

CHAPTER 27

OLIGOPOLY

We have now investigated two important forms of market structure: pure competition, where there are typically many small competitors, and pure monopoly, where there is only one large firm in the market. However, much of the world lies between these two extremes. Often there are a number of competitors in the market, but not so many as to regard each of them as having a negligible effect on price. This is the situation known as **oligopoly**.

The model of monopolistic competition described in Chapter 24 is a special form of oligopoly that emphasizes issues of product differentiation and entry. However, the models of oligopoly that we will study in this chapter are more concerned with the strategic interactions that arise in an industry with a small number of firms.

There are several models that are relevant since there are several different ways for firms to behave in an oligopolistic environment. It is unreasonable to expect one grand model since many different behavior patterns can be observed in the real world. What we want is a guide to some of the possible patterns of behavior and some indication of what factors might be important in deciding when the various models are applicable.

For simplicity, we will usually restrict ourselves to the case of two firms; this is called a situation of **duopoly**. The duopoly case allows us to capture many of the important features of firms engaged in strategic interaction without the notational complications involved in models with a larger number of firms. Also, we will limit ourselves to investigation of cases in which each firm is producing an identical product. This allows us to avoid the problems of product differentiation and focus only on strategic interactions.

27.1 Choosing a Strategy

If there are two firms in the market and they are producing a homogeneous product, then there are four variables of interest: the price that each firm charges and the quantities that each firm produces.

When one firm decides about its choices for prices and quantities it may already know the choices made by the other firm. If one firm gets to set its price before the other firm, we call it the **price leader** and the other firm the **price follower**. Similarly, one firm may get to choose its quantity first, in which case it is a **quantity leader** and the other is a **quantity follower**. The strategic interactions in these cases form a **sequential game**.¹

On the other hand, it may be that when one firm makes its choices it doesn't know the choices made by the other firm. In this case, it has to guess about the other firm's choice in order to make a sensible decision itself. This is a **simultaneous game**. Again there are two possibilities: the firms could each simultaneously choose prices or each simultaneously choose quantities.

This classification scheme gives us four possibilities: quantity leadership, price leadership, simultaneous quantity setting, and simultaneous price setting. Each of these types of interaction gives rise to a different set of strategic issues.

There is also another possible form of interaction that we will examine. Instead of the firms competing against each other in one form or another they may be able to **collude**. In this case the two firms can jointly agree to set prices and quantities that maximize the sum of their profits. This sort of collusion is called a **cooperative game**.

27.2 Quantity Leadership

In the case of quantity leadership, one firm makes a choice before the other firm. This is sometimes called the **Stackelberg model** in honor of the

¹ We will examine game theory in more detail in the next chapter, but it seems appropriate to introduce these specific examples here.

first economist who systematically studied leader-follower interactions.²

The Stackelberg model is often used to describe industries in which there is a dominant firm, or a natural leader. For example, IBM is often considered to be a dominant firm in the computer industry. A commonly observed pattern of behavior is for smaller firms in the computer industry to wait for IBM's announcements of new products and then adjust their own product decisions accordingly. In this case we might want to model the computer industry with IBM playing the role of a Stackelberg leader, and the other firms in the industry being Stackelberg followers.

Let us turn now to the details of the theoretical model. Suppose that firm 1 is the leader and that it chooses to produce a quantity y_1 . Firm 2 responds by choosing a quantity y_2 . Each firm knows that the equilibrium price in the market depends on the total output produced. We use the inverse demand function $p(Y)$ to indicate the equilibrium price as a function of industry output, $Y = y_1 + y_2$.

What output should the leader choose to maximize its profits? The answer depends on how the leader thinks that the follower will react to its choice. Presumably the leader should expect that the follower will attempt to maximize profits as well, given the choice made by the leader. In order for the leader to make a sensible decision about its own production, it has to consider the follower's profit-maximization problem.

The Follower's Problem

We assume that the follower wants to maximize its profits

$$\max_{y_2} p(y_1 + y_2)y_2 - c_2(y_2).$$

The follower's profit depends on the output choice of the leader, but from the viewpoint of the follower the leader's output is predetermined—the production by the leader has already been made, and the follower simply views it as a constant.

The follower wants to choose an output level such that marginal revenue equals marginal cost:

$$MR_2 = p(y_1 + y_2) + \frac{\Delta p}{\Delta y_2} y_2 = MC_2.$$

The marginal revenue has the usual interpretation. When the follower increases its output, it increases its revenue by selling more output at the

² Heinrich von Stackelberg was a German economist who published his influential work on market organization, *Marktform und Gleichgewicht*, in 1934.

market price. But it also pushes the price down by Δp , and this lowers its profits on all the units that were previously sold at the higher price.

The important thing to observe is that the profit-maximizing choice of the follower will depend on the choice made by the leader. We write this relationship as

$$y_2 = f_2(y_1).$$

The function $f_2(y_1)$ tells us the profit-maximizing output of the follower as a function of the leader's choice. This function is called the **reaction function** since it tells us how the follower will react to the leader's choice of output.

Let's derive a reaction curve in the simple case of linear demand. In this case the (inverse) demand function takes the form $p(y_1 + y_2) = a - b(y_1 + y_2)$. For convenience we'll take costs to be zero.

Then the profit function for firm 2 is

$$\pi_2(y_1, y_2) = [a - b(y_1 + y_2)]y_2$$

or

$$\pi_2(y_1, y_2) = ay_2 - by_1y_2 - by_2^2.$$

We can use this expression to draw the **isoprofit lines** in Figure 27.1. These are lines depicting those combinations of y_1 and y_2 that yield a constant level of profit to firm 2. That is, the isoprofit lines are comprised of all points (y_1, y_2) that satisfy equations of the form


$$ay_2 - by_1y_2 - by_2^2 = \bar{\pi}_2.$$

Note that profits to firm 2 will increase as we move to isoprofit lines that are further to the left. This is true since if we fix the output of firm 2 at some level, firm 2's profits will increase as firm 1's output decreases. Firm 2 will make its maximum possible profits when it is a monopolist; that is, when firm 1 chooses to produce zero units of output.

For each possible choice of firm 1's output, firm 2 wants to choose its own output to make its profits as large as possible. This means that for each choice of y_1 , firm 2 will pick the value of y_2 that puts it on the isoprofit line furthest to the left, as illustrated in Figure 27.1. This point will satisfy the usual sort of tangency condition: the slope of the isoprofit line must be vertical at the optimal choice. The locus of these tangencies describes firm 2's reaction curve, $f_2(y_1)$.

To see this result algebraically, we need an expression for the marginal revenue associated with the profit function for firm 2. It turns out that this expression is given by

$$MR_2(y_1, y_2) = a - by_1 - 2by_2.$$



Derivation of a reaction curve. This reaction curve gives the profit-maximizing output for the follower, firm 2, for each output choice of the leader, firm 1. For each choice of y_1 the follower chooses the output level $f_2(y_1)$ associated with the isoprofit line farthest to the left.

(This is easy to derive using calculus. If you don't know calculus, you'll just have to take this statement on faith.) Setting the marginal revenue equal to marginal cost, which is zero in this example, we have

$$a - by_1 - 2by_2 = 0,$$

which we can solve to derive firm 2's reaction curve:

$$y_2 = \frac{a - by_1}{2b}.$$

This reaction curve is the straight line depicted in Figure 27.1.

The Leader's Problem

We have now examined how the follower will choose its output *given* the choice of the leader. We turn now to the leader's profit-maximization problem.

Presumably, the leader is also aware that its actions influence the output choice of the follower. This relationship is summarized by the reaction

function $f_2(y_1)$. Hence when making its output choice it should recognize the influence that it exerts on the follower.

The profit-maximization problem for the leader therefore becomes

$$\begin{aligned} & \max_{y_1} p(y_1 + y_2)y_1 - c_1(y_1) \\ & \text{such that } y_2 = f_2(y_1). \end{aligned}$$

Substituting the second equation into the first gives us

$$\max_{y_1} p[y_1 + f_2(y_1)]y_1 - c_1(y_1).$$

Note that the leader recognizes that when it chooses output y_1 , the total output produced will be $y_1 + f_2(y_1)$: its own output *plus* the output produced by the follower.

When the leader contemplates changing its output it has to recognize the influence it exerts on the follower. Let's examine this in the context of the linear demand curve described above. There we saw that the reaction function was given by

$$f_2(y_1) = y_2 = \frac{a - by_1}{2b}. \quad (27.1)$$

Since we've assumed that marginal costs are zero, the leader's profits are

$$\pi_1(y_1, y_2) = p(y_1 + y_2)y_1 = ay_1 - by_1^2 - by_1y_2. \quad (27.2)$$

But the output of the follower, y_2 , will depend on the leader's choice via the reaction function $y_2 = f_2(y_1)$.

Substituting from equation (27.1) into equation (27.2) we have

$$\begin{aligned} \pi_1(y_1, y_2) &= ay_1 - by_1^2 - by_1f_2(y_1) \\ &= ay_1 - by_1^2 - by_1 \frac{a - by_1}{2b}. \end{aligned}$$

Simplifying this expression gives us

$$\pi_1(y_1, y_2) = \frac{a}{2}y_1 - \frac{b}{2}y_1^2.$$

The marginal revenue for this function is

$$MR = \frac{a}{2} - by_1.$$

Setting this equal to marginal cost, which is zero in this example, and solving for y_1 gives us


$$y_1^* = \frac{a}{2b}.$$

In order to find the follower's output, we simply substitute y_1^* into the reaction function,

$$\begin{aligned} y_2^* &= \frac{a - by_1^*}{2b} \\ &= \frac{a}{4b}. \end{aligned}$$

These two equations give a total industry output of $y_1^* + y_2^* = 3a/4b$.

The Stackelberg solution can also be illustrated graphically using the isoprofit curves depicted in Figure 27.2. (This figure also illustrates the Cournot equilibrium which will be described in section 27.5.) Here we have illustrated the reaction curves for both firms and the isoprofit curves for firm 1. The isoprofit curves for firm 1 have the same general shape as the isoprofit curves for firm 2; they are simply rotated 90 degrees. Higher profits for firm 1 are associated with isoprofit curves that are lower down since firm 1's profits will increase as firm 2's output decreases.



Stackelberg equilibrium. Firm 1, the leader, chooses the point on firm 2's reaction curve that touches firm 1's lowest possible isoprofit line, thus yielding the highest possible profits for firm 1.

Firm 2 is behaving as a follower, which means that it will choose an output along its reaction curve, $f_2(y_1)$. Thus firm 1 wants to choose an

output combination on the reaction curve that gives it the highest possible profits. But the highest possible profits means picking that point on the reaction curve that touches the *lowest* isoprofit line, as illustrated in Figure 27.2. It follows by the usual logic of maximization that the reaction curve must be tangent to the isoprofit curve at this point.

27.3 Price Leadership

Instead of setting quantity, the leader may instead set price. In order to make a sensible decision about how to set its price, the leader must forecast how the follower will behave. Accordingly, we must first investigate the profit-maximization problem facing the follower.

The first thing we observe is that in equilibrium the follower must always set the same price as the leader. This follows from our assumption that the two firms are selling identical products. If one charged a different price from the other, all of the consumers would prefer the producer with the lower price, and we couldn't have an equilibrium with both firms producing.

Suppose that the leader has set a price p . We will suppose that the follower takes this price as given and chooses its profit-maximizing output. This is essentially the same as the competitive behavior we investigated earlier. In the competitive model, each firm takes the price as being outside of its control because it is such a small part of the market; in the price-leadership model, the follower takes the price as being outside of its control since it has already been set by the leader.

The follower wants to maximize profits:

$$\max_{y_2} py_2 - c_2(y_2).$$


This leads to the familiar condition that the follower will want to choose an output level where price equals marginal cost. This determines a supply curve for the follower, $S(p)$, which we have illustrated in Figure 27.3.

Turn now to the problem facing the leader. It realizes that if it sets a price p , the follower will supply $S(p)$. That means that the amount of output the leader will sell will be $R(p) = D(p) - S(p)$. This is called the **residual demand curve** facing the leader.

Suppose that the leader has a constant marginal cost of production c . Then the profits that it achieves for any price p are given by:

$$\pi_1(p) = (p - c)[D(p) - S(p)] = (p - c)R(p).$$

In order to maximize profits the leader wants to choose a price and output combination where marginal revenue equals marginal cost. However, the marginal revenue should be the marginal revenue for the *residual* demand curve—the curve that actually measures how much output it will be able to



Price leader. The demand curve facing the leader is the market demand curve minus the follower's supply curve. The leader equates marginal revenue and marginal cost to find the optimal quantity to supply, y_L^* . The total amount supplied to the market is y_T^* and the equilibrium price is p^* .

sell at each given price. In Figure 27.3 the residual demand curve is linear; therefore the marginal revenue curve associated with it will have the same vertical intercept and be twice as steep.

Let's look at a simple algebraic example. Suppose that the inverse demand curve is $D(p) = a - bp$. The follower has a cost function $c_2(y_2) = y_2^2/2$, and the leader has a cost function $c_1(y_1) = cy_1$.

For any price p the follower wants to operate where price equals marginal cost. If the cost function is $c_2(y_2) = y_2^2/2$, it can be shown that the marginal cost curve is $MC_2(y_2) = y_2$. Setting price equal to marginal cost gives us

$$p = y_2.$$

Solving for the follower's supply curve gives $y_2 = S(p) = p$.

The demand curve facing the leader—the residual demand curve—is

$$R(p) = D(p) - S(p) = a - bp - p = a - (b + 1)p.$$

From now on this is just like an ordinary monopoly problem. Solving for p as a function of the leader's output y_1 , we have

$$p = \frac{a}{b + 1} - \frac{1}{b + 1}y_1. \quad (27.3)$$

This is the inverse demand function facing the leader. The associated marginal revenue curve has the same intercept and is twice as steep. This means that it is given by

$$MR_1 = \frac{a}{b+1} - \frac{2}{b+1}y_1.$$

Setting marginal revenue equal to marginal cost gives us the equation

$$MR_1 = \frac{a}{b+1} - \frac{2}{b+1}y_1 = c = MC_1.$$

Solving for the leader's profit-maximizing output, we have

$$y_1^* = \frac{a - c(b+1)}{2}.$$

We could go on and substitute this into equation (27.3) to get the equilibrium price, but the equation is not particularly interesting.

27.4 Comparing Price Leadership and Quantity Leadership

We've seen how to calculate the equilibrium price and output in the case of quantity leadership and price leadership. Each model determines a different equilibrium price and output combination; each model is appropriate in different circumstances.

One way to think about quantity setting is to think of the firm as making a capacity choice. When a firm sets a quantity it is in effect determining how much it is able to supply to the market. If one firm is able to make an investment in capacity first, then it is naturally modeled as a quantity leader.

On the other hand, suppose that we look at a market where capacity choices are not important but one of the firms distributes a catalog of prices. It is natural to think of this firm as a price setter. Its rivals may then take the catalog price as given and make their own pricing and supply decision accordingly.

Whether the price-leadership or the quantity-leadership model is appropriate is not a question that can be answered on the basis of pure theory. We have to look at how the firms actually make their decisions in order to choose the most appropriate model.

27.5 Simultaneous Quantity Setting

One difficulty with the leader-follower model is that it is necessarily asymmetric: one firm is able to make its decision before the other firm. In some

situations this is unreasonable. For example, suppose that two firms are *simultaneously* trying to decide what quantity to produce. Here each firm has to forecast what the other firm's output will be in order to make a sensible decision itself.

In this section we will examine a one-period model in which each firm has to forecast the other firm's output choice. Given its forecast, each firm then chooses a profit-maximizing output for itself. We then seek an equilibrium in forecasts—a situation where each firm finds its beliefs about the other firm to be confirmed. This model is known as the **Cournot model**, after the nineteenth-century French mathematician who first examined its implications.³

We begin by assuming that firm 1 expects that firm 2 will produce y_2^e units of output. (The e stands for *expected* output.) If firm 1 decides to produce y_1 units of output, it expects that the total output produced will be $Y = y_1 + y_2^e$, and output will yield a market price of $p(Y) = p(y_1 + y_2^e)$. The profit-maximization problem of firm 1 is then

$$\max_{y_1} p(y_1 + y_2^e)y_1 - c(y_1).$$

For any given belief about the output of firm 2, y_2^e , there will be some optimal choice of output for firm 1, y_1 . Let us write this functional relationship between the *expected output* of firm 2 and the *optimal choice* of firm 1 as

$$y_1 = f_1(y_2^e).$$

This function is simply the reaction function that we investigated earlier in this chapter. In our original treatment the reaction function gave the follower's output as a function of the leader's choice. Here the reaction function gives one firm's optimal choice as a function of its *beliefs* about the other firm's choice. Although the interpretation of the reaction function is different in the two cases, the mathematical definition is exactly the same.

Similarly, we can derive firm 2's reaction curve:

$$y_2 = f_2(y_1^e),$$

which gives firm 2's optimal choice of output for a given expectation about firm 1's output, y_1^e .

Now, recall that each firm is choosing its output level *assuming* that the other firm's output will be at y_1^e or y_2^e . For arbitrary values of y_1^e and y_2^e this won't happen—in general firm 1's *optimal* level of output, y_1 , will be different from what firm 2 *expects* the output to be, y_1^e .

Let us seek an output combination (y_1^*, y_2^*) such that the optimal output level for firm 1, assuming firm 2 produces y_2^* , is y_1^* and the optimal output

³ Augustin Cournot (pronounced "core-no") was born in 1801. His book, *Researches into the Mathematical Principles of the Theory of Wealth*, was published in 1838.

level for firm 2, assuming that firm 1 stays at y_1^* , is y_2^* . In other words, the output choices (y_1^*, y_2^*) satisfy

$$y_1^* = f_1(y_2^*)$$

$$y_2^* = f_2(y_1^*).$$

Such a combination of output levels is known as a **Cournot equilibrium**. In a Cournot equilibrium, each firm is maximizing its profits, given its beliefs about the other firm's output choice, and, furthermore, those beliefs are confirmed in equilibrium: each firm optimally chooses to produce the amount of output that the other firm expects it to produce. In a Cournot equilibrium neither firm will find it profitable to change its output once it discovers the choice actually made by the other firm.

An example of a Cournot equilibrium is given in Figure 27.2. The Cournot equilibrium is simply the pair of outputs at which the two reaction curves cross. At such a point, each firm is producing a profit-maximizing level of output given the output choice of the other firm.

27.6 An Example of Cournot Equilibrium

Recall the case of the linear demand function and zero marginal costs that we investigated earlier. We saw that in this case the reaction function for firm 2 took the form

$$y_2 = \frac{a - by_1^e}{2b}.$$

Since in this example firm 1 is exactly the same as firm 2, its reaction curve has the same form:

$$y_1 = \frac{a - by_2^e}{2b}.$$

Figure 27.4 depicts this pair of reaction curves. The intersection of the two lines gives us the Cournot equilibrium. At this point each firm's choice is the profit-maximizing choice, given its beliefs about the other firm's behavior, and each firm's beliefs about the other firm's behavior are confirmed by its *actual* behavior.

In order to calculate the Cournot equilibrium algebraically, we look for the point (y_1, y_2) where each firm is doing what the other firm expects it to do. We set $y_1 = y_1^e$ and $y_2 = y_2^e$, which gives us the following two equations in two unknowns:

$$y_1 = \frac{a - by_2}{2b}$$

$$y_2 = \frac{a - by_1}{2b}.$$




Figure
27.4

Cournot equilibrium. Each firm is maximizing its profits, given its beliefs about the other firm's output decision. The Cournot equilibrium is at (y_1^*, y_2^*) , where the two reaction curves cross.

In this example, both firms are identical, so each will produce the same level of output in equilibrium. Hence we can substitute $y_1 = y_2$ into one of the above equations to get

$$y_1 = \frac{a - by_1}{2b}.$$

Solving for y_1^* , we get

$$y_1^* = \frac{a}{3b}.$$

Since the two firms are identical, this implies that

$$y_2^* = \frac{a}{3b}$$

as well, and the total industry output is

$$y_1^* + y_2^* = \frac{2a}{3b}.$$

27.7 Adjustment to Equilibrium

We can use Figure 27.4 to describe a process of adjustment to equilibrium. Suppose that at time t the firms are producing outputs (y_1^t, y_2^t) , which are not necessarily equilibrium outputs. If firm 1 expects that firm 2 is going to continue to keep its output at y_2^t , then next period firm 1 would want to choose the profit-maximizing output given that expectation, namely $f_1(y_2^t)$. Thus firm 1's choice in period $t + 1$ will be given by

$$y_1^{t+1} = f_1(y_2^t).$$

Firm 2 can reason the same way, so firm 2's choice next period will be

$$y_2^{t+1} = f_2(y_1^t).$$

These equations describe how each firm adjusts its output in the face of the other firm's choice. Figure 27.4 illustrates the movement of the outputs of the firms implied by this behavior. Here is the way to interpret the diagram. Start with some operating point (y_1^t, y_2^t) . Given firm 2's level of output, firm 1 optimally chooses to produce $y_1^{t+1} = f_1(y_2^t)$ next period. We find this point in the diagram by moving horizontally to the left until we hit firm 1's reaction curve.

If firm 2 expects firm 1 to continue to produce y_1^{t+1} , its optimal response is to produce y_2^{t+1} . We find this point by moving vertically upward until we hit firm 2's reaction function. We continue to move along the "staircase" to determine the sequence of output choices of the two firms. In the example illustrated, this adjustment process converges to the Cournot equilibrium. We say that in this case the Cournot equilibrium is a **stable equilibrium**.

Despite the intuitive appeal of this adjustment process, it does present some difficulties. Each firm is assuming that the other's output will be fixed from one period to the next, but as it turns out, both firms keep changing their output. Only in equilibrium is one firm's expectation about the other firm's output choice actually satisfied. For this reason, we will generally ignore the question of how the equilibrium is reached and focus only on the issue of how the firms behave in the equilibrium.

27.8 Many Firms in Cournot Equilibrium

Suppose now that we have several firms involved in a Cournot equilibrium, not just two. In this case we suppose that each firm has an expectation about the output choices of the other firms in the industry and seek to describe the equilibrium output.

Suppose that there are n firms and let $Y = y_1 + \cdots + y_n$ be the total industry output. Then the “marginal revenue equals marginal cost condition” for firm i is

$$p(Y) + \frac{\Delta p}{\Delta Y} y_i = MC(y_i).$$

If we factor out $P(Y)$ and multiply the second term by Y/Y , we can write this equation as

$$p(Y) \left[1 + \frac{\Delta p}{\Delta Y} \frac{Y}{p(Y)} \frac{y_i}{Y} \right] = MC(y_i).$$

Using the definition of elasticity of the aggregate demand curve and letting $s_i = y_i/Y$ be firm i 's share of total market output, this reduces to

$$p(Y) \left[1 - \frac{s_i}{|\epsilon(Y)|} \right] = MC(y_i). \quad (27.4)$$

We can also write this expression as

$$p(Y) \left[1 - \frac{1}{|\epsilon(Y)|/s_i} \right] = MC(y_i).$$

This looks just like the expression for the monopolist except for the s_i term. We can think of $\epsilon(Y)/s_i$ as being the elasticity of the demand curve facing the firm: the smaller the market share of the firm, the more elastic the demand curve it faces.

If its market share is 1—the firm is a monopolist—the demand curve facing the firm is the market demand curve, so the condition just reduces to that of the monopolist. If the firm is a very small part of a large market, its market share is effectively zero, and the demand curve facing the firm is effectively flat. Thus the condition reduces to that of the pure competitor: price equals marginal cost.

This is one justification for the competitive model described in Chapter 22. If there are a large number of firms, then each firm's influence on the market price is negligible, and the Cournot equilibrium is effectively the same as pure competition.

27.9 Simultaneous Price Setting

In the Cournot model described above we have assumed that firms were choosing their quantities and letting the market determine the price. Another approach is to think of firms as setting their prices and letting the market determine the quantity sold. This model is known as **Bertrand competition**.⁴

⁴ Joseph Bertrand, also a French mathematician, presented his model in a review of Cournot's work.

When a firm chooses its price, it has to forecast the price set by the other firm in the industry. Just as in the case of Cournot equilibrium we want to find a pair of prices such that each price is a profit-maximizing choice given the choice made by the other firm.

What does a Bertrand equilibrium look like? When firms are selling identical products, as we have been assuming, the Bertrand equilibrium has a very simple structure indeed. It turns out to be the competitive equilibrium, where price equals marginal cost!

First we note that price can never be less than marginal cost since then either firm would increase its profits by producing less. So let us consider the case where price is greater than marginal cost. Suppose that both firms are selling output at some price $\hat{p}$ greater than marginal cost. Consider the position of firm 1. If it lowers its price by any small amount ϵ and if the other firm keeps its price fixed at $\hat{p}$, all of the consumers will prefer to purchase from firm 1. By cutting its price by an arbitrarily small amount, it can steal all of the customers from firm 2.

If firm 1 really believes that firm 2 will charge a price $\hat{p}$ that is greater than marginal cost, it will always pay firm 1 to cut its price to $\hat{p} - \epsilon$. But firm 2 can reason the same way! Thus any price higher than marginal cost cannot be an equilibrium; the only equilibrium is the competitive equilibrium.

This result seems paradoxical when you first encounter it: how can we get a competitive price if there are only two firms in the market? If we think of the Bertrand model as a model of competitive bidding it makes more sense. Suppose that one firm “bids” for the consumers’ business by quoting a price above marginal cost. Then the other firm can always make a profit by undercutting this price with a lower price. It follows that the only price that each firm cannot rationally expect to be undercut is a price equal to marginal cost.

It is often observed that competitive bidding among firms that are unable to collude can result in prices that are much lower than can be achieved by other means. This phenomenon is simply an example of the logic of Bertrand competition.

27.10 Collusion

In the models we have examined up until now the firms have operated independently. But if the firms collude so as to jointly determine their output, these models are not very reasonable. If collusion is possible, the firms would do better to choose the output that maximizes total industry profits and then divide up the profits among themselves. When firms get together and attempt to set prices and outputs so as to maximize total industry profits, they are known as a **cartel**. As we saw in Chapter 24, a

cartel is simply a group of firms that jointly collude to behave like a single monopolist and maximize the sum of their profits.

Thus the profit-maximization problem facing the two firms is to choose their outputs y_1 and y_2 so as to maximize total industry profits:

$$\max_{y_1, y_2} p(y_1 + y_2)[y_1 + y_2] - c_1(y_1) - c_2(y_2).$$

This will have the optimality conditions

$$p(y_1^* + y_2^*) + \frac{\Delta p}{\Delta Y} [y_1^* + y_2^*] = MC_1(y_1^*)$$

$$p(y_1^* + y_2^*) + \frac{\Delta p}{\Delta Y} [y_1^* + y_2^*] = MC_2(y_2^*).$$

The interpretation of these conditions is interesting. When firm 1 considers expanding its output by Δy_1 , it will contemplate the usual two effects: the extra profits from selling more output and the reduction in profits from forcing the price down. But in the second effect, it now takes into account the effect of the lower price on both its own output and the output of the other firm. This is because it is now interested in maximizing total industry profits, not just its own profits.

The optimality conditions imply that the marginal revenue of an extra unit of output must be the same no matter where it is produced. It follows that $MC_1(y_1^*) = MC_2(y_2^*)$, so that the two marginal costs will be equal in equilibrium. If one firm has a cost advantage, so that its marginal cost curve always lies below that of the other firm, then it will necessarily produce more output in equilibrium in the cartel solution.

The problem with agreeing to join a cartel in real life is that there is always a temptation to cheat. Suppose, for example, that the two firms are operating at the outputs that maximize industry profits (y_1^*, y_2^*) and firm 1 considers producing a little more output, Δy_1 . The marginal profits accruing to firm 1 will be

$$\frac{\Delta \pi_1}{\Delta y_1} = p(y_1^* + y_2^*) + \frac{\Delta p}{\Delta Y} y_1^* - MC_1(y_1^*). \quad (27.5)$$

We saw earlier that the optimality condition for the cartel solution is

$$p(y_1^* + y_2^*) + \frac{\Delta p}{\Delta Y} y_1^* + \frac{\Delta p}{\Delta Y} y_2^* - MC_1(y_1^*) = 0.$$

Rearranging this equation gives us

$$p(y_1^* + y_2^*) + \frac{\Delta p}{\Delta Y} y_1^* - MC_1(y_1^*) = -\frac{\Delta p}{\Delta Y} y_2^* > 0. \quad (27.6)$$

The last inequality follows since $\Delta p/\Delta Y$ is negative, since the market demand curve has a negative slope.

Inspecting equations (27.5) and (27.6) we see that

$$\frac{\Delta\pi_1}{\Delta y_1} > 0.$$

Thus, if firm 1 believes that firm 2 will keep its output fixed, then it will believe that it can increase profits by increasing its own production. In the cartel solution, the firms act together to restrict output so as not to “spoil” the market. They recognize the effect on joint profits from producing more output in either firm. But if each firm believes that the other firm will stick to its output quota, then each firm will be tempted to increase its own profits by unilaterally expanding its output. At the output levels that maximize joint profits, it will always be profitable for each firm to unilaterally increase its output—if each firm expects that the other firm will keep its output fixed.

The situation is even worse than that. If firm 1 believes that firm 2 will keep its output fixed, then it will find it profitable to increase its own output. But if it thinks that firm 2 will increase its output, then firm 1 would want to increase its output first and make its profits while it can!

Thus, in order to maintain an effective cartel, the firms need a way to detect and punish cheating. If they have no way to observe each other's output, the temptation to cheat may break the cartel. We'll return to this point a little later.

To make sure that we understand the cartel solution, let's calculate it for the case of zero marginal costs and the linear demand curve we used in the Cournot case.

The aggregate profit function will be

$$\pi(y_1, y_2) = [a - b(y_1 + y_2)](y_1 + y_2) = a(y_1 + y_2) - b(y_1 + y_2)^2,$$

so the marginal revenue equals marginal cost conditions will be


$$a - 2b(y_1^* + y_2^*) = 0,$$

which implies that

$$y_1^* + y_2^* = \frac{a}{2b}.$$

Since marginal costs are zero, the division of output between the two firms doesn't matter. All that is determined is the total level of industry output.

This solution is shown in Figure 27.5. Here we have illustrated the isoprofit curves for each of the firms and have highlighted the locus of common tangents. Why is this line of interest? Since the cartel is trying to maximize total industry profits, it follows that the marginal profits from having either firm produce more output must be the same—otherwise it would pay to have the more profitable firm produce more output. This in



A cartel. If industry profits are maximized, then the marginal profit from producing more output in either firm must be the same. This implies that the isoprofit curves must be tangent to one another at the profit-maximizing levels of output.

turn implies that the slopes of the isoprofit curves must be the same for each firm; that is, that the isoprofit curves must be tangent to each other. Hence the output combinations that maximize total industry profits—the cartel solution—are those that lie along the line illustrated in Figure 27.5.

Figure 27.5 also illustrates the temptation to cheat that is present at the cartel solution. Consider, for example, the point where the two firms split the market equally. Think about what would happen if firm 1 believed that firm 2 would keep its output constant. If firm 1 increased its output and firm 2 kept constant output, then firm 1 would move to a lower isoprofit curve—which means that firm 1 would increase its profits. This is exactly the story told in the algebra above. If one firm thinks that the other's output will remain constant, it will be tempted to increase its own output and thereby make higher profits.

27.11 Punishment Strategies

We have seen that a cartel is fundamentally unstable in the sense that it is always in the interest of each of the firms to increase their production above that which maximizes aggregate profit. If the cartel is to operate

successfully, some way must be found to “stabilize” the behavior. One way to do this is for firms to threaten to punish each other for cheating on the cartel agreement. In this section, we investigate the size of punishments necessary to stabilize a cartel.

Consider a duopoly composed of two identical firms. If each firm produces half the monopoly amount of output, total profits will be maximized and each firm will get a payoff of, say, π_m . In an effort to make this outcome stable, one firm announces to the other: “If you stay at the production level that maximizes joint industry projects, fine. But if I discover you cheating by producing more than this amount, I will punish you by producing the Cournot level of output forever.” This is known as a **punishment strategy**.

When will this sort of threat be adequate to stabilize the cartel? We have to look at the benefits and costs of cheating as compared to those of cooperating. Suppose that cheating occurs, and the punishment is carried out. Since the optimal response to Cournot behavior is Cournot behavior (by definition), this results in each firm receiving a per-period profit of, say, π_c . Of course, the Cournot payoff, π_c is less than the cartel payoff, π_m .

Let us suppose that the two firms are each producing at the collusive, monopoly level of production. Put yourself in the place of one of the firms trying to decide whether to continue to produce at your quota. If you produce more output, deviating from your quota, you make profit π_d , where $\pi_d > \pi_m$. This is the standard temptation facing a cartel member described above: if each firm restricts output and pushes the price up, then each firm has an incentive to capitalize on the high price by increasing its production.

But this isn’t the end of the story because of the punishment for cheating. By producing at the cartel amount, each firm gets a steady stream of payments of π_m . The present value of this stream starting today is given by

$$\text{Present value of cartel behavior} = \pi_m + \frac{\pi_m}{r}.$$

If the firm produces more than the cartel amount, it gets a one-time benefit of profits π_d , but then has to live with the breakup of the cartel and the reversion to Cournot behavior:

$$\text{Present value of cheating} = \pi_d + \frac{\pi_c}{r}.$$

When will the present value of remaining at the cartel output be greater than the present value of cheating on the cartel agreement? Obviously when

$$\pi_m + \frac{\pi_m}{r} > \pi_d + \frac{\pi_c}{r},$$

which can also be written as

$$r < \frac{\pi_m - \pi_c}{\pi_d - \pi_m}.$$

Note that the numerator of this fraction is positive, since the monopoly profits are larger than the Cournot profits, and the denominator is positive, since deviation is even more profitable than sticking with the monopoly quota.

The inequality says that as long as the interest rate is sufficiently small, so that the prospect of future punishment is sufficiently important, it will pay the firms to stick to their quotas.

The weakness of this model is that the threat to revert to Cournot behavior forever is not very believable. One firm certainly may believe that the other will punish it for deviating, but “forever” is a long time. A more realistic model would consider shorter periods of retaliation, but the analysis then becomes much more complex. In the next chapter, we discuss some models of “repeated games” that illustrate some of the possible behaviors.

EXAMPLE: Price Matching and Competition

We have seen that there is always a temptation for each member of a cartel to produce more than its quota. In order to maintain a successful cartel, some way must be found to police members’ behavior by some form of punishment for deviations from the joint profit-maximizing output. In particular this means that firms must be able to keep track of the prices and production levels of the other firms in the cartel.

One easy way to acquire information about what the other firms in your industry are charging is to use your customers to spy on the other firms. It is common to see retail firms announce that they will “beat any price.” In some cases, such an offer may indicate a highly competitive retail environment. But in other cases, this same policy can be used to gather information about other firms’ prices in order to maintain a cartel.

Suppose, for example, that two firms agree, either explicitly or implicitly to sell a certain model of refrigerator for \$700. How can either of the stores be sure that the other firm isn’t cheating on their agreement and selling the refrigerator for \$675? One way is to offer to beat any price a customer can find. That way, the customers report any attempts to cheat on the collusive arrangement.

EXAMPLE: Voluntary Export Restraints

During the 1980s, the Japanese automobile companies agreed to a “voluntary export restraint (VER).” This meant that they would “voluntarily” reduce the exports of their automobiles to the United States. The typical U.S. consumer thought that this was a great victory for U.S. trade negotiators.

But if you think about this for a minute, things look quite different. In our examination of oligopoly we have seen that the problem facing firms in an industry is how to *restrict* output in order to support higher prices and discourage competition. As we've seen, there will always be a temptation to cheat on production agreements; every cartel must find a way to detect and prevent this cheating. It is especially convenient for the firms if a third party, such as the government, can serve this role. This is exactly the role that the U.S. government played for the Japanese auto makers!

According to one estimate Japanese imported cars were about \$2500 more expensive in 1984 than they would have been without the VERs. Furthermore, the higher prices of imported cars allowed American producers to sell their automobiles at about \$1000 more than they would have otherwise.⁵

Due to these higher prices the U.S. consumers paid about \$10 billion more for Japanese cars in 1985–86 than they would have otherwise. This money has gone directly into the pockets of the Japanese automobile producers. Much of this additional profit appears to have been invested in increasing productive capabilities, which allowed the Japanese auto producers to reduce the cost of producing new cars in subsequent years. The VERs did succeed in saving American jobs; however, it appears that the cost per job saved was about \$160,000 per year.

If the goal of the VER policy was simply to increase the health of the American automobile industry, there was a much simpler way to do this: just impose a \$2500 tariff on each imported Japanese car. This way the revenues due to the restriction of trade would accrue to the U.S. government rather than to the Japanese automobile industry. Rather than send \$10 billion abroad during 1985–86, the U.S. government could have spent the money on projects designed to increase the long-term health of the U.S. auto industry.

27.12 Comparison of the Solutions

We have now examined several models of duopoly behavior: quantity leadership (Stackelberg), price leadership, simultaneous quantity setting (Cournot), simultaneous price setting (Bertrand), and the collusive solution. How do they compare?

In general, collusion results in the smallest industry output and the highest price. Bertrand equilibrium—the competitive equilibrium—gives us the highest output and the lowest price. The other models give results that are in between these two extremes.

⁵ Robert Crandall, "Import Quotas and the Automobile Industry: the Costs of Protectionism," *The Brookings Review*, Summer, 1984.

A variety of other models are possible. For example, we could look at a model with differentiated products where the two goods produced were not perfect substitutes for each other. Or we could look at a model where the firms make a sequence of choices over time. In this framework, the choices that one firm makes at one time can influence the choices that the other firm makes later on.

We have also assumed that each firm knows the demand function and the cost functions of the other firms in the industry. In reality these functions are never known for sure. Each firm needs to estimate the demand and cost conditions facing its rivals when it makes its own decisions. All of these phenomena have been modeled by economists, but the models become much more complex.

Summary

1. An oligopoly is characterized by a market with a few firms that recognize their strategic interdependence. There are several possible ways for oligopolies to behave depending on the exact nature of their interaction.
2. In the quantity-leader (Stackelberg) model one firm leads by setting its output, and the other firm follows. When the leader chooses an output, it will take into account how the follower will respond.
3. In the price-leader model, one firm sets its price, and the other firm chooses how much it wants to supply at that price. Again the leader has to take into account the behavior of the follower when it makes its decision.
4. In the Cournot model, each firm chooses its output so as to maximize its profits given its beliefs about the other firm's choice. In equilibrium each firm finds that its expectation about the other firm's choice is confirmed.
5. A Cournot equilibrium in which each firm has a small market share implies that price will be very close to marginal cost—that is, the industry will be nearly competitive.
6. In the Bertrand model each firm chooses its price given its beliefs about the price that the other firm will choose. The only equilibrium price is the competitive equilibrium.
7. A cartel consists of a number of firms colluding to restrict output and to maximize industry profit. A cartel will typically be unstable in the sense that each firm will be tempted to sell more than its agreed upon output if it believes that the other firms will not respond.

REVIEW QUESTIONS

1. Suppose that we have two firms that face a linear demand curve $p(Y) = a - bY$ and have constant marginal costs, c , for each firm. Solve for the Cournot equilibrium output.
2. Consider a cartel in which each firm has identical and constant marginal costs. If the cartel maximizes total industry profits, what does this imply about the division of output between the firms?
3. Can the leader ever get a lower profit in a Stackelberg equilibrium than he would get in the Cournot equilibrium?
4. Suppose there are n identical firms in a Cournot equilibrium. Show that the absolute value of the elasticity of the market demand curve must be greater than $1/n$. (Hint: in the case of a monopolist, $n = 1$, and this simply says that a monopolist operates at an elastic part of the demand curve. Apply the logic that we used to establish that fact to this problem.)
5. Draw a set of reaction curves that result in an unstable equilibrium.
6. Do oligopolies produce an efficient level of output?