CHAPTER 26

MONOPOLY BEHAVIOR

In a competitive market there are typically several firms selling an identical product. Any attempt by one of the firms to sell its product at more than the market price leads consumers to desert the high-priced firm in favor of its competitors. In a monopolized market there is only one firm selling a given product. When a monopolist raises its price it loses some, but not all, of its customers.

In reality most industries are somewhere in between these two extremes. If a gas station in a small town raises the price at which it sells gasoline and it loses most of its customers, it is reasonable to think that this firm must behave as a competitive firm. If a restaurant in the same town raises its price and loses only a few of its customers, then it is reasonable to think that this restaurant has some degree of monopoly power.

If a firm has some degree of monopoly power it has more options open to it than a firm in a perfectly competitive industry. For example, it can use more complicated pricing and marketing strategies than a firm in a competitive industry. Or it can try to differentiate its product from the products sold by its competitors to enhance its market power even further. In this chapter we will examine how firms can enhance and exploit their market power.

26.1 Price Discrimination

We have argued earlier that a monopoly operates at an inefficient level of output since it restricts output to a point where people are willing to pay more for extra output than it costs to produce it. The monopolist doesn't want to produce this *extra* output, because it would force down the price that it would be able to get for *all* of its output.

But if the monopolist could sell different units of output at different prices, then we have another story. Selling different units of output at different prices is called **price discrimination**. Economists generally consider the following three kinds of price discrimination:

First-degree price discrimination means that the monopolist sells different units of output for different prices *and* these prices may differ from person to person. This is sometimes known as the case of **perfect price discrimination**.

Second-degree price discrimination means that the monopolist sells different units of output for different prices, but every individual who buys the same amount of the good pays the same price. Thus prices differ across the units of the good, but not across people. The most common example of this is bulk discounts.

Third-degree price discrimination occurs when the monopolist sells output to different people for different prices, but every unit of output sold to a given person sells for the same price. This is the most common form of price discrimination, and examples include senior citizens' discounts, student discounts, and so on.

Let us look at each of these to see what economics can say about how price discrimination works.

26.2 First-Degree Price Discrimination

Under **first-degree price discrimination**, or **perfect price discrimination**, each unit of the good is sold to the individual who values it most highly, at the maximum price that this individual is willing to pay for it.

Consider Figure 26.1, which illustrates two consumers' demand curves for a good. Think of a reservation price model for demand where the individuals choose integer amounts of the goods and each step in the demand curve represents a change in the willingness to pay for additional units of the good. We have also illustrated (constant) marginal cost curves for the good.

A producer who is able to perfectly price discriminate will sell each unit of the good at the highest price it will command, that is, at each consumer's reservation price. Since each unit is sold to each consumer at his or her reservation price for that unit, there is no consumers' surplus generated in

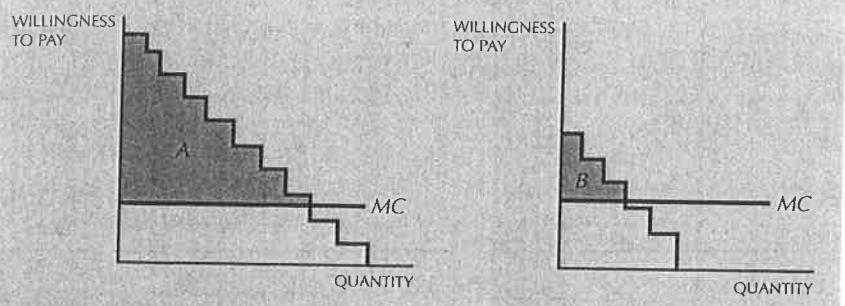

First-degree price discrimination. Here are two consumers' demand curves for a good along with the constant marginal cost curve. The producer sells each unit of the good at the maximum price it will command, which yields it the maximum possible profit.

Figure 26.1

this market; all the surplus goes to the producer. In Figure 26.1 the colored areas indicate the *producer's surplus* accruing to the monopolist. In an ordinary competitive market setting these areas would represent *consumers' surplus*, but in the case of perfect price discrimination, the monopolist is able to appropriate this surplus for itself.

Since the producer gets all the surplus in the market, it wants to make sure that the surplus is as large as possible. Put another way, the producer's goal is to maximize its profits (producer's surplus) subject to the constraint that the consumers are just willing to purchase the good. This means that the outcome will be Pareto efficient, since there will be no way to make both the consumers and the producer better off: the producer's profit can't be increased, since it is already the maximal possible profit, and the consumers' surplus can't be increased without reducing the profit of the producer.

If we move to the smooth demand curve approximation, as in Figure 26.2, we see that a perfectly price-discriminating monopolist must produce at an output level where price equals marginal cost: if price were greater than marginal cost, that would mean that there is someone who is willing to pay more than it costs to produce an extra unit of output. So why not produce that extra unit and sell it to that person at his or her reservation price, and thus increase profits?

Just as in the case of a competitive market, the sum of producer's and consumers' surpluses is maximized. However, in the case of perfect price discrimination the producer ends up getting *all* the surplus generated in the market!

We have interpreted first-degree price discrimination as selling each unit at the maximum price it will command. But we could also think of it as selling a fixed amount of the good at a "take it or leave it" price. In the

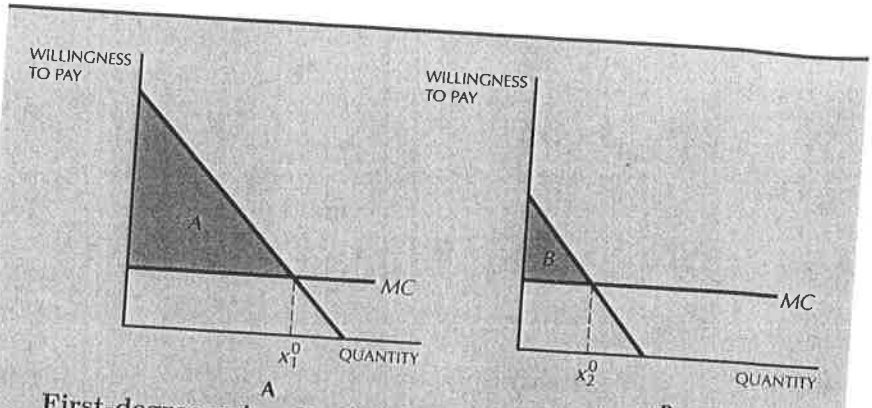

Figure 26.2 **First-degree price discrimination with smooth demand curves.** Here are two consumers' smoothed demand curves for a good along with the constant marginal cost curve. Here the producer maximizes profits by producing where price equals marginal cost, just as in the case of a competitive market.

case illustrated in Figure 26.2, the monopolist would offer to sell x_1^0 units of the good to person 1 at a price equal to the area under person 1's demand curve and offer to sell x_2^0 units of the good to person 2 at a price equal to the area under person 2's demand curve B. As before, each person would end up with zero consumer's surplus, and the entire surplus of $A+B$ would end up in the hands of the monopolist.

Perfect price discrimination is an idealized concept—as the word "perfect" might suggest—but it is interesting theoretically since it gives us an example of a resource allocation mechanism other than a competitive market that achieves Pareto efficiency. There are very few real-life examples of perfect price discrimination. The closest example would be something like a small-town doctor who charges his patients different prices, based on their ability to pay.

EXAMPLE: First-degree Price Discrimination in Practice

As mentioned earlier, first-degree price discrimination is primarily a theoretical concept. It's hard to find real-world examples in which every individual is charged a different price. One possible example would be cases where prices are set by bargaining, as in automobile sales or in antique markets. However, these are not ideal examples.

Southwest Airlines recently introduced a system called Ding that attempts something rather close to first-degree price discrimination.[1] The

[1] See Christopher Elliott, "Your Very Own Personal Air Fare," *New York Times*, August 9, 2005.

system uses the Internet in a clever way. The user installs a program on her computer and the airline sends special fare offers to the user periodically. The fares are announced with a "ding" sound, hence the system name. According to one analyst, the fares offered by Ding were about 30 percent lower than comparable fares.

But will these low fares persist? One might also use such a system to offer higher fares. However, that possibility seems unlikely given the intensely competitive nature of the airline industry. It's easy to switch back to standard ways of buying tickets if prices start creeping up.

26.3 Second-Degree Price Discrimination

Second-degree price discrimination is also known as the case of **nonlinear pricing**, since it means that the price per unit of output is not constant but depends on how much you buy. This form of price discrimination is commonly used by public utilities; for example, the price per unit of electricity often depends on how much is bought. In other industries bulk discounts for large purchases are sometimes available.

Let us consider the case depicted earlier in Figure 26.2. We saw that the monopolist would *like* to sell an amount x_1^0 to person 1 at price $A+$ cost and an amount x_2^0 to person 2 at price $B+$ cost. To set the right prices, the monopolist has to *know* the demand curves of the consumers; that is, the monopolist has to know the exact willingness to pay of each person. Even if the monopolist knows something about the statistical distribution of willingness to pay—for example, that college students are willing to pay less than yuppies for movie tickets—it might be hard to tell a yuppie from a college student when they are standing in line at the ticket booth.

Similarly, an airline ticket agent may know that business travelers are willing to pay more than tourists for their airplane tickets, but it is often difficult to tell whether a particular person is a business traveler or a tourist. If switching from a grey flannel suit to Bermuda shorts would save $500 on travel expenses, corporate dress codes could change quickly!

The problem with the first-degree price discrimination example depicted in Figure 26.2 is that person 1—the high-willingess-to-pay person—can *pretend* to be person 2, the low-willingess-to-pay person. The seller may have no effective way to tell them apart.

One way to get around this problem is to offer two different price-quantity packages in the market. One package will be targeted toward the high-demand person, the other package toward the low-demand person. It can often happen that the monopolist can construct price-quantity packages that will induce the consumers to choose the package meant for them; in economics jargon, the monopolist constructs price-quantity packages that give the consumers an incentive to **self select**.

In order to see how this works, Figure 26.3 illustrates the same kind of demand curves used in Figure 26.2, but now laid on top of each other. We've also set marginal cost equal to zero in this diagram to keep the argument simple.

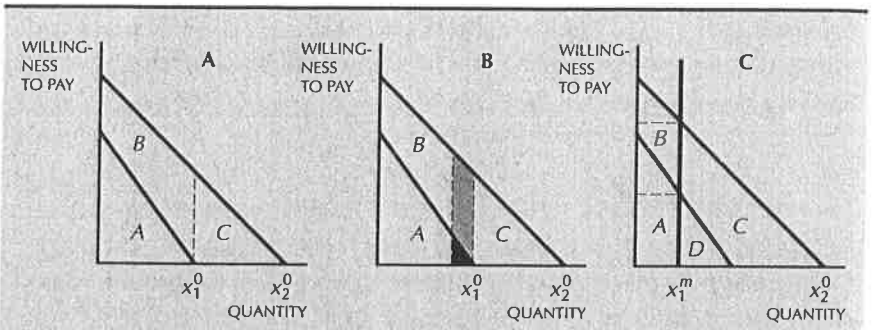

Figure 26.3

Second-degree price discrimination. These are the demand curves of two consumers; the producer has zero marginal cost by assumption. Panel A illustrates the self-selection problem. Panel B shows what happens if the monpolist reduces the output targeted for consumer 1, and panel C illustrates the profit-maximizing solution.

As before, the monopolist would like to offer x_1^0 at price A and to offer x_2^0 at price $A + B + C$. This would capture all the surplus for the monopolist and generate the most possible profit. Unfortunately for the monopolist, these price-quantity combinations are not compatible with self-selection. The high-demand consumer would find it optimal to choose the quantity x_1^0 and pay price A; this would leave him with a surplus equal to area B, which is better than the zero surplus he would get if he chose x_2^0.

One thing the monopolist can do is to offer x_2^0 at a price of $A+C$. In this case the high-demand consumer finds it optimal to choose x_2^0 and receive a gross surplus of $A + B + C$. He pays the monopolist $A + C$, which yields a net surplus of B for consumer 2—just what he would get if he chose x_1^0. This generally yields more profit to the monopolist than it would get by offering only one price-quantity combination.

But the story doesn't end here. There's yet a further thing the monopolist can do to increase profits. Suppose that instead of offering x_1^0 at price A to the low-demand consumer, the monopolist offers a bit less than that at a price slightly less than A. This reduces the monopolist's profits on person 1 by the small colored triangle illustrated in Figure 26.3B. But note that since person 1's package is now less attractive to person 2, the

monopolist can now charge *more* to person 2 for x_2^0! By reducing x_1^0, the monopolist makes area A a little smaller (by the dark triangle) but makes area C bigger (by the triangle plus the light trapezoid area). The net result is that the monopolist's profits increase.

Continuing in this way, the monopolist will want to reduce the amount offered to person 1 up to the point where the profit lost on person 1 due to a further reduction in output just equals the profit gained on person 2. At this point, illustrated in Figure 26.3C, the marginal benefits and costs of quantity reduction just balance. Person 1 chooses x_1^m and is charged A; person 2 chooses x_2^0 and is charged $A + C + D$. Person 1 ends up with a zero surplus and person 2 ends up with a surplus of B—just what he would get if he chose to consume x_1^m.

In practice, the monopolist often encourages this self-selection not by adjusting the *quantity* of the good, as in this example, but rather by adjusting the *quality* of the good. The quantities in the model just examined can be re-interpreted as qualities, and everything works as before. In general, the monopolist will want to reduce the quality offered to the low end of its market so as not to cannibalize sales at the high end. Without the high-end consumers, the low-end consumers would be offered higher quality, but they would still end up with zero surplus. Without the low-end consumers, the high-end consumers would have zero surplus, so it is beneficial to the high-end consumers to have the low-end consumers present. This is because the monopolist has to cut the price to the high-end consumers to discourage them from choosing the product targeted to the low-end consumers.

EXAMPLE: Price Discrimination in Airfares

The airline industry has been very successful at price discrimination (although industry representatives prefer to use the term "yield management.") The model described above applies reasonably well to the problem faced by airlines: there are essentially two types of consumers, business travelers and individual travelers, who generally have quite different willingnesses to pay. Although there are several competing airlines in the U.S. market, it is quite common to see only one or two airlines serving specific city pairs. This gives the airlines considerable freedom in setting prices.

We have seen that the optimal pricing policy for a monopolist dealing with two groups of consumers is to sell to the high-willingness-to-pay market at a high price and offer a reduced-quality product to the market with the lower willingness to pay. The point of the reduced-quality product is to dissuade those with a high willingness to pay from purchasing the lower priced good.

The way the airlines implement this is to offer an "unrestricted fare" for business travel and a "restricted fare" for non-business travel. The

restricted fare often requires advanced purchase, a Saturday-night stayover, or other such impositions. The point of these impositions, of course, is to be able to discriminate between the high-demand business travelers and the more price sensitive individual travelers. By offering a "degraded" product—the restricted fares—the airlines can charge the customers who require flexible travel arrangements considerably more for their tickets.

Such arrangements may well be socially useful; without the ability to price discriminate, a firm may decide that it is optimal to sell *only* to the high-demand markets.

Another way that airlines price discriminate is with first-class and coach-class travel. First-class travelers pay substantially more for their tickets, but they receive an enhanced level of service: more space, better food, and more attention. Coach-class travelers, on the other hand, receive a lower level of service on all these dimensions. This sort of quality discrimination has been a feature of transportation services for hundreds of years. Witness, for example, this commentary on railroad pricing by Emile Dupuit, a nineteenth century French economist:

> It is not because of the few thousand francs which would have to be spent to put a roof over the third-class carriage or to upholster the third-class seats that some company or other has open carriages with wooden benches ... What the company is trying to do is prevent the passengers who can pay the second-class fare from traveling third class; it hits the poor, not because it wants to hurt them, but to frighten the rich ... And it is again for the same reason that the companies, having proved almost cruel to the third-class passengers and mean to the second-class ones, become lavish in dealing with first-class customers. Having refused the poor what is necessary, they give the rich what is superfluous.[2]

The next time you fly coach class, perhaps it will be of some solace to know that rail travel in nineteenth century France was even more uncomfortable!

EXAMPLE: Prescription Drug Prices

A month's supply of the antidepressant Zoloft sells for $29.74 in Austria, $32.91 in Luxembourg, $40.97 in Mexico, and $64.67 in the United States. Why the difference? Drug makers, like other firms, charge what the market

[2] Translation by R. B. Ekelund in "Price Discrimination and Product Differentiation in Economic Theory: An Early Analysis," *Quarterly Journal of Economics*, 84 (1970), 268–78.

will bear. Poorer countries can't pay as much as richer ones, so drug prices tend to be lower.

But that's not the whole story. Bargaining power also differs dramatically from country to country. Canada, which has a national health plan, often has lower drug prices than the United States, where there is no centralized provider of health care.

It has been proposed that drug companies be forced to charge a single price worldwide. Leaving aside the thorny question of enforcement, we might well ask what the consequences of such a policy would be. Would the world overall end up with lower prices or higher prices?

The answer depends on the relative size of the market. A drug for malaria would find most of its demand in poor countries. If forced to charge a single price, drug companies would likely sell such a drug at a low price. But a drug for diseases that afflicted those in wealthy countries would likely sell for a high price, making it too expensive for those in poorer areas.

Typically, moving from price discrimination to a single-price regime will raise some prices and lower others, making some people better off and some people worse off. In some cases, a product may not be supplied at all to some markets if a seller is forced to apply uniform pricing.

26.4 Third-Degree Price Discrimination

Recall that this means that the monopolist sells to different people at different prices, but every unit of the good sold to a given group is sold at the same price. Third-degree price discrimination is the most common form of price discrimination. Examples of this might be student discounts at the movies, or senior citizens' discounts at the drugstore. How does the monopolist determine the optimal prices to charge in each market?

Let us suppose that the monopolist is able to identify two groups of people and can sell an item to each group at a different price. We suppose that the consumers in each market are not able to resell the good. Let us use $p_1(y_1)$ and $p_2(y_2)$ to denote the inverse demand curves of groups 1 and 2, respectively, and let $c(y_1 + y_2)$ be the cost of producing output. Then the profit-maximization problem facing the monopolist is

$$\max_{y_1, y_2} \; p_1(y_1)y_1 + p_2(y_2)y_2 - c(y_1 + y_2).$$

The optimal solution must have

$$MR_1(y_1) = MC(y_1 + y_2)$$

$$MR_2(y_2) = MC(y_1 + y_2).$$

That is, the marginal cost of producing an extra unit of output must be equal to the marginal revenue in *each* market. If the marginal revenue in

market 1 exceeded marginal cost, it would pay to expand output in market 1, and similarly for market 2. Since marginal cost is the same in each market, this means of course that marginal revenue in each market must also be the same. Thus a good should bring the same increase in revenue whether it is sold in market 1 or in market 2.

We can use the standard elasticity formula for marginal revenue and write the profit-maximization conditions as

$$p_1(y_1)\left[1 - \frac{1}{|\epsilon_1(y_1)|}\right] = MC(y_1 + y_2)$$

$$p_2(y_2)\left[1 - \frac{1}{|\epsilon_2(y_2)|}\right] = MC(y_1 + y_2),$$

where $\epsilon_1(y_1)$ and $\epsilon_2(y_2)$ represent the elasticities of demand in the respective markets, evaluated at the profit-maximizing choices of output.

Now note the following. If $p_1 > p_2$, then we must have

$$1 - \frac{1}{|\epsilon_1(y_1)|} < 1 - \frac{1}{|\epsilon_2(y_2)|},$$

which in turn implies that

$$\frac{1}{|\epsilon_1(y_1)|} > \frac{1}{|\epsilon_2(y_2)|}.$$

This means that

$$|\epsilon_2(y_2)| > |\epsilon_1(y_1)|.$$

Thus the market with the higher price must have the lower elasticity of demand. Upon reflection, this is quite sensible. An elastic demand is a price-sensitive demand. A firm that price discriminates will therefore set a low price for the price-sensitive group and a high price for the group that is relatively price insensitive. In this way it maximizes its overall profits.

We suggested that senior citizens' discounts and student discounts were good examples of third-degree price discrimination. Now we can see why they have discounts. It is likely that students and senior citizens are more sensitive to price than the average consumer and thus have more elastic demands for the relevant region of prices. Therefore a profit-maximizing firm will price discriminate in their favor.

EXAMPLE: Linear Demand Curves

Let us consider a problem where the firm faces two markets with linear demand curves, $x_1 = a - bp_1$ and $x_2 = c - dp_2$. Suppose for simplicity that marginal costs are zero. If the firm is allowed to price discriminate,

it will produce where marginal revenue equals zero in each market—at a price and output combination that is halfway down each demand curve, with outputs $x_1^* = a/2$ and $x_2^* = c/2$ and prices $p_1^* = a/2b$ and $p_2^* = c/2d$.

Suppose that the firm were forced to sell in both markets at the same price. Then it would face a demand curve of $x = (a + c) - (b + d)p$ and would produce halfway down this demand curve, resulting in an output of $x^* = (a + c)/2$ and price of $p^* = (a + c)/2(b + d)$. Note that the total output is the same whether or not price discrimination is allowed. (This is a special feature of the linear demand curve and does not hold in general.)

However, there is an important exception to this statement. We have assumed that when the monopolist chooses the optimal single price it will sell a positive amount of output in each market. It may very well happen that at the profit-maximizing price, the monopolist will sell output to only one of the markets, as illustrated in Figure 26.4.

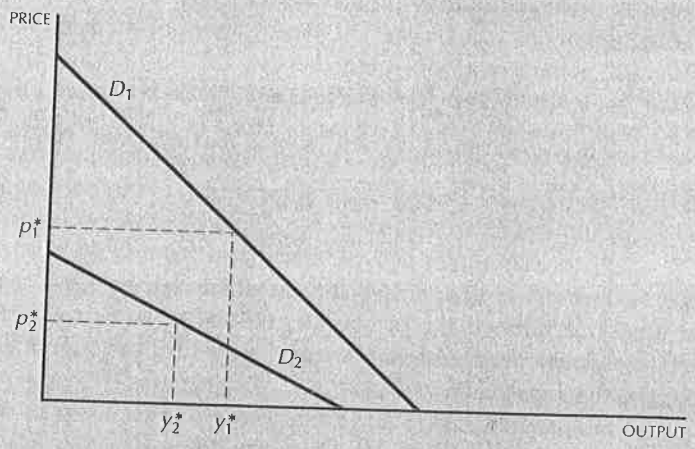

Price discrimination with linear demands. If the monopolist can charge only one price, it will charge p_1^*, and sell only to market 1. But if price discrimination is allowed, it will also sell at price p_2^* to market 2.

Figure 26.4

Here we have two linear demand curves; since marginal cost is assumed to be zero, the monopolist will want to operate at a point where the elasticity of demand is -1, which we know to be halfway down the market demand curve. Thus the price p_1^* is a profit-maximizing price—lowering the price any further would reduce revenues in market 1. If the demand in market 2 is very small, the monopolist may not want to lower its price any further in order to sell to this market: it will end up selling only to the larger market.

In this case, allowing price discrimination will unambiguously increase total output, since the monopolist will find it in its interest to sell to both markets if it can charge a different price in each one.

EXAMPLE: Calculating Optimal Price Discrimination

Suppose that a monopolist faces two markets with demand curves given by

$$D_1(p_1) = 100 - p_1$$
$$D_2(p_2) = 100 - 2p_2.$$

Assume that the monopolist's marginal cost is constant at \$20 a unit. If it can price discriminate, what price should it charge in each market in order to maximize profits? What if it can't price discriminate? Then what price should it charge?

To solve the price-discrimination problem, we first calculate the inverse demand functions:

$$p_1(y_1) = 100 - y_1$$
$$p_2(y_2) = 50 - y_2/2.$$

Marginal revenue equals marginal cost in each market yields the two equations:

$$100 - 2y_1 = 20$$
$$50 - y_2 = 20.$$

Solving we have $y_1^* = 40$ and $y_2^* = 30$. Substituting back into the inverse demand functions gives us the prices $p_1^* = 60$ and $p_2^* = 35$.

If the monopolist must charge the same price in each market, we first calculate the total demand:

$$D(p) = D_1(p_1) + D_2(p_2) = 200 - 3p.$$

The inverse demand curve is

$$p(y) = \frac{200}{3} - \frac{y}{3}.$$

Marginal revenue equals marginal cost gives us

$$\frac{200}{3} - \frac{2}{3}y = 20,$$

which can be solved to give $y^* = 70$ and $p^* = 43\frac{1}{3}$.

In accord with the discussion in the previous section, it is important to check that this price generates non-negative demands in each market. However, it is easily checked that this is the case.

EXAMPLE: Price Discrimination in Academic Journals

Most written scholarly communication takes place in academic journals. These journals are sold by subscription to libraries and to individual scholars. It is very common to see different subscription prices being charged to libraries and individuals. In general, we would expect that the demand by libraries would be much more inelastic than demand by individuals, and, just as economic analysis would predict, the prices for library subscriptions are typically much higher than the prices for individual subscriptions. Often library subscriptions are 2 to 3 times more expensive than subscriptions to individuals.

More recently, some publishers have begun to price discriminate by geography. During 1984, when the U.S. dollar was at an all-time high as compared to the English pound, many British publishers began to charge different prices to U.S. subscribers than to European subscribers. It would be expected that the U.S. demand would be more inelastic. Since the dollar price of British journals was rather low due to the exchange rate, a 10 percent increase in the U.S. price would result in a smaller percentage drop in demand than a similar increase in the British price. Thus, on grounds of profit maximization, it made sense for the British publishers to raise the prices of their journals to the group with the lower elasticity of demand—the U.S. subscribers. According to a 1984 study, North American libraries were charged an average of 67 percent more for their journals than U.K. libraries, and 34 percent more than anyone else in the world.[3]

Further evidence for price discrimination can be found by examining the pattern of price increases. According to a study by the University of Michigan Library, "... publishers have carefully considered their new pricing strategy. There seems to be a direct correlation ... between patterns of library usage and the magnitude of the pricing differential. The greater the use, the larger the differential."[4]

By 1986 the exchange rate had turned in favor of the pound, and the dollar prices of the British journals had increased significantly. Along with the price increase came some serious resistance to the higher prices. The concluding sentences of the Michigan report are illustrative: "One expects that a vendor with a monopoly on a product will charge according to demand. What the campus as a customer must determine is whether it will continue to pay up to 114% more than its British counterparts for the identical product."

[3] Hamaker, C. and Astle, D., "Recent Pricing Patterns in British Journal Publishing," *Library Acquisitions: Practice and Theory*, 8, 4 (Spring 1984), 225–32.

[4] The study was conducted by Robert Houbeck for the University of Michigan Library, and published in Vol. 2, No. 1 of the *University Library Update*, April 1986.

26.5 Bundling

Firms often choose to sell goods in **bundles**: packages of related goods offered for sale together. A noteworthy example is a bundle of software, sometimes known as a "software suite." Such a bundle might consist of several different software tools—a word processor, a spreadsheet, and a presentation tool—that are sold together in one set. Another example is a magazine: this consists of a bundle of articles that could, in principle, be sold separately. Similarly, magazines are often sold via subscription—which is just a way of bundling separate issues together.

Bundling can be due to cost savings: it is often less expensive to sell several articles stapled together than it is to sell each of them separately. Or it may be due to complementarities among the goods involved: software programs sold in bundles often work together more effectively than off-the-shelf programs.

But there can also be reasons involving consumer behavior. Let's consider a simple example. Suppose that there are two classes of consumers and two different software programs, a word processor and a spreadsheet. Type A consumers are willing to pay \$120 for the word processor and \$100 for the spreadsheet. Type B consumers have the opposite preferences: they are willing to pay \$120 for the spreadsheet and \$100 for the word processor. This information is summarized in Table 26.1.

Table 26.1 **Willingness to pay for software components.**

Type of consumer	Word processor	Spreadsheet
Type A consumers	120	100
Type B consumers	100	120

Suppose that you are selling these products. For simplicity, let us assume that the marginal cost is negligible so that you only want to maximize revenue. Furthermore, make the conservative assumption that the willingess to pay for the bundle consisting of the word processor and the spreadsheet is just the sum of the willingesses to pay for each component.

Now consider the profits from two different marketing policies. First, suppose that you sell each item separately. The revenue maximizing policy is to set a price of \$100 for each piece of software. If you do this, you will sell two copies of the word processor and two copies of the spreadsheet, and receive a total revenue of \$400.

But what if you bundle the items together? In this case, you could sell *each* bundle for $220, and receive a net revenue of $440. The bundling strategy is clearly more attractive!

What is going on in this example? Recall that when you sell an item to several different people, the price is determined by the purchaser who has the *lowest* willingess to pay. The more diverse the valuations of the individuals, the lower the price you have to charge to sell a given number of items. In this case bundling the word processor and the spreadsheet reduces the dispersion of willingess to pay—allowing the monopolist to set a higher price for the bundle of goods.

EXAMPLE: Software Suites

Microsoft, Lotus, and other software manufacturers have taken to bundling much of their applications software. For example, in 1993 Microsoft offered a spreadsheet, word processor, presentation tool, and database as the "Microsoft Office" package at a suggested retail price of $750. (The discounted "street price" was about $450.) If bought separately, the individual software applications would total $1,565! Lotus offered its "Smart Suite" at essentially the same price; its separate components sold for a total of $1,730.

According to an article by Steve Lohr in the October 15, 1993, *New York Times*, 50 percent of Microsoft's applications software was sold in bundles, and generated revenue of over $1 billion a year.

These software suites fit the bundling model well. Tastes for software are often very heterogeneous. Some people use a word processor every day and use a spreadsheet only occasionally. Other people have the reverse pattern of software use. If you wish to sell a spreadsheet to a large number of users, you have to sell it at a price that will be attractive to an occasional user. Similarly with the word processor: it is the willingness to pay of the *marginal* user that sets the market price. By bundling the two products together, the dispersion of willingnesses to pay is reduced and total profits can increase.

This is not to say that bundling is the whole story in software suites; other phenomena are also at work. The individual components of the suites are guaranteed to work well together; they are complementary goods in this respect. Furthermore, the success of a piece of software tends to depend strongly on how many people use it, and bundling software helps to build market share. We will investigate this phenomenon of **network externalities** in Chapter 36.

26.6 Two-Part Tariffs

Consider the pricing problem facing the owners of an amusement park. They can set one price for tickets to get into the park and another price for

the rides. How should they set these two prices if they want to maximize profits? Note that the demand for access and the demand for rides are interrelated: the price that people are willing to pay to get into the park will depend on the price that they have to pay for the rides. This kind of two-part pricing scheme is known as a **two-part tariff**.[5]

Other applications of two-part tariffs abound: Polaroid sells its camera for one price and its film for another. People who are deciding whether or not to purchase the camera presumably consider the price of the film. A company that makes razor blades sells the razor for one price and the blades for another—again the price they set for the blades influences the demand for razors and vice versa.

Let us consider how to solve this pricing problem in the context of the original example: the so-called Disneyland Dilemma. As usual we will make some simplifying assumptions. First, we assume that there is only one kind of ride in Disneyland. Second, we assume that people only desire to go to Disneyland for the rides. Finally, we assume that everyone has the same tastes for rides.

In Figure 26.5 we have depicted the demand curve and the (constant) marginal cost curve for rides. As usual the demand curve slopes down—if Disney sets a high price for each ride, fewer rides will be taken. Suppose that they set a price of p^*, as in Figure 26.5, that leads to a demand for x^* rides. How much will they be able to charge for admission to the park, given that the rides cost p^*?

The total willingness to pay for x^* rides is measured by the consumers' surplus. Hence the most that the owners of the park can charge for admission is the area labeled "consumer's surplus" in Figure 26.5. The total profits to the monopolist is this area plus the profit on the rides, $(p^* - MC)x^*$.

It is not hard to see that total profits are maximized when price equals marginal cost: we've seen before that this price gives the largest possible consumer plus producer surplus. Since the monopolist gets to charge people their consumers' surplus, setting price equal to marginal cost and the entry fee to the resulting consumer's surplus is the profit-maximizing policy.

Indeed, this is the policy that Disneyland, and most other amusement parks follow. There is one price for admission, but then the attractions inside are free. It appears that the marginal cost of the rides is less than the transactions cost of collecting a separate payment for them.

26.7 Monopolistic Competition

We have described a monopolistic industry as being one in which there is a single large producer. But we've been somewhat vague about exactly what

[5] See the classic article by Walter Oi, "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly," *Quarterly Journal of Economics*, 85 (1971), 77–96.

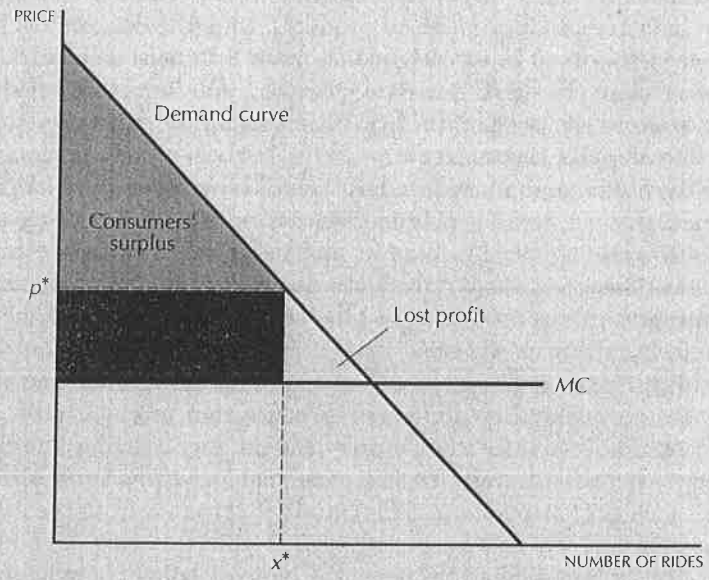

Disneyland Dilemma. If the owners of the park set a price of p^*, then x^* rides will be demanded. The consumers' surplus measures the price that they can charge for admission to the park. The total profits of the firm are maximized when the owners set price equal to marginal cost.

Figure 26.5

comprises an industry. One definition of an industry is that it consists of all firms that produce a given product. But then what do we mean by product? After all, there is only one firm that produces Coca-Cola—does that mean that this firm is a monopolist?

Clearly the answer is no. The Coca-Cola firm still has to compete with other producers of soft drinks. We should really think of an industry as being the set of firms that produce products that are viewed as close substitutes by consumers. Each firm in the industry can produce a unique product—a unique brand name, say—but consumers view each of the brands as being substitutes to some degree.

Even though a firm may have a legal monopoly on its trademarks, and brand names, so that other firms can't produce *exactly* the same product, it is usually possible for other firms to produce *similar* products. From the viewpoint of a given firm, the production decisions of its competitors will be a very important consideration in deciding exactly how much it will produce and what price it can charge.

Thus the demand curve facing a firm will usually depend on the output decisions and the prices charged by other firms that produce similar products. The slope of the demand curve facing the firm will depend on

how similar the other firms' products are. If a large number of the firms in the industry produce *identical* products, then the demand curve facing any one of them will be essentially flat. Each firm must sell its product for whatever price the other firms are charging. Any firm that tried to raise its price above the prices of the other firms selling identical products would soon lose all of its customers.

On the other hand, if one firm has the exclusive rights to sell a particular product, then it may be able to raise its price without losing all of its customers. Some, but not all, of its customers may switch to competitors' products. Just how many customers switch depends on how similar the customers think the products are—that is, on the elasticity of the demand curve facing the firm.

If a firm is making a profit selling a product in an industry, and other firms are not allowed to perfectly reproduce that product, they still may find it profitable to enter that industry and produce a similar but distinctive product. Economists refer to this phenomenon as **product differentiation**—each firm attempts to differentiate its product from the other firms in the industry. The more successful it is at differentiating its product from other firms selling similar products, the more monopoly power it has—that is, the less elastic is the demand curve for the product. For example, consider the soft drink industry. In this industry there are a number of firms producing similar, but not identical products. Each product has its following of consumers, and so has some degree of market power.

An industry structure such as that described above shares elements of both competition and monopoly; it is therefore referred to as **monopolistic competition**. The industry structure is monopolistic in that each firm faces a downward-sloping demand curve for its product. It therefore has some market power in the sense that it can set its own price, rather than passively accept the market price as does a competitive firm. On the other hand the firms must compete for customers in terms of both price and the kinds of products they sell. Furthermore, there are no restrictions against new firms entering into a monopolistically competitive industry. In these aspects the industry is like a competitive industry.

Monopolistic competition is probably the most prevalent form of industry structure. Unfortunately, it is also the most difficult form to analyze. The extreme cases of pure monopoly and pure competition are much simpler and can often be used as first approximations to more elaborate models of monopolistic competition. In a detailed model of a monopolistically competitive industry, much depends on the specific details of the products and technology, as well as on the nature of the strategic choices available to firms. It is unreasonable to model a monopolistically competitive industry in the abstract, as we have done with the simpler cases of pure competition and pure monopoly. Rather, the institutional details of the particular industry under consideration must be examined. We will describe some methods that economists use to analyze strategic choice in the next two

chapters, but a detailed study of monopolistic competition will have to wait for more advanced courses.

We can, however, describe an interesting feature of the free entry aspect of monopolistic competition. As more and more firms enter the industry for a particular kind of product, how would we expect the demand curve of an incumbent firm to change? First, we would expect the demand curve to shift inward since we would expect that at each price, it would sell fewer units of output as more firms enter the industry. Second, we would expect that the demand curve facing a given firm would become more elastic as more firms produced more and more similar products. Thus entry into an industry by new firms with similar products will tend to shift the demand curves facing existing firms to the left and make them flatter.

If firms continue to enter the industry as long as they expect to make a profit, equilibrium must satisfy the following three conditions:

1. Each firm is selling at a price and output combination on its demand curve.

2. Each firm is maximizing its profits, given the demand curve facing it.

3. Entry has forced the profits of each firm down to zero.

These facts imply a very particular geometrical relationship between the demand curve and the average cost curve: the demand curve and the average cost curve must be tangent to each other.

The argument is illustrated in Figure 26.6. Fact 1 says that the output and price combination must be somewhere on the demand curve, and fact 3 says that the output and price combination must also be on the average cost curve. Thus the operating position of the firm must be at a point that lies on both curves. Could the demand curve cross the average cost curve? No, because then there would be some point on the demand curve above the average cost curve—but this would be a point yielding *positive* profits.[6] And by fact 2, the zero profit point is a profit maximum.

Another way to see this is to examine what would happen if the firm depicted in Figure 26.6 charged any price other than the break-even price. At any other price, higher or lower, the firm would lose money, while at the break-even price, the firm makes zero profits. Thus the break-even price is the profit-maximizing price.

There are two worthwhile observations about the monopolistically competitive equilibrium. First, although profits are zero, the situation is still Pareto inefficient. Profits have nothing to do with the efficiency question: when price is greater than marginal cost, there is an efficiency argument for expanding output.

[6] If $p > c(y)/y$, then simple algebra shows that $py - c(y) > 0$.

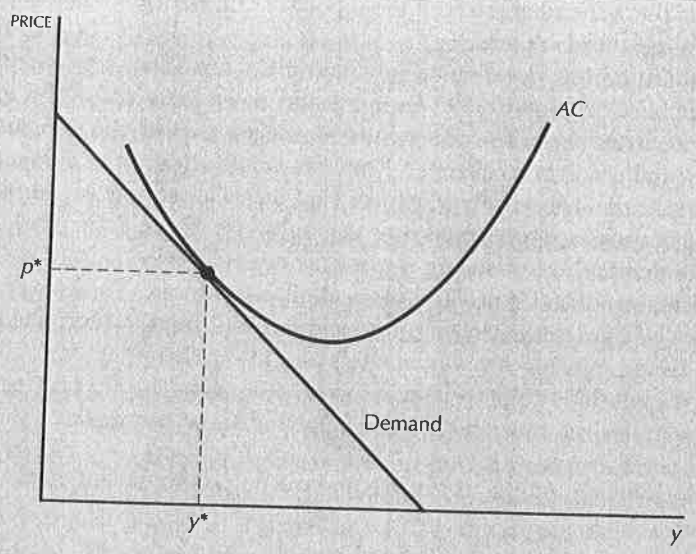

Figure 26.6 **Monopolistic competition.** In a monopolistically competitive equilibrium with zero profits, the demand curve and the average cost curve must be tangent.

Second, it is clear that firms will typically be operating to the left of the level of output where average cost is minimized. This has sometimes been interpreted as saying that in monopolistic competition there is "excess capacity." If there were fewer firms, each could operate at a more efficient scale of operation, which would be better for consumers. However, if there were fewer firms there would also be less product variety, and this would tend to make consumers worse off. Which of these effects dominates is a difficult question to answer.

26.8 A Location Model of Product Differentiation

In Atlantic City there is a boardwalk that stretches along the beach. Some ice cream vendors with pushcarts want to sell ice cream on the boardwalk. If one vendor is going to be given the concession to sell ice cream on the boardwalk, where should he locate?[7]

Suppose that consumers are distributed evenly along the beach. From a social point of view, it makes sense to locate the ice cream vendor so that

[7] The discussion here is based on the classic model of Harold Hotelling, "Stability in Competition," *Economic Journal*, March 1929.

the total distance walked by all the consumers is minimized. It is not hard to see that this optimal location is halfway along the boardwalk.

Now suppose that two ice cream vendors are allowed. Suppose that we fix the price that they are able to charge for their ice cream and just ask where they should locate in order to minimize the total distance walked. If each consumer walks to the ice cream vendor nearest him, we should put one vendor a quarter of the way along the boardwalk and one vendor three-quarters of the way along the boardwalk. The consumer halfway along the boardwalk will be indifferent between the two ice cream vendors; each has a market share of one-half of the consumers. (See Figure 26.7A.)

But do the ice cream vendors have an incentive to stay in these locations? Put yourself in the position of vendor L. If you move a little bit to the right, you will steal some of the other vendor's customers and you won't lose any of your own. By moving to the right, you will still be the closest vendor to all the customers to your left and you will still be closer to the customers on your right. You will therefore increase your market share and your profits.

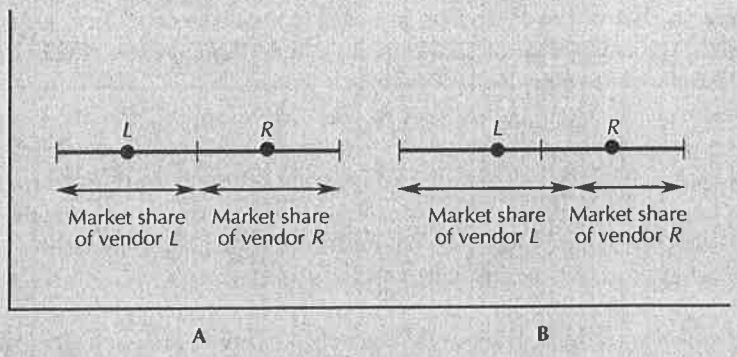

Competition in location. Panel A shows the socially optimal location pattern; L locates one-quarter of the way along the line and R locates three-quarters of the way along. But each vendor will find it in its private interest to move toward the middle. The only equilibrium location is for both vendors to be in the middle, as shown in Panel B.

Figure 26.7

But vendor R can reason the same way—by moving to the left, he will steal some of the other vendor's customers and not lose any of his own! This shows that the socially optimal location patterns are not an equi-

librium. The only equilibrium is for both vendors to sell in the middle of the boardwalk, as shown in Figure 26.7B. In this case, competition for customers has resulted in an *inefficient* location pattern.

The boardwalk model can serve as a metaphor for other sorts of product-differentiation problems. Instead of the boardwalk, think of the choice of music varieties by two radio stations. At one extreme we have classical music and at the other we have heavy metal rock. Each listener chooses the station that appeals more to his tastes. If the classical station plays music that is a bit more toward the middle of the taste spectrum, it won't lose the classical clients, but it will gain a few of the middlebrow listeners. If the rock station moves a bit toward the middle, it won't lose any of its rock lovers but will get a few of the middlebrow listeners. In equilibrium, both stations play the same sort of music and the people with more extreme tastes are unhappy with both of them!

26.9 Product Differentiation

The boardwalk model suggest that monopolistic competition will result in too little product differentiation: each firm will want to make its product similar to that of the other firm in order to steal the other firm's customers. Indeed, we can think of markets in which there is too much imitation relative to what seems to be optimal.

However, it doesn't always work this way. Suppose that the boardwalk is *very* long. Then each ice cream vendor would be perfectly happy sitting near each end of the boardwalk. If their market areas don't overlap, nothing is to be gained from moving closer to the middle of the boardwalk. In this case, neither monopolist has an incentive to imitate the other, and the products are about as different as they can get.

It is possible to produce models of monopolistic competition where there is *excessive* product differentiation. In such models, each firm attempts to make consumers think that its product is different from the products of its competitors so as to create some degree of market power. If the firms succeed in convincing the consumers that their product has no close substitutes, they will be able to charge a higher price for it than they would otherwise be able to do.

This leads each producer to invest heavily in creating a distinctive brand identity. Laundry soap, for example, is a pretty standardized commodity. Yet manufacturers invest huge amounts of money in advertisements that claim cleaner clothes, better smell, a better marriage, and and a generally happier life if you choose their brand rather than a competitor's. This "product positioning" is much like the ice cream vendors locating far away from each other in order to avoid head-to-head competition.

There are critics who have argued that such excessive investment in product positioning is wasteful. Perhaps this is true in some cases, but then

again, "excessive variety" may simpley be a consequence of encouraging firms to provide consumers with a variety of products from which to choose.

26.10 More Vendors

We have shown that if there are two vendors whose market areas overlap, and each seller sells the same price, they will both end up located at the "middle" of the boardwalk. What happens if there are more than two vendors who compete in their location?

The next easiest case is that of three vendors. This case gives rise to a rather peculiar outcome: there may be *no* equilibrium location pattern! To see this, look at Figure 26.8. If there are three vendors located on the boardwalk, there must be one located between the other two. As before, it pays each of the "outside" vendors to move towards the middle vendor since they can steal some of its customers without losing any of their own. But if they get *too* close to the other vendor, it pays it to jump immediately to the right of its right-hand competitor or immediately to the left of its left-hand competitor to steal *its* market. No matter what the location pattern, it pays someone to move!

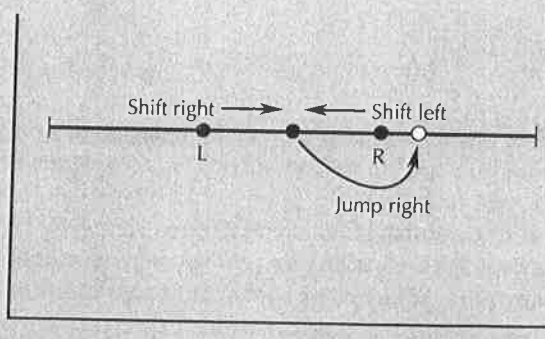

No equilibrium. There is no pure strategy equilibrium in the Hotelling model with 3 firms since for any configuration, at least one firm wants to change location.

Figure 26.8

Luckily, this "perverse" result only holds in the case of three competitors. If there are four or more competitors, an equilibrium location pattern will generally emerge.

Summary

1. There will typically be an incentive for a monopolist to engage in price discrimination of some sort.

2. Perfect price discrimination involves charging each customer a different take-it-or-leave-it price. This will result in an efficient level of output.

3. If a firm can charge different prices in two different markets, it will tend to charge the lower price in the market with the more elastic demand.

4. If a firm can set a two-part tariff, and consumers are identical, then it will generally want to set price equal to marginal cost and make all of its profits from the entry fee.

5. The industry structure known as monopolistic competition refers to a situation in which there is product differentiation, so each firm has some degree of monopoly power, but there is also free entry so that profits are driven to zero.

6. Monopolistic competition can result in too much or too little product differentiation in general.

REVIEW QUESTIONS

1. Will a monopoly ever provide a Pareto efficient level of output on its own?

2. Suppose that a monopolist sells to two groups that have constant elasticity demand curves, with elasticity ϵ_1 and ϵ_2. The marginal cost of production is constant at c. What price is charged to each group?

3. Suppose that the amusement park owner can practice perfect first-degree price discrimination by charging a different price for each ride. Assume that all rides have zero marginal cost and all consumers have the same tastes. Will the monopolist do better charging for rides and setting a zero price for admission, or better by charging for admission and setting a zero price for rides?

4. Disneyland also offers a discount on admissions to residents of Southern California. (You show them your zip code at the gate.) What kind of price discrimination is this? What does this imply about the elasticity of demand for Disney attractions by Southern Californians?