

How to Price

*A Guide to
Pricing Techniques and
Yield Management*

OZ SHY

This page intentionally left blank

How to Price

A Guide to Pricing Techniques and Yield Management

Over the past four decades, business and academic economists, operations researchers, marketing scientists, and consulting firms have increased their interest in and research on pricing and revenue management. This book attempts to introduce the reader to a wide variety of research results on pricing techniques in a unified, systematic way at varying levels of difficulty. The book contains a large number of exercises and solutions and therefore can serve as a main or supplementary course textbook, as well as a reference guide for pricing consultants, managers, industrial engineers, and writers of pricing software applications. Despite a moderate technical orientation, the book is accessible to readers with a limited knowledge in these fields as well as to readers who have had more training in economics. Most pricing models are first demonstrated by numerical and calculus-free examples and then extended for more technically oriented readers.

Oz Shy is a Research Professor at WZB – Social Science Research Center in Berlin, Germany, and a Professor of Economics at the University of Haifa, Israel. He received a BA degree from the Hebrew University of Jerusalem and a PhD from the University of Minnesota. His previous books are *Industrial Organization: Theory and Applications* (1996) and *The Economics of Network Industries* (Cambridge University Press, 2001). Professor Shy has published more than 40 journal and book articles in the areas of industrial organization, network economics, and international trade, and he serves on the editorial boards of *International Journal of Industrial Organization*, *Journal of Economic Behavior & Organization*, and *Review of Network Economics*. He has taught at the State University of New York, Tel Aviv University, University of Michigan, Stockholm School of Economics, and Swedish School of Economics at Helsinki.

How to Price

A Guide to Pricing Techniques and Yield Management

Oz Shy

WZB – Social Science Research Center, Berlin, Germany
and
University of Haifa, Israel



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press

The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521887595

© Oz Shy 2008

This publication is in copyright. Subject to statutory exception and to the provision of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published in print format 2008

ISBN-13 978-0-511-38912-2 eBook (Adobe Reader)

ISBN-13 978-0-521-88759-5 hardback

Cambridge University Press has no responsibility for the persistence or accuracy of urls for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

For Sarah, Daniel, and Tianlai

Contents

Preface xi

1	Introduction to Pricing Techniques	1
1.1	Services, Booking Systems, and Consumer Value	2
1.2	Overview of Pricing Techniques	5
1.3	Revenue Management and Profit Maximization	9
1.4	The Role Played by Capacity	10
1.5	YM, Consumer Welfare, and Antitrust	12
1.6	Pricing Techniques and the Use of Computers	13
1.7	The Literature and Presentation Methods	14
1.8	Notation and Symbols	14
2	Demand and Cost	19
2.1	Demand Theory and Interpretations	20
2.2	Discrete Demand Functions	24
2.3	Linear Demand Functions	26
2.4	Constant-elasticity Demand Functions	30
2.5	Aggregating Demand Functions	34
2.6	Demand and Network Effects	39
2.7	Demand for Substitutes and Complements	42
2.8	Consumer Surplus	45
2.9	Cost of Production	52
2.10	Exercises	56
3	Basic Pricing Techniques	59
3.1	Single-market Pricing	60
3.2	Multiple Markets without Price Discrimination	67
3.3	Multiple Markets with Price Discrimination	79
3.4	Pricing under Competition	89
3.5	Commonly Practiced Pricing Methods	99
3.6	Regulated Public Utility	104
3.7	Exercises	110

4	Bundling and Tying	115
4.1	Bundling	117
4.2	Tying	131
4.3	Exercises	145
5	Multipart Tariff	151
5.1	Two-part Tariff with One Type of Consumer	152
5.2	Two-part Tariff with Multiple Consumer Types	159
5.3	Menu of Two-part Tariffs	165
5.4	Multipart Tariff	171
5.5	Regulated Public Utility	176
5.6	Exercises	178
6	Peak-load Pricing	181
6.1	Seasons, Cycles, and Service-cost Definitions	183
6.2	Two Seasons: Fixed-peak Case	185
6.3	Two Seasons: Shifting-peak Case	190
6.4	General Computer Algorithm for Two Seasons	194
6.5	Multi-season Pricing	194
6.6	Season-interdependent Demand Functions	201
6.7	Regulated Public Utility	205
6.8	Demand, Cost, and the Lengths of Seasons	214
6.9	Exercises	223
7	Advance Booking	227
7.1	Two Booking Periods with Two Service Classes	232
7.2	Multiple Periods with Two Service Classes	238
7.3	Multiple Booking Periods and Service Classes	245
7.4	Dynamic Booking with Marginal Operating Cost	248
7.5	Network-based Dynamic Advance Booking	250
7.6	Fixed Class Allocations	254
7.7	Nested Class Allocations	258
7.8	Exercises	262
8	Refund Strategies	265
8.1	Basic Definitions	267
8.2	Consumers, Preferences, and Seller's Profit	270
8.3	Refund Policy under an Exogenously Given Price	274
8.4	Simultaneous Price and Refund Policy Decisions	280
8.5	Multiple Price and Refund Packages	288
8.6	Refund Policy under Moral Hazard	290
8.7	Integrating Refunds within Advance Booking	293
8.8	Exercises	294

9	Overbooking	297
9.1	Basic Definitions	299
9.2	Profit-maximizing Overbooking	305
9.3	Overbooking of Groups	313
9.4	Exercises	322
10	Quality, Loyalty, Auctions, and Advertising	325
10.1	Quality Differentiation and Classes	326
10.2	Damaged Goods	332
10.3	More on Pricing under Competition	335
10.4	Auctions	343
10.5	Advertising Expenditure	352
10.6	Exercises	355
11	Tariff-choice Biases and Warranties	359
11.1	Flat-rate Biases	360
11.2	Choice in Context and Extremeness Aversion	362
11.3	Other Consumer Choice Biases	366
11.4	Warranties	369
11.5	Exercises	375
12	Instructor and Solution Manual	377
12.1	To the Reader	377
12.2	Manual for Chapter 2: Demand and Cost	378
12.3	Manual for Chapter 3: Basic Pricing Techniques	382
12.4	Manual for Chapter 4: Bundling and Tying	387
12.5	Manual for Chapter 5: Multipart Tariff	391
12.6	Manual for Chapter 6: Peak-load Pricing	395
12.7	Manual for Chapter 7: Advance Booking	402
12.8	Manual for Chapter 8: Refund Strategies	406
12.9	Manual for Chapter 9: Overbooking	411
12.10	Manual for Chapter 10: Quality, Loyalty, Auctions, and Advertising	414
12.11	Manual for Chapter 11: Tariff-choice Biases and Warranties	417
	References	421
	Index	431

Preface

What This Book Will NOT Teach You

The key to successful profit-maximizing pricing is *knowing your potential customers*. If a firm does not manage to learn the characteristics of all its potential customer types, such as consumers' willingness to pay, the firm will not be able to properly price its products and services.

This book will *not* teach you how to identify the characteristics of your consumers. Although several econometric techniques for identifying these characteristics are described in Chapter 2, a comprehensive analysis of this subject is beyond the scope of this book. The two main reasons for not attempting to include these techniques in this book are (a) consumers' preferences in general, and willingness to pay in particular, vary all the time when new competing products, services, and brands are introduced to the market, which means that (b) the most efficient way of learning about customers is by *trial and error*, or simply experimenting with different tariffs while recording how consumers respond. That is, as this book shows, successful pricing techniques should not only be profitable, they should also induce consumers to *reveal* their characteristics.

What This Book Attempts to Teach You

Revenue and profit are affected by a wide variety of observable and unobservable parameters. Therefore, even if various pricing techniques are well chosen and properly used, there is still no guarantee that the firm will be profitable. However, despite the high degree of uncertainty, if one takes the approach that pricing with some reasoning cannot be inferior (profitwise) to implementing arbitrary pricing strategies, then it is hoped this book will provide you with the right intuition and with a wide variety of tools under which sellers can enhance their profits. During the past 40 years, business and academic economists, operations researchers, marketing scientists, and consulting firms have increased their interest in and research on pricing and revenue management. This book attempts to introduce the reader to a wide variety of their research results in a unified systemic way, but at varying levels of difficulty. Traditionally, the different disciplines manifested different views on pricing techniques; however, recently the attitudes toward pricing in these

disciplines have exhibited a sharp convergence that recognizes price discrimination and market segmentation as an important part of the design of profitable pricing techniques. It is hoped that the present book contributes to this convergence process.

Motivation for Writing This Book

Yield and revenue management (or profit management, as it should be called) is commonly taught in business schools, where very often teachers simply combine it with a marketing course. Revenue management is also taught in special courses and seminars for employees of the airline and hotel industries. Most of these special courses tend to be nontechnical. All this means that the analytical work on yield management, which was written mainly by scientists in the field of operations research, cannot be diffused to the general audience. Such a diffusion is not always needed, however, given that large companies tend to rely on software packages and automated reservation systems.

On the other side of the campus, the economics profession has managed over the years to develop a large number of theories on profitable pricing techniques. Most of these techniques are based on price discrimination. Other theories come from extensive research conducted by economists during the 1970s and 1980s on optimal regulation and deregulation of public utilities. Often, the economics approach goes somewhat further than the operations research approach by considering the strategic response of rival firms competing in the same market.

The purpose of this book is to combine the relevant theories from economics (mainly from microeconomics, industrial organization, and regulation) with some operations research, and to make it accessible to students and practitioners who have limited knowledge in these fields. On the other hand, readers who have had more training in economics will easily find more advanced material. Knowledge of calculus is not needed for the major part of this book, because calculus techniques are not very useful for handling discrete data, which a computer can manipulate. However, more mathematically trained readers should be able to find various topics and extensions that are based on calculus. To summarize, this book attempts to introduce the formal analysis of revenue management and pricing techniques by bridging the knowledge gained from economic theory and operations research. This book is also designed as a reference guide for pricing consultants and managers as well as computer programmers who are equipped with the appropriate technical knowledge.

Computer Applications

Professional price practitioners may want to simulate the studied pricing techniques on a computer to ultimately bring these techniques to practical use. For this reason, I have attempted to sketch some algorithms according to which programmers can write simple macros. These macros can be easily written using popular spreadsheet software and thus do not always require sophisticated programming. Of course, some readers may feel more comfortable writing in formal programming languages. The reader is invited to visit the Web site www.how-to-price.com to observe how these short macros can be implemented on the Web using the JavaScript language. Clearly, limited space does not allow me to write complete algorithms. But I hope that the logic behind the suggested algorithms would benefit the potential programmer by serving as a benchmark for more sophisticated pricing software. For convenience, the algorithms in this book are written to resemble algorithms in Pascal (a computer programming language designed in 1970 for teaching students structured programming).

To the Instructor

The instructor will find sufficient material to fill at least a one-semester course, if not an entire year. This book uses lots of calculus-free models, so it can be used without calculus if needed. An instructor's manual is provided in Chapter 12, where I also provide abbreviated solutions for all exercises. I urge the instructor to read this manual before writing the course syllabus because for each chapter, I provide some suggestions regarding which topics should suit students with different backgrounds.

Basically, the book can be divided into three parts. Although topics from all chapters are interrelated, Chapters 2 through 5 may be classified as pricing techniques (mostly for static and stationary markets). Chapters 6 through 9 roughly fall under the category of yield and revenue management as they analyze dynamic markets under capacity constraints. Chapters 10 and 11 offer a variety of topics related to both pricing and revenue management.

Each chapter ends with several exercises. These exercises attempt to motivate students to understand and memorize the basic definitions associated with the various theories developed in that chapter. The solution to all these exercises are provided in Chapter 12. Providing all the solutions to students has its pros and cons. However, I have found that students who go over these solutions perform much better on the exams than do students who are not exposed to the solutions. As a result, instead of placing the solution manual on the Internet (as I have done for my other books), I have integrated the solutions into the book itself.

This book is clearly on the technical side. However, most topics in this book are covered at multiple levels of difficulty. Hence, numerical examples should appeal to the less technical reader, whereas the general formulations and computer algorithms should appeal to more technical readers and researchers. Topics from this book can be arranged as a one-semester course for advanced undergraduate and graduate students in economics, as well as for those in some advanced MBA programs that go beyond the purely descriptive case-based method. Students of industrial engineering should also be able to grasp most of the material.

Errors, Typos, and Errata Files

My experience with my first two books (Shy 1996, 2001) has been that it is nearly impossible to publish a completely error-free book. Writing a book very much resembles writing a large piece of software because literally all software packages contain some bugs that the author could not predict. In addition, 80% of the time is devoted to debugging the software after the basic code has been written. I will therefore make an effort to publish all errors known to me on my Web site: www.ozshy.com.

Typesetting and Acknowledgments

This book was typeset by the author using the $\text{\LaTeX 2}_{\epsilon}$ document preparation software developed by Leslie Lamport (a special version of Donald Knuth's $\text{\TeX}$ program) and modified by the $\text{\LaTeX 3}$ Project Team. For most parts, I used $\text{\MikTeX}$, developed by Christian Schenk, as the main compiler.

Staffan Ringbom, Swedish School of Economics at Helsinki and HECER, has offered many suggestions, ideas, and comments that greatly improved the exposition and the content of this book. In addition, Staffan was the first to teach this book in a university environment and to collect some comments directly from students. I also would like to thank the Social Science Research Center Berlin (WZB) for providing me with the best possible research environment, which enabled me to complete this book in only two years.

During the preparation of the manuscript, I was very fortunate to work with Scott Parris of Cambridge University Press, to whom I owe many thanks for managing the project in the most efficient way. Scott has been fond of this project for several years, and his interest in this topic encouraged me to go ahead and write this book. Finally, I thank Barbara Walthall of Aptara, Inc. and the entire Cambridge University Press team for the fast production of this book.

Berlin, Germany (May 2007)

www.ozshy.com

Chapter 1

Introduction to Pricing Techniques

1.1	Services, Booking Systems, and Consumer Value	2
1.1.1	Service definitions	
1.1.2	Dynamic reservation systems	
1.1.3	Consumer value	
1.2	Overview of Pricing Techniques	5
1.2.1	Why is price discrimination needed?	
1.2.2	Classifications of market segmentation	
1.2.3	Classifications of price discrimination	
1.3	Revenue Management and Profit Maximization	9
1.4	The Role Played by Capacity	10
1.4.1	Price-based YM under capacity constraints	
1.4.2	Quantity-based YM versus price-based YM	
1.5	YM, Consumer Welfare, and Antitrust	12
1.6	Pricing Techniques and the Use of Computers	13
1.7	The Literature and Presentation Methods	14
1.8	Notation and Symbols	14

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

Value of the brand = Reference value

+ “Positive” differentiation values – “Negative” differentiation values.

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

be filled by high-valuation consumers (nonstudents in the present example). Each consumer is willing to pay \$10, so the revenue $\$10 \times 250 = \2500 is maximal.

The above discussion demonstrates that the stock of capacity is crucial for the determination of revenue and profit-maximizing prices. But, clearly capacity constraints can only be temporary because the service operator can always invest and expand her service capacity in the long run. Using the present example, the concert hall can expand or build new halls to accommodate a larger audience. All this leads us to the following conclusions:

- (a) Pricing strategy in the short run may differ from pricing in the long run.
- (b) A complete short-run and long-run pricing strategy must also include a plan for investing in additional capacity.

The above decisions must be made by any utility company. For example, electricity companies must decide on prices based on how much electricity they can generate as measured by the number of Kw/H (kilowatts per hour) the currently operable generators can produce. However, in the long run, an electricity company can vary its electricity generation capacity by purchasing additional generators or by switching to nuclear technologies, for example. The tight connection between pricing and capacity decisions actually defines the classical peak-load pricing problem to be analyzed in Chapter 6.

1.4.2 Quantity-based YM versus price-based YM

Yield management gained momentum (in fact, was initiated) as a result of the 1978 deregulation of the airline industry in the United States (followed by a similar deregulation in Europe in 1997). Newly emerging airlines, such as People's Express in the United States in the early 1980s, undercut the airfare charged by the established airlines by more than 60%. Established airlines were left with excess capacity (empty seats) on each route served by a new entrant. Consequently, during the 1980s all established airlines began allocating the seating capacity of each flight according to different fare classes. The practice of class allocation is commonly referred to as *quantity-based* YM.

The “art” of conducting proper YM is not so much how to divide the capacity among the different fare classes, but how to restrict the low-fare classes so that passengers with high willingness to pay will continue to buy the high-fare tickets. Such restrictions include advance purchase, nonrefundability, and Saturday night stay, as well as the more visible market segmentation techniques involving the division of service into classes (first class, business class, and economy class).

In this book, we do not make much use of the formal distinction between price-based YM and quantity-based YM, and this is for two reasons. First, price decisions and quantity decisions are related. For example, if an airline reservation system closes the booking of economy-class tickets, this may look like a quantity decision,

but actually this decision is equivalent to raising all airfares to match the airfare for business class. Second, the book is organized according to topics (subjects) rather than according to whether price or quantity techniques are used.

Academic economists have always been interested in profit-maximizing pricing techniques, long before the airline industry was deregulated. For this reason, most of the pricing models in this book are taken directly from economic theory. In contrast, a large number of YM theorists have been working mainly on capacity/quantity allocation techniques (quantity-based YM). These operations research techniques are also applied to inventory control problems, commonly referred to as *supply chain management* (SCM). Only recently, have academic economists been combining the choice of price into models in which consumers make advance reservations before the contracted service is scheduled to be delivered.

1.5 YM, Consumer Welfare, and Antitrust

This book is about pricing techniques firms can use to enhance their revenues and profits. At first thought, one might be tempted to say that when profits and revenue are up, consumer welfare is reduced. However, as it turns out this is not necessarily the case. There are situations under which firms can use certain profit-maximizing pricing techniques that also increase consumer welfare.

In particular, it is now well known (see Varian 1985) that price discrimination could, under certain circumstances, enhance consumer welfare. To see an example, let us return to the example displayed in Table 1.1. Consider two scenarios: (a) If price discrimination is prohibited or simply impossible to implement, we have already shown that the firm should charge a uniform price of \$10, thereby servicing only nonstudents and, yielding a revenue of \$3000. (b) Suppose now that price discrimination between students and nonstudents becomes feasible. Then, suppose that the firm announces that students are eligible for a \$6 discount on a ticket (upon presentation of a valid ID). The resulting revenue level is $(\$10 - \$6)200 + \$10 \times 300 = \$3800 > \$3000$. Comparing consumer welfare in the absence of price discrimination with the level when price discrimination is implemented, nonstudents are indifferent as they pay the same price. However, students are strictly better off with price discrimination as they are able to purchase a ticket at \$4, which is \$1 less than their maximum willingness to pay.

The key element in this example is that uniform pricing results in the exclusion of low-valuation consumers from the market. In contrast, price discrimination “invites” the students to enter the market. The entrance of low-valuation consumers constitutes a net gain in social welfare. Thus, a *necessary* condition for price discrimination to enhance social welfare is the inclusion of newly served consumers who are not served when price discrimination is not used.

Unlike textbooks in microeconomics and industrial organization, this book does not analyze social welfare. In rare cases, we will compute social welfare when

YM via price discrimination is used by a regulator (such as in public utility pricing). Further, we will not discuss and analyze the antitrust implications of each pricing technique, mainly because competition bureaus do not have general rules and guidelines to judge pricing techniques directly, unless these techniques reduce competition. Thus, reading antitrust law could be misleading as strict interpretation implies that most of these techniques are simply illegal. Just to take an example, Section 2 of the Clayton Act of 1914 amended by the Robinson-Patman Act of 1936, states that

It shall be unlawful for any person engaged in commerce, in the course of such commerce, either directly or indirectly, to discriminate in price between different purchasers of commodities of like grade and quality, . . . where the effect of such discrimination may be substantially to lessen competition or tend to create a monopoly in any line of commerce, or to injure, destroy, or prevent competition. . . .

Thus, Section 2 explicitly states that price discrimination should not be considered illegal unless price discrimination substantially decreases competition.

Finally, perhaps the main reason general rules regarding the use of each pricing strategy studied in this book do not exist is that nowadays most competition courts apply the so-called *rule of reason* as opposed to the *per se rule*. In plain language, this means that each case is judged individually and the only concern of the court is whether the action taken by a firm weakens price competition.

1.6 Pricing Techniques and the Use of Computers

This book tries something new, at least in comparison with standard textbooks on economics and business. This book provides a wide variety of computer algorithms that can provide the core for building the software needed for the computation of profit-maximizing prices. The algorithms are written in a language closely resembling the well-known Pascal programming language that makes it easy to follow the basic logic behind the algorithms. Clearly, the idea of using computers for selecting prices is not novel. In fact, there are many software companies selling services to hotel chains and airlines. Therefore, the only attempt here is to demonstrate how economic theory can be embedded into simple computer algorithms.

As we mentioned earlier in this chapter, computers cannot substitute for human intuition in determining profit-maximizing prices. Most companies still determine prices by intuition combined with trial and error. Moreover, at least at the time of writing this book, computers cannot determine which pricing techniques should be used in each market. Loosely speaking, computers cannot think. All computers can do is process a large number of computations at a much greater speed and with greater accuracy than what humans can do. Because of this feature, computers can be used to verify whether a particular intuition happens to be correct or false.

Having said all that, it should be pointed out that with the increase in speed and reduction in the cost of running computers, there is a growing tendency among researchers to try different methods of explorations with the use of computers (see, for example, Wolfram 2002). Computers can simply search large databases and experiment with different price structures. The simple algorithms in this book can serve as examples of what kinds of “programming loops” are needed for searching for the “right” prices. Price practitioners who use such algorithms should also design additional algorithms that test the results against alternative price mechanisms.

1.7 The Literature and Presentation Methods

As I mentioned before in this introduction, similar to my earlier two books, the presentation in this book is based on two beliefs of mine: First, high-level math is not always needed to present a full argument. For example, a model with two or three states of nature can easily replace a continuous density function. More importantly, my second belief (although many researchers may disagree) is that a simple model is not necessarily less general than a complicated model, or a model that uses high-level math.

For these reasons, and because the presentation level in scientific journals differs substantially from the presentation of this book, I was not able to fully use scientific literature for writing this book. Therefore, this book is not intended to survey the vast literature on yield management. The reason is that to make the models accessible to undergraduate students as well as to general pricing practitioners, I had to design my own models rather than use someone else's. I guess this is the right place to formally apologize to all those researchers whose work is not cited. Clearly, the choice of which paper to cite or not cite is not based on any value judgment, but on convenience and relevance to the simplified version used in this book. Readers seeking a comprehensive reading list of the scientific literature on YM should consult recent books by Talluri and van Ryzin (2004), and Ingold, Yeoman, and McMahon (2001), as well as a literature survey by McGill and van Ryzin (1999). On pricing in general, Monroe (2002), Nagle and Holden (2002), McAfee (2005, Ch. 11), and Winer (2005) provide comprehensive studies of pricing techniques as well as extensive discussions on all aspects related to pricing, including behavioral and psychological approaches.

1.8 Notation and Symbols

The book tries to minimize the use of mathematical symbols. For the sake of completeness, Table 1.3 contains all the symbols used in this book.

Notation is classified into two groups: *parameters*, which are numbers that are treated as exogenous by the agents described in the model, and *variables*, which

are endogenously determined. Thus, the purpose of every theoretical model is to define an equilibrium concept that yields a unique solution for these variables for given values of the model's parameters.

For example, production costs and consumers' valuations of products are typically described by parameters (constants), which are estimated in the market by econometricians and are taken exogenously by the theoretical economist. In contrast, quantity produced and quantity consumed are classical examples of variables that are endogenously determined, meaning that they are solved within the model itself.

We now set the rule for assigning notation to parameters and variables. *Parameters are denoted either by **Greek** letters or by **uppercase English** letters. In contrast, variables are denoted by **lowercase English** letters.* Table 1.4 lists the notation used for denoting parameters throughout this book. Finally, Table 1.5 lists the notation used for denoting variables throughout this book.

Symbols

$=$	equal by derivation
$\stackrel{\text{def}}{=}$	equal by definition
$\approx$	approximately equal
$\implies$	implies that
$\iff$	if and only if
Σ	sum [Sigma]
Δ	a change in a variable/parameter [Delta]
$\%\Delta$	percentage change in a variable/parameter
∂	partial derivative
$\in$	is an element of the set
E	expectation operator (expected value of...)
$\times$ or $\cdot$	simple multiplication operators
$!$	factorial, for example, $3! = 1 \times 2 \times 3 = 6$
$\lfloor x \rfloor$	floor of x , for example, $\lfloor 3.16 \rfloor = \lfloor 3.78 \rfloor = 3$
$\lceil x \rceil$	ceiling of x , for example, $\lceil 3.16 \rceil = \lceil 3.78 \rceil = 4$
$\text{func}(\text{var})$	is a function of the variable, for example, $f(x)$
$\leftarrow$	assignment operation in a computer algorithm, $x \leftarrow 2$
$\text{Pr}\{\text{event}\}$	probability of an event: $[0, 1]$
$\mathbb{R}$	positive or negative real numbers: $(-\infty, +\infty)$
$\mathbb{R}^+$	nonnegative real numbers: $[0, +\infty)$
$\mathbb{R}^{++}$	strictly positive real numbers: $(0, +\infty)$
$\mathbb{Z}$	integer numbers: $\dots, -2, -1, 0, 1, 2, \dots$
$\mathbb{N}^+$	natural numbers: $0, 1, 2, 3, \dots$
$\mathbb{N}^{++}$	strictly positive natural numbers: $1, 2, 3, \dots$
$\emptyset$	empty set (a set containing no elements)
$\{, , , \}$	set of elements (order does not matter)
$(, , ,)$	vector (order does matter)
LHS, RHS	Right-hand side and left-hand side of an equation

Table 1.3: Symbols.

Parameters

Notation	Type	Interpretation
ϕ	$\mathbb{R}^+$	fixed or sunk production cost [phi]
μ_o	$\mathbb{R}^+$	marginal operating/production cost [mu]
μ_k	$\mathbb{R}^+$	marginal cost of capital/capacity/reservation
K	$\mathbb{R}^+, \mathbb{N}^+$	available capacity level, or amount of capital
F	$\mathbb{N}^+$	number of firms in a given industry
α	$\mathbb{R}^{++}$	demand (shift) parameter [alpha]
β	$\mathbb{R}^+$	demand parameter (slope or exponent) [beta]
π	$[0, 1]$	probability ($0 \leq \pi \leq 1$) [pi]
π^i	$[0, 1]$	probability of a booking request for class i , $i \in \mathcal{B}$
π^0	$[0, 1]$	probability of not booking ($\pi^0 \stackrel{\text{def}}{=} 1 - \sum_i \pi^i$)
τ	$\mathbb{N}^+$	a particular time period (e.g., $t = \tau$) [tau]
T	$\mathbb{N}^{++}$	number of periods/seasons, or the last period/season
D	$\mathbb{R}^{++}$	duration of a season
ε	$\mathbb{R}^{++}$	a small number [epsilon]
M	$\mathbb{N}^{++}$	# of consumer types, # markets ($\ell = 1, 2, \dots, M$)
V_ℓ	$\mathbb{R}^+$	consumer ℓ 's willingness to pay for a product/service
N	$\mathbb{N}^{++}$	number of consumers (N_ℓ of type ℓ) (N_i in group i)
$U_\ell(\cdot)$	func.	consumer ℓ 's utility function (of V_ℓ , p , etc.)
δ	$\mathbb{R}^{++}$	differentiation (or switching) cost [delta]
ψ	$\mathbb{R}^+$	penalty level [psi]
P	$\mathbb{R}^+$	exogenously given price charged by a firm
G	$\mathbb{N}^{++}$	grid (computer algorithms' precision of price change)
B	$\mathbb{N}^+$	maximum allowable booking level, or # block rates
$\mathcal{B}$	set	booking classes, set of goods, set of quality levels, for example, $\mathcal{B} = \{A, B, C, \dots, i, \dots\}$, $i \in \mathcal{B}$, etc.

Table 1.4: General notation for parameters (uppercase English and Greek letters).

Variables

Notation	Type	Interpretation
d_t	0, 1	decision in period t , $d_t = 1$ (accept), $d_t = 0$ (reject)
x	$\mathbb{R}^+$	(expected) revenue of a firm
c	$\mathbb{R}^+$	(expected) total cost borne by a firm
y	$\mathbb{R}$	(expected) profit of a firm ($y \stackrel{\text{def}}{=} x - c$)
p or f	$\mathbb{R}^+$	endogenously determined price/fee set by a firm
up	$\mathbb{R}^+$	markup on marginal cost ($up \stackrel{\text{def}}{=} p - \mu$)
q	$\mathbb{R}^+$	quantity produced or quantity demanded
$\bar{q}$	$\mathbb{R}^+$	aggregate industry output ($\bar{q} \stackrel{\text{def}}{=} \sum_j q_j$)
e	func.	elasticity function [$e(q) \stackrel{\text{def}}{=} (\Delta q / \Delta p)(p/q)$]
$\check{e}$	func.	arc elasticity function
gcs	$\mathbb{R}^+$	gross consumer surplus
ncs	$\mathbb{R}$	net consumer surplus ($= gcs - \text{expenditure}$)
b	$\mathbb{N}^+$	booking level (number of confirmed reservations)
b_t^i	$\mathbb{N}^+$	period t cumulative booking level for class $i \in \mathcal{B}$
k_t	$\mathbb{N}^+$	period t remaining capacity ($k_t = K - \sum_i b_t^i$)
p^i	$\mathbb{R}^+$	price of a service of class $i \in \mathcal{B}$
p_t	$\mathbb{R}^+$	period t price offer (by a consumer or set by a seller)
r	$\mathbb{R}^+$	refund level ($r \leq p$)
cn	$\mathbb{R}^+$	cancellation fee
n	$\mathbb{R}^+$	number of consumers who buy/book the product/service
s	$\mathbb{R}^+$	number of consumers who show up
ds	$\mathbb{R}^+$	number of consumers who are denied service
g	$\mathbb{N}^+$	number of (booked) consumer groups
a	$\mathbb{R}^+$	advertising expenditure

Indexing variables

i	$\mathbb{N}^+$	general, product/service types, booking classes & groups
j	$\mathbb{N}^+$	general firms in a given industry
ℓ	$\mathbb{N}^+$	consumer types
t	$\mathbb{N}^+$	time period or season (e.g., $t = 0, 1, 2, \dots, T$)

Table 1.5: General notation for variables (lowercase English letters) and indexing.

Chapter 2

Demand and Cost

2.1	Demand Theory and Interpretations	20
2.1.1	Definitions	
2.1.2	Interpreting goods and services	
2.1.3	The elasticity and revenue functions	
2.2	Discrete Demand Functions	24
2.3	Linear Demand Functions	26
2.3.1	Definition	
2.3.2	Estimation of linear demand functions	
2.3.3	Elasticity and revenue for linear demand	
2.4	Constant-elasticity Demand Functions	30
2.4.1	Definition and characterization	
2.4.2	Estimation of constant-elasticity demand functions	
2.4.3	Elasticity and revenue for constant-elasticity demand	
2.5	Aggregating Demand Functions	34
2.5.1	Aggregating single-unit demand functions	
2.5.2	Aggregating continuous demand functions	
2.6	Demand and Network Effects	39
2.7	Demand for Substitutes and Complements	42
2.8	Consumer Surplus	45
2.8.1	Consumer surplus: Discrete demand functions	
2.8.2	Consumer surplus: Continuous demand functions	
2.9	Cost of Production	52
2.10	Exercises	56

The key to a successful implementation of any yield management strategy is getting to know the consumers. Large firms invest tremendous amounts of money on research seeking to characterize their own customers as well as potential consumers. In economic theory, the most useful instrument for characterizing consumer behavior is the demand function. A demand function shows the quantity demanded by an

individual, a group, or all the consumers in a given market, as a function of market prices, and some other variables.

Knowing the demand structure is a necessary condition for proper selection of profit-maximizing actions by the firm. But it is not a sufficient condition because the firm must also take cost-of-production considerations into account. Therefore, price decision makers within a firm should properly study the structure of the cost of the service or the product sold by their own firms. They should also distinguish among the different types of costs, particularly between costs associated with a marginal expansion of output and costs associated with investing in infrastructure, research and development (R&D), and capacity.

For this reason, we devote an entire chapter to study the most widely used demand and cost structures. Some readers, particularly readers who took microeconomics courses at a second-year undergraduate level, may be familiar with this material. In general, readers can skip this chapter and use it as a reference for the various concepts whenever necessary.

2.1 Demand Theory and Interpretations

In most cases, knowing consumers' demands constitutes the key information on which producers, service providers, and sellers in general base their profit-maximizing pricing and marketing techniques. However, the concepts of demand, products, and services may be given a wide variety of interpretations. In this section, we discuss some of these interpretations as it is important that firms identify the precise type of demand they are facing before they design their pricing mechanisms.

2.1.1 Definitions

DEFINITION 2.1

- (a) The **demand function** $q(p)$ shows the quantity demanded at any given price, $p \geq 0$, by a single consumer or a group of consumers.
- (b) The **inverse demand function** $p(q)$ shows the maximum amount that an individual or a group of individuals is willing to pay at any given consumption level, $q \geq 0$. Mathematically, the function $p(q)$ is the inverse of the function $q(p)$.

Notice that Definition 2.1 is general enough to be applied to different levels of market aggregation in the sense that it can be applied to individuals as well as to different-sized markets. The technique for how to combine individuals' demand functions into a single market demand function will be studied in Section 2.5.

Definition 2.1 is incomplete in the sense that it makes the quantity demanded depend on the price/fee only. However, as the reader may be well aware of, demand is also influenced by a wide variety of other factors, such as prices of other

goods and services that consumers may view as substitutes or complements, income levels, time of delivery, bundling and tying with other goods and services, advertising, social conformity and nonconformity, social pressure, network effects, environmental concerns, brand loyalty, and more. Clearly, all these factors are very important, and most of them are incorporated in this book.

2.1.2 Interpreting goods and services

This book analyzes profit-maximizing pricing techniques. For this reason, producers and service providers should fully understand and very often define the nature of the products and services they supply. Furthermore, it is important that sellers understand how these products and services are used by consumers and how consumers perceive the benefit they gain from consuming these goods. For this reason, we now attempt to characterize goods and services according to several criteria.

Frequency of purchase: Flow versus stock goods

The distinction between flow and stock goods is based on the frequency of purchase. Theoretically, pure stock goods are those that can be stored indefinitely. Diamonds, gold, silver, and artworks are good examples. However, often we view mechanical and electrical appliances also as stock goods despite the fact that they tend to be replaced every few years. Flow goods are generally perishable goods for which the quantity demanded is measured by units of consumption during a certain time period. By *perishable*, we mean goods that cannot be stored for a long time, if at all. Therefore, perishable products must be repeatedly purchased. Food items, perhaps, provide the best examples of flow goods. Some food items can be stored for only a week, whereas others (such as boxed food) can be stored for six months or even longer.

Because the majority of pricing models presented in this book apply to services, we should mention that services are generally regarded as highly perishable, which means they are nonstorable. The reason for this is that most services are time dependent, so postponements and delays may result in a partial or total utility loss to consumers. For example, traveling today via ground, air, or sea transportation, or a hotel room for tonight, may be regarded as totally different services from traveling and a hotel room tomorrow. Reading or watching the news today constitute a totally different service from getting tomorrow's news. Of course, this is not always the case as for some services, such as changing the engine oil in your car, postponing the service for a limited time will not matter to you very much. Despite the fact that most services are perishable, not all services are flow goods in the sense that some services, such as a trip to the Galápagos Islands, may be purchased once in a lifetime. In contrast, a bus trip to work is definitely a flow good, as it is repeated on a daily basis. A ski trip can also be repeated on a yearly basis or for some people, can be a once-in-a-lifetime event.

Back to products, in the “new” information economy, information goods consume a large portion of individuals’ budgets. We interpret *information goods* in the broad sense of the term to include books, software, encyclopedias, databases, music, and video. We tend to treat these information products as stock goods. Not only can these products be stored for a long period of time, they can also be duplicated without any deterioration if they are stored in digital formats. Of course, storage devices such as magnetic tapes, digital disks, and diskettes, are themselves perishable and therefore require some maintenance or replacement. We should point out that in some sense, information goods can also be viewed as services simply because there is no benefit from storing them. News on current prices in the stock markets, or any other type of news, can be viewed as perishable goods that consumers must purchase repeatedly.

Quantity of purchase and willingness to pay

Generally, there are two major interpretations for demand schedules. The first interpretation involves consumers who increase the number of units purchased when the price drops (holding other parameters affecting demand constant). This interpretation is commonly found in introductory textbooks that are used in university microeconomics courses. With the risk of finding many counterexamples, we can say that this interpretation is more suitable for markets of flow goods, where frequent purchases make the quantity purchased highly sensitive to short-run and small changes in prices. Figure 2.1(left) illustrates a downward-sloping inverse demand function for one individual consumer. Thus, the consumer depicted on Figure 2.1(left) demands $q = 2$ units at the price of $p = \$30$, $q = 3$ units at $p = \$25$, and so on.




Figure 2.1: Illustration of inverse demand functions by individuals. *Left:* Downward-sloping demand. *Right:* Single-unit demand.

The second interpretation, which will be used extensively in this book, applies to markets composed of a large number of individuals. Each consumer is assumed

to buy *at most one unit* of the product/service, and will not buy additional units even when the price drops. Figure 2.1(right) exhibits a demand function where the consumer does not purchase the product ($q = 0$) at prices exceeding \$20 ($p > \20). However, as the price drops to \$20 or below ($p \leq \20), the consumer demands exactly *one unit* ($q = 1$) and does not buy more units, even as the price drops to zero. We should point out that consumers with this type of demand are not homogeneous in the sense that each consumer may have a different level of willingness to pay for a unit of consumption. That is, we could also plot a similar demand function whose maximum willingness to pay is \$40 and not \$20 as for the consumer plotted in Figure 2.1(right). These differences may be generated by differences in income, value of time, and the utility generated from the services of the product or the service. Under this interpretation, the market demand function represents a summation of the individuals whose willingness to pay exceeds the market price. We refer the reader to Section 2.5 for a demonstration of how market demand functions can be derived from groups of individuals whose demand functions are not price sensitive, as illustrated in Figure 2.1(right).

2.1.3 The elasticity and revenue functions

The elasticity function is derived from the demand function and maps the quantity purchased to a very useful number that we call the *elasticity at a point* on the demand. The elasticity measures how fast quantity demanded adjusts to a small change in price. Formally, we define the price elasticity of demand by

$$e(q) \stackrel{\text{def}}{=} \left(\frac{\Delta q}{\Delta p} \right) \left(\frac{p}{q} \right) = \frac{\text{percentage change of } q}{\text{percentage change of } p} = \frac{\% \Delta q}{\% \Delta p}. \quad (2.1)$$

DEFINITION 2.2

At a particular level of consumption q , the demand

- is called **elastic** if $e(q) < -1$ (or, $|e(q)| > 1$),
- is called **inelastic** if $-1 < e(q) < 0$, (or, $|e(q)| < 1$),
- and has a **unit elasticity** if $e(q) = -1$ (or, $|e(q)| = 1$).

The inverse demand function shows the maximum amount a consumer is willing to pay per unit of consumption at a given consumption level q . The *revenue* function shows the amount of revenue collected by a seller at a particular price–quantity combination. Formally, we define the revenue function by

$$x(q) \stackrel{\text{def}}{=} p(q) \cdot q, \quad (2.2)$$

which is the price multiplied by the corresponding quantity demanded.

Finally, the *marginal revenue* function is the change in revenue resulting from a “slight” increase in quantity demanded. Formally we define the marginal revenue function by $\Delta x / \Delta q$, for “small” increments of q , as given by Δq .

2.2 Discrete Demand Functions

Data on demand are discrete in nature because the number of observations in the form of data points that can be collected is always finite. Therefore, when a researcher plots the raw data, the demand function consists of discrete points in the price–quantity space. We refer to graphs representing the raw data as *discrete* demand functions, and distinguish them from the continuous demand functions we also analyze in this chapter. In fact, in Sections 2.3 and 2.4, we demonstrate how to generate *continuous* linear and constant-elasticity demand functions from discrete demand by estimating the continuous functions directly from the raw data. Although most textbooks in economics use continuous demand functions, this book focuses mainly on discrete demand functions. The reason is that, by construction, computers generally handle computations based on a finite amount of data. This means that the computations based on raw data, or based on data generated by computer reservation systems, must be handled using discrete algebra (as opposed to calculus).

The first two rows of Table 2.1 display the raw data for a discrete demand function. Rows 3–6 in this table display various computations derived directly

p	\$35	\$30	\$25	\$25	\$20	\$20	\$15	\$15	\$10
q	1	2	3	4	5	6	7	8	9
$e(q)$	−7.00	−3.0	n/d	−1.25	n/d	−0.67	n/d	−0.38	n/a
$\check{e}(q)$	−4.33	−2.2	n/d	−1	n/d	−0.54	n/d	−0.29	n/a
$x(q)$	\$35	\$60	\$75	\$100	\$100	\$120	\$105	\$120	90
$\frac{\Delta x(q)}{\Delta q}$	\$25	\$15	\$25	\$0	\$20	−\$15	\$15	−\$30	n/a

Table 2.1: Discrete demand function $\langle p, q \rangle$, and the corresponding price elasticity $e(q)$, arc price elasticity $\check{e}(q)$, total revenue $x(q)$, and marginal revenue $\frac{\Delta x(q)}{\Delta q}$. Note: n/d means not defined (division by 0), n/a means data not available.

from the raw data. Figure 2.2 plots the raw data provided by Table 2.1. The third row of Table 2.1 computes the price elasticity defined by (2.1). To demonstrate two examples, we compute the price elasticity at $q = 1$ and $q = 6$ units of consumption to be

$$e(1) = \frac{2 - 1}{\$30 - \$35} \frac{\$35}{1} = -7 \quad \text{and} \quad e(6) = \frac{7 - 6}{\$15 - \$20} \frac{\$20}{6} = -0.67. \quad (2.3)$$

Observe that we are unable to compute the price elasticity at $q = 9$, because we have no data on the price that would induce the consumer(s) to buy $q = 10$ units. For this reason, we marked $e(9)$ as n/a in Table 2.1.

Inspecting the elasticity computations given by (2.3) reveals that the elasticity formula is based on the evaluation of the price and the corresponding quan-




Figure 2.2: Discrete demand function based on data given in Table 2.1.

tity “before” the change takes place. For example, $e(6)$ in (2.3) is evaluated at the pair $\langle p, q \rangle = \langle \$20, 6 \rangle$, although there is no reason why this evaluation should not be made at $\langle p, q \rangle = \langle \$15, 7 \rangle$, which is the demand point “after” the changes $\Delta p = \$15 - \20 and $\Delta q = 7 - 6$ take place. The arc elasticity function, denoted by $\check{e}(q)$, corrects this problem somewhat by evaluating the elasticity at the midpoints $\bar{p} = (\$15 + \$20)/2 = \$17.5$ and $\bar{q} = (6 + 7)/2 = 6.5$ instead of at the “start” and “end” points of this change. Therefore, redoing the computation (2.3) for the arc elasticity yields

$$\check{e}(1) = \frac{2 - 1}{\$30 - \$35} \frac{\$32.5}{0.5} = -4.33 \quad \text{and} \quad \check{e}(6) = \frac{7 - 6}{\$15 - \$20} \frac{\$17.5}{0.5} = -0.54, \quad (2.4)$$

which is considered to be a more accurate measure of elasticity.

The revenue collected by sellers (which equals consumers’ expenditure on this good) and the marginal revenue function are simply defined by

$$x \stackrel{\text{def}}{=} p(q)q \quad \text{and} \quad \frac{\Delta x(q)}{\Delta q}. \quad (2.5)$$

For example, using the data in Table 2.1, $\frac{\Delta x(1)}{\Delta q} = x(2) - x(1) = \$60 - \$35 = \25 , and $\frac{\Delta x(6)}{\Delta q} = x(7) - x(6) = \$105 - \$120 = -\15 .

Inspecting Table 2.1 reveals that the maximum revenue that can be extracted from consumers is \$120, which is accomplished by setting the price either to $p = \$20$ and selling $q = 6$ units, or to $p = \$15$ and selling $q = 8$ units. Moreover, the price elasticities $e(q)$ and $\check{e}(q)$ are close to -1 around the revenue-maximizing output levels. This is not a coincidence, and as shown in Sections 2.3 and 2.4, for continuous demand functions, revenue is always maximized at the output level

where both elasticities take a value of $e(q) = \check{e}(q) = -1$. Observe that regular and arc elasticities are always equal for continuous demand functions because their evaluation is based on infinitesimal changes.

2.3 Linear Demand Functions

2.3.1 Definition

A linear demand is a special type of the general demand function characterized by Definition 2.1. Its special property is that it is drawn as, or fitted to be represented by, a straight line. Formally, the general formulas for the inverse and the direct demand functions are defined by

$$p(q) \stackrel{\text{def}}{=} \alpha - \beta q \quad \text{or} \quad q(p) \stackrel{\text{def}}{=} \frac{\alpha}{\beta} - \frac{1}{\beta} p, \quad (2.6)$$

where the parameters α and β may be estimated using econometric techniques, as we briefly demonstrate in Section 2.3.2. Figure 2.3 plots the linear inverse demand function given by the formula on the left-hand side of (2.6). Note that part of the




Figure 2.3: Inverse linear demand function.

demand is not drawn in Figure 2.3. That is, for any price exceeding the intercept α , the (inverse) demand becomes vertical at $q = 0$. In other words, the demand coincides with the vertical axis for prices in the range $p > \alpha$.

For the inverse demand function, the parameter $\alpha > 0$ is called the demand intercept, whereas the parameter $\beta \geq 0$ is called the slope of the inverse demand curve. We should mention that in rare cases, $-\beta$ may be found to have a positive slope for some price range. Inverting the inverse demand function formulated on the left-hand side of (2.6) generates the direct demand function on the right-hand side of (2.6), with an intercept α/β and a slope of $-1/\beta$.

2.3.2 Estimation of linear demand functions

Firms can collect data on the demand facing their products and services by experimenting with different prices and recording the quantity demanded at every price. Alternatively, firms can conduct consumer surveys. What is common to both methods is that the data collected are discrete (as opposed to continuous). In other words, observations generally consist of data points that are a collection of the pairs $\langle p, q \rangle$. Table 2.2 provides an example of data points on quantity demanded at various prices.

ℓ (observation no.)	1	2	3	4	5	Mean
p_ℓ (price)	\$30	\$25	\$20	\$15	\$10	\$20.00
q_ℓ (quantity)	2	3	6	7	9	5.40

Table 2.2: Data points for estimating linear demand.

The discrete data on price and quantity observations provided by Table 2.2 may often be hard to use for analyzing pricing techniques. That is, some pricing tactics are easier to configure when the demand is represented by a continuous function. We therefore would like to fit a formula that would approximate the behavior of the consumer whose preferences are represented by the data given in Table 2.2. Figure 2.4 illustrates how a linear regression line can be fit to approximate the linear demand function. In fact, using the data in Table 2.2, regressing q on p yields the direct demand function

$$q = 12.6 - 0.36p, \quad \text{or, in an inverted form,} \quad p = 35 - 2.7777q. \quad (2.7)$$




Figure 2.4: Fitting a linear regression line.

Note that regressing q (the dependent variable) on p (an independent variable) is different from regressing p on q . In fact, regressing p on q yields $p =$

$34.63855 - 2.71084 q$, which is somewhat different from the inverse demand function (2.7) obtained from the estimation of the direct demand function. Unfortunately, there is no easy way to reconcile the two estimation methods. Given that our computations rely on frequent inversions of direct demand functions into inverse demand functions, and the other way around, in this chapter we always estimate the direct demand function equation, and then derive the inverse demand function from the estimated direct demand equation.

The estimated demand function (2.7) is drawn in Figure 2.4. The linear regression yielding the demand schedule (2.7) from the data given in Table 2.2 can be computed on basically any personal computer that runs a popular spreadsheet program, or a commonly used statistical package. Most widely used spreadsheet programs allow users to type in the observed data points then highlight the relevant rows or columns and obtain the intercept and slope of the demand function. In fact, some of these packages even draw the regression line with the data points, just as we did in Figure 2.4. Therefore, there is no need for us to go into detail about how to derive the formulas for obtaining the intercept and the slope. For the sake of completeness, we merely add a few remarks on the regression fitting technique. In general, regression is a statistical technique of fitting a functional relationship among economic variables by cataloging the observed variables, and then using well-known formulas to extract the exogenous parameters of the functional relationship. The most commonly used regression method is known as ordinary least squares (a method that minimizes the sum of squared errors). Let M be the number of observations in hand ($M = 5$ in the example given in Table 2.2). Then, define the mean price (or average price) and mean quantity demanded (or average quantity) by

$$\bar{p} \stackrel{\text{def}}{=} \frac{p_1 + p_2 + \cdots + p_M}{M} = \frac{\sum_{\ell=1}^M p_{\ell}}{M} \quad \text{and} \quad \bar{q} \stackrel{\text{def}}{=} \frac{\sum_{\ell=1}^M q_{\ell}}{M}. \quad (2.8)$$

Under the ordinary least squares method, the slope of the direct demand function $q(p) = \gamma - \delta p$ and its intercept with the horizontal q -axis are obtained from

$$\delta = -\frac{\sum_{\ell=1}^M (q_{\ell} - \bar{q})(p_{\ell} - \bar{p})}{\sum_{\ell=1}^M (p_{\ell} - \bar{p})^2} \quad \text{and} \quad \gamma = \bar{q} + \delta \bar{p}. \quad (2.9)$$

Again, it may not be practical to use (2.9) directly to compute that $\delta = 0.36$ and $\gamma = 12.6$ as given in (2.7), because linear regressions can be computed on most spreadsheet programs, statistical packages, and even some freely available Web pages.

2.3.3 Elasticity and revenue for linear demand

The elasticity and revenue functions defined in Section 2.1.3 for the general case can be easily derived for the linear demand function given by (2.6). Using calculus,

we obtain

$$e(q) = \frac{dq(p)}{dp} \frac{p}{q} = \left(-\frac{1}{\beta}\right) \left(\frac{\alpha - \beta q}{q}\right) = 1 - \frac{\alpha}{\beta q}. \quad (2.10)$$

Therefore, the demand has a unit elasticity at the consumption level $q = \alpha/(2\beta)$. Consequently, according to Definition 2.2, the demand is elastic at output levels $q < \alpha/(2\beta)$ and is inelastic for $q > \alpha/(2\beta)$. The elasticity regions for the linear demand case are illustrated in Figure 2.3.

The total revenue function associated with the linear demand function (2.6) is derived as follows:

$$x(q) = p(q)q = \alpha q - \beta q^2. \quad (2.11)$$

Hence, using calculus, the marginal revenue function is given by

$$\frac{dx(q)}{dq} = \frac{d[p(q)q]}{dq} = \alpha - 2\beta q. \quad (2.12)$$

The total and marginal revenue functions (2.11) and (2.12) are drawn in Figure 2.5.




Figure 2.5: Total and marginal revenue functions for linear demand.

Comparing the marginal revenue function (2.12) with the inverse demand function (2.6) reveals that the marginal revenue function has the same intercept α and twice the negative slope of the inverse demand function. This comparison can be easily visualized by comparing Figure 2.3 with Figure 2.5, which illustrates a linear marginal revenue function that intersects the quantity axis at $\alpha/2\beta$, which equals half of α/β (the quantity at the intersection with the demand function). As also shown in Figure 2.5, revenue is maximized when the firm sells $q = \alpha/(2\beta)$ units where the demand elasticity is $e(\alpha/2\beta) = -1$, or in absolute value $|e(\alpha/2\beta)| = 1$.

Finally, you probably have noticed already that the demand elasticity and the marginal revenue functions are related. That is, Figures 2.3 and 2.5 illustrate that $\Delta x/\Delta q = 0$ when $e(q) = 1$, and $\Delta x/\Delta q > 0$ whenever $|e(q)| > 1$. In fact, as it turns