Part II
Monopoly Power in Theory and Practice

In Part II, we consider the pure monopoly problem in much more detail than the simple, textbook case presented in Chapter 2. In particular, we consider a range of price and non-price tactics the might be used by firms that operate as a pure monopolist. There are two reasons to study such tactics. First, there are a number of special cases in which firms might actually have an effective monopoly. In many regions, for example, there is just one ski lift operator within a radius of fifty miles and only one amusement park serving an even greater area. Second, and more importantly, the tactics we discuss such as quantity discounts and bundling are also available to imperfect competitors who are not monopolists. However, because it is much easier to understand the role of such measures in this more competitive environment *after* seeing them used by a single firm with no strategically linked rival, we introduce these concepts here.

We consider three basic techniques that a firm with a downward-sloping demand curve may use to improve its product over that achieved by the standard, uniform-pricing monopolist. In Chapters 5 and 6, we explore various price discrimination schemes. This includes both linear (market segmentation) and nonlinear (two-part tariff and quantity discount) pricing strategies.

In Chapter 7 we introduce the idea of product design as a means of enhancing profit. Beyond demonstrating that design issues are important, this permits us to introduce the concept of horizontal differentiation and its spatial representation as first formalized by Hotelling (1929). It also permits us to introduce the alternative concept of vertical product differentiation along the lines initially presented by Mussa and Rosen (1978). Although these models are more technically demanding, they each have an accessible underlying intuition. More importantly, both horizontal and vertical differentiation are crucial concepts that we will employ repeatedly in the settings strategic interaction among a small number of firms that form the framework for the later chapters, 9 to 25. Hence, giving those concepts a formal representation that can be applied in a wide variety of settings is an important building block for later work. Chapter 7 also includes an empirical application based on the study by Stavins (1989) of price-discrimination techniques in airline ticketing. Although we have introduced discriminatory practices in a monopoly setting, everyday experience offers considerable evidence that such tactics are not limited to the rare case of a pure monopolist. This study helps to clarify the role of price discrimination in more competitive settings and gives a more realistic understanding of such tactics.

Microsoft sells much of it software as a bundle, e.g., the package of programs including Word, Excel, and PowerPoint available in its Office package. Similarly, the owner of a Hewlett–Packard inkjet printer is required or tied in to also use of a Hewlett–Packard inkjet printer cartridge. Such bundling and tying practices are the focus of Chapter 8. We demonstrate how such practices may again be used to increase the firm's profit. Here again, however, everyday experience offers abundant evidence that these practices are not limited to pure or even near monopoly cases. Therefore, we again examine the role of such practices in more competitive environments.

5

Price Discrimination and Monopoly: Linear Pricing

A central part of the growing debate over U.S. healthcare policy has been the high price of prescription drugs. It was this concern that led to the expansion of Medicare to include coverage for prescription drugs starting in 2006. It is this same issue that has led many Americans to buy their prescription drugs in other countries, most notably Canada. Yet for that to happen legally requires additional legislation. The existing statutes—enacted during the Clinton Administration—only permit such imports if the Secretary of Health and Human Services certifies that the drug "pose no additional risk" to consumers, and no such certifications have been issued. As of this writing, legislation to remove the certification requirement has so far failed to win passage in both houses of the U.S. Congress.

Whatever the risk of Canadian prescriptions, it is indisputable that there are significant differences in prescription drug prices between the United States and Canada. Graham and Robson (2000) collected detailed 1999 price information for 45 brand name prescription drugs, collectively covering approximately 25 percent of the total prescriptions written in the United States. From this sample they calculated that Canadian retail prices were far less than American ones, with the median discount approximately 46 percent. For one drug in their sample, this discount was 95 percent. In a related study Graham and Tabler (2001) analyzed the retail prices charged in 2001 by a randomly selected set of pharmacies for three patented drugs in three Canadian and three neighboring American areas.

Table 5.1 provides summary information on brand name drug prices that further confirms that these are generally lower in Canada than in the United States, with normal discounts running in the range of 50 percent. At the same time, however, there is considerable evidence that generic drug prices are much lower in the U.S. than in Canada. Studies by both the U.S. Food and Drug Administration (FDA) and the Canadian Patented Medicine Prices Review Board (PMPRB) find that generic drugs sell in the U.S. at prices that are on average 35 percent less than the comparable charge in Canada.[1]

What explains this pattern of drug prices? This is an especially pressing question in the case of branded drugs because these are generally made by the same firms regardless of whether they are sold in the U.S. or Canada. What is it that makes it profitable for these firms to

[1] Patented Medicine Prices Review Board, "Canadian and Foreign Price Trends," http://www.pmprb-cepmb.gc.ca (July, 2006) and Food and Drug Administration, "Generic Drug Prices in the US Are Lower Than Prices in Canada," White Paper (November, 2003).

Table 5.1 Comparison of prescription drug prices (U.S. dollars)

	Celebrex® 200 mg		*Lipitor® 40 mg*		*Paxil® 20 mg*	
	Mean	*Standard deviation*	*Mean*	*Standard deviation*	*Mean*	*Standard deviation*
Washington	86.26	5.66	110.01	8.97	82.47	3.86
British Columbia	33.17	2.37	52.83	3.50	40.75	2.54
North Dakota and Minnesota	78.08	5.70	107.75	7.03	78.63	6.08
Manitoba	32.36	1.60	52.43	1.52	39.80	2.00
New York	88.57	7.59	117.69	5.44	85.06	4.39
Ontario	34.82	1.96	55.52	2.09	42.62	2.03

Source: Graham and Tabler (2001) *Prescription Drug Prices in Canada and the United States: Part 3 Retail Price Distribution*, Public Policy Sources no. 50, Fraser Institute

set higher prices in the U.S. market? And what makes it possible for this differential to be sustained?

These questions go to the heart of the analysis in this and the following two chapters. Charging different prices to different consumers for the same good is referred to as price discrimination. We want to know what makes this profitable and when it is feasible. Essentially we are asking what tactics firms can use to implement price discrimination in a way that increases profits relative to charging the same price to everyone. Remember though that any increased profit must come either from a reduction in consumer surplus, improved market efficiency, or some combination of the two. From a policy perspective, it matters a great deal as to which of these is the case. Hence, we will also want to explore the welfare implications of price discrimination. Finally, it is worthwhile noting that discriminatory prices can also affect market competition. This occurs when the buyers are not final consumers but instead, retailers such as drug stores. If large drug store chains are charged different wholesale prices than are small, independent pharmacies, then retail competition between these two groups will not be conducted on a level playing field. In these next three chapters, we address each of these issues in turn.

5.1 FEASIBILITY OF PRICE DISCRIMINATION

A firm with market power faces a downward sloping demand curve, so if the firm charges the same price to each consumer—the standard case of non-discriminatory pricing—the revenue it gets from selling an additional unit of output is less than the price charged. In order to sell the additional unit the firm must lower its price not only to the consumer who buys the additional unit, but to all its other consumers as well. Having to lower price to all its customers in order to gain an additional consumer limits the monopolist's incentive to serve more consumers. As a result the textbook monopoly undersupplies its product relative to the efficient outcome.

However, non-discriminatory pricing is not just a source of potential inefficiency. It is also a constraint on the firm's ability to extract consumer surplus, particularly from those consumers willing to pay a lot for its product. If we allow the monopolist to price discriminate

we shall see that this is a powerful technique that permits the firm to earn considerably more profit. We will also find that in some cases price discrimination may induce the monopolist to sell more output and so come closer to the competitive market outcome. Price discrimination can sometimes make a monopolized market more efficient.

While a monopolist can increase profit through price discrimination, it is important to realize that price discrimination is not always easily accomplished. There is a reason why the standard textbook case assumes that each customer pays the same price. To discriminate successfully the monopolist must overcome two main obstacles. The first of these is identifying who is who on the demand curve. The second is the problem of arbitrage.

In considering the identification problem it is useful to recall a common assumption in the textbook monopoly model. This assumption is that the monopolist has somehow learned what is the quantity demanded at each price—otherwise it would not know its marginal revenue curve and, hence, would not be able to determine the profit-maximizing output. Let's examine more carefully just what this assumption means in practice.

For some products such as bicycles, TVs, DVD players, or haircuts a single consumer will purchase at most one unit of the good over a given period of time. The firm's demand curve is then an explicit ordering of consumers by their reservation prices—the top price each is willing to pay. For these goods, knowledge of the demand curve means that the monopoly firm knows that the top part of the demand curve is made up of those consumers willing to pay a relatively large amount for the one unit they will purchase whereas the bottom part of the demand curve is made up by those willing to pay only a little. For other products, however, such as movies, CDs, refreshments, and tennis lessons, what is happening on the demand curve is slightly more complex. This is because each individual consumer can be induced to purchase more than one unit of the good if the price is sufficiently low. Hence, for these goods the demand curve reflects not only differences in the willingness to pay across consumers but also differences in the willingness to pay as any one consumer buys more of the product. When the monopolist practices uniform pricing, these distinctions are not relevant. In that case the assumption that the firm knows its demand curve means only that it knows how willingness to pay for the good *in the overall market* varies with the quantity of the good sold.

To be able to practice price discrimination the monopolist must learn or know more information about consumers than is assumed in the standard model. The monopolist must know how the market demand curve has been constructed from the individual consumer demand curves. In other words, the monopolist must know how different kinds of consumers differ in their demands for its good. This is easier for some sellers than for others. For example, tax accountants effectively sell one unit of their services to each client in any given year. Further, they know exactly how much their clients earn and, more importantly, how much they save their clients by way of reduced tax liabilities. They can certainly use this information to identify the customers' willingness-to-pay. Similarly, a car dealer typically sells one car to a customer. The dealer may be able to identify those buyers with the greatest or least willingness to pay by asking potential buyers where they live or work or shop. The same is often true for realtors, dentists, and lawyers.

Sellers of retail merchandise, however, face a more anonymous market. Various schemes such as varying the price depending on time of purchase, "early-bird" specials or Saturday morning sales, or offering coupons that take time to collect, can help retailers identify "who's who" on their demand curve. Nevertheless, the identification problem is still difficult to overcome. Moreover, even if weekend sales or coupon schemes do successfully identify the firm's different consumers, such schemes may be too costly to implement.

Even when a monopolist can solve the identification problem, there is still a second obstacle to price discrimination, arbitrage. To discriminate successfully, the monopolist must be able to prevent those consumers who are offered a low price from reselling their purchases to other consumers to whom the monopolist wants to offer a high price. Again, this will be

Reality Checkpoint

Old Wines in a New Format

On May 16, 2005, the five-year struggle of two entrepreneurial women, Eleanor Heald of Michigan and Juanita Swedenburg of Virginia, finally came to a happy end. Eleanor Heald and her husband, Ray, are wine aficionados living in Troy, Michigan. They love to get together with friends at wine tasting parties, and they contributed numerous articles about wine and specific vintners to magazines such as *Practical Winery and Vineyard*. Juanita Swedenburg ran a small winery in Middleburg, Virginia. Until her death in early 2004, she too had the active support of her husband, Wayne.

Since both Ms. Heald is an avid wine consumer and Ms. Swedenburg was a producer, one might think that what brought them together was Ms. Heald's purchase of some of Ms. Swedenburg's output. Actually, the truth is almost the opposite. The reason is that prior to the Supreme Court ruling of that mid-May day, Ms. Heald was forbidden by Michigan law to order wine directly from any winery located outside of Michigan. Like New York, Massachusetts, and about 12 other states Michigan had laws that prohibited such shipments.

In general, wine in the United States is marketed via a three-tier system. There are the wineries themselves that produce the wine. Then there are the wholesalers who buy from the wineries. Finally, there are the retailers who purchase supplies from the distributors and sell to final consumers. The winery level has seen fantastic growth since the 1970s as Americans have increasingly turned to wine as a beverage of choice. There are now over 2,000 individual commercial wineries in the U.S. Many of these are, however, very small firms shipping only a few thousand cases of very individualized wine per year. As a result of this small volume, it is cost prohibitive for these firms to sell indirectly via through the wholesale and retail tiers—especially as wholesale market structure has become an increasingly concentrated one, which, in many regions, is dominated by just one or two firms. For these many small wineries the Internet has been a lifesaver allowing them to ship directly to consumers—at least in some states. Others prohibit all such direct buying. States like Michigan and New York permitted such shipments but only if the winery was located in the state. It was this discrimination against out-of-state wineries that upset both Ms. Heald and Ms. Swedenburg. Ms. Heald could not get wine from outside Michigan and Ms. Swedenburg could not ship her wine to some interested buyers in New York. The two filed separate lawsuits claiming that these two states were violating the constitution. The cases wound their way through the court system and were finally combined into one case by the Supreme Court. Of course, Michigan and New York—backed by their own wineries, retailers, and wholesalers—and the states with similar laws fought back. Their principal defense was that the laws were meant to prevent minors from getting easy access to alcohol. Yet since the laws did not prevent direct shipments from wineries within the state, the Supreme Court found this argument unpersuasive. On that fateful day, it struck down the Michigan and New York laws and, by implication, all similar ones in other states.

Source: K. McLaughlin "Will Buying Wine Get Easier?" *Wall Street Journal*, May 17, 2005, p. A1; and *Granholm v. Heald*, 544 U.S. 460 (2005).

more easily accomplished for some goods and services than for others. Medical, legal, and educational services are not easily resold. One consumer can't sell her appendectomy to another! Similarly, a senior citizen cannot easily resell a discounted movie theater ticket to a teenager. For other markets, particularly consumer durables such as bicycles and automobiles, resale—or sale across different markets—is difficult to prevent. This is an important part of the drug pricing story noted at the start of this chapter. Pharmaceutical companies can only price discriminate successfully if they can keep the American and Canadian markets separate, in other words, only if they can prevent arbitrage.

To sum up, we expect firms with monopoly power to try to price discriminate. In turn, this implies that we should expect that these firms will want to identify the different types of consumers who buy their goods and to prevent resale or consumer arbitrage among them. The ability to do this and the best strategy for achieving price discrimination will vary from firm to firm and from market to market. We now turn to the practice of price discrimination and investigate some of the more popularly practiced techniques. The tradition in economics has been to classify these techniques into three broad classes: first-degree, second-degree, and third-degree price discrimination.[2] More recently, these types of pricing schemes have been referred to respectively as personalized pricing, menu pricing and group pricing.[3] In this chapter we focus on third-degree price discrimination or group pricing.

5.2 THIRD-DEGREE PRICE DISCRIMINATION OR GROUP PRICING

Third-degree price discrimination or group pricing is defined by three key features. First, there is some easily observable characteristic such as age, income, geographic location, or education status by which the monopolist can group consumers in terms of their willingness to pay for its product. Second, the monopolist can prevent arbitrage across the different groups. In the prescription drug case with which we started this chapter the issue would be to prevent the re-importing of prescription drugs initially exported from the United States to Canada. Finally, third-degree price discrimination requires that the monopolist quotes the same price per unit to all consumers within a particular group and consumers in each group then decide how much to purchase at their quoted price.

Group pricing reflects price discrimination because for the same good the price quoted to one group of consumers is not the same as the price quoted to another group. This type of pricing policy is the one most commonly found in economics textbooks and is referred to in the industrial organization literature as *linear pricing*—hence the title of this chapter. Consumers within a group are free to buy as much as they like at the quoted price, so that the average price per unit paid by each consumer is the same as the marginal price for the last unit bought.

The world is full of examples of third-degree price discrimination. Senior discounts and "kids are free" programs are both examples. An interesting case that is particularly familiar

[2] Price discrimination is a fascinating topic and its interest to economists goes well beyond the field of industrial organization. The distinction between first-, second-, and third-degree discrimination follows the work of Pigou (1920). A more modern treatment appears in Phlips (1983). Varian (1989) offers an excellent summary.

[3] These terms were first coined by Shapiro and Varian (1999).

Table 5.2 Schedule of annual membership fees for the American Economic Association (U.S. dollars)

Regular members with annual incomes of $47,000 or less	64
Regular members with annual incomes above $47,000 but no more than $62,000	77
Regular members with annual incomes above $62,000	90
Junior members (available to registered students—student status must be certified)	32
Family member (persons living at the same address as a regular member, additional membership without a subscription to AEA publications)	13

to economists is the fee schedule for membership in the American Economic Association, the major professional organization for economists in the United States. Payment of the fee entitles a member to receive professional announcements, newsletters, and three very important professional journals, the *American Economic Review*, the *Journal of Economic Perspectives*, and the *Journal of Economic Literature*, each of which is published quarterly.

The 2006 fee schedule is shown below in Table 5.2. As can readily be seen, the aim is to price discriminate on the basis of income. A particularly interesting feature of this scheme is that the Association makes no attempt to check the veracity of the income declared by a prospective member. What they appear to rely upon is that economists will be either honest or even boastful in reporting their income. In addition, the Association must also hope to avoid the arbitrage problem whereby junior faculty members who pay a low subscription fee resell to senior faculty members who pay a high one. Here again, we can only report on casual observation, and on this basis, such reselling actually appears to be rare so that the arbitrage problem seems to be effectively non-existent.

The practice of the American Economic Association is not unique. Many academic journals charge a different price to institutions such as university libraries than to individuals. A subscription rate to the *Journal of Economics and Management Strategy*, for example, is $45 for an individual but $115 for an institution.

Airlines are particularly adept at applying third-degree price discrimination. It has sometimes been suggested that the number of different fares charged to economy class passengers on a particular flight is approximately equal to the number of passengers! A common feature of this type of price discrimination is that it is implemented by restrictions on the characteristics of the ticket. These include constraints upon the time in advance by which the flight must be booked, whether flights can be changed, the number of days between departure and return, whether the trip involves staying over a Saturday night and so on. We return to the airline case later in this chapter.

Other examples of third-degree price discrimination are restaurant "early bird specials" and supermarket discounts to shoppers who clip coupons. Similarly, department stores that lower their apparel prices at the end of the season are attempting to charge a different price based on the observable characteristic of the time of purchase.[4] Segmenting consumers by time of purchase is also evident in other markets. Consumers typically pay more to see a

[4] Discounting over time in a systematic fashion runs the risk that if consumers know prices will fall in the future, they will delay their purchases. If the number of customers that postpones is "too" large, seasonal discounts will not be a very good strategy.

film at a first-run theater when the film is newly released than to see it at a later date at a second-run cinema or, still later, as a rented DVD at home.

An essential feature of all third-degree price discrimination schemes is that the monopolist has some easily observed characteristic that serves as a good proxy for differences in consumer willingness to pay. This characteristic can be used effectively to divide the market into two or more groups, each of which will be charged a different price. The monopolist must next be able to ensure that resale of the product by those who are offered a low price to those who are offered a high one is not feasible. Consider the airlines again. The requirement to stay over a Saturday night effectively discriminates between those consumers who are traveling on business and those who are not.

Once the different consumer groups have been identified and separated, the general rule that characterizes third-degree price discrimination is easily stated. *Consumers for whom the elasticity of demand is low should be charged a higher price than consumers for whom the elasticity of demand is relatively high.*

5.3 IMPLEMENTING THIRD-DEGREE PRICE DISCRIMINATION OR GROUP PRICING

The logic underlying our pricing rule is fairly straightforward. Here, we illustrate it with a simple example and defer formal presentation of the general case to a Derivation Checkpoint (Discriminatory and Nondiscriminatory Pricing) later in the chapter. Suppose then that the publishers of J. K Rowling's final volume in the Harry Potter series, *Harry Potter and the Deathly Hallows*, estimate that inverse demand for this book in the United States is $P_U = 36 - 4Q_U$ and in Europe is $P_E = 24 - 4Q_E$. In each case, prices are measured in dollars and quantities in millions of books sold at publication of the first edition of the book. Marginal cost is assumed to be the same in each market and equal to \$4 per book. The publisher also incurs other costs such as cover design and promotion, but we treat these as fixed and independent of sales volume. Therefore, we can ignore them in our current analysis.

As a first case, assume that the publisher treats the two markets as a single, integrated market. To work out the profit maximizing price, the publisher will first need to calculate aggregate market demand at any price P. This means that they will need to add the two market demand curves *horizontally*. In the United States we have $P = 36 - 4Q_U$ which can be inverted to give $Q_U = 9 - P/4$ provided, of course, that $P \leq \$36$. In Europe we have $P = 24 - 4Q_E$ so that $Q_E = 6 - P/4$. This gives us the aggregate demand equation:

$$\begin{aligned} Q &= Q_U + Q_E = 9 - P/4 \qquad \text{for } \$36 \geq P \geq \$24 \\ Q &= Q_U + Q_E = 15 - P/2 \qquad \text{for } P < \$24 \end{aligned} \tag{5.1}$$

We can write this in the more normal inverse form as:

$$\begin{aligned} P &= 36 - 4Q \qquad \text{for } \$36 \geq P \geq \$24 \\ P &= 30 - 2Q \qquad \text{for } P < \$24 \end{aligned} \tag{5.2}$$

This demand relationship is illustrated in Figure 5.1. The kink in the aggregate demand function at a price of \$24 and a quantity of 3 million arises because at any price above \$24 books will be sold only in the United States whereas once the price drops below \$24 both

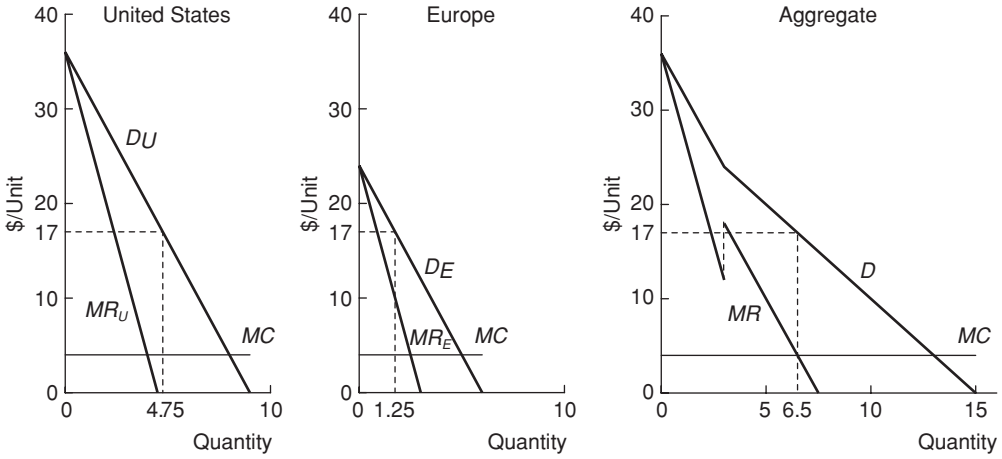

Figure 5.1 Non-disciminatory pricing: constant marginal cost
The firm identifies aggregate demand and the associated marginal revenue. It chooses total output where marginal revenue equals marginal cost and the non-discriminatory price from the aggregate demand function. Output in each market is the market clearing output.

markets are active. The marginal revenue function associated with this demand function satisfies the usual "twice as steep" rule:

$$\begin{aligned} MR &= 36 - 8Q \qquad \text{for } Q \leq 3 \\ MR &= 30 - 4Q \qquad \text{for } Q > 3 \end{aligned} \tag{5.3}$$

This is also illustrated in Figure 5.1. The jump in the marginal revenue function at a quantity of 3 million arises because when price falls from just above \$24 to just below \$24 the inactive European market becomes active. That is, when the price falls to just below \$24, it brings in a new set of consumers.

We are now is a position to calculate the profit maximizing price, aggregate quantity and quantity in each market. Equating marginal revenue with marginal cost assuming that both markets are active we have $30 - 4Q = 4$ so that $Q^* = 6.5$ million. From the aggregate demand curve this gives a price of $P^* = \$17$. It follows that 4.75 million books will be sold in the United States and 1.75 million books in Europe. Aggregate profit (ignoring all the fixed and other set-up costs) is $(17 - 4)*6.5 = \$84.5$ million.

That this pricing strategy is not the best that the monopolist can adopt is actually clear from Figure 5.1. At the equilibrium we have just calculated, the marginal revenue on the last book sold in Europe is greater than marginal cost whereas marginal revenue on the last book sold in the United States is less than marginal cost. Transferring some of the books sold in the United States to the European market will, therefore, lead to an increase in profit.

Let us be more explicit. A necessary condition for profit maximization under third-degree price discrimination is that marginal revenue must equal marginal cost in *each* market that the monopolist serves. If this were not the case in a particular market, then the last unit sold in that market would be generating either more or less in cost than it is earning in revenue. Cutting back or increasing total production in that market would therefore raise profits. If marginal cost in serving each market is identical, as in our case, then the rule implies that

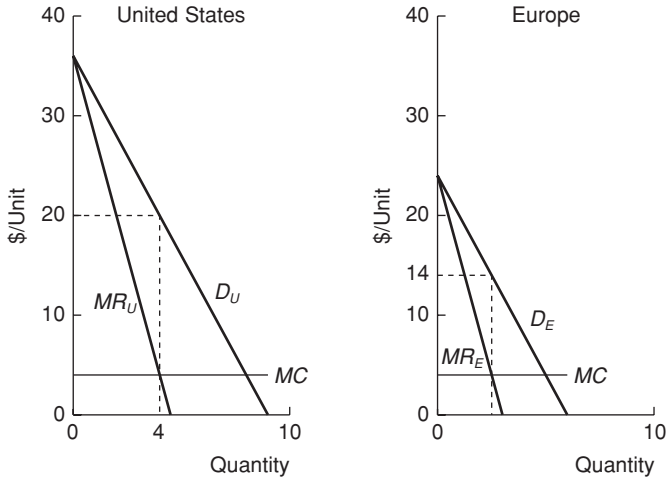

Figure 5.2 Third-degree price discrimination or group pricing: constant marginal cost
The firm sets output where marginal revenue equals marginal cost in each market and sets the market clearing price in each market.

marginal revenue be the same on the last unit sold in each market. If this condition does not hold, the monopolist can raise revenue and profit with no increase in production (and hence, no increase in costs), simply by shifting sales from the low marginal revenue market to the high one.

The application of these rules to our example is illustrated in Figure 5.2. Recall that demand in the United States market is $P_U = 36 - 4Q_U$ and in Europe is $P_E = 24 - 4Q_E$. This means that marginal revenue in the United States is $MR_U = 36 - 8Q_U$ and in Europe is $MR_E = 24 - 8Q_E$. Now apply the rule that marginal revenue equals marginal cost in each market. This gives a profit maximizing output in the United States of $Q_U^* = 4$ million books at a price of $P_U^* = \$20$ and in Europe a profit-maximizing output of $Q_E^* = 2.5$ million books at a price of $P_E^* = \$14$. Profit from sales in the Unites States is \$64 million and in Europe is \$25 million, giving aggregate profit (again ignoring all the fixed and other set-up costs) of \$89 million, an increase of \$4.5 million over the non-discriminatory profit.

How does this outcome relate to the elasticity rule that we presented above? An important property of linear demand curves is that the elasticity of demand falls smoothly from infinity to zero as we move down the demand curve. This means that, for any price less than \$24 (and greater than zero) the elasticity of demand in the United States market is lower than in the European market. (You can check this by evaluating the demand elasticity in the two markets at this or any other particular price.) Our rule then states that we should find a higher price in the United States than in Europe. This, of course, is precisely the result that our example gave.

How would our analysis be affected if marginal cost were not constant? The same basic principles apply with one important change. If marginal production costs are not constant, we cannot treat the two markets independently since whatever output the monopolist chooses to supply to the United States, for example, affects the marginal cost of supplying Europe. So the different markets have to be looked at together. Nevertheless, we still have simple rules that guide the monopolist's pricing decisions in these markets.

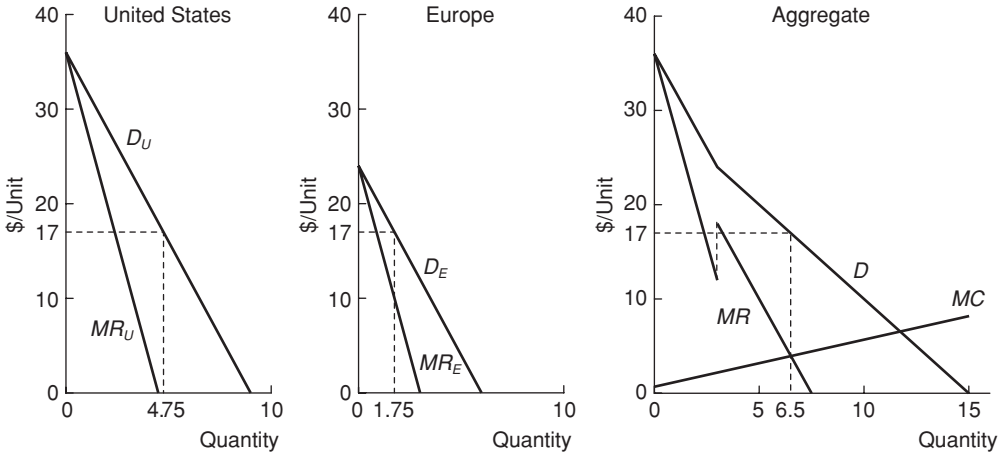

Figure 5.3 Non-discriminatory pricing with non-constant marginal cost
The firm identifies aggregate demand and the associated marginal revenue. It chooses total output where marginal revenue equals marginal cost and the non-discriminatory price from the aggregate demand function. Output in each market is the market clearing output.

To illustrate this point, suppose that the publisher of *Harry Potter and the Deathly Hallows* has a single printing facility that produces books for both the United States and European markets and that marginal cost is given by $MC = 0.75 + Q/2$, where Q is the total number of books printed.

Figure 5.3 illustrates the profit-maximizing behavior if the monopolist chooses not to price discriminate. The basic analytical steps in this process are as follows:

1. Calculate aggregate market demand as above.
2. Identify the marginal revenue function for this aggregate demand function. From our example, if $Q > 3$ so that both markets are active, this is $MR = 30 - 4Q$.
3. Equate marginal revenue with marginal cost to determined aggregate output. So we have $0.75 + Q/2 = 30 - 4Q$ giving $Q^* = 6.5$ million books.
4. Identify the equilibrium price from the aggregate demand function. Since both markets are active, the relevant part of the aggregate demand function is $P = 30 - 2Q$, giving an equilibrium price of $P^* = \$17$.
5. Calculate demand in each market at this price: 4.75 million books in the United States and 1.75 million books in Europe.

Now suppose that the monopolist chooses to price discriminate. This outcome is illustrated in Figure 5.4. The underlying process is clearly different, and the steps in implementing profit maximizing price discrimination are as follows:

1. Derive marginal revenue in each market and add these *horizontally* to yield an allocation of output across the two markets with the same marginal revenue. Marginal revenue in the United States is $MR = 36 - 8Q_U$ for any marginal revenue less than \$36 and in Europe is $MR = 24 - 8Q_E$ for any marginal revenue below \$24. Inverting these gives $Q_U = 4.5 - MR/8$ and $Q_E = 3 - MR/8$. Summing these gives an aggregated marginal revenue:

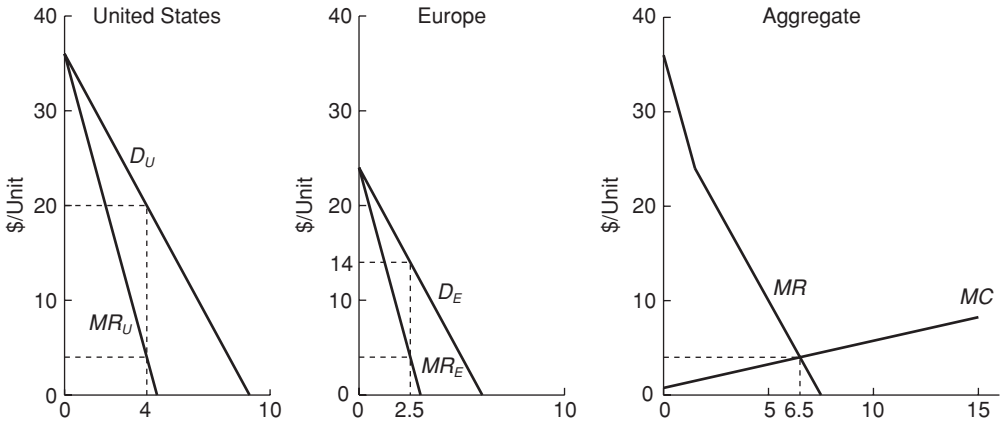

Figure 5.4 Third-degree price discrimination or group pricing with non-constant marginal cost
The firm calculates aggregate marginal revenue and equates this with marginal cost. Output in each market equates marginal revenue with aggregate marginal cost. Price in each market is the market-clearing price.

$$\begin{aligned} Q = Q_U + Q_E &= 4.5 - MR/8 \qquad \text{for } Q \le 1.5 \\ Q = Q_U + Q_E &= 7.5 - MR/4 \qquad \text{for } Q > 1.5 \end{aligned} \tag{5.4}$$

This can be inverted to give aggregate marginal revenue its more usual form

$$\begin{aligned} MR &= 36 - 8Q \qquad \text{for } Q \le 1.5 \\ MR &= 30 - 4Q \qquad \text{for } Q > 1.5 \end{aligned} \tag{5.5}$$

Note how this step differs from the non-discriminatory case. In the latter both markets are treated as one, so we start with aggregate demand and derive its associated marginal revenue. In the discriminatory pricing case, by contrast, the markets are supplied separately, with the profit maximizing condition that $MC = MR$ in both markets so we need aggregate marginal revenue, not aggregate demand.

2. Equate aggregate marginal revenue with marginal cost to identify the equilibrium aggregate quantity *and* marginal revenue. So we have $30 - 4Q = 0.75 + 2/Q$ giving $Q^* = 6.5$. As a result, the equilibrium marginal revenue is \$4, which is equal to the marginal cost of the last unit produced.
3. Identify the equilibrium quantities in each market by equating individual market marginal revenue with the equilibrium marginal revenue and marginal cost. In the United States this gives $36 - 8Q_U = 4$ or $Q_U^* = 4$ million books and in Europe $24 - 8Q_U = 4$ or $Q_E^* = 2.5$ million books.
4. Identify the equilibrium price in each market from the individual market demand functions, giving a price of \$20 in the United States and \$14 in Europe.

The foregoing procedure is again derived from two simple rules that guide the monopolist's pricing decisions with third-degree price discrimination. These rules apply no matter the shape of the monopolist's marginal cost function. The rules are:

Derivation Checkpoint

Discriminatory and Non-discriminatory Pricing

Suppose that a monopolist supplies two groups of consumers with inverse demand for each group is given by:

$$P_1 = A_1 - B_1Q_1;\ P_2 = A_2 - B_2Q_2$$

In these demand functions we assume that $A_1 > A_2$ so that group 1 is the "high demand" group whose demand is the less elastic at any given price. Inverting the inverse demands gives the direct demands at some price P:

$$Q_1 = (A_1 - P)/B_1;\ Q_2 = (A_2 - P)/B_2$$

and so aggregate demand is:

$$Q = Q_1 + Q_2 = \frac{A_1B_2 + A_2B_1}{B_1B_2} - \frac{B_1 + B_2}{B_1B_2}P$$

Of course, this holds only for any price less that A_2. Invert this to get the aggregate inverse demand for the two groups, again for any price less than A_2 yields:

$$P = \frac{A_1B_2 + A_2B_1}{B_1 + B_2} - \frac{B_1B_2}{B_1 + B_2}Q$$

The marginal revenue associated with this aggregate demand is:

$$MR = \frac{A_1B_2 + A_2B_1}{B_1 + B_2} - 2\frac{B_1B_2}{B_1 + B_2}Q$$

We can simplify matters a bit by assuming, without loss of generality, that marginal cost is zero. So solving $MR = 0$ for Q gives the equilibrium aggregate output with uniform pricing:

$$Q^U = \frac{A_1B_2 + A_2B_1}{2B_1B_2}$$

Substituting Q^* into the price equation gives the equilibrium uniform price

$$P^U = \frac{A_1B_2 + A_2B_1}{2(B_1 + B_2)}$$

Substituting this price into the individual demands then gives equilibrium output in each market

$$Q_1^U = \frac{(2A_1 - A_2)B_1 + A_1B_2}{2B_1(B_1 + B_2)};\ Q_2^U = \frac{(2A_2 - A_1)B_2 + A_2B_1}{2B_2(B_1 + B_2)}$$

With third-degree price discrimination the firm sets marginal revenue equal to marginal cost each group. From the demand curves, we know that the marginal revenues are:

$$MR_1 = A_1 - 2B_1Q_1;\ MR_2 = A_1 - 2B_2Q_2$$

It follows immediately that the equilibrium outputs for each group are

$$Q_1^D = \frac{A_1}{2B_1};\ Q_2^D = \frac{A_2}{2B_2}$$

Comparison with (d.7) confirms that $Q_1^D < Q_1^U$ and $Q_2^D > Q_2^U$. In other words, third-degree price discrimination diverts output from the high-demand market to the low-demand market increasing price in the former and lowering price in the latter. You can also confirm that $Q_1^D + Q_2^D = Q^U$. In other words, when demands are linear *aggregate output is identical with uniform pricing and with third-degree price discrimination or group pricing*.

1. Marginal revenue must be equalized in each market.
2. Marginal revenue must equal marginal cost, where marginal cost is measured at the *aggregate* output level.

There is one further interesting point that is worth noting regarding the contrast between uniform pricing (no price discrimination) and third-degree price discrimination. When demand is linear and both markets are active under both pricing schemes *aggregate demand is identical with the two pricing policies*. This is proved formally in the Derivation Checkpoint: Discriminatory and Nondiscriminatory Pricing. The intuition is simple to see. When both markets are active aggregate marginal revenue is identical with the two pricing policies (we are below the discontinuity in *MR* in Figure 5.3). So equating aggregate marginal revenue with aggregate marginal cost must give the same aggregate output. The reason that third-degree price discrimination is more profitable in this case is because the aggregate output is allocated more profitably across the two markets—to ensure that marginal revenue on the last unit sold in each market is equal.

We complete our discussion of third-degree price discrimination in this section by making explicit the relationship between the price set and the elasticity of demand in any specific market segment. Our review of monopoly and market power in Chapters 2 and 3 explained how we could express the firm's marginal revenue in any market in terms of price and the point elasticity of demand at that price. Specifically, marginal revenue in market i is given by $MR_i = P_i\left(1 - \frac{1}{\eta_i}\right)$ where η_i is (the negative of) the elasticity of firm i's demand. (See Derivation Checkpoint: The Calculus of Competition, and Problem 4, both in Chapter 2.) The larger is η_i the more elastic is demand in this market. Now recall that third-degree price discrimination requires that the profit-maximizing aggregate output must be allocated such that marginal revenue is equalized across each market (and, of course, equal to marginal cost). For example, if there are two markets then says that $MR_1 = MR_2$. Substituting from the equations above, we then know that

$$MR_1 = P_1\left(1 - \frac{1}{\eta_1}\right) = MR_2 = P_2\left(1 - \frac{1}{\eta_2}\right).$$

We can solve this for the ratio of the two prices to give:

$$\frac{P_1}{P_2} = \frac{(1 - 1/\eta_2)}{(1 - 1/\eta_1)} = \frac{\eta_1\eta_2 - \eta_1}{\eta_1\eta_2 - \eta_2}. \tag{5.6}$$

From this it is clear that price will indeed be lower in the market with the higher elasticity of demand. The intuition is that prices must be lower in those markets in which consumers are sensitive to price. Such price sensitivity means that raising the price will lose too many customers and this loss more than offsets any gain in surplus per customer. To put it differently, when consumers are price sensitive the strategy of lowering price can actually raise the monopolist's total surplus because it brings in many additional purchases. We encourage you to reinterpret the various examples with which we motivated our analysis in terms of demand elasticities. For example, is it reasonable to think that business travelers will have a lower elasticity of demand for air travel at a particular time than vacation travelers?

5.1 Practice Problem

The manager of a local movie theater believes that demand for a film depends on when the movie is shown. Early moviegoers who go to films before 5 p.m. are more sensitive to price than are evening moviegoers. With some market research the manager discovers that the demand curves for daytime (D) and evening (E) moviegoers are $Q_D = 100 - 10P_D$ and $Q_E = 140 - 10P_E$ respectively. The marginal cost of showing a movie is constant and equal to \$3 per customer no matter when the movie is shown. This includes the costs of ticketing and aning.

a. What is the profit maximizing pricing policy if the manager charges the same price for daytime and evening attendance? What is attendance in each showing and what is aggregate profit per day?
b. Now suppose that the manager adopts a third-degree price discrimination scheme, setting a different day and evening price. What are the profit maximizing prices? What is attendance at each session? Confirm that aggregate attendance is as in a. What is aggregate profit per day?

5.4 PRODUCT VARIETY AND THIRD-DEGREE PRICE DISCRIMINATION OR GROUP PRICING

We have thus far defined price discrimination as occurring whenever a firm sells an identical product to two or more buyers at different prices. But what if the products are not identical? Ford, for example, offers several hundred (perhaps even several thousand) varieties of the Ford Taurus with slightly different features. Procter & Gamble offers a wide range of toothpastes in different tastes, colors, and claimed medicinal qualities. Kellogg offers dozens of breakfast cereals that vary in terms of grain, taste, consistency, and color.

Many examples of what looks like third-degree price discrimination or group pricing arise when the seller offers such *differentiated* products. For example, books are first released as expensive hardcover editions and only later as cheap paperbacks. Hotels in a ski area are more expensive in winter than in summer. First class air travel costs more than coach. The common theme of these examples is that they all involve variations of a basic product. This is a phenomenon that we meet every day in buying restaurant meals, refrigerators, haircuts, and many other goods and services. In each of these situations, what we observe is a firm selling different varieties of the same good—distinguished by color, or material, or design. As a brief reflection on the typical restaurant menu will reveal, what we also usually observe is that the different varieties are aimed at different groups and sell at different prices.

In considering these as applications of price discrimination we have to be careful. After all, the cost incurred in producing goods of different types, such as hardback and paperback books, or first class versus coach flights, is different. Phlips (1983) provides perhaps the best definition of third-degree price discrimination or group pricing once we allow for product differentiation:

> Price discrimination should be defined as implying that two varieties of a commodity are sold [by the same seller] to two buyers at different *net* prices, the net price being the price (paid by the buyer) corrected for the cost associated with the product differentiation. (Phlips, 1983, p. 6)

Using this definition, it would not be discriminatory to charge $750 extra for a car with antilock brakes if it costs $750 extra to assemble a car with such brakes. By contrast, the difference in price between a coach class fare of $450 and a first class fare of $8,000 for service between Boston and London must be seen as almost entirely reflecting price discrimination because the additional cost of providing first class service per passenger is well below the $7,550 difference in price. In other words, price discrimination among different versions of the same good exists only if the difference in prices is not justified by differences in underlying costs: which is what Phlips means by the *net price*.

Consideration of product variety leads to a very important question. Does offering different varieties of a product enhance the monopolist's ability to charge different net prices? That is, does a firm with market power increase its ability to price discriminate by offering different versions of its product? As we shall see, the general answer is, yes.

We can obtain at least some insight into this issue by recalling the two problems that successful discrimination must overcome, namely, identification and arbitrage. In order to price discriminate, the firm must determine who is who on its demand curve and then be able to prevent resale between separate consumers. By offering different versions or models of its

Reality Checkpoint

"Seventeen Tickets for *Seven Guitars*: Price Discrimination on Broadway"

In New York, about 25,000 people, on average, attend Broadway shows each night. As avid theatergoers know, prices for these tickets have been rising inexorably. The top price for Broadway shows has risen 31 percent since 1998. However, due to various discounts offered through coupons, two-for-one deals, special student prices, and the TKTS booth in Times Square, the actual price paid has gone up by only 24 percent.

Why so much discounting? The value of a seat in a theater, like a seat on an airplane, is highly perishable. Once the show starts or the plane takes off, a seat is worth next to nothing. So, it's better to fill the seat at a low price than not fill it at all.

Stanford economist, Phillip Leslie, investigated Broadway ticket price discrimination using detailed data for a 1996 Broadway play, *Seven Guitars*. Over 140,000 people saw this play, and they bought tickets in 17 price categories. While some of the difference was due to seat quality—opera versus mezzanine versus balcony—a large amount of price differentials remained even after quality adjustments. The average difference of two tickets chosen at random on a given night was about 40 percent of the average price. This is comparable to the price variation in airline tickets.

Leslie used advanced econometric techniques to estimate the values that different income groups put on the various categories of tickets. He found that Broadway producers do a pretty good job, in general, at maximizing revenue. He found the average price set for *Seven Guitars* was about $55 while, according to Mr. Leslie's estimates, the value that would maximize profit was a very close $60. His data also indicated that the optimal uniform price would be a little over $50. Again, price discrimination is less about the average price charged and more about varying the price in line with the consumer's willingness to pay. In this connection, Leslie found that optimal price discrimination drew in over 6 percent more patrons than would optimal uniform pricing.

Source: P. Leslie "Price Discrimination in Broadway Theatre," *Rand Journal of Economics*, 35 (Autumn, 2004), 520–41.

product the monopolist may be able to solve these two problems. Different consumer types may buy different versions of product and therefore reveal who they are through their purchase decision. Moreover since different customers are purchasing different varieties, the problem of resale is considerably reduced.

As an example of the potential for product differentiation to enhance profit, consider an airline that we will call Northwest Airlines (NA), operating direct passenger flights between Boston and Amsterdam. NA knows that there are three types of customers for these flights: those who prefer to travel first class, those who wish to travel business class and those who are reconciled to having to travel coach. One part of the arbitrage problem is easily solved of course: in order to sit in a first class seat you need a first class ticket. However, there is another aspect to this problem. If the difference in price is great enough relative to the valuation a consumer places on a higher class of travel, a business class traveler, for example, might choose to fly coach. For simplicity, we assume that this arbitrage, or self-selection problem does not arise. That is, we assume that first class passengers prefer not to travel rather than sit in business or coach and business class passengers similarly will not consider coach travel—they place sufficiently high values on the differences in quality between the types of seat that they will not trade down. (See end-of-chapter problem 5 for an example of this case.)[5]

NA's market research indicates that daily demand for first class travel on this route is $P_F = 18{,}500 - 1{,}000Q_F$, for business class travel is $P_B = 9{,}200 - 250Q_B$ and for Coach travel is $P_C = 1{,}500 - 5Q_C$. The marginal cost is estimated to be \$100 for a coach passenger, \$200 for a business class passenger and \$500 for a first class passenger.

The profit maximizing third-degree price discrimination scheme for differentiated products of this type satisfies essentially the same rules as for homogeneous products. Simply put, NA should identify the quantity that equates marginal revenue with marginal cost for each class of seat and then identify the equilibrium price from the relevant demand function. For first class passengers this requires $MR_F = 18{,}500 - 2{,}000Q_F = 500$, or $Q_F^* = 9$. The resulting first class fare is $P_F^* = \$9{,}500$. In business class we have $MR_B = 9{,}200 - 500Q_B = 200$, or $Q_B^* = 18$ and $P_B^* = \$4{,}700$. Finally, in Coach we have $1{,}500 - 10Q_C = 100$, giving $Q_C^* = 140$ and $P_C^* = \$800$.

The example we have just presented resolved the arbitrage problem by assuming that different types of traveler are committed to particular classes of travel. Of course, this may not always be the case. For example, the downturn in economic activity through 2003 encouraged many businesses to seek ways to cut costs. In particular, business travelers increasingly are required by their companies to fly coach. It remains the case that these types of travelers are willing to pay more (though not as much more as before) for air travel than casual or vacation travelers. Now, however, the airline's ability to exploit the difference in willingness to pay faces a potentially severe arbitrage problem.

To see this more clearly, let's simplify the problem and suppose that the airline has just two types of customers, business people and vacationers. Business people are known to have a high reservation price, or willingness to pay, for a return ticket, which we will denote as V^B. Vacationers, by contrast, have a low reservation price, denoted as V^V. By assumption, $V^B > V^V$, and the airline would obviously like to exploit this difference by charging business customers a high price and vacationers a low one. However, the airline cannot simply

[5] There is also the possibility that coach or business travelers would want to trade up. The equilibrium prices that we derive in the example preclude such a possibility.

impose this distinction. A policy of explicitly charging business customers more than vacationers would quickly lead to every customer claiming to be on holiday and not on business. To be sure, the airline could try to identify which passengers really are on holiday, but this would be costly and likely to alienate customers.

If this were the end of our story, it would appear that the airline has no choice but to sell its tickets at a single, uniform price. It would then face the usual textbook monopoly dilemma. A high price will earn a large surplus from every customer that buys a ticket but clearly leads to a smaller, mostly business set of passengers. In contrast, a low price will encourage many more people to fly but, unfortunately, leave the company with little surplus from any one consumer.

Suppose, however, that business and holiday travelers differ in another respect as well as in their motives for flying. To be specific, suppose that business travelers want to complete their trip and return home within three days, whereas vacationers want to be away for at least one week. Suppose also that the airline learns (through surveys and other market research) that business travelers would pay a premium beyond a normal ticket price if they could be guaranteed a return flight within their preferred three-day span. In this case, product differentiation by means of offering two differentiated tickets—one with a minimum time away of a week and another with no minimum stay—will enable the airline to extract considerable surplus from each type of consumer.

The complete strategy would be as follows: First, set a low price of V^V for tickets requiring a minimum of one week before returning. Since holiday travelers do not mind staying away seven days, and since the ticket price does not exceed their reservation price, they will willingly purchase this ticket. Since such travelers are paying their reservation price, the airline has extracted their entire consumer surplus and converted it into profit for itself.[6]

Second, the airline should set a price as close to V^B as possible for flights with no minimum stay. The limit on its ability to do this will be such factors as the cost of paying for a hotel for extra nights, the price of alternative transportation capable of returning individuals in three days and related considerations. Denote the dollar value of these other factors as M. Business people wanting to return quickly will gladly pay a premium over the one week price V^V up to the value of M, so long as their total fare is less than V^B. (The precise condition is $V^V + M < V^B$.) Using such a scheme enables the airline to extract considerable surplus from business customers, while simultaneously extracting the entire surplus from vacationers.

In short, even if the airline cannot squeeze out the entire consumer surplus from the market, it can nevertheless improve its profits greatly by offering two kinds of tickets. This is undoubtedly the reason that the practice just described is so common among airlines and other transportation companies (see inset). Such companies offer different varieties of their product as a means of having their customers self-select into different groups. Automobile and appliance manufacturers utilize a similar strategy—offering different product lines meant to appeal to consumers of different incomes or otherwise different willingness to pay. Stiglitz (1977) labels such mechanisms as *screening devices* because they screen or separate customers precisely along the relevant dimension of willingness to pay.

[6] An alternative and frequently used distinction is to require that the traveler stay over a Saturday night in order to qualify for a cheap fare. Presumably a corporation will not want to finance the lodgings of its employees when they are not on company business. Further, business travelers will typically want to spend weekends with family and loved ones. On both counts, the Saturday night requirement works as a self-separating device.

Reality Checkpoint

You Can't Go Before You Come Back

It is not uncommon to find that a coach fare to fly out on Tuesday and return quickly on Thursday costs well over twice the coach fare to fly out on Tuesday and return a week later. So, for travelers wanting to return in two days, an obvious strategy is to buy two round-trip tickets—one say that departs on Thursday the 10th and returns on Thursday the 17th, and another that departs on Tuesday the 15th and returns on Tuesday the 22nd. The passenger can use the outgoing half of the first ticket on Tuesday the 15th and then fly back on the return flight of the second ticket that flies on the 17th. Unfortunately for such savvy travelers—and for the students and other needy consumers who could use the unused portions of each flight—the airlines are alert to such practices. In particular, when a passenger checks in for a flight, the airline checks to see if the passenger has an unused portion of a return flight. If so, the fee is automatically adjusted to the higher fare. The airlines have a great incentive to make sure that those who are willing to pay a substantial premium to return in two days really do pay it.

Source: "Why It Doesn't Pay to Change Planes or Plans," *London Daily Telegraph*, March 11, 2000, p. 27.

A rather curious kind of screening is illustrated by Wolfram Research, manufacturers of the *Mathematica*® software package. In making their student version of the software, Wolfram disables a number of functions that are available in the full academic or commercial versions. In 2007 Wolfram offered the full version of *Mathematica*® at around $2,495, the academic version at around $1,095 and the student version at around $140. There is little doubt that this is a case involving substantial differences in net prices.

The motivation behind this screening by means of product differentiation seems equally clear. Wolfram realizes that some customers do not need—or at least do not want to pay very much for—the full version of their software. Wolfram markets the low-priced version of *Mathematica*® for these consumers, and then sells the extended version to customers with a high willingness to pay for the improved product. Note that the two products must really differ in some important respect (to consumers at least). If Wolfram did not reduce the capabilities of the student version it would have to worry about arbitrage between the two customer groups, with students buying for their professors!

The Wolfram example just described is a type of screening referred to by marketing experts as "crimping the product." Deneckere and McAffee (1996) argue that crimping, or deliberately damaging a product to enhance the ability to price discriminate, has been a frequent practice of manufacturers throughout history. Among the examples that they cite are (1) IBM's Laser Printer E, an intentionally slower version of the company's higher-priced top-of-the-line laser printer and (2) simple cooking wine, which is ordinary table wine with so much salt added that it is undrinkable. Some people have even argued that the U.S. Post Office deliberately reduces the quality of its standard, first class service so as to raise demand for its two-day priority and overnight mail services.

Each of these examples is a clear case of a difference in net prices. The lower-quality product sells for a lower price, yet—because it starts as a high-quality product and then requires the further cost of crimping—the lower-quality product is actually more expensive to make.

Why do firms crimp a high-quality product to produce a low-quality one instead of simply producing a low-quality one in the first place? The most obvious answer relates to costs of production. Given that a firm with monopoly power such as Wolfram knows that there are consumers of different types willing to buy different varieties of its product, the firm must decide how these consumer types can be supplied with products "close" to those that they most want at least cost. It may well be cheaper to produce the student version of *Mathematica*® by crimping the full version rather than to set up a separate production line dedicated to manufacturing different versions of the software package.

The final type of product differentiation that we consider in this chapter is differentiation by *location of sale*.[7] In many cases a product for sale in one location is not the same as the otherwise identical product for sale in another location. A prescription drug such as Lipitor® for sale in Wisconsin is not identical to the same prescription drug for sale in New York State. Even with the advent of sophisticated Internet search engines, a new automobile for sale in one state is not identical to the same new automobile for sale in another state.

To illustrate why this type of product differentiation can lead to price discrimination, suppose that there is a company, Boston Sea Foods (BSF), which sells a proprietary brand of clam chowder. BSF knows that demand for its chowder in Boston is $P_B = A - BQ_B$ and in Manhattan is $P_M = A - BQ_M$ where quantities are measured in thousands of pints. In other words, the firm believes that these two markets have identical demands. BSF has constant marginal costs of c per thousand pints of chowder. Transport costs to reach the Boston market are negligible but it costs BSF an amount t to transport a thousand pints of chowder to Manhattan.

How does BSF maximize its profits from these two markets, given that BSF employs linear pricing? BSF should apply the rules that we have already developed. It should equate marginal revenue with marginal cost in each market. In the Boston market this requires that $A - 2BQ_B = c$, so that $Q_B^* = (A - c)/2B$ and the Boston price is $P_B^* = (A + c)/2$. In the Manhattan market we have, by contrast, $A - 2BQ_M = c + t$, so that $Q_M^* = (A - c - t)/2B$ and the Manhattan price is $P_M^* = (A + c + t)/2$.

Why is this outcome an example of third-degree price discrimination? Recall our definition of price discrimination with differentiated products. For there to be *no* such discrimination any difference in price should be equal to the difference in the costs of product differentiation. In our BSF example it costs BSF t per thousand pints to send chowder from Boston to Manhattan but the difference in price in the two markets is only $t/2$. In other words, BSF is price discriminating by absorbing 50 percent of the transport costs of sending its chowder to Manhattan.

What about the arbitrage problem in the BSF example? Manhattanites might want to buy their chowder directly in the Boston market but it is economic for them to do so only if they have access to a transport technology that is at least 50 percent cheaper than that employed by BSF, a very tall order other than for those who choose to vacation in Boston.

Returning to our prescription drug example in Table 5.1, one possible explanation for the difference in prices in the three United States regions might be differences in costs of supplying these three regions. Another, of course, would be differences in demands in the three regions arising from differences in these regions' demographics or incomes.

[7] We return to spatial differentiation in more detail in Chapter 7.

5.2 Practice Problem

NonLegal Seafoods (NS) sells its excellent clam chowder in Boston, New York, and Washington. NS has estimated that the demands in these three markets are respectively $Q_B = 10{,}000 - 1{,}000P_B$, $Q_{NY} = 20{,}000 - 2{,}000P_{NY}$ and $Q_W = 15{,}000 - 1{,}500P_W$ where quantities are pints of clam chowder per day. The marginal cost of making a pint of clam chowder in their Boston facility is \$1. In addition it costs \$1 per pint to ship the chowder to New York and \$2 per pint to ship to Washington.

a. What are the profit maximizing prices that NS should set in these three markets? How much chowder is sold per day in each market?
b. What profit does NS make in each market?

5.5 THIRD-DEGREE PRICE DISCRIMINATION OR GROUP PRICING AND SOCIAL WELFARE

The term "price discrimination" suggests inequity and, from a social perspective, sounds like a "bad thing." Is it? To answer this question we must recall the economist's approach to social welfare and the problem raised by the standard monopoly model. Economists view arrangements as less than socially optimal whenever there are potential trades that could make both parties better off. This is the reason that a standard monopoly is sub-optimal. The textbook monopolist practicing uniform pricing restricts output. At the margin, consumers value the product *more* than it costs the monopolist to produce it. A potentially mutually beneficial trade exists but under uniform pricing such a trade will not occur.

The question that arises with third-degree price discrimination is whether such discrimination worsens or reduces this monopoly distortion. The intuitive reason why third-degree discrimination may reduce efficiency relative to the uniform pricing case is essentially that such a policy amounts to uniform pricing within two or more separate markets. It thus runs the risks of compounding the output-reducing effects of monopoly power.

We can be more specific regarding the welfare effects of third-degree price discrimination by drawing on the work of Schmalensee (1981). This is illustrated for the case of two markets in Figure 5.5. In this Figure, P_1 and P_2 are the profit maximizing discriminatory prices—obtained by equating marginal revenue with marginal cost in each market—while P_U is the optimal non-discriminatory price. Market 2 is referred to as the strong market since the discriminatory price is higher than the uniform price while market 1 is the weak market. ΔQ_1 and ΔQ_2 are respectively the difference between the discriminatory output and the non-discriminatory output in the weak and the strong market. It follows, of course, that $\Delta Q_1 > 0$ and $\Delta Q_2 < 0$.

Our normal definition of welfare is the sum of consumer plus producer surplus. Using this definition, an upper limit on the increase in surplus that follows from third-degree price discrimination in Figure 5.5 is the area G minus the area L. This gives us the following equation. (In writing equation (5.7) we have used the property that $\Delta Q_2 < 0$.)

$$\Delta W \leq G - L = (P_U - MC)\Delta Q_1 + (P_U - MC)\Delta Q_2 = (P_U - MC)(\Delta Q_1 + \Delta Q_2) \tag{5.7}$$

Extending this analysis to n markets, we have

$$\Delta W \leq (P_U - MC)\sum_{i=1}^{n} \Delta Q_i \tag{5.8}$$

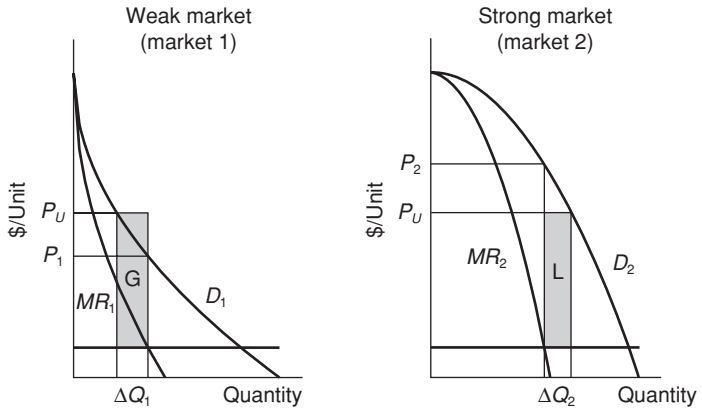

Figure 5.5 Welfare effects of third-degree price discrimination or group pricing
The upper limit on the welfare gain is area G and the lower limit on welfare loss is area L. The upper limit of the net welfare impact is $G - L$ and is positive only if aggregate output is greater with discriminatory pricing than with non-discriminatory pricing.

It follows from equation (5.8) that for $\Delta W \geq 0$ it is necessary that $\sum_{i=1}^{n} \Delta Q_i \geq 0$. In other words, a necessary condition for third-degree price discrimination to increase welfare is that it increases total output.

We know from the Harry Potter example and from the more general Derivation Checkpoint (see inset) that when demands in the various markets are linear total output is identical with discriminatory and non-discriminatory pricing. It follows that with linear demands third-degree price discrimination reduces total welfare. The increase in profit is more than offset by the reduction in consumer surplus. Schmalensee then states:

> If one thinks that demand curves are about as likely to be concave as convex . . . [this] . . . might lead one to the conclusion that monopolistic third-degree price discrimination should be outlawed. (1981, p. 246)

However, before jumping to the suggested conclusion, we need to note an important *caveat*. The qualification is that our analysis implicitly assumes that the same markets are served with and without price discrimination. This may very well not be the case. In particular, one property of price discrimination is that it can make it profitable to serve markets that would not be served with non-discriminatory prices. If this is the case, then the additional welfare from the new markets that third-degree price discrimination introduces more than offsets any loss of welfare in the markets that were previously being served.[8]

A simple example serves to make this point. Suppose that monthly demand for a patented AIDS treatment is $P_N = 100 - Q_N$ in North America but $P_S = \alpha 100 - Q_S$ in Sub-Saharan Africa with $\alpha < 1$, on the assumption that African consumers have a lower demand because their income is much smaller. We also assume that the marginal cost of producing a month's treatment is constant at $c = 20$ per unit and that transport costs to the African market are negligible.

[8] See also Shih et al. (1988) for a formal discussion of output effects under third-degree price discrimination.

Now assume that the patent holder either does not or cannot price discriminate across the two markets. As before, we start by inverting the demand functions to give $Q_N = 100 - P$ and $Q_S = \alpha 100 - P$. If the price is low enough to attract buyers in both markets then aggregate demand is: $Q = (1 + \alpha)100 - 2P$ or $P = (1 + \alpha)50 - Q/2$, and marginal revenue is $MR = (1 + \alpha)50 - Q$. Equating marginal revenue with marginal cost $c = 20$, gives the equilibrium output, $Q = (1 + \alpha)50 - 20 = 30 + \alpha 50$, and price $P = 35 + 25\alpha$.

Now recall our assumption that both markets are active without price discrimination. For this assumption to hold it must be that the equilibrium price when there is no discrimination is less than the maximum price—$\alpha 100$—that Sub-Saharan African consumers are willing and able to pay. That is, for our assumption to hold it must be the case that $35 + 25\alpha < \alpha 100$. In turn, this implies that for both markets to be active with no price discrimination it is necessary that $\alpha > 35/75$ or $\alpha > 0.466$. In other words, for the Sub-Saharan African market to be served it is necessary that the maximum willingness to pay for AIDS drugs in that market be about 47 percent of the maximum willingness to pay in North America.

Moreover, even if $\alpha > 0.466$ the Sub-Saharan African market may not be served. From the patent-holding firm's perspective, it is not quite enough that the maximum willingness to pay exceeds the price charged ***if*** it serves both markets. This is because the monopolist always has the option of choosing a higher price and serving only the North American market. In the end-of-chapter problem 6, you are asked to show that $\alpha > 0.531$ for it to be profitable for the firm to serve both markets when price discrimination is for some reason prohibited.

Practice Problem

5.3 Return to practice problem 5.1 and confirm that total welfare is greater with non-discriminatory pricing than with third-degree price discrimination.

Summary

We started this chapter with a discussion of prescription drug price differentials that seem not to be related to costs. In a well-functioning market, such differentials can only occur if there is something that separates the two groups of consumers buying at different prices. We showed that a firm with monopoly power that supplies consumers of different types can increase its profits if the firm can figure out a way to separate its consumers into these types and charge different prices to the different types. Our analysis has concentrated on third-degree price discrimination or group pricing, in which the firm offers different prices to different groups of consumer, but leaves it up to the consumers to determine how much they will purchase at the quoted prices. This is often referred to as linear pricing.

In order to implement third-degree price discrimination the firm has to solve two problems. First, it needs some observable characteristic by which it can identify the different groups of consumers: the identification problem. Secondly, the firm must be able to prevent consumers who pay a low price from selling to consumers offered a high price: the arbitrage problem. Provided that both problems can be overcome, there is then a simple principle that guides the monopolist in setting prices. Set a high price in markets in which elasticity of demand is low and a low price in markets in which elasticity of demand is high. When the firm makes a single homogeneous product this implies that different groups of consumers will be paying different prices for the same good. If the firm sells differentiated products, it implies that the prices of different varieties will vary by something other than the difference in their marginal production costs.

While third-degree price discrimination or group pricing is undoubtedly profitable, it is less clear that it is socially desirable. Again there is a simple principle that can guide us. For third-degree price discrimination to increase social

welfare it is necessary, but not sufficient, that it lead to an increase in output. This makes intuitive sense. After all, we know that under uniform pricing or non-discrimination a monopolist makes profit by restricting output. If price discrimination leads to increased output it might reduce the monopoly distortion. This is, however, a tall order, usually requiring some very restrictive conditions regarding the shapes of the demand functions in the different markets. For example, it is a condition that is *never* satisfied when demands are linear and the same markets are served with and without price discrimination.

The qualification regarding the same markets being served is, however, quite important. In particular, group price discrimination has the beneficial effect of encouraging the monopolist to serve markets that would otherwise have been left unserved. For example, markets populated by very low-income groups might not be supplied if the monopolist were not able to set discriminatory prices. When price discrimination leads the monopolist to serve additional markets, the likelihood that it increases social welfare is greatly increased.

We conclude by noting one limitation of focusing on third-degree price discrimination. Restricting the monopolist to simple, linear forms of price discrimination is qualitatively the same as allowing it to charge a monopoly price in each of its separable markets. Yet we know that in any given market, charging a monopoly price reduces the surplus. The monopolist knows this too and therefore, cannot help but wonder if a more complicated, i.e., a non-linear pricing strategy might permit the monopolist to capture more of the potential surplus as profit. It is to this question that we turn in the next chapter.

Problems

1. *True* or *false*: Price discrimination always increases economic efficiency, relative to what would be achieved by a single, uniform monopoly price.

2. A nearby pizza parlor offers pizzas in three sizes: small, medium, and large. Its corresponding price schedule is: \$6, \$8, and \$10. Do these data indicate that the firm is price discriminating? Why or why not?

3. A monopolist has two sets of customers. The inverse demand for one set may be described by $P = 200 - X$. For the other set, the inverse demand is $P = 100 - 2X$. The monopolist faces constant marginal cost of 40.
   a. Show that the monopolist's total demand, if the two markets are treated as one is:

      $X = 0$; $P \geq 200$
      $X = 200 - P$; $100 < P \leq 200$
      $X = 300 - (3/2)P$; $0 \leq P \leq 100$

   b. Show that the monopolist's profit maximizing price is $P = 120$ if both groups are to be charged the same price. At this price, how much is sold to members of Group 1 and how much to members of Group 2? What is the consumer surplus of each group? What are total profits?

4. Now suppose that the monopolist in #3 can separate the two groups and charge separate, profit-maximizing prices to each group
   a. What will these prices be? What is consumer surplus? What are total profits?
   b. If total surplus is consumer surplus plus profit, how has price discrimination affected total surplus?

5. Suppose that Coca-Cola uses a new type of vending machine that charges a price according to the outside temperature. On "hot" days —defined as days in which the outside temperature is 25 degrees Celsius or higher— demand for vending machine soft drinks is: $Q = 300 - 2P$. On "cool" days—when the outside temperature is below 25 degrees Celsius —demand is: $Q = 200 - 2P$. The marginal cost of a tinned soft drink is 20 cents.
   a. What price should the machine charge for a soft drink on "hot" days? What price should it charge on "cool" days?
   b. Suppose that half of the days are "hot" and the other half are "cool." If Coca-Cola uses a traditional machine that is programmed to charge the same price regardless of the weather, what price should it set?
   c. Compare Coca-Cola's profit from a weather-sensitive machine to the traditional, uniform pricing machine.

6. Return to the final example of section 5.4, in which the demand for AIDS drugs was $Q_N = 100 - P$ in North America and $Q_S = \alpha 100 - P$ in Sub-Saharan Africa. Show that with marginal cost = 20 for such drugs, it must be the case that $\alpha > 0.531$ if the drug manufacturer is to serve both markets while charging the same price in each market. (*Hint*: Calculate the total profit if it serves only North America and then calculate the total profit if it serves both markets. Then determine the value of α for which the profit from serving both markets is at least as large.)

7. Frank Buckley sells his famous bad tasting but very effective cough medicine in Toronto and Montreal. The demand functions in these two urban areas, respectively, are: $P_T = 18 - Q_T$ and $P_M = 14 - Q_M$. Buckley's plant is located in Kingston, Ontario, which is roughly midway between the two cities. As a result, the cost of producing and delivering cough syrup to each town is: $2 + 3Q_i$ where $i = T, M$.
   a. Compute the optimal price of Buckley's cough medicine in Toronto and Montreal if the two markets are separate;
   b. Compute the optimal price of Buckley's medicine if Toronto and Montreal are treated as a common market.

8. The Mount Sunburn Athletic Club has two kinds of tennis players, Acers and Netters, in its membership. A typical Ace has a weekly demand for hours of: $Q_A = 6 - P$. A typical Netter has a weekly demand of: $Q_N = 3 - P/2$. The marginal cost of a court is zero and there are one thousand players of each type. If the MSAB charges the same price per hour regardless of who plays, what price should it charge if it wishes to maximize club revenue?

References

Deneckere, R. and R. P. McAffee. 1996. "Damaged Goods." *Journal of Economics and Management Strategy* 5 (Summer): 149–74.

Graham, J. and B. Robson. 2000. "Prescription Drug Prices in Canada and the United States—Part 1: A Comparative Survey." *Public Policy Sources* 22. Fraser Institute.

Graham, J. and T. Tabler. 2001. "Prescription Drug Prices in Canada and the United States—Part 3: Retail Price Distribution." *Public Policy Sources* 50. Fraser Institute

Leslie, P. 2004. "Price Discrimination in Broadway Theatre." *Rand Journal of Economics* 35 (Autumn): 520–41.

Phlips, L. 1983. *The Economics of Price Discrimination*. Cambridge: Cambridge University Press.

Pigou, A. C., 1920. *The Economics of Welfare*. London: Macmillan.

Schmalensee, R. 1981. "Output and Welfare Implications of Monopolistic Third-degree Price Discrimination," *American Economic Review* 71 (March): 242–7.

Shapiro, C. and H. R. Varian. 1999. *Information Rules*. Boston: Harvard Business School Press.

Shih, J., and C. Mai, and J. Liu. 1988. "A General Analysis of the Output Effect under Third-degree Price Discrimination," *Economic Journal* 98 (March): 149–58.

Varian, H. 1989. "Price Discrimination." In R. Schmalansee and R. Willig, eds., *The Handbook of Industrial Organization. Vol. 1*. Amsterdam: North-Holland, 597–654.