

6

PRICING QUALITY

In the previous chapter we introduced the concept of indirect price discrimination, in which all consumers are offered the same set of prices, but they make choices such that “strong market” consumers pay a higher total amount to the seller (as by definition they are willing to do), and “weak market” customers, but *only* weak market customers, pay less, albeit an amount that still adds to the profit of the seller (i.e. they are paying a price above the marginal cost of the service they receive). In this chapter we discuss a second method of indirect price discrimination, where the seller offers goods and services of different quality levels to consumers, charging more for the higher quality options. The practice is commonly observed: automobile, appliance, and consumer electronics manufacturers produce different models of varying luxury and price, couriers offer different speeds of delivery for different prices, airlines offer different classes of seats, and clothing companies produce premium and discount brands. In the cultural sector, performing arts venues “scale the house,” charging different amounts for different quality seats, publishers have hardcover and paperback versions of books, record companies produce “deluxe” editions of CDs along with standard fare, and cinemas charge different amounts according to the day of the week and time of day. The practice is common; in this chapter we see if it is possible to think about setting prices for different qualities in a systematic way.

We begin with recognition that people have different willingness to pay for higher quality. When a publisher sets different prices for hardcover and paperback books, it knows that some customers will choose one and some will choose the other. We know different preferences are out there – what

58 Pricing quality

makes this situation different from direct price discrimination is that we cannot easily group people by preferences. Indeed, many individuals will be considered strong consumers in some markets and weak consumers in others. This may be the case even within a genre; for example, I might have a favorite author whose books I always want to acquire in hardcover, but I am content with paperback books for other authors. For some touring musicians it might be important to me to get the best seats I can, even at a very high price, while for other musicians I am perfectly fine sitting in the balcony.

Can the seller simply follow the rules we set in [Chapter 4](#) for direct price discrimination? For any particular artist performing in a venue, why not estimate the demand for the lower quality seats, the demand for the higher quality seats, and for each type look to where marginal revenue equals marginal cost? As a thought experiment it is a good place to start. But this situation is different from direct price discrimination: the two sets of demands are interdependent. The demand for high-quality seats depends, in part, on what is being offered in the low-quality section. If the low-quality section is really not *that* bad, and if the price of low-quality seats is quite low, then some people who otherwise might buy a high-quality seat will pick the low-quality offering instead. With direct price discrimination, say between seniors and non-seniors, customers do not get to choose their age when purchasing a ticket – they are in one category or the other. But with quality differentials, individuals do have a choice to make in terms of which quality level is best for them.

To begin, let's start with a case where the seller is offering tickets to a concert at two different quality levels. We will take the quality differential as a given for now; later in this chapter we will look at how sellers strategically alter the quality levels of their products. For a seller to ask whether it has optimally set the price differential between high and low quality, it can examine the effects of changes at the margin.

First, let's consider the lower priced, lower quality seats. If this were the only type of seat available, then we know that the optimal price is where marginal revenue (the change in total revenues from slightly lowering the price in order to sell one more seat) equals marginal cost (the cost of seating one additional customer; for performing arts venues this is probably close to zero). In the case of two different qualities of seats on offer, the principle of optimizing at the margin continues to hold, but the analysis is slightly more complex. Consider what happens if the price of low-quality seats is lowered. First, there is the effect on revenue from gaining additional customers who otherwise would not have bought a ticket at all, offset by the fact that all

low-quality seats are now selling at a lower price. That is the usual way we have been thinking about marginal revenue. But in this case there is an additional effect: some customers who had previously been willing to purchase a high-quality seat might switch to buying a low-quality seat now that the price is lower. A customer might have preferred an orchestra seat when they were priced at \$25 and the balcony was \$18, but if the balcony price falls to \$12 the customer might switch from orchestra to balcony. And, presuming the marginal cost of seating a customer in one type of seat is the same as the marginal cost of seating him in any other seat, customers switching from high-quality to low-quality seats represent a loss in revenue to the seller. Taking the sum of these effects, the low-quality seats are optimally priced when, at the margin, the gain from slightly lowering the price, through bringing customers into the venue who otherwise would not have attended, is just equal to (1) the loss from lowering the price for all those already willing to pay the going rate for low-quality seats, plus (2) the loss from customers who previously would have purchased a high-quality seat now switching to purchasing a low-quality seat. If we do not have equality between these marginal gains and losses, such that lowering the price of low-quality seats would bring more gains than losses, then the price should be lowered. In that case, we would eventually find a price where the equality holds: as we lower price, the marginal revenue from bringing into the venue new customers is falling, and there will be an increasing number of customers who previously would have purchased high-quality switching to low-quality seats. When the marginal gains from lowering the low-quality price even further just equal the marginal losses, we have found the optimal low-quality price.

Now let's turn to the price of high-quality seats; how does the seller know it is set optimally? Again, look to the margin. If the price of high-quality seats were lowered slightly, the effects would be (1) those who were previously willing to purchase high-quality seats would continue to do so, although now at a lower price than before, which represents a loss to the seller, (2) some customers who previously were purchasing low-quality seats will change their minds and purchase a high-quality seat now that the price is lower, which is a gain, and (3) some customers who previously would not have purchased any ticket at all will be induced to purchase a high-quality ticket, which is a gain. If the high-quality ticket is priced optimally, then with a change at the margin it should be the case that the loss (1) would be just offset by the gains (2) plus (3), such that there is no net gain to be had from lowering the price. If it turns out that there *is* in fact a net gain to be had from lowering price, the seller should do so.

Another way to think of the above analysis is as follows. If the difference in quality levels is taken as a given – say, for example, the difference in quality between seats in the orchestra and seats in the balcony – then the seller wants to set balcony seats at a price low enough to entice the weak market consumer, but not so low that too many strong market customers who would otherwise have bought seats in the orchestra now decide that the balcony offers the best value. And the seller wants to set the orchestra price high enough that it captures the high willingness to pay from the strong market consumer, without being so high that, again, the strong market consumers either switch to preferring balcony seats, or decide not to attend at all. Even if strong and weak market customers have very different reservation prices, the price differential cannot become *too* large, as that will cut too deeply into demand for the higher quality option. There are two differentials in play: the differential in quality levels, and the differential in prices. They are linked; a significant difference in prices can only be sustainable when there is a big enough difference in quality. Otherwise, too many strong market consumers will simply opt for the low-quality option. If the difference in the quality of experience of sitting in the orchestra or the balcony is small, that limits the degree to which prices can differ, since too high a price differential will simply induce nearly everyone to purchase a balcony ticket.

The fact that the amount of the price differential is constrained, even if willingness to pay in the strong market is *much* higher than in the weak market, is a part of the explanation for why so many arts presenters are nonprofit organizations. Consider this example: suppose the total amount that customers would be willing to pay for a season of opera, adding the reservation prices of all potential consumers, is high enough to cover the costs of producing the season. In other words, the opera season is worth doing purely for the attendees alone, even before considering any benefits to the wider community. But there is no way that price discrimination can capture the full amount that different consumers would, in theory, be willing to pay. Segmenting the market between students, seniors, and those paying full fare, using two-part pricing for additional goods and services within the opera hall, and using price differentials for different quality of seats cannot capture the fact that, for some patrons, the opera season is worth thousands of dollars. A few wealthy opera-lovers would be willing to pay a very large amount to ensure the production of an opera season, but there is no way to price tickets to capture that. But the nonprofit form allows for “voluntary” price discrimination in the form of donations. Donors know that their contributions to the opera company will be devoted to opera, as by law nonprofit organizations cannot

simply distribute net earnings to owners or managers. Charitable donations are the only way the opera company can capture the very high value that some patrons place on its art, such that costs can be covered. While this book is devoted to the myriad ways in which arts organizations can increase their revenues through strategic pricing, in some cases what can be raised through even quite sophisticated pricing schemes is not enough to capture the entire value that patrons place on the art, or to cover total costs. We turn to the question of whether nonprofit (and public) arts organizations ought to price for profit-maximization in [Chapter 10](#).

To this point we have assumed that the quality differentials that are offered by the arts presenter are given. Orchestra and balcony seats simply offer a different experience for the consumer, and cinemas cannot do much about the fact that Tuesday nights are less popular for attending the movies than Friday nights. But there are many situations where the quality differential is not simply a given, but can be adjusted by the seller. Sometimes firms will *deliberately* lower the quality of the inferior good in order to maintain demand for the premium product.

Suppose a courier offers customers two levels of service – overnight delivery and two-day delivery – with a higher price for overnight delivery. If customers, through experience, began to find that it was often the case when they purchased two-day delivery that the package was in fact delivered the very next day, they have much less incentive to purchase the higher priced guaranteed-overnight delivery the next time they have a time-sensitive package to send. Why pay for overnight when paying for two-day delivery gets the package to the recipient overnight in any case? As a consequence, the courier has an incentive to delay delivery of two-day packages until the second day even if the package has arrived in the destination city overnight; in such a case the package can be held in the storage facility until the next day. This maintains an effective quality differential (if you pay for two-day delivery, two-day delivery is all you will receive) that enables the courier to maintain the premium price for overnight shipping. And this is so regardless of cost differentials. If a two-day delivery was purchased on Monday, but the package arrives in the destination city early on Tuesday morning, there is no cost differential to the courier in delivering the package to the recipient on Tuesday or on Wednesday. Indeed, it costs the courier *more* to deliver on Wednesday since that means the package takes up space in the storage facility for a day. But it is worthwhile as a means of enabling price discrimination, and ensuring that strong market buyers who really do need fast shipment keep paying the premium price.