 + q_2$, where q_1 is your output and q_2 is your rivals. The board of directors has directed you to choose an output level that will *maximize* the firm's profit. How many widgets should your firm produce in order to achieve the profit-maximizing goal? Moreover you must present your strategy to the board of directors and explain to them why producing this amount of widgets is the profit-maximizing strategy.

5. You are still a manager of a small widget-producing firm. Now however there are 14 such firms (including yours) in the industry. Each firm is identical; each one produces the same product and has the same costs of production. Your firm, as well as each one of the other firms, has the same total cost function, namely: total cost = 200 + $50q$ where q is the output of an individual firm. The price at which you can sell your widgets is determined by market demand, which has been estimated as: $P = 290 - \frac{1}{3}Q$ where Q is the sum of all the individual firms producing in this industry. So, for example, if 120 widgets are produced in the industry then the market-clearing price will be 250 whereas if 300 widgets are produced then the market-clearing price will be 190. The Board of Directors has directed you to choose an output level that maximizes the firm's profit. You have an incentive to maximize profits because your job and salary depend upon the profit performance of this company. Moreover you should also be able to present your profit-maximizing strategy to the Board of Directors and explain to them why producing this amount maximizes the firm's profit.

6. The inverse market demand for fax paper is given by $P = 400 - 2Q$. There are two firms who produce fax paper. Each firm has a unit cost of production equal to 40, and they compete in the market in quantities. That is, they can choose any quantity to produce, and they make their quantity choices simultaneously.
   a. Show how to derive the Cournot–Nash equilibrium to this game? What are firms' profits in equilibrium?
   b. What is the monopoly output, i.e. the one that maximizes total industry profit? Why isn't producing one half the monopoly output a Nash equilibrium outcome?

7. Return to problem 6, but suppose now that firm 1 has a cost advantage. Its unit cost is constant and equal to 25 whereas firm 2 still has a unit cost of 40. What is the Cournot outcome now? What are the profits for each firm?

8. We can use the Cournot model to derive an equilibrium industry structure. For this purpose, we will define an equilibrium as that structure

9781405176323_4_009.qxd 10/19/07 8:10 PM Page 222

in which no firm has an incentive to leave or enter the industry. If a firm leaves the industry, it enters an alternative competitive market in which case it earns zero (economic) profit. If an additional firm enters the industry when there are already n firms in it, the new firm's profit is determined by the Cournot equilibrium with $n + 1$ firms. For this problem, assume that each firm has the cost function: $C(q) = 256 + 20q$. Assume further that market demand is described by: $P = 100 - Q$.

a. Find the long-run equilibrium number of firms in this industry.
b. What industry output, price, and firm profit levels will characterize the long-run equilibrium?

References

Cournot, A. 1836. *Researches into the Mathematical Principles of the Theory of Wealth*. Paris: Hachette (English Translation by N. T. Bacon, New York: Macmillan, 1897).

Friedman, J. 1977. *Oligopoly Theory*. Amsterdam: North-Holland.

Marion, B. W., W. F. Mueller, R. W. Cotterill, F. E. Geithman, and J. R. Schmelzer. 1979. *The Food Retailing Industry Market Structure, Profits and Prices*. New York: Praeger.

Marvel, H. 1989. "Concentration and Price in Gasoline Retailing." In Leonard Weiss, ed., *Concentration and Price*. Cambridge, MA: MIT Press.

Mas-Colell, A., M. D. Whinston, and J. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Rasmusen, E. 2007. *Games and Information: An Introduction to Game Theory*. 4th edition. Cambridge: Blackwell.

Schelling, T. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Tedlow, R. 1996. *New and Improved: The Story of Mass Marketing in America*. 2nd edition. Boston: Harvard Business School Press.