Jean Tirole

The Theory of Industrial Organization

The MIT Press
Cambridge, Massachusetts
London, England

Introduction

Price and Nonprice Competition

In an oligopolistic market structure, a firm no longer encounters a passive environment. Therefore, we need to incorporate the strategic interactions of various decision makers in our models. To do this, we will make extensive use of the theory of noncooperative games.

Firms can use many instruments to compete in a market. Grossly simplifying, we can classify these instruments according to the speed at which they can be altered. In the short run, price is often the main instrument that a firm can change easily (other instruments include advertising and sales-force effort). Therefore, we begin our analysis with price competition in the context of rigid cost structures and product characteristics. In the somewhat longer run, cost structures and product characteristics can be altered, either together or separately. Production techniques can be rearranged and improved upon; capacity can be increased. Product characteristics (quality, product design, delivery delay, location of outlets, and so forth) can be changed. The consumers' perception of the product, which influences the demand function, can be modified by advertising. Ultimately, there is the decision of whether or not to enter or stay in the market (a "0-1" choice). Finally, in the long run, the product characteristics and the cost structures may be changed, not only by simple adjustments within the existing set of products and feasible costs, but also by a modification of this set. Research and development allows firms to expand their choice sets. "Process innovation" alters the technological production possibilities, and "product innovation" affords the creation of new products.

We can very crudely schematize the different stages of competition as in figure 1.

Chapter 5 deals with short-run price competition, examines the Bertrand paradox (in which two or more identical firms producing a homogeneous good with a

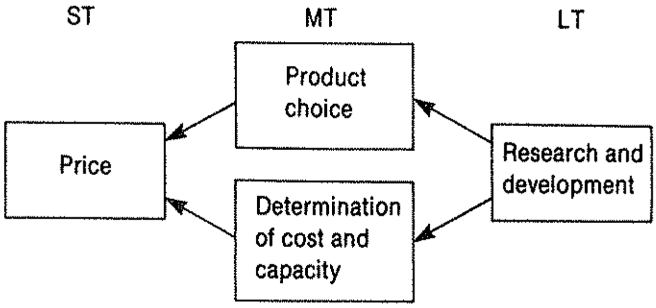

Figure 1

constant-returns-to-scale technology in equilibrium sell at marginal cost and make no profit), discusses why the Bertrand conclusion is disturbing, and suggests three factors that, in practice, smooth price competition. The effect of capacity constraints is studied later in the chapter. Chapter 6 looks at repeated price competition, and chapter 7 introduces product differentiation. The last three chapters deal with barriers to entry, accommodation, predation, and exit (chapters 8 and 9) and with competition in research and development and the adoption of new technologies (chapter 10).

Noncooperative Games and Strategic Behavior

We shall model oligopolistic behavior as noncooperative games in which each firm behaves in its own self-interest. We are especially interested in the equilibria of these games. Nash equilibrium is the basic solution concept in game theory. A set of actions[1] is in Nash equilibrium if, given the actions of its rivals, a firm cannot increase its own profit by choosing an action other than its equilibrium action. For example, take two firms (the analysis generalizes trivially to n firms). Firm i ($i = 1, 2$) earns profit $\Pi^i(a_i, a_j)$, where a_i is the action of firm i and a_j is the action of its rival. We say that a pair of feasible actions is in Nash equilibrium if, for all i and any feasible action a_i,

$$\Pi^i(a_i^*, a_j^*) \geqslant \Pi^i(a_i, a_j^*). \tag{1}$$

The strategies we study here are pure strategies; each firm chooses a simple action. We could also consider mixed strategies where each firm chooses randomly from a set of actions. Of course, in order for firm i to be willing to choose randomly from this set of different actions, all actions in it must yield the same profit (or expected profit, if firm j also plays a mixed strategy), and this profit must be optimal relative to the set of feasible actions a_i.

Nash equilibrium generalizes naturally to dynamic situations and to problems of incomplete information. We shall consider first the dynamic concept of Nash equilibrium (designated "perfect" in the jargon of game theory). This idea becomes particularly important as soon as there are many time periods and any intertemporal dependency of profits or feasible action sets; that is, when the actors make choices in period t that affect their objective functions or their set of feasible choices in a future time period $t + t'$, where $t' > 0$. To determine the consequences of actions taken in t, the players must forecast what will happen in $t + t'$ given the state of the game at the beginning of that period (which is influenced by their actions in t). To calculate these expectations, each player assumes that all other players will play an optimal strategy in $t + t'$. Therefore, the solution of a dynamic game is "backward looking." For example, in a two-period game, we start by solving the second-period Nash equilibrium as a function of the state of the game at the beginning of the second period (that is, on the basis of what happened during the first period). This means that the players can determine the future consequences of their first-period actions, because their first-period actions determine which second-period equilibrium will ensue; in a sense, given their first-period actions, the remainder of the game is a foregone conclusion. Therefore, the players choose their first-period actions with an eye toward their consequences in *both* periods. Thus, it suffices to determine the Nash equilibrium of the corresponding game in which players take only first-period actions but with the same set of consequences as in the original two-period game. All this may appear rather abstract, but it will become much clearer with some examples.

The Nash concept also generalizes to situations of asymmetric information. For instance, a firm may *ex ante* have one of two cost structures and be the only party to know which of these two is realized. The other parties must then figure out how this firm plays optimally for each possible realization of the cost structure. The no-

1. For expositional simplicity, we blur the distinction between action and strategy here. See the Game Theory User's Manual for more details.

tion of Bayesian equilibrium shows precisely how Nash equilibrium may be extended to this type of situation. Finally, in dynamic games with asymmetric information, the notions of perfect equilibrium and Bayesian equilibrium can be combined to further extend the relevance of Nash equilibrium.

Because most problems of industrial organization can be solved with a handful of basic game-theoretic concepts, it is recommended that the reader develop at least a casual familiarity with game theory. Although most of the arguments in part II can be understood at an intuitive level, the reader will benefit from a formal acquaintance with the concept of Nash equilibrium and its extensions, in the same way that optimization techniques clarify the study of the exercise of monopoly power. The reader may find the Game Theory User's Manual (chapter 11) useful in this respect.

Does noncooperative game theory remain relevant in situations in which firms appear to collude? In industrial organization, as in other fields, collusion and noncooperative behavior are not inconsistent. First, an altruistic party's objective function may embody the objectives of another party. In such a case, the first party's own interest is to make decisions that help the other party. (Here, altruism means cooperative actions taken purely for reason of self-interest.) Second, in the absence of altruism, parties facing conflicts may wish to change the rules of the game they are playing if this game has disastrous consequences for them. Signing a contract is a way of doing so. For instance, duopolists may agree to share the market in order to avoid cutthroat competition. However, signing a contract is formally only a part of a bigger noncooperative game. These two reasons why collusion can emerge from self-interested behavior may have limited relevance in IO. First, firms are rarely thought of as altruistic. Second, signing collusive contracts to prevent competition is often illegal. A third and more important reason is that, in a dynamic context, a firm may want to "pull its punches" because an aggressive action would trigger a rational reaction or retaliation from its opponents. (This will be emphasized in chapter 6 and, to a lesser extent, in chapter 8.) Again, the collusion is only apparent; it results from optimal noncooperative behavior. (This type of collusion is sometimes called *tacit* collusion.)

Reaction Functions: Strategic Complements and Substitutes

Consider a simultaneous-move game between (for simplicity) two firms. Assume that each action belongs to the real line and that the profit functions $\Pi^i(a_i, a_j)$ are twice continuously differentiable in the actions. The (necessary) first-order condition for a Nash equilibrium is that for each firm i

$$\Pi^i_i(a_i^*, a_j^*) = 0, \tag{2}$$

where a subscript denotes a partial derivative (e.g., $\Pi^i_i \equiv \partial\Pi^i/\partial a_i$). The second-order condition is that $a_i = a_i^*$ yields a local maximum:

$$\Pi^i_{ii}(a_i^*, a_j^*) \leqslant 0. \tag{3}$$

Assume that each firm's profit function is strictly concave in its own action everywhere: $\Pi^i_{ii}(a_i, a_j) < 0$ for all (a_i, a_j). Then the second-order condition is satisfied and, furthermore, the first-order condition given in equation 2 is sufficient for a Nash equilibrium. A Nash equilibrium is then given by a system of two equations with two unknowns (equation 2).

Let us define $R_i(a_j)$ as the best action for firm i given that firm j chooses a_j:

$$\Pi^i_i(R_i(a_j), a_j) = 0. \tag{4}$$

$a_i = R_i(a_j)$ is unique from our strict-concavity assumption,[2] and is called firm i's reaction to a_j. A Nash equilibrium is a pair of actions (a_1^*, a_2^*) such that $a_1^* = R_1(a_2^*)$ and $a_2^* = R_2(a_1^*)$. In such an equilibrium, each firm reacts optimally to the other firm's anticipated action.

A crucial element of part II is the sign of the slope of reaction functions for the various strategic variables we consider. This slope is obtained by differentiating equation 4:

$$R_i'(a_j) = \frac{\Pi^i_{ij}(R_i(a_j), a_j)}{-\Pi^i_{ii}(R_i(a_j), a_j)}. \tag{5}$$

2. We will assume that it exists and is an interior solution. In other words, going to the boundary of the feasible set of actions (e.g., $-\infty$ or $+\infty$) is not optimal for firm i.

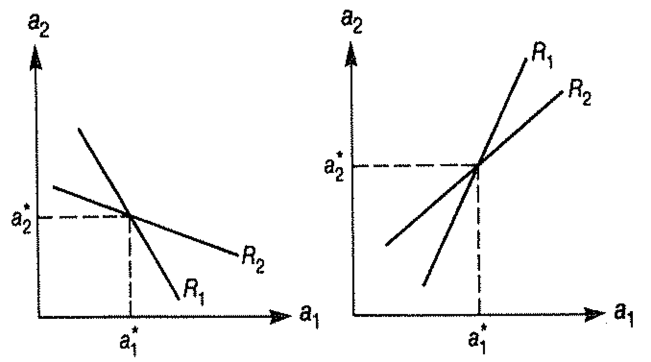

Strategic substitutes ($\Pi^i_{ij} < 0$) Strategic complements ($\Pi^i_{ij} > 0$)

Figure 2

We thus have $\text{sign}(R_i') = \text{sign}(\Pi^i_{ij})$. Π^i_{ij} is the cross-partial derivative of firm i's profit function, i.e., the derivative of its marginal profit with respect to its opponent's action. The reaction curve is upward sloping if $\Pi^i_{ij} > 0$ and downward sloping if $\Pi^i_{ij} < 0$. Following Bulow, Geanakoplos, and Klemperer,[3] we will also consider the actions of the two firms to be strategic complements if $\Pi^i_{ij} > 0$ and strategic substitutes if $\Pi^i_{ij} < 0$.[4] As we shall see further on, prices are often strategic complements, and capacities are often strategic substitutes.

The construction of the reaction functions in a simultaneous-move game, performed in figure 2, is no more than a technical and illustrative device. By definition of simultaneous choices, a firm chooses its action before observing that of its opponent. Hence, it has no possibility of reacting. Reaction functions depict what a firm would do if it were to learn of a change in its opponent's action (which it does not). Points other than the Nash point on the reaction curves are never observed.

In contrast, reaction functions have real economic content in dynamic (sequential) games. For instance, if firm i chooses a_i first and firm j observes this choice before choosing a_j, firm i can use the function R_j to compute how a change in its behavior affects its opponent's behavior.

3. J. Bulow, J. Geanakoplos, and P. Klemperer, "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy* 93 (1985): 488–511.

4. This terminology is inspired by demand theory. Two goods are complements for a consumer if a decrease in the price of one good makes the other good more attractive to the consumer. Here, a decrease in a_j induces a decrease in a_i if $\Pi^i_{ij} > 0$, and conversely for substitutes.

Short-Run Price Competition

The study of price competition—a fundamental part of oligopoly theory—is one of its weakest links. It so happens that the most obvious natural formalization yields a result that is sometimes unconvincing. Deeper reflection shows that this formalization is economically naive, and alternative approaches come to mind.

In this chapter we assume that firms "meet only once" in the market. They simultaneously and noncooperatively charge a price. The Bertrand paradox, discussed in section 5.1, states that under these circumstances even oligopolists behave like competitive firms—that is, the number of firms in the industry is irrelevant to the study of price behavior. Section 5.2 offers an overview of the three alternative approaches that will be developed below and in the next two chapters. Section 5.3 introduces one such approach which is associated with decreasing returns to scale or capacity constraints; it studies foundations for the rival model to the Bertrand paradigm, the Cournot model of competition in quantities. The Cournot model assumes that firms pick quantities rather than prices, and that an auctioneer chooses the price to equate supply and demand. This model has been justly criticized on the grounds that no such auctioneer exists and that firms ultimately choose prices. The point of section 5.3 is that Cournot competition may be thought of as a two-stage game in which firms first choose capacities (or, more generally, scale variables) and then compete through prices. Section 5.4 reviews the main properties of the Cournot paradigm. Section 5.5 discusses concentration indices. The supplementary section completes sections 5.3 and 5.4 with a discussion of capacity-constrained price competition and other aspects of the Cournot model.

5.1 The Bertrand Paradox

To simplify, let us take the case of a duopoly. The analysis generalizes straightforwardly to the case of n firms.

Assume that two firms produce identical goods which are "nondifferentiated" in that they are perfect substitutes in the consumers' utility functions. Consequently, consumers buy from the producer who charges the lowest price. If the firms charge the same price, we must make an assumption about the distribution of consumers between them. We assume that each firm faces a demand schedule equal to half of the market demand at the common price (the half is not a crucial assumption). Further, we assume that the firm always supplies the demand it faces (this assumption is not crucial here). The market demand function is $q = D(p)$. Each firm incurs a cost c per unit of production. Therefore, the profit of firm i is

$$\Pi^i(p_i, p_j) = (p_i - c)D_i(p_i, p_j), \tag{5.1}$$

where the demand for the output of firm i, denoted D_i, is given by

$$D_i(p_i, p_j) = \begin{cases} D(p_i) & \text{if } p_i < p_j \\ \frac{1}{2}D(p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j. \end{cases}$$

The aggregate profit,

$$\min_{p_i}(p_i - c)D(p_i),$$

cannot exceed the monopoly profit,

$$\Pi^m = \max_p(p - c)D(p).$$

Each firm can guarantee itself a non-negative profit by charging a price above the marginal cost. Hence, any reasonable prediction must yield

$$0 \leqslant \Pi^1 + \Pi^2 \leqslant \Pi^m.$$

The firms choose their prices both simultaneously and noncooperatively. Simultaneity means that each firm has not yet observed the other firm's price when choosing its own. Rather, a firm anticipates it. We assume that it does so correctly. A Nash equilibrium in prices—sometimes referred to as a Bertrand equilibrium—is a pair of prices (p_1^*, p_2^*) such that each firm's price maximizes that firm's profit given the other firm's price. Formally, for all $i = 1$, 2 and for all p_i,

$$\Pi^i(p_i^*, p_j^*) \geqslant \Pi^i(p_i, p_j^*).$$

The Bertrand (1883) paradox states that the unique equilibrium has the two firms charge the competitive price: $p_1^* = p_2^* = c$. The proof is as follows: Consider, for example,

$$p_1^* > p_2^* > c.$$

Then firm 1 has no demand, and its profit is zero. On the other hand, if firm 1 charges

$$p_1 = p_2^* - \varepsilon$$

(where ε is positive and "small"), it obtains the entire market demand, $D(p_2^* - \varepsilon)$, and has a positive profit margin of

$$p_2^* - \varepsilon - c.$$

Therefore, firm 1 cannot be acting in its own best interest if it charges p_1^*. Now suppose that

$$p_1^* = p_2^* > c.$$

The profit of firm 1 is

$$D(p_1^*)(p_1^* - c)/2.$$

If firm 1 reduces its price slightly to $p_1^* - \varepsilon$, its profit becomes

$$D(p_1^* - \varepsilon)(p_1^* - \varepsilon - c),$$

which is greater for small ε. In this situation, the market share of the firm increases in a discontinuous manner. Because no firm will charge less than the unit cost c (the lowest-price firm would make a negative profit), we are left with one or two firms charging exactly c. To show that both firms do charge c, suppose that

$$p_1^* > p_2^* = c.$$

Then firm 2, which makes no profit, could raise its price slightly, still supply all the demand, and make a positive profit—a contradiction.

The conclusions of this simple model are the following:

(i) that firms price at marginal cost, and

(ii) that firms do not make profits.

These conclusions suggest that the monopoly results of chapter 1 are very special. Even a duopoly would suffice to restore competition. We call this the Bertrand *paradox* because it is hard to believe that firms in industries with few firms never succeed in manipulating the market price to make profits.[1]

In the asymmetric case (say, where firm i has constant unit cost c_i, where $c_1 < c_2$), conclusions i and ii do not hold. Indeed, the following can be shown (up to some technical considerations, see exercise 5.1):

(iii) that both firms charge price $p = c_2$ (actually, firm 1 charges an ε below c_2 to make sure it has the whole market), and

(iv) that firm 1 makes a profit of $(c_2 - c_1)D(c_2)$, and firm 2 makes no profit (as long as $c_2 \leqslant p^m(c_1)$, where $p^m(c_1)$ maximizes $(p-c_1)D(p)$; otherwise, firm 1 charges $p^m(c_1)$).

Thus, firm 1 charges above marginal cost and makes a positive profit, and the Bertrand equilibrium is no longer welfare-optimal. But, again, the conclusion is a bit strained. Firm 1 makes very little profit if c_2 is close to c_1, and firm 2 makes no profit at all.

*Exercise 5.1** Prove conclusions iii and iv.

5.2 Solutions to the Bertrand Paradox: An Introduction

We can resolve the Bertrand paradox by relaxing any one of the three crucial assumptions of the model. Each of these generalizations brings more realism to the problem of price determination. The first is studied in section 5.3; the other two are taken up in chapters 6 and 7. In the present section we will sketch these possible solutions.

5.2.1 The Edgeworth Solution

Edgeworth (1897) solved the Bertrand paradox by introducing capacity constraints, by which firms cannot sell more than they are capable of producing. To understand this idea, suppose that firm 1 has a production capacity smaller than $D(c)$. Is $(p_1^*, p_2^*) = (c, c)$ still an equilibrium price system? At this price, both firms make zero profit. Suppose that firm 2 increases its price slightly. Firm 1 then faces demand $D(c)$, which it cannot satisfy. Then, rationing dictates that some consumers must resort to firm 2. Firm 2 has a (residual) nonzero demand at a price greater than its marginal cost and, therefore, makes positive profit. Consequently, the Bertrand solution is no longer an equilibrium.

To solve explicitly for the equilibrium, we must introduce a more specific assumption concerning the manner in which consumers are rationed. As a general rule, in models with capacity constraints, firms make positive profit and the market price is greater than the marginal cost. The crucial question now is whether or not this property is relevant: Won't firms accumulate capital *ex ante* until they are capable of satisfying the entire market demand at marginal cost? The answer is No. To accumulate capital is expensive, and it is not in the self-interest of the firms to do so if such behavior yields zero gross profit (without netting out capital costs).[2]

Using capacity constraints to justify noncompetitive prices is quite reasonable in some applications. For example, imagine the case of two hotels in a small town. In the short run, these hotels cannot adjust the number of beds (capacity). It is useless for them to get involved in cut-throat price competition if they are incapable of satisfying market demand individually. In the longer run they do not increase their capacity very much, because they expect keen competition in a situation of collective overcapacity. One can also consider the case where the production of a good requires a certain delay. Then, the available quantities for sale in the very short run cannot be adjusted at all, and therefore they act as capacity constraints at the time of price competition.

The existence of a rigid capacity constraint is a special case of a decreasing-returns-to-scale technology. In our previous example, a firm has marginal cost that is equal to c up to the capacity constraint and is then equal to infinity. More generally, the marginal cost may increase with the output. Except in special cases (such as the hotel example), a firm usually has some leeway to increase its production beyond its "efficient level"—extra machines can be rented, existing ones can be utilized beyond their efficient intensity of use, inputs can be provided on short

1. Another paradox of the model is that one wonders why firms bother to enter at all if they do not make any profit. Along the same lines, suppose that the firms face a fixed cost of entering the market. Then, if one firm enters, the other firm will not follow suit, however small the fixed cost. Thus, if one believes in the existence of at least a small fixed cost of production or of entry, the market is likely to yield a monopoly.

2. See sections 5.3 and 5.7.

notice, and workers can work overtime. The cost of producing these extra units exceeds that for the inframarginal ones, but is, in general, not infinite.

5.2.2 The Temporal Dimension

The second crucial assumption underlying the Bertrand paradox is the "timing" of the game, which does not always seem to reflect economic reality as it stands. To see this, consider a crucial condition for the Bertrand solution. In particular, why is $p_1 = p_2 > c$ not an equilibrium? The answer is that firm 1, for example, would benefit from a slight decrease in its price (i.e., to $p_2 - \varepsilon$) and from its resulting takeover of the entire market. What would happen then? Nothing, given Bertrand's crucial condition that players are assumed to play only once. Firm 2 would lose all its customers, and would make zero profit because it would not react. In reality, firm 2 would probably decrease its price in order to regain its share of the market. If we introduce this temporal dimension and the possibility of reaction, it is no longer clear that firm 1 would benefit from decreasing its price below p_2. Firm 1 would have to compare the short-run gain (the increase of its market share) to the longer-run loss in a price war. Chapter 6 shows that more collusive behavior than in the Bertrand equilibrium can be sustained by the threat of future losses in a price war.

5.2.3 Product Differentiation

An important assumption in the Bertrand analysis is the perfect substitutability of the firms' products. Consumers are indifferent between the goods at equal price and, thus, buy from the lowest-priced producer. This creates a pressure on price, which is somewhat relaxed when the firms' products are not quite identical (see chapter 2 for a description of a few differentiation spaces). Then, in general, the firms do not charge their marginal cost. For instance, think of two firms selling the same good but located at different places. Suppose that firm 1 charges $p_1 = c$. Firm 2, by charging $p_2 = c + \varepsilon$ (for ε small), keeps at least some consumers who are located near it. For these consumers, the price differential is more than offset by the difference in transportation cost. Hence, the zero-profit price system ($p_1 = c$, $p_2 = c$) is no longer an equilibrium. (An extreme case of product differentiation occurs when the demands for the firms' products are unrelated. Each firm then charges its monopoly price.) Price competition with differentiated products is analyzed in chapter 7.

5.2.4 What to Make of the Bertrand Analysis

Bertrand competition is interesting because it depicts a polar case. It represents what we have in mind when we think of sharp small-number competition. In general, of course, oligopoly pricing will lead to an outcome intermediate between the Bertrand one and the outcome of the other polar case (the monopoly situation). Most of our analysis of price rivalry will concern the determination of the factors that induce tough or soft competition.

5.4 Traditional Cournot Analysis

We now analyze the one-stage game in which firms choose their quantities (understand their capacities) simultaneously. We will use either the general reduced form for the profit function, $\Pi^i(q_i, q_j)$, or the more specific exact Cournot form:

$$\Pi^i(q_i, q_j) = q_i P(q_i + q_j) - C_i(q_i)$$

(see the caveats on this exact form in section 5.3).

Each firm maximizes its profit given the quantity chosen by the other firm. Assuming that the profit function Π^i is strictly concave in q_i and twice differentiable, we get

$$q_i = R_i(q_j), \tag{5.4}$$

where R_i is firm i's reaction curve:

$$\Pi^i_i(R_i(q_j), q_j) = 0.$$

Recall from the introduction to part II that if we assume that firm i's marginal profit is decreasing with the other firm's quantity, then the reaction functions are downward sloping. The equilibrium quantities are depicted in figure 5.5 by the intersection of the two reaction curves. Of course, such an intersection need not be unique; in that case we would have multiple equilibria.

To be a bit more specific, let us consider the first-order condition for profit maximization for the exact Cournot form:

$$\Pi^i_i = P(q_i + q_j) - C_i'(q_i) + q_i P'(q_i + q_j) = 0. \tag{5.5}$$

This has a simple interpretation. The first two terms yield

undersold. This confirms the idea that a high capacity threatens one's rivals and forces them to price aggressively.

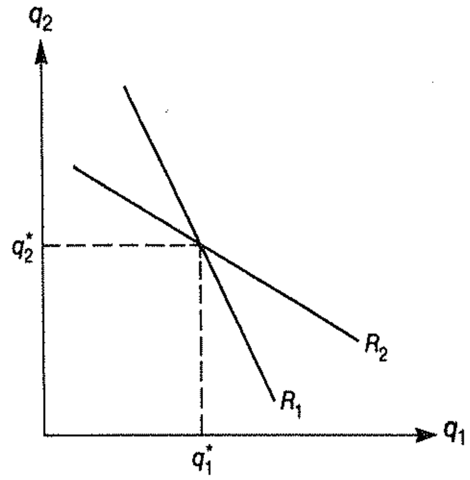

Figure 5.5

the profitability of an extra unit of output, which is equal to the difference between price and marginal cost. The third term represents the effect of this extra unit on the profitability of inframarginal ones. The extra units create a decrease in price P', which affects the q_i units already produced. Equation 5.5 is similar to the formulas obtained for a competitive firm and a monopoly. For a competitive firm there is no third term, because the firm is too small to affect the market price; for a monopoly, q_i is equal to the output of the industry.

The preceding comparison actually illustrates the negative *externality* between the firms: When choosing its output, firm i takes into account the adverse effect of the market-price change on its own output, rather than the effect on aggregate output. Hence, each firm will tend to choose an output that exceeds the optimal output from the industry's point of view[14] (but, of course, not from a welfare point of view). Thus, *the market price will be lower than the monopoly price, and the aggregate profit will be lower than the monopoly profit.* Another interesting consequence of equation 5.5 is that Cournot equilibrium does not equalize marginal costs except in the symmetric case. Not only is too little produced, but *the industry's cost of production is not minimized.*

Equation 5.5 can be rewritten as

$$L_i = \frac{\alpha_i}{\varepsilon}, \tag{5.6}$$

where

$$L_i \equiv \frac{P - C_i'}{P}$$

is the Lerner index (price-cost margin) for firm i;

$$\alpha_i \equiv \frac{q_i}{Q}$$

is firm i's market share ($Q \equiv q_i + q_j$), and

$$\varepsilon \equiv -\frac{P'}{P}Q$$

is the elasticity of demand. Thus, the Lerner index is proportional to the firm's market share and inversely proportional to the elasticity of demand. This index is positive—that is, firms sell at a price exceeding marginal cost. Thus, the Cournot equilibrium is not socially efficient.

A technical note about the concavity of the firm's objective function and the sign of the cross-partial derivative: From equation 5.5, we obtain

$$\Pi^i_{ii} = 2P' + q_i P'' - C_i'' \tag{5.7}$$

and

$$\Pi^i_{ij} = P' + q_i P''. \tag{5.8}$$

Recall that $P' < 0$. For the objective function to be concave ($\Pi^i_{ii} < 0$), it suffices that the firm's cost be convex ($C_i'' \geqslant 0$) and that the inverse-demand function be concave ($P'' \leqslant 0$). The latter assumption suffices for quantities to be strategic substitutes ($\Pi^i_{ij} < 0$). These two assumptions are met, for instance, for linear demand ($P'' = 0$) and constant returns to scale ($C_i'' = 0$). For more on the concavity of the objective function and the existence of a Cournot equilibrium see section 5.7.

The Cournot equilibrium is easily derived in the case of linear demand and cost. Suppose that $D(p) = 1 - p$ (or $P(Q) = 1 - Q$) and $C_i(q_i) = c_i q_i$. Then the reaction functions are

$$q_i = R_i(q_j) = \frac{1 - q_j - c_i}{2}.$$

Hence, the Cournot equilibrium is given by

$$q_i = \frac{1 - 2c_i + c_j}{3},$$

14. This is obtained by replacing $q_i P'$ in equation 5.5 with $(q_i + q_j)P'$.

and the profit is

$$\Pi^i = \frac{(1 - 2c_i + c_j)^2}{9}.$$

A firm's output decreases with its marginal cost. More interestingly, it increases with its competitor's marginal cost; this is because a higher c_j leads firm j to produce less, which raises the residual demand faced by firm i, encouraging firm i to produce more.

That a firm's output decreases with its marginal cost and increases with its competitor's marginal cost can be obtained for more general demand and cost functions as long as the following two conditions are satisfied: (a) the reaction curves are downward sloping (quantities are strategic substitutes) and (b) the reaction curves cross only once (there exists a unique Cournot equilibrium) and the slope of R_2 in the (q_1, q_2) space is smaller in absolute value than the slope of R_1.[15]

It is easily shown that an increase in a firm's marginal cost shifts the firm's reaction curve down. To prove this, recall from chapter 1 that a monopolist's price (respectively, quantity) increases (respectively, decreases) with the firm's marginal cost. But in duopoly, for a given output q_j, firm i is a monopolist on the residual-demand curve $P(\cdot + q_j)$. Hence, the proof of chapter 1 applies, and firm i's optimal output given q_j is a decreasing (more precisely, a nonincreasing) function of firm i's marginal cost. This result is completely general; conditions such as (a) or (b) are not required. (The reader is advised to go through the argument again as an exercise.)

Figure 5.6 depicts reaction curves satisfying (a) and (b), and shows the effect of an increase in firm 1's marginal cost. Firm 1's equilibrium output is reduced, while firm 2's increases.

The results generalize straightforwardly to the case of n firms. Let

$$Q \equiv \sum_{i=1}^{n} q_i.$$

Equation 5.5 then becomes

$$P(Q) - C_i'(q_i) + q_i P'(Q) = 0. \tag{5.9}$$

Firm i's Lerner index is still equal to the ratio of its market

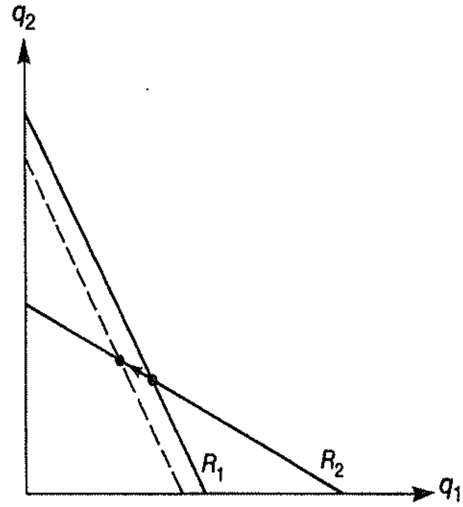

Figure 5.6
Effect of an increase in firm 1's marginal cost.

share to the elasticity of demand. For instance, for the symmetric case with linear cost and demand,

$P(Q) = 1 - Q$

and $C_i(q_i) = cq_i$

for all i (with $c < 1$), equation 5.9 becomes

$$1 - Q - c - q_i = 0. \tag{5.10}$$

The equilibrium is symmetric for this symmetric model: $Q = nq$, where q is the output per firm. Hence, we obtain

$$q = \frac{1 - c}{n + 1}. \tag{5.11}$$

The market price is

$$p = 1 - nq = c + \frac{1 - c}{n + 1}, \tag{5.12}$$

and each firm's profit is

$$\Pi = \frac{(1 - c)^2}{(n + 1)^2}. \tag{5.13}$$

The market price and each firm's profit decrease with the number of firms. Furthermore, because the market price decreases with n, so does the aggregate profit $n\Pi$. Indeed, *when the number of firms becomes very large ($n \to \infty$), the market price tends to the competitive price c.* Thus, a Cournot

15. This "stability condition" is treated in chapter 8. A sufficient condition is that, for $i = 1, 2$, $|R_i'| < 1$—that is, a decrease in one firm's production yields a decrease in aggregate output even if its rival reacts (optimally) to the decrease in production. For instance, for a linear demand, $|R_i'| = \frac{1}{2}$.

equilibrium with a large number of firms is approximately competitive. This is natural, because each firm has only a small influence on the price and thus acts almost like a price taker.

See section 5.7 for more on the convergence to the competitive equilibrium, and for a discussion of the existence and uniqueness of a Cournot equilibrium.

*Exercise 5.3** There are three identical firms in the industry. The demand is $1 - Q$, where $Q = q_1 + q_2 + q_3$. The marginal cost is zero.

(i) Compute the Cournot equilibrium.

(ii) Show that if two of the three firms merge (transforming the industry into a duopoly), the profit of these firms decreases. Explain.

(iii) What happens if all three firms merge?

(iv)** If the firms were competing in prices and sold differentiated products, would a merger between two of them be profitable? (Work at an intuitive level, and assume that prices are strategic complements.)

*Exercise 5.4** Consider a duopoly producing a homogeneous product. Firm 1 produces one unit of output with one unit of labor and one unit of raw material. Firm 2 produces one unit of output with two units of labor and one unit of raw material. The unit costs of labor and raw material are w and r. The demand is $p = 1 - q_1 - q_2$, and the firms compete in quantities.

(i) Compute the Cournot equilibrium.

(ii) Show that firm 1's profit is not affected by the price of labor (over some range). To prove this elegantly, use the envelope theorem. Explain.

*Exercise 5.5** This exercise illustrates the strategic considerations faced by a multimarket firm. It is inspired by the more general theory of Bulow et al. (1985).

There are two firms in a market. They produce perfect substitutes at cost $C(q) = q^2/2$. The demand is $p = 1 - (q_1 + q_2)$.

(i) Compute the Cournot equilibrium.

(ii) Suppose now that firm 1 has the opportunity to sell the same output on another market as well. The quantity sold on this market is x_1, so firm 1's cost is $(q_1 + x_1)^2/2$. The demand on the second market is $p = a - x_1$. Consider the Cournot game in which firm 1 chooses q_1 and x_1 and firm 2 chooses q_2 simultaneously. Show that $q_1 = (2 - a)/7$ and $q_2 = (5 + a)/21$ over the relevant range of a. Show that for $a = \frac{1}{4}$ a small increase in a hurts firm 1. (Use the envelope theorem.) Interpret.