Chapter 1
Introduction to Pricing Techniques

This book focuses on pricing techniques that enable firms to enhance their profits. This book, however, cannot provide a complete recipe for success in marketing a certain product as this type of recipe, if it existed, would depend on a very large number of factors that cannot be analyzed in a single book. However, what this book does offer is a wide variety of pricing methods by which firms can enhance their revenue and profit. Such pricing strategies constitute part of the field called *yield management*. As explained and discussed in Section 1.3, throughout this book we will be using the term *yield management* (YM) to mean profit management and profit maximization, as opposed to the more commonly used term *revenue management* (RM).

1.1 Services, Booking Systems, and Consumer Value

Before we discuss pricing techniques, we wish to characterize the "output" that firms would like to sell. Therefore, Subsection 1.1.1 defines and characterizes the type of services and goods for which YM turns out to be most useful as a profit-enhancing set of tools. Clearly, this book emphasizes services that constitute around 70% of the gross domestic product of a modern economy. Subsection 1.1.2 identifies dynamic industry characteristics that make YM pricing techniques highly profitable. These characteristics highlight the role of the timing under which the potential consumers approach the sellers for the purpose of booking and purchasing the services sellers provide. Subsection 1.1.3 discusses the difficulty in determining consumer value and willingness to pay for services and products.

1.1.1 Service definitions

YM pricing techniques will not enhance the profit of *every* seller of goods and services. YM pricing techniques are particularly profitable for selling services, for the following reasons:

- *Nonstorability*: Services are time dependent and are therefore nonstorable. This feature is essential as otherwise service providers could transfer unused capacity from one service date to another. For example, airline companies cannot transfer unsold seats from one aircraft to another. Hotel managers cannot "save" vacant rooms for future sales.

- *Advance purchase/booking*: Time of purchase need not be the same as the service delivery time. In this book we demonstrate how reservation systems can be designed to enhance profits from the utilization of a given capacity level. For example, we show how airline companies can exploit consumer heterogeneity with respect to their ability to commit to buying services.

- *No-shows and cancellations*: Consumers who book in advance may not show up and may even cancel their reservation. Service providers should be able to segment the market according to how much refund (if any) is given upon no-shows.

- *Service classes*: The service can be provided in different quality classes. Market segmentation is profitable whenever the difference in price between, say, first and second class exceeds the difference in marginal costs.

The first item on the list is essential for the practice of YM to be profitable. The second item is not essential but definitely helps to generate extra revenue from segmenting the market according to the time reservations are made. The third item on

the list also applies to physical products (as opposed to services) because sellers often practice refund policies for goods in the form of monetary refunds and product replacements.

1.1.2 Dynamic reservation systems

As it turns out, the procedure under which consumers buy or reserve a service can be viewed as part of the service itself. Moreover, another characteristic of the type of many of the services analyzed in this book is that consumers make their reservations at different time periods. More precisely, some consumers reserve the service long before the service is scheduled to be delivered. Others make last-minute reservations.

In the absence of full refunds on purchased services, an early reservation reflects a commitment on the part of the consumer. Service providers can exploit different levels of willingness to commit by offering discounts to those consumers who are willing to make an early commitment, and charge higher prices for a last-minute booking.

The airline industry was perhaps the first industry to fully computerize reservation systems. It was also the first to systematically discriminate in price according to when bookings are made. During the late 1980s, these computerized reservation systems (CRS) were perfected and became fully dynamic so that capacity allocations could be revised according to which types of reservations were already made in addition to which reservation types would be expected to emerge before the service delivery time.

This discussion and the analysis provided in this book should help us understand the following observed phenomena, for example:

(a) Why travelers sitting in the same economy class on the same flight pay different airfares. Why people who stay at identical hotel room sizes end up paying different prices.

(b) Why capacity underutilization is often observed, such as empty seats on an aircraft and vacant hotel rooms.

Roughly speaking, the answer to (a) is that profit is enhanced when passengers and consumers pay near their maximum willingness to pay. Therefore, as long as consumers are heterogeneous with respect to their willingness to pay, proper use of YM always results in having people paying different prices for what appears to be an identical service. This is implemented via market segmentation, discussed in Subsection 1.2.2.

The answer to (b) is that because service providers seek to maximize profit, it may become profitable not to sell the entire capacity but to leave some capacity in case consumers with high willingness to pay show up at the last minute. If they

don't, then sellers are left with unused capacity. However, the reader may be wondering at this point whether profit is indeed maximized and may be asking why service providers do not at the last minute sell unused capacity at low prices, thereby avoiding empty seats and vacant rooms. The answer is simple. If consumers observe that a certain service provider sells discounted tickets at the last minute, they may be deterred from making early reservations. Thus, service providers may suffer from a bad reputation if they often practice last-minute discounts. This is known as the *sellers' commitment problem*. That is, in the short run, service providers may find it profitable to sell last-minute unbooked capacity at a lower price just to fill up the entire capacity. However, long-run considerations, such as reputation effect, prevent such practices.

1.1.3 Consumer value

The main point that this book attempts to stress is that sellers earn much higher profit if they set prices according to consumer value as opposed to basing all pricing decisions on unit cost only. It is not rare to hear managers state that their profits are generated by charging consumers a certain fixed markup above unit cost. In most cases, such cost-based pricing techniques fail to extract a large part of what consumers are actually willing to pay.

In view of this discussion, the "conflict" between buyers and sellers, particularly if the two parties allow bargaining to take place, is manifested in the following two rules:

Rule for sellers: Make an effort to set the price according to buyers' value and not according to cost.

Rule for buyers: Bargain, if you can, for prices closer to marginal cost.

Note that the rule for sellers becomes essential for services produced at near-zero marginal costs, such as those provided on the Internet.

With a few exceptions, throughout this book it is assumed that sellers know the consumers' value and willingness to pay for the services and products they sell. The firms may not know the exact valuation of a specific consumer, but it is assumed that they know the distribution of the willingness to pay among different consumer groups. Clearly, firms should exert a lot of effort to unveil these valuations. For demand functions, the next chapter shows how this can be done by running regressions on data on past sales collected from the market. However, because these data are not always available, firms may resort to market surveys. Market surveys are less reliable because consumers don't always understand the question they are asked, and even if they do, they may understate their willingness to pay.

In cases in which the seller faces competition from other firms selling similar products and services, consumers may base their willingness to pay on the prices charged by the competing firms, that is, by placing a reference value for the product

or service. In this situation, the seller must carefully study and compare the features of products and services offered by his or her competitors with the features of the product or service he or she offers. In fact, as often argued in this book, the seller should attempt to differentiate his or her brand from competing brands, by adding more features, including his or her services. Clearly, a lack of features relative to competing brands would necessitate a price reduction. After translating the observed differences between brands into their monetary equivalent, a seller should determine consumer value by

$$\begin{aligned}\text{Value of the brand} &= \text{Reference value} \\ &+ \text{“Positive” differentiation values} - \text{“Negative” differentiation values.}\end{aligned}$$

The above formula relies on the assumption that all consumers agree on the pluses and minuses of each brand, which need not always be the case – for example, in markets where the brands are horizontally differentiated (as opposed to vertically differentiated).

Finally, there are other factors that affect consumers' willingness to pay for a certain brand, including:

Switching costs: If the seller is an established firm with a large number of returning customers, the seller can add to the price the cost consumers would pay to switch to a competing brand. If the seller is a new entrant, the seller may want to reduce the price to subsidize consumer switching costs; see the analysis in Section 3.4.

Essential input: Sellers can augment the price in cases in which the product/service serves as an essential input to goods and services produced by buyers. Some economists refer to this type of action as the “holdup problem.”

Location costs: When reference prices are used, the cost of shipping or the location of the service should be reflected in the price, or shared by the parties.

1.2 Overview of Pricing Techniques

YM pricing techniques are not cost based. On the contrary, the key to successful YM is to make different consumers pay different prices for what seem to be identical services. The key to profit-enhancing pricing plans is the ability to engage in price discrimination via what economists call market segmentation. *Price discrimination* prevails if different (groups of) consumers pay different prices for what appears to be the same or a similar service or good. *Market segmentation* prevails whenever firms manage to divide the market into subgroups of consumers in which consumers belonging to different groups end up paying different prices.

1.2.1 Why is price discrimination needed?

An inexperienced reader may wonder why price discrimination is so important and ask why a strategy whereby all consumers are charged the same price is generally not profit maximizing? The answer to this question is that the practice of price discrimination enables service providers to enlarge their customer base and to create new markets. Consider the following example taken from a market for classical orchestra performances. Table 1.1 displays the willingness to pay by students and nonstudents. Suppose that each potential consumer considers buying at most one

	Students	Nonstudents
Max. Price	\$5	\$10
Number	200	300

Table 1.1: Maximum willingness to pay by students and nonstudents

ticket for a specific concert. As Table 1.1 indicates, each student will not pay more than \$5 for a ticket, whereas a nonstudent will not pay more than \$10.

First suppose that the concert hall is restricted to offering all tickets at the same price to all consumers. Then, a profit-maximizing single price can be set to a high level of \$10, thereby serving nonstudents only. Alternatively, the provider can set a low price of \$5, in which case both consumer groups will buy tickets. Ignoring costs, a high price would generate a revenue of $\$10 \times 300 = \$3000$, whereas a low price would generate a revenue of $\$5 \times (200 + 300) = \$2500$. Clearly, in this example the concert hall would set the price equal to \$10 per ticket and sell only to nonstudents.

Suppose now that the concert hall announces that all consumers who show a valid student ID are entitled to a \$5 discount from the price printed on the ticket. Under this policy, nonstudents pay the full price of \$10, whereas students end up paying \$5 for a ticket. The total revenue is given by $\$10 \times 300 + \$5 \times 200 = \$4000$, which is greater than \$3000, which is the maximal revenue generated by a single uniform pricing strategy.

Three major conclusions can be drawn from this simple example. First, as noted in Varian (1989), the key step to revenue maximization is to avoid average pricing (in our example, prices between, but not equal to, \$5 and \$10). Second, setting more than one price will increase revenue only if market segmentation is feasible. To make market segmentation feasible, the service provider must possess the physical means for avoiding *arbitrage*. In the present example, it is the student ID card that prevents arbitrage because, if checked, it prevents students from selling discounted tickets to nonstudents. If student cards are not required, all students will buy some extra tickets and sell them to nonstudents for a profit at any price between \$5.01 and \$9.99. The third conclusion to be drawn from this example is that a *discount* does not mean lower revenue. Here, revenue increases precisely because

student cards make it possible to lower the price for low-valuation consumers. In fact, later on we will analyze a similar strategy in which damaging a good (artificially lowering the quality of the service) can also enhance sales revenue.

1.2.2 Classifications of market segmentation

The above discussion was intended to convince the reader that market segmentation is necessary for the success of any price discrimination strategy. Broadly speaking, a market can be segmented along the following dimensions:

- *Consumer identifiable characteristics*: Charging different prices according to age group, profession, affiliation, location, type of delivery, and means of payment.
- *Quality*: Selling high-quality versions of the product/service to high-income buyers, and low-quality versions to low-income buyers. Segmentation of this type is possible only if the desire for higher quality increases with income. Note that firms often reduce quality (damaging the good/service) to keep differential pricing.
- *Bundling and tying*: Bundling refers to volume discounts. Segmentation of this type is possible only if consumers have different demand elasticity with respect to the quantity they purchase. Tying refers to selling packages of different goods at a single price. This market segmentation is profitable when consumer preferences for the different goods are negatively correlated.
- *Delivery time and delay*: The seller segments the market according to consumers' willingness to pay for how fast the product or service is provided or delivered. This segmentation is feasible provided that those consumers who urgently need the product or service are willing to pay a higher amount than those who don't mind waiting.
- *Components*: Sellers can segment the market by mixing different components and providing a different number of components comprising the system to be used by the buyer. This strategy is commonly observed in the software industry, where a piece of software is sold in standard, pro, and professional versions.
- *Advance booking and refunds*: Sellers can segment the market based on consumers' willingness to commit to showing up at the time the service is scheduled to be delivered. Market segmentation is achieved by charging lower prices either to those who reserve in advance or to those who seek less refund on a no-show. Conversely, those who seek to obtain a full refund on a no-show are charged a higher price.

As this book will make clear, these classifications are not mutually exclusive. To the contrary, many types of the above-listed segmentations are often combined into

a single pricing strategy. For example, book publishers tend to sell books with a hard cover during the first year of publication. Then, the same book is printed with a soft cover and sold at a lower price. Thus, consumers' willingness to pay for the first printing (fast delivery) seems to be correlated with the quality of the binding.

1.2.3 Classifications of price discrimination

Traditionally, academic economists (see, for example, Varian 1989) classify price discrimination according to first, second, and third degrees as follows:

- *First degree*: Consumers may be charged different prices so that the price of each unit they buy equals each consumer's maximum willingness to pay.
- *Second degree*: Each consumer faces the same price schedule, but the schedule involves different prices for different amounts of the good purchased. This practice is sometimes referred to as bundling (quantity discounts).
- *Third degree*: The seller segments the market into different consumer groups (with identifiable characteristics) that are charged different per-unit prices. This practice is referred to as market segmentation.

In this book, we will *not* be making much use of these classifications because the goal of this book is to characterize the proper pricing strategy to be able to segment the market, rather than just targeting a specific type of price discrimination taken from the above list. That is, from a practical point of view, the firm should be attempting to ensure that the chosen pricing techniques will indeed lead to the desired market segmentation, and that the resulting segmentation is the most profitable segmentation among all the feasible market segmentations. Moreover, the problem with the above classifications (according to first, second, and third degrees) is that these three classifications are *not* mutually exclusive. For example, second- and third-degree price discrimination can be implemented by, say, offering students different bundles from those offered to customers who cannot present student identification cards. For this reason, we deviate from the traditional classifications and follow the entry on price discrimination in Wikipedia, which suggests the following classifications based on a seller's ability to segment a market:

- *Complete discrimination*: Basically, the same as the first-degree price discrimination described above. Each consumer purchases where the marginal benefit equals the consumer-specific price.
- *Direct segmentation*: The seller segments the market into different consumer groups (with identifiable characteristics).
- *Indirect segmentation*: The seller offers variations of the product based on quality, quantity, delivery time, bundled service, and so on. The proper use of this

technique leads to self-selection of consumers according to their nonidentifiable characteristics.

The reader should note that there is a fundamental difference between direct and indirect segmentation. Direct segmentation is clearly more profitable but requires the ability and knowledge to group consumers according to age, gender, geographic location, profession, prior consumption record, and so on. However, if this knowledge is not available (or illegal under nondiscrimination laws), sellers must resort to the less profitable indirect segmentation, which relies on selecting product and service variations instead of directly selecting different consumer groups. Finally, complete segmentation is clearly the most profitable; however, it is unlikely to become feasible (and more likely to be illegal) as it requires the seller to obtain full characterization of each consumer separately.

1.3 Revenue Management and Profit Maximization

Students of economics generally fail to understand why academic and nonacademic business people use the terms *yield* and *revenue management* as the goal of their pricing strategy. This is because economics students are always taught that firms should attempt to maximize profit and that revenue maximization does not imply profit maximization in the presence of strictly positive marginal costs.

However, as it turns out very often, profit-maximizing pricing strategies are sometimes better understood in the context of revenue maximization rather than by attempting to maximize profits even when all production costs are taken into account. In addition, with the ongoing information revolution and the fast penetration of the Internet as the main source of information, yield and revenue management can in many cases lead to profit-maximizing prices, mainly because most costs of producing information are sunk whereas the cost of duplicating information services could be negligible. But even if we ignore information products, there are some industries that operate under significant capacity constraints, such as the airline and hotel industries. In these industries, most costs are sunk and indeed the marginal costs can be ignored as long as the firms operate below their full capacity.

In view of this discussion, this book uses the term *yield management* to mean the utilization of profit-maximizing pricing techniques. Therefore, we will generally avoid mentioning the commonly used term *revenue management*, although recently it seems that the use of RM is gradually replacing the use of YM. Historically, YM was associated with early problems that treated price and capacity as fixed and maximized "yield" or utilization of capital. This book, however, interprets the term *yield* as profit.

To demonstrate why profit maximization differs from revenue maximization, Table 1.2 displays the willingness to pay for a meal by three consumer groups: students, civil servants, and members of parliament. When marginal cost is zero,

	Students	Civil Servants	Parliament Members
Maximal Price:	\$5	\$8	\$10
# Consumers:	200	100	100
Marginal Cost		Profit Levels	
\$0	**2000**	1600	1000
\$4	400	**800**	600
\$7	−800	200	**300**

Table 1.2: The effect of marginal cost on the choice of profit-maximizing price. *Note*: Boldface figures are profit levels under the profit-maximizing price.

profit equals revenue. Under zero marginal cost, profit/revenue is $\$5(200+100+100) = \$2000$ when the price is lowered to \$5. Raising the price to \$8 and \$10 lowers the profit/revenue levels. Now, if the marginal cost equals 4, profit does not equal revenue. Clearly, the revenue-maximizing price has already shown to be \$5. However, it can be shown that the profit-maximizing price is \$8, yielding a profit of $(8-4)(100+100) = 800$. As Table 1.2 indicates, any other price generates lower profit levels. Finally, for a high marginal cost, Table 1.2 reveals that the profit-maximizing price is \$10, resulting in a profit level of $(\$10-\$7)100 = \$300$.

1.4 The Role Played by Capacity

Capacity constraints play a key role in yield management. First, if the service provider (seller) uses various pricing techniques as the sole strategic variable (price-based YM), these prices must depend on the amount of available capacity. This is discussed in Subsection 1.4.1. In contrast, if the seller fixes the prices according to the estimated maximum willingness to pay by the potential consumers, or if prices are fixed by competing sellers, profit can be maximized by allocating different capacities according to the different fare classes (quantity-based YM); see the discussion in Subsection 1.4.2.

1.4.1 Price-based YM under capacity constraints

To see why capacity matters, let us recall our concert hall example displayed in Table 1.1. That example showed that under unlimited capacity, price discrimination via market segmentation between students and nonstudents enhances sales to the entire potential consumer population. Now, suppose that we add a restriction to Table 1.1 whereby no more than 250 people can be seated in one performance. Such a restriction may be imposed by a regulator such as the fire department or could be structural, such as the size of the concert hall itself. Clearly, under this capacity constraint, market segmentation is not profitable as the entire capacity can