PART II

PRICE-BASED RM

Chapter 5

DYNAMIC PRICING

5.1 Introduction and Overview

In this chapter, we look at settings in which prices rather than quantity controls are the primary variables used to manage demand. While the distinction between quantity and price controls is not always sharp (for instance, closing the availability of a discount class can be considered equivalent to raising the product's price to that of the next highest class), the techniques we look at here are distinguished by their explicit use of price as the control variable and their explicit modeling of demand as a price-dependent process.

In terms of business practice, varying prices is often the most natural mechanism for revenue management. In most retail and industrial trades, firms use various forms of dynamic pricing—including personalized pricing, markdowns, display and trade promotions, coupons, discounts, clearance sales, and auctions and price negotiations (request for proposals and request for quotes—RFP/RFQ processes)—to respond to market fluctuations and uncertainty in demand. Exactly how to make such price adjustments in a way that maximizes revenues (or profits, in the case where variable costs are involved) is the subject of this chapter.

Dynamic pricing is as old as commerce itself. Firms and individuals have always resorted to price adjustments (such as haggling at the bazaar) in an effort to sell their goods at a price that is as high as possible yet acceptable to customers. However, the last decade has witnessed an increased application of scientific methods and software systems for dynamic pricing, both in the estimation of demand functions and the optimization of pricing decisions.

5.1.1 Price versus Quantity-Based RM

Some industries use price-based RM (retailing), whereas others use quantity-based RM (airlines). Even in the same industry, firms may use a mixture of price- and quantity-based RM. For instance, many of the RM practices of the new low-cost airlines more closely resemble dynamic pricing than the quantity-based RM of the traditional carriers. What explains these differences?

It is hard to give a definitive answer, and indeed Chapter 8 is devoted to different theoretical explanations of RM practice. But in essence, it boils down to a question of the extent to which a firm is able to vary quantity or price in response to changes in market conditions. This ability, in turn, is determined by the commitments a firm makes (to price or quantity), its level of flexibility in supplying products or services, and the costs of making quantity or price changes.

Consider airlines, for example. While arguably less true today than in the past, airlines normally commit to prices for their various fare products in advance of taking bookings. This is due to advertising constraints (such as the desire to publish fares in print media and fare tariff books), distribution constraints, and a desire to simplify the task of managing prices. For these marketing and administrative reasons, most airlines advertise and price fare products on an aggregate origin-destination market level, for a number of flights over a given interval of time, and do not price on a departure-by-departure basis. This limits their ability to use price to manage the demand on any given departure, demand that varies considerably by flight and is quite uncertain at the time of the price posting. At the same time, the supply of the various classes is almost perfectly flexible between the products (subject to the capacity constraint of the flight), since all fare products sold in the same cabin of service share a homogeneous seat capacity. It is this combination of price commitments together with flexibility on the supply side that make quantity-based RM an attractive tactic in the airline industry. Hotels, cruise ships, and rental cars—other common quantity-based RM industries—share many of these same attributes.

In other cases, however, firms have more price flexibility than quantity flexibility. In apparel retailing, for example, firms commit to order quantities well in advance of a sales season—and may even commit to certain stocking levels in each store. Often, it is impossible (or very costly) to reorder stock or reallocate inventory from one store to another. At the same time, it is easier (though not costless) for most retailers to change prices, as this may require only changing signage and making data entries into a point-of-sale system. Online retailers in particular enjoy tremendous price flexibility because changing prices is almost costless.

Business-to-business sales are often conducted through a RFP/RFQ process, which allows firms to determine prices on a transaction-by-transaction basis. In all these situations, price-based RM is therefore a more natural practice. Of course, the context could dictate a different choice even in these industries. For example, if a retailer commits to advertised prices in different regional markets yet retains a centralized stock of products, it might then choose to manage demand by tactically allocating its supply to these different regions—a quantity-based RM approach.

However, given the choice between price- and quantity-based RM, one can argue that price-based RM is the preferred option. The argument is as follows (see Gallego and van Ryzin [199]). Quantity-based RM operates by rationing the quantity sold to different products or to different segments of customers. But rationing, by its very nature, involves reducing sales by *limiting* supply. If one has price flexibility, however, rather than reducing sales by *limiting supply,* we can reduce sales by *increasing price.* This achieves the same quantity-reducing function as rationing, but does it more profitably because by increasing price we both reduce sales *and* increase revenue at the same time. In short, price-based "rationing" is simply a more profitable way to limit sales than quantity-based rationing.

In practice, of course, firms rarely have the luxury of choosing price and quantity flexibility. Therefore, practical business constraints dictate which tactical response—price- or quantity-based RM (or a mixture of both)—is most appropriate in any given business context.

5.1.2 Industry Overview

To give a sense of the scope of activity in the area of dynamic pricing, we next review pricing innovations in a few industries.

5.1.2.1 Retailing

Retailers, especially in apparel and other seasonal-goods sectors, have been at the forefront in deploying science-based software for pricing, driven primarily by the importance of pricing decisions to retailers' profits. For example, Kmart alone wrote off $400 million due to markdowns in one quarter of 2001, resulting in a 40% decline in its net income [194].

Several software firms specializing in RM in retailing have recently emerged. Most of this software is currently oriented toward optimizing markdown decisions. Demand models fit to historical point-of-sale data together with data on available inventory serve as inputs to optimization models that recommend the timing and magnitude of markdown decisions.

Major retailers—including Gymboree, J. C. Penney, L. L. Bean, Liz Claiborne, Safeway, ShopKo, and Walgreen's—are experimenting with this new generation of software [194, 214, 270, 379]. Many have reported significant improvements in revenue from using pricing models and software. For example, ShopKo reported a 24% improvement in gross margins as a result of using its model-based pricing software [270] and other retailers report gains in gross margins of 5% to 15% [194]. Academic studies based on retail data have also documented significant improvements in revenues using model-based markdown recommendations [70, 247].

5.1.2.2 Manufacturing

Scientific approaches to pricing are gaining acceptance in the manufacturing sector as well. For example, Ford Motor Co. reported a high-profile implementation of pricing-software technology to support pricing and discounts for its products [135]. The project, started in 1995, focused on identifying features that customers were most willing to pay for and changing salesforce incentives to focus on profit margins rather than unit-sale volumes. Ford then applied pricing models developed by an outside consulting firm to optimize prices and dealer and customer incentives across its various product lines. In 1998, Ford reported that the first five U.S. sales regions using this new pricing approach collectively beat their profit targets by $1 billion, while the 13 that used their old methods fell short of their targets by about $250 million [135].

5.1.2.3 E-business

E-commerce has also had a strong influence on the practice of pricing [529]. Companies such as eBay and Priceline have demonstrated the viability of using innovative pricing mechanisms that leverage the capabilities of the Internet. E-tailers can discount and markdown on the fly based on customer loyalty and click-stream behavior. Since a large e-tailer like Amazon.com has to make a large number of such pricing decisions based on real-time information, automating decision making is a natural priority. The success of these e-commerce companies—inconsistent and volatile as it may appear at times—is at least partly responsible for the increased interest among traditional retailers in using more innovative approaches to pricing.

On the industrial side, e-commerce pricing has been influenced by the growth of business-to-business (B2B) exchanges and other innovations in using the Internet to gain trading efficiencies. While this sector too has had its ebbs and flows, it has produced an astounding variety of new pricing and trading mechanisms, some of which are use regularly for the sale products such as raw materials, generic commodity items and excess inventory. For example, Freemarkets has had significant success in providing software and service for industrial-procurement auctions, and as of this writing claims to have facilitated over $30 billion in trade since its founding in 1999. Covisint—an exchange jointly funded by Daimler-Chrysler, Ford Motor Company, and General Motors—while slow to develop, looks nevertheless to become a permanent feature of the auto-industry procurement market. Most infrastructure software for B2B exchanges—sold by firms such as Ariba, i2, IBM, and Commerce One—also has various forms of dynamic pricing capabilities built in.

For all these reasons, e-commerce has given price-based RM a significant boost in recent years.

5.1.3 Examples of Dynamic Pricing

We next examine three specific example of dynamic pricing and the qualitative factors driving price changes in each case.

5.1.3.1 Style-Goods Markdown Pricing

Retailers of style and seasonal goods use markdown pricing to clear excess inventory before the end of the season. This type of price-based RM is most prevalent in apparel, sporting goods, high-tech, and perishable-foods retailing. The main incentive for price reductions in such cases is that goods perish or have low salvage values once the sales season is over; hence, firms have an incentive to sell inventory while they can, even at a low price, rather than salvage it.

But apart from inventory considerations, there are other proposed explanations for markdown pricing. One explanation, proposed by Lazear [332] (see Examples 8.11 and 8.12) and investigated empirically in Pashigan [415] and Pashigan and Bowen [414], is that retailers are uncertain about which products will be popular with customers. Therefore, firms set high prices for all items initially. Products that are popular are the ones for which customers have high reservation prices, so these sell out at the high initial price. The firm then identifies the remaining items as low-reservation-price products and marks them down. In this explanation, markdown pricing serves as a form of demand learning.

A second explanation for markdowns is that customers who purchase early have higher willingness to pay, either because they can use the product for a full season (a bathing suit at the start of summer) or because there is some cache to being the first to own it (a new dress style or electronic gadget). Markdown pricing then serves as a segmentation mechanism to separate price-insensitive customers from those price-sensitive customers willing to defer consumption to get a lower price.

Warner and Birsky [554] give yet another explanation, with empirical evidence, for markdown pricing. On holidays and during peak-shopping periods (such as before Christmas), customers can search for the lowest prices more efficiently because they are actively engaged in search, making many shopping trips over a concentrated period of time. Even those customers who normally do not spend much time searching for the best price change their behavior during these peak shopping periods and become more vigilant. The result is that demand during peak periods is more price-sensitive and retailers respond by running "sales" during these periods.

5.1.3.2 Discount Airline Pricing

Not all dynamic pricing involves price reductions, however. As we mentioned earlier, discount airlines use primarily price-based RM, but with prices often going up over time. These airlines (some examples are easyJet and Ryanair in Europe and JetBlue in the U.S.) typically offer only one type of ticket on each flight, a non-refundable, one-way fare without advance-purchase restrictions. However, they offer these tickets at different prices for different flights, and moreover, during the booking period for each flight, vary prices dynamically based on capacity and demand for that specific departure. To quote from one practitioner of this type of dynamic pricing (Easyjet website, 2003):

> The way we structure our fares is based on supply and demand and prices usually increase as seats are sold on every flight. So, generally speaking, the earlier you book, the cheaper the fare will be. Sometimes, however, due to market forces our fares may be reduced further. Our booking system continually reviews bookings for all future flights and tries to predict how popular each flight is likely to be.

Figure 5.1 shows the evolution of prices for a particular European discount airline flight as a function of the number of weeks prior to departure. Note that prices are highest in the last few weeks prior to departure.

There are some fundamental differences between air travel and style- and seasonal-goods products that explain this increasing price pattern. For one, the value of air travel to customers does not necessarily go down as the deadline approaches. Conversely, the value of a ticket earlier on is lower for customers as customers multiply the value by the probability that they will indeed use the ticket (especially for a non-refundable ticket). Somewhat related to these points, additionally, although customers purchase tickets at different points of time, all customers consume

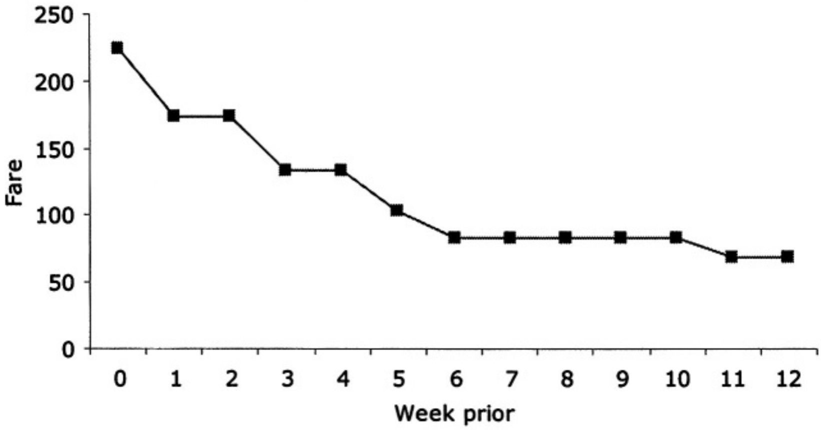

Figure 5.1. **Prices as a function of weeks prior to departure at a European low-cost discount air carrier.**

the product (fly the flight) at the same time. So two factors come into play. Customers who purchased early may get upset to see prices drop while they are still holding a reservation; indeed, many airlines give a price guarantee to refund the difference if there is a price drop (to encourage passengers to book early), making it costly for the firms to lower prices. And in the travel business, high-valuation high-uncertainty customers tend to purchase closer to the time of service. Hence, demand is less price-sensitive close to the time of service.

5.1.3.3 Consumer-Packaged Goods Promotions

In contrast to markdown and discount airline pricing, promotions are short-run, temporary price reductions. Promotions are the most common form of price-based RM in the consumer packaged-goods (CPG) industry (soap, diapers, coffee, yogurt, and so on).

The fact that customers purchase CPG products repeatedly has important implications for pricing and promotions. Specifically, customers are aware of past prices and past promotions, so running promotions too frequently may condition customers to view the brand as a frequently discounted product, cutting into brand equity in the long run. Because customers are aware of past prices, promotions impact their subjective "reference price"—or sense of the "fair" price—for products. And customers may *stockpile* products, so short-run increases in demand due to promotions may come at the expense of reduced future demand.

The institutional structure of promotions is also more complicated. There are three parties involved—manufacturers, retailers, and end customers. Promotions are run either by a manufacturer as discounts to retailers (trade promotions), which may or may not be passed on to the customers by the retailers (retailer pass-thru), or by retailers (retail promotions or consumer promotions). In some forms of promotion (e.g., mail-in coupons) manufacturers gives a discount directly to the end customer.

The motivations of the manufacturer and the retailer are different as well. While a manufacturer is interested in increasing sales or profits for its brand, retailers are interested in overall sales or profits for a category constituting multiple brands from multiple manufacturers. For a retailer, discounting a particular brand may increase sales for that brand but dilute overall category profits as customers switch from high-margin brands to the discounted brand. So in designing optimal promotions structures, one has to consider complex incentive compatibility constraints.

5.1.4 Modeling Dynamic Price-Sensitive Demand

Any dynamic-pricing model requires a model of how demand—either individual or aggregate—responds to changes in price. The basic theory of consumer choice and the resulting market-response models are covered in Chapter 7. We draw on these models in this chapter.

However, in dynamic-pricing problems some additional factors must be considered. The first concerns how individual customers behave over time—what factors influence their purchase decisions and how sophisticated their decision-making process is, and so on. The second concerns the state of market conditions—specifically the level of competition and the size of the customer population. We next look at each of these assumptions qualitatively.

5.1.4.1 Myopic- Versus Strategic-Customer Models

One important demand-modeling assumption concerns the level of sophistication of customers. Most of the models we consider in this chapter assume *myopic customers*—those who buy as soon as the offered price is less than their willingness to pay. Myopic customers do not adopt complex buying strategies, such as refusing to buy in the hope of lower prices in the future. They simply buy the first time the price drops below their willingness to pay. Models that incorporate *strategic customers,* in contrast, allow for the fact that customers will optimize their own purchase behavior in response to the pricing strategies of the firms.

Of course, the strategic-customer model is more realistic. However, such a demand model makes the pricing problem essentially a strategic game between the customers and the firm, and this significantly complicates the estimation and analysis of optimal pricing strategies—often making the problem intractable. In contrast, the myopic-customer model is much more tractable and hence is more widely used. The issue in practice is really a matter of how "bad" the myopic assumption is in any given context. In many situations, customers are sufficiently spontaneous in making decisions that one can ignore their strategic behavior. Moreover, customers often do not have the sufficient time or information to behave very strategically. However, the more expensive and durable the purchase, the more important it becomes to model strategic customer behavior (for example, automobile buyers waiting to purchase at the end of a model year).

One common defense of the myopic assumption is the following. The forecasting models that use observations of past customer behavior in a sense reflects the effects of our customers' strategic behavior. For example, if the customers who are most price-sensitive tend to adopt a strategy of postponing their purchases until end-of-season clearance sales, then the estimated price sensitivity in these later periods will tend to appear much higher than in earlier periods. Therefore, even though we do not model the strategic behavior directly, our forecasting models indirectly capture the correct price response.

This view is plausible if the pricing strategies obtained from a model are roughly similar to past policies, so that they can be viewed as "perturbations" or "fine tuning" of a historical pricing strategy—a strategy that customers have already factored into their behavior. On the other hand, if optimized pricing recommendations are radically different in structure from past pricing strategies, then it is reasonable to expect that customers will adjust their buying strategies in response. If this happens, the predictions of myopic models that are fit to historical data may be very bad indeed.

Yet even when the myopic approach works (in the sense of correctly predicting price responses), it runs the risk of reinforcing "bad equilibrium" pricing strategies. For example, a myopic model fit to past data may reconfirm the "optimality" of lowering prices significantly at the end of a sales season or running periodic holiday sales because it estimates, based on historical data, that demand is especially price-sensitive in these periods. But this price sensitivity may be due to the fact that customers have learned not to buy at other times, because they know prices will be cut at the end of the season or during holidays. If the firm was to adopt a constant price strategy—and customers were convinced that the firm was sticking to this strategy—then the observed price sensitivity might shift. The resulting equilibrium might be more profitable, but it is one that the firm would not discover using a myopic-customer model.

Despite these limitations and potential pitfalls of the myopic model, it is practical, is widely used, and provides useful insight into dynamic pricing. We therefore focus on the myopic case for the most part in this chapter. However, we consider strategic customers in Section 5.5.2 below and in considerably more depth in Chapter 6, where we look at auctions, the analysis of which is entirely based on strategic-customer models.

5.1.4.2 Infinite- Versus Finite-Population Models

Another important assumption in demand modeling is whether the population of potential customers is finite or infinite. Of course, in reality, every population of customers is finite; the question is really a matter of whether the number and type of customers that have already bought changes one's estimate of the number or type of future customers.

In an infinite-population model, we assume that we are sampling *with replacement* when observing customers. As a result, the distribution of the number of customers and the distribution of their willingness to pay is not affected by the past history of observed demand. This is often termed the *nondurable-goods assumption* in economics because we can view this as a case where customers immediately consume their purchase and then reenter the population of potential customers (say, for a can of Coke). This assumption is convenient analytically because one does not need to retain the history of demand (or a suitable sufficient statistic) as a state variable in a pricing-optimization problem.

The finite-population model assumes a random process *without replacement.* That is, there are a finite (possibly random) number of customers with heterogeneous willingness to pay values. If one of the customers in the population purchases, the customer is removed from the population of potential customers, and therefore future purchases only occur from the remaining customers. This is termed the *durable-goods assumption* in economics because we can consider it as a case where the good being purchased is consumed over a long period of time (for example, an automobile) and hence once a customer purchases, he effectively removes himself from the population of potential customers.

For example, suppose we assume a price $p(t)$ is offered in period t and all customers who value the item at more than $p(t)$ purchase in period t (myopic behavior). Then, under a finite-population model, we know that after period t, the remaining customers all have valuations less than $p(t)$. In particular, the future distribution of willingness to pay is conditioned on the values being less than $p(t)$. As a result, in formulating a dynamic-pricing problem, we have to keep track of past pricing decisions and their effect on the residual population of customers.

Which of these models is most appropriate depends on the context. While often the infinite-population model is used simply because it is easier to deal with analytically, the key factors in choosing one model over the other are the number of potential customers relative to the number that actually buy and the type of good (durable versus nondurable). Specifically, the infinite-population model is a reasonable approximation when there is a large population of potential customers and the firm's demand represents a relatively small fraction of this population because in such cases the impact of the firm's past sales on the number of customers and the distribution of their valuations is negligible. It is also reasonable for consumable goods. However, if the firm's demand represents a large fraction of the potential pool of customers or if the product is a durable good, then past sales will have a more significant impact on the statistics of future demand, and the finite-population assumption is more appropriate.

Qualitatively, the two models lead to quite different pricing policies. Most notably, finite-population models typically lead to *price skimming* as an optimal strategy, in which prices are lowered over time in such a way that high-valuation customers pay higher prices earlier while low-valuation customers pay lower prices in later periods. Effectively, this creates a form of second-degree price discrimination, segmenting customers with different values for the good and charging differential prices over time. In infinite-population models, there is no such price-skimming incentive. Provided the distribution of customer valuations does not shift over time, the same price that yields a high revenue in one period will yield a high revenue in later periods, and thus a firm has no incentive to deviate from this revenue-maximizing price.

5.1.4.3 Monopoly, Oligopoly, and Perfect-Competition Models

Another key assumption in dynamic-pricing models concerns the level of competition the firm faces. Many pricing models used in RM practice are *monopoly models,* in which the demand a firm faces is assumed to depend only on its own price and not on the price of its competitors. Thus, the model does not explicitly consider the competitive reaction to a price change. Again, one makes this assumption primarily for tractability and is not always realistic.

As with the myopic-customer model, the monopoly model can be partly justified on empirical grounds—namely, that an observed historical price response has embedded in it the effects of competitors' responses to the firm's pricing strategy. So for instance, if a firm decides to lower its price, the firm's competitors might respond by lowering their prices. With market prices lower, the firm and its competitors see an increase in demand. The observed increase in demand is then measured empirically and treated as the "monopoly" demand response to the firm's price change in a dynamic-pricing model—even though competitive effects are at work.

Again, while such a view is pragmatic and reflects the conventional wisdom behind the pricing models used in practice, there are some dangers inherent in it, paralleling those of the myopic-customer model. The price-sensitivity estimates may prove wrong if the optimized strategy deviates significantly from past strategies because then the resulting competitive response may be quite different from the historical response. Also, the practice runs the risk of reinforcing "bad" equilibrium responses. Despite these risks, monopoly models have still proved to be valuable for decision support.

It is worth noting that oligopoly models, in which the equilibrium-price response of competitors is explicitly modeled and computed, also have their pitfalls. Most notably, the assumption that firms behave rationally (or quasi-rationally, if heuristics are used in place of optimal strategies) may result in a poor predictor of their actual price response. These potential modeling errors together with the increased complexity of analyzing oligopoly models—and the difficulty in collecting competitor data to estimate the models accurately—has made them less popular in practice. Shugan [468] provides a good summary of this point of view; he notes that "the strong approximating assumption of no competitive response is sometimes better than the approximating assumption of pre-existing optimal behavior." However, properly designed and validated, oligopoly models can provide valuable insights on issues of pricing strategy.

Finally, one can also consider perfectly competitive models—in which many competing firms supply an identical commodity. As described in Section 8.2, the output of each firm is assumed to be small relative the market size, and this, combined with the fact that each firm is offering identical commodities, means that a firm cannot influence market prices.[1] Therefore, each firm is essentially a *price taker*—able to sell as much as it wants at the prevailing market price but unable to sell anything at higher prices. Despite the importance of perfect-competition models in economic theory, the assumption that firms have no pricing power means that the results are not that useful for price-based RM. Nevertheless, they do play a role in quantity-based RM. For example, one can interpret the capacity-control models of Chapters 2 and 3 as stemming from competitive, price-taking models; firms take the price for their various products as given (set by competitive market forces), and control only the quantity they supply (the availability or allocation) at these competitive prices. As our focus in the chapter is on price-based RM, we do not consider this model of competition further in this chapter.

5.2 Single-Product Dynamic Pricing Without Replenishment

The first problem we look at is dynamic pricing of a single product over a finite sales horizon given a fixed inventory at the start of the sales horizon. We assume that the firm is a monopolist, customers are myopic, and there is no replenishment of inventory.

The models are representative of the type used in style and seasonal goods retail RM. For such retailers, production and ordering cycles are typically much larger than the sales season, and the main challenge is to determine the price path of a particular style at a particular store location, given a fixed set of inventory at the beginning of the season.

At one level, such models are simplistic: they consider only a single product in isolation and assume customers are myopic, and therefore demand is a function solely of time and the current price (although other factors such as inventory depletion are sometimes included). They therefore ignore competition, the impact of substitution, and the possible strategic behavior of customers over time. Despite these simplifications, the models provide good rough-cut approximations and are useful in practice. In addition, by decomposing the problem and treating products independently, it is possible to solve such models efficiently even when there are hundreds of thousands of product-location combinations. Finally, even with the simplifying assumptions, the analysis can

[1] This is in contrast to the Cournot model of quantity competition discussed in Section 8.4, in which there are only a small number of firms whose quantity decisions do affect the market price. Roughly speaking, Cournot competition approaches perfect competition as the number of firms in the industry tends to infinity.

still become complex if we allow stochastic demand and put constraints on prices.

Since we consider only a single product, there is a single (scalar) price decision at each time t, denoted $p(t)$, which induces a unique (scalar) demand rate $d(t,p)$. The set of allowable prices is denoted Ω_p, and Ω_d denotes the set of achievable demand rates. We assume that these functions satisfy the regularity conditions in Assumptions 7.1, 7.2, and 7.3 unless otherwise specified. These include several regularity properties, which we summarize here:

- The demand functions are continuously differentiable and strictly decreasing, $d'(t,p) < 0$, on Ω_p. Hence, they have an inverse, denoted $d(t,p)$.

- The demand functions are bounded above and below and tend to zero for sufficiently high prices—namely,

$$\inf_{p \in \Omega_p} d(t,p) = 0.$$

- The revenue functions $r(t,p) = pd(t,p)$ (equivalently $r(t,d) = dp(t,d)$) are finite for all $p \in \Omega_p$ and have a finite maximizer interior to Ω_p.

- The marginal revenue as a function of demand, d, defined by

$$J(t,d) \equiv \frac{\partial}{\partial d} r(t,d) = p(t,d) + dp'(t,d),$$

  is strictly decreasing in d. (Assumption 7.2)

Readers who are not familiar with demand functions are encouraged to review Section 7.3 for more discussion of these and other related properties of demand functions. As discussed in Section 7.3, the demand function can also be expressed as $d(t,p) = N_t(1 - F(t,p))$, where N_t is the market-size parameter and $F(t,p)$ is the fraction of the market with willingness to pay less than p. We let $x(t)$ denote the inventory at time $t = 1, \ldots, T$, where T is the number of periods in the sale horizon. The initial inventory is $x(0) = C$.

5.2.1 Deterministic Models

The simplest deterministic pricing model is formulated in discrete time as follows. Given an initial inventory $x(0) = C$, select a sequence of prices $p(t)$ (inducing demand rates of $d(t, p(t))$) that maximize total revenues. Formulating the problem in terms of the demand rates $d(t)$, the optimal rates $d^*(t)$ must solve

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} r(t, d(t)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} d(t) \le C \\ & d(t) \ge 0. \end{aligned} \tag{5.1}$$

Let π^* be the Lagrange multiplier on the inventory constraint, and recall that $J(t,d) = \frac{\partial}{\partial d} r(t,d)$ denotes the marginal revenue. Then the first-order necessary conditions for the optimal rates $d^*(t)$ and multiplier π^* are

$$J(t, d^*(t)) = \pi^*, \tag{5.2}$$

subject to the complementary slackness condition

$$\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0 \tag{5.3}$$

and the multiplier nonnegativity constraint $\pi^* \ge 0$. Under Assumption 7.2, $J(t,d)$ is decreasing in d and so $r(t,d)$ is concave; hence, these conditions are also sufficient.

The optimality conditions are quite intuitive. The Lagrange multiplier π^* has the interpretation as the marginal opportunity cost of capacity. The condition $J(t, d^*(t)) = \pi^*$ says that the marginal revenue should equal the marginal opportunity cost of capacity in each period. This makes sense because if marginal revenues and costs are not balanced, we can increase revenues by reallocating sales (by adjusting prices) from a period of low marginal revenue to a period of higher marginal revenue. Finally, the complementary slackness condition says that the opportunity cost cannot be positive if there is an excess of stock. If the opportunity cost is zero ($\pi^* = 0$), then if we maximize revenue without a constraint in every period (pricing to the point where marginal revenue is zero), we will still not exhaust the supply. This means it can be optimal—even in the absence of any costs for capacity—not to sell all the available supply.

Note that this problem is essentially equivalent to the problem of optimal third-degree price discrimination (see Section 8.3.3.2) if we consider customers in each period t to be different segments who are offered discriminatory prices $p(t)$. Another way of viewing the above argument is that the firm, faced with a capacity constraint, decides how much to sell in each period, and its optimal allocation of capacity occurs when the marginal revenue in all the periods are the same. The following example illustrates the idea:

Example 5.1 Consider a two-period selling horizon, where during the first period demand is given by $d_1 = -p_1 + 100$ and in period 2 demand is given by $d_2 = -2p_2 + 120$. (Customers in the second period are more price-sensitive than those in the first period.) Purchase behavior is assumed to be myopic. Considered separately, the revenue-maximizing price for the first period (maximizing $r_1 = p_1(-p_1 + 100)$) is given by $p_1^* = 50$ and $d_1^* = 50$, and in the second period by $p_2^* = 30, d_2^* = 60$ (maximizing $r_2 = p_2(-2p_2 + 120)$).

Intertemporal effects come into play if the firm has only a limited number of items to sell (less than 50+60). Suppose the firm's capacity is 40. How should it divide the sale between the two periods?

Note that here, $J(1, d_1) = -2d_1 + 100$ and $J(2, d_2) = -d_2 + 60$. Consider the table of marginal values, Table 5.1, at various allocations and the corresponding revenues. The total revenue is maximized at the point where the marginal values for the two periods are approximately the same (when $d_1 = 27, d_2 = 13$), conforming to our intuition; if they were not equal, the firm would reallocate capacity to the higher marginal-value period.

Table 5.1. **Allocations of capacity between periods 1 and 2 and the marginal values and total revenue.**

d_1	d_2	$J(1, d_1)$	$J(2, d_2)$	r
22	18	56	42	2634
23	17	54	43	2646.5
24	16	52	44	2656
25	15	50	45	2662.5
26	14	48	46	2666
27	**13**	**46**	**47**	**2666.5**
28	12	44	48	2664
29	11	42	49	2658.5
30	10	40	50	2650
31	9	38	51	2638.5
32	8	36	52	2624
33	7	34	53	2606.5

To see qualitatively how prices will change over time, we can write the optimality condition (5.2) as

$$\frac{p^*(t) - \pi^*}{p^*(t)} = \frac{1}{|\epsilon(t, p^*)|},$$

where $\epsilon(t, p)$ is the elasticity of demand in period t, defined by

$$\epsilon(t, p) \equiv \frac{p}{d(t, p)} \frac{\partial d(t, p)}{\partial p}.$$

See Section 7.3.1.3 for a further discussion of price elasticity. Thus, more elastic demand in period t implies a lower optimal price $p^*(t)$.

For example, if customers that buy toward the end of the sales horizon are more price-sensitive than those that buy early, then optimal prices will decline over time. If customers early on are price-sensitive, and those buying later are less price-sensitive, then optimal prices will increase over time. This observation offers one explanation for why in some industries (such as apparel retailing) prices tend to decline over time, while in others (such as airlines) prices increase over time. Chapter 8 provides additional explanations for intertemporal price patterns.

5.2.1.1 Computational Approaches

Problem (5.1) is a rather simple nonlinear program to solve. Each value π implies a value $d^*(t)$ by (5.2). If the value π is too low, these demand rates will be too high, and the constraint $\sum_{t=1}^{T} d^*(t) \leq C$ will be violated. If π is too high, total demand will not exhaust supply, and (5.3) will be violated. Of course, if $\pi = 0$ results in a total demand that is less than C, then this is the optimal dual value. Using these rules, it is straightforward to derive a search procedure to find the optimal π^*.

Another computational approach is to apply a greedy allocation algorithm, based on the observation that the marginal revenues in all periods are equal at optimality. Specifically, discretize the capacity C into M units of size δ each, so that $C = M\delta$. The greedy algorithm then proceeds by allocating demand in discrete amounts δ so as to equalize the marginal revenue:

STEP 0 (Initialize): Initialize solution $d(t) = 0$, $t = 1, \ldots, T$. Initialize counter $k = 0$.

STEP 1 (Evaluate marginal revenues): IF $\max_t\{J(t, d(t))\} > 0$, THEN DO:
Increment the demand of this highest marginal revenue period t^*:

$$d(t^*) \leftarrow d(t^*) + \delta.$$

ELSE, IF $\max\{J(t, d(t))\} \leq 0$ STOP (Current solution optimal).

STEP 2 (Check capacity constraint and repeat): IF $k = M$, STOP;
ELSE $k \leftarrow k + 1$ and GOTO STEP 1.

This algorithm takes $O(M \log T)$ time and is quite simple to program. Provided the marginal revenue is decreasing in each period, this greedy procedure produces an optimal (discretized) solution. (See Federgruen and Groenvelt [182].) The following example illustrates the algorithm:

Example 5.2 Consider a two-period problem with inverse-demand functions $p(1,d_1) = 10 - d_1$ and $p(2,d_2) = 10 - 2d_2$. The corresponding marginal revenue functions are

$$J(1,d_1) = 10 - 2d_1 \quad \text{and} \quad J(2,d_2) = 10 - 4d_2.$$

There are $C = 6$ units of capacity and we let the increment $\delta = 1$. The algorithm then proceeds as shown in Table 5.2. At the start, both periods have the same marginal revenue of 10. We break ties arbitrarily by assigning demand to period 1, so we assign the first unit to period 1. After this assignment, the marginal revenue in period 1 drops to 8 while the marginal revenue in period 2 is still 10, so we assign the next unit to period 2. The process continues as shown in Table 5.2, assigning units to the period with highest marginal revenue until all six units are used up. The algorithm terminates with $d_1^* = 4$ and $d_2^* = 2$; all six units are allocated and the marginal revenues are equalized $J(1,d_1^*) = J(2,d_2^*) = 2$.

Table 5.2. Example of the marginal-allocation algorithm.

k	d_1	d_2	$J(1,d_1)$	$J(2,d_2)$
0	0	0	10	10
1	1	0	8	10
2	1	1	8	6
3	2	1	6	6
4	3	1	4	6
5	3	2	4	2
6	4	2	2	2

5.2.1.2 Solution in the Time-Homogenous Case

A few additional observations can be made from this model when demand is time-homogenous, i.e., $d(t,p) = d(p)$ for all t. In this case, the optimal price p^*, given by $J(p^*) = \pi^*$, is the same in each period. This shows that prices fluctuate from period to period in the deterministic model (5.1) only as a result of changes in the demand function over time.

The optimal static price will either be the price that causes the supply to run out exactly at the end of the horizon (if $\pi^* > 0$) or the price at which the unconstrained revenue is maximized (if $\pi^* = 0$—that is, the revenue-maximizing price). Specifically, let p^0 be defined to be the value at which marginal revenue is zero, $J(p^0) = 0$, called the *revenue-maximizing price.* Let $\bar{p}$ denote the value at which $d(\bar{p}) = C/T$, which we call the *stock-clearing price.* Then

$$p^* = \max\{p^0, \bar{p}\},$$

so the optimal solution reduces to using the maximum of the revenue-maximizing price and the stock-clearing price. Simply, one cannot do better than pricing at p^0 at all times. If $Td(p^0) \leq C$, this price is feasible because demand is less than supply. If not, $Td(p^0) > C$, and demand at p^0 exceeds supply. We then have to raise the price, and $\bar{p} > p^0$ is the highest price at which we can still manage to sell all C units.

5.2.1.3 Discrete Prices

Often, in practice, we would like to choose prices from a discrete set. For example, prices close to convenient whole dollar amounts (such as \$24.99 or \$149.99), or fixed percentage markdowns (such as 25% off or 50% off) are often used because they are familiar to customers and easy to understand. In such cases, it may be desirable as a matter of policy to constrain prices to a finite set of k discrete price points, so that $p(t) \in \Omega_p$, where $\Omega_p = \{p_1, \ldots, p_k\}$. Equivalently, the sales rate $d(t)$ is constrained to a discrete set $d(t) \in \Omega_d(t)$ (time-varying in this case if the demand function is time-varying), where $\Omega_d(t) = \{d_1(t), \ldots, d_k(t)\}$, and $d_i(t) = d(t, p_i)$ denotes the sales rate at time t when using the price p_i.

The discreteness of the prices imposes technical complications when attempting to solve the dynamic pricing problem (5.1) because the problem is no longer continuous or convex. However, one can overcome this difficulty by relaxing the problem to allow the use of convex combinations of the discrete prices (or demand rates). In most periods, the optimal solution will be to use only one of the discrete prices; in the remaining periods, the solution has the interpretation of allocating a fraction of time to each of several prices.

To see this, define a vector of new variables $\alpha_i(t)$ for each t, $\alpha(t) = (\alpha_1(t), \ldots, \alpha_k(t))$, which represent convex weights: they are nonnegative and sum to one. Next, in each period replace the variable $d(t)$ with the convex combination

$$d(t) = \sum_{i=1}^{k} \alpha_i(t) d_i(t),$$

and replace the constraint $d(t) \in \Omega_d(t)$ with the constraint

$$\alpha(t) \in W \equiv \{\alpha \in \Re^k : \sum_{i=1}^{k} \alpha_i = 1, \alpha \geq 0\}.$$

The optimization problem is then

$$\max_{\boldsymbol{\alpha}(t)\in W} \quad \sum_{t=1}^{T}\sum_{i=1}^{k} r_i(t)\alpha_i(t) \tag{5.4}$$

$$\text{s.t.} \quad \sum_{t=1}^{T}\sum_{i=1}^{k} \alpha_i(t)d_i(t) \le C,$$

where $r_i(t) = p_i d_i(t)$ is the revenue rate at price p_i. This is a linear program in the variables $\boldsymbol{\alpha}(t)$, so it is easy to solve numerically.

To relate the solution to the unconstrained price case, introduce a dual variable π^* on the capacity constraint as before. The optimal solution $\boldsymbol{\alpha}^*(t)$ in each period is then characterized by solving

$$\max_{\boldsymbol{\alpha}(t)\in W} \left\{ \sum_{i=1}^{k} \alpha_i(t)(r_i(t) - \pi^* d_i(t)) \right\}, \tag{5.5}$$

where $\pi^* \ge 0$ and $\boldsymbol{\alpha}^*(t)$ are convex weights satisfying the complementary slackness condition

$$\pi^* \left(\sum_{t=1}^{T}\sum_{i=1}^{k} \alpha_i^*(t)d_i(t) - C \right) = 0. \tag{5.6}$$

Since the objective function of (5.5) is linear in $\alpha(t)$, if there is a unique index i^* for which $r_{i^*}(t) - \pi^* d_{i^*}(t)$ is greatest, then the optimal solution is simply $\alpha_{i^*}(t) = 1$, which corresponds to using the discrete price p_{i^*}. If there is more than one such value i^*, then there will be multiple solutions to (5.5), and determining which is optimal can be resolved by appealing to the complementary slackness condition (5.6). Of course, such a choice could result in a fractional solution in which $\alpha_i(t) > 0$ for two or more values i. However, this can be interpreted as saying that we should use the price i for a fraction $\alpha_i(t)$ of period t. Hence, the solution of (5.4) can be converted in practice into a discrete-price recommendation. The following example illustrates the calculation.

Example 5.3 Consider a two week selling season in which there is a linear-demand function $d(1,p) = 100 - p$ in week 1 and a demand function $d(2,p) = 100 - 1.4p$ in week 2. The firm is constrained to offer prices in the set {40, 50, 70}. The demand and revenues are then given in Table 5.3. Solving the linear program (5.4) for different value of the initial inventory *C*, we obtain the results in Table 5.4. For example, with an initial inventory of 50, the solution has $\alpha_{70}(1)$ and $\alpha_{50}(2) = 0.64$ and $\alpha_{70}(2) = 0.36$. This corresponds to pricing at \$70 for all of week 1 and 36% of week 2 then lowering the price to \$50 for the remainder of week 2. Similarly, when the initial inventory is 70, the solution calls for pricing at \$70 for half of week 1 and then lowering the price to \$50 for the remainder of the selling season. At very high levels of inventory (110 and 120), it is optimal to charge a price of \$50 in week 1 and a price of \$40 in week 2.

Table 5.3. **Example of discrete prices and revenues.**

p	$d(1,p)$	$r(1,p)$	$d(2,p)$	$r(2,p)$
20	60	2,400	44	1,760
30	50	2,500	30	1,500
50	30	2,100	2	140

Table 5.4. **Solution of a linear program for the discrete-price example.**

Inv. (C)	$\alpha_{40}(1)$	$\alpha_{50}(1)$	$\alpha_{70}(1)$	$\alpha_{40}(2)$	$\alpha_{50}(2)$	$\alpha_{70}(2)$	*Total Sold*
50	0.00	0.00	1.00	0.00	0.64	0.36	50
60	0.00	0.00	1.00	0.00	1.00	0.00	60
70	0.00	0.50	0.50	0.00	1.00	0.00	70
80	0.00	1.00	0.00	0.00	1.00	0.00	80
90	0.00	1.00	0.00	0.71	0.29	0.00	90
100	0.00	1.00	0.00	1.00	0.00	0.00	94
110	0.00	1.00	0.00	1.00	0.00	0.00	94
120	0.00	1.00	0.00	1.00	0.00	0.00	94

5.2.1.4 Maximum Concave Envelope

In the discrete-price problem, certain discrete prices may never be optimal to use and can in fact be eliminated from the problem. Indeed, suppose that for a give price p_j there exist convex weights $\alpha_i(t)$ such that

$$\begin{aligned} \sum_{i=1}^{k} \alpha_i(t) r_i(t) &> r_j(t) \\ \sum_{i=1}^{k} \alpha_i(t) d_i(t) &\leq d_j(t). \end{aligned} \tag{5.7}$$

Then the price p_j is never optimal at time t. Intuitively, this follows since a convex combination of other prices produces strictly higher revenue yet consumes no more capacity than using p_j. This is in fact the same notion of efficiency described in Section 2.6.2 for the discrete-choice model of demand in the single-resource capacity-control problem. All such inefficient prices j can be eliminated from consideration at time t. The remaining efficient prices define the *maximum concave envelope* of the pairs of values $\{(d_i(t), r_i(t)) : i = 1, \ldots, k\}$ as shown in Figure 5.2.

5.2.1.5 Inventory-Depletion Effect

Another practical factor affecting dynamic pricing in many retailing contexts is the adverse effects of low inventory levels. This is sometimes

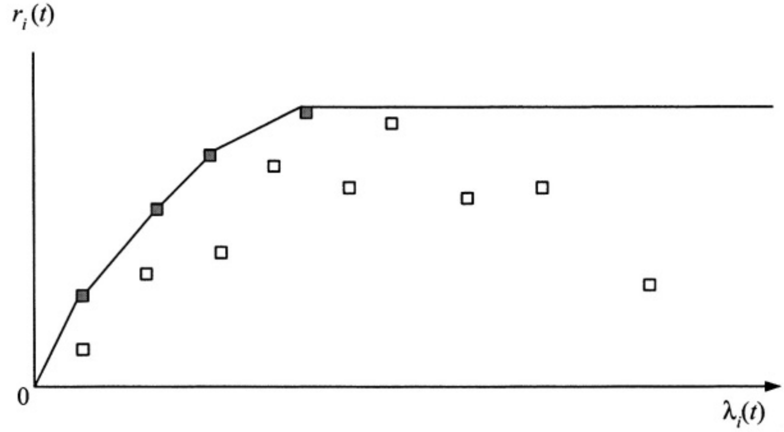

Figure 5.2. The maximum concave envelope produced by discrete prices (scatter plot of pairs $(r_i(t), d_i(t))$: efficient points are shaded).

referred to in retailing as a *broken-assortment effect.* For example, if the inventory-pricing model is applied at an aggregate item level, where an item contains several SKUs—such as color-size combinations in apparel retailing—then when inventories run low, certain SKUs may be out of stock even though there is a positive inventory for the item as a whole (for example, if a color or size runs out). The resulting reduction in alternatives naturally reduces the sales rate at any given price. Indeed, empirical studies have confirmed a positive correlation between inventory levels and sales rates [65].

These inventory-depletion effects can be modeled by making the demand rate a function of inventory as well as of price and time, so that the demand rate becomes $d(t, p(t), x(t))$. We can use a variety of functional forms to represent this inventory-depletion effect. For example, one proposed model is the following multiplicative form [480]:

$$\hat{d}(t, x(t)) = d(t)g(x(t)), \tag{5.8}$$

where $g(\cdot)$ is a depletion-effect term. We will call $d(t)$ the *unadjusted sales rate* (the rate of sales if inventory were unlimited) and $\hat{d}(t, x(t))$ the *adjusted sales rate* (the rate adjusted for inventory-depletion effects). One choice for g is

$$g(x) = 1 - \gamma \max\{0, 1 - x/x_0\},$$

where x_0 is the minimum *full-fixture inventory* and $0 \leq \gamma \leq 1$ is a sensitivity parameter. Both x_0 and γ can be estimated from historical data. Note that $g(x)$ is concave in x.

Another possible form is

$$g(x) = e^{-\gamma \max\{0, 1 - x/x_0\}},$$

where γ and x_0 have the same interpretation (see Smith and Achabal [480]).

For this model with inventory depletion one must keep track of the inventory at each time t in the optimization problem. For example, assuming the multiplicative inventory-depletion model of (5.8) and formulating the problem in terms of the unadjusted sales rate $d(t)$, the inventory evolves according to the state equation

$$x(t+1) = x(t) - d(t)g(x(t)),$$

and the revenue-maximization problem can be formulated as

$$\begin{aligned} \max_{d(t)\geq 0} \quad & \sum_{t=1}^{T} r(t, d(t))g(x(t)) && (5.9) \\ \text{s.t.} \quad & x(t+1) = x(t) - d(t)g(x(t)), \quad t = 1, \ldots, T \\ & x(T) \geq 0, \\ & x(0) = C, \end{aligned}$$

where $r(t, d(t)) = p(t, d(t))d(t)$ is the unadjusted revenue-rate function.

While somewhat more complex than the case without inventory-depletions effects, this is still a relatively simple nonlinear program to solve because the objective function is separable and the constraints are linear. (The objective function, however, is not necessarily jointly concave even if $r(t, d(t))$ and $g(x)$ are both concave.)

One qualitative impact of this inventory-depletion phenomenon is that optimal prices may decline over time even though the unadjusted revenue-rate function is time-invariant. (Recall that in the problem without inventory-depletion effects, a time-invariant revenue-rate function implied a time-invariant optimal price.) For example, Smith and Achabal [480] show, for the continuous-time version of this model, that if the unadjusted revenue-rate function is constant and the inventory-depletion effect is multiplicative, then optimal prices decline over time in such a way that the adjusted sales rate $g(x(t))d(t)$ is constant; that is, as inventory depletion reduces demand, the optimal prices fall to exactly compensate for the drop in sales due to inventory depletion.

5.2.1.6 A Retail Markdown Application

Here we look at the study of Heching et al. [247] that applied deterministic pricing models of the sort discussed above to analyze markdown

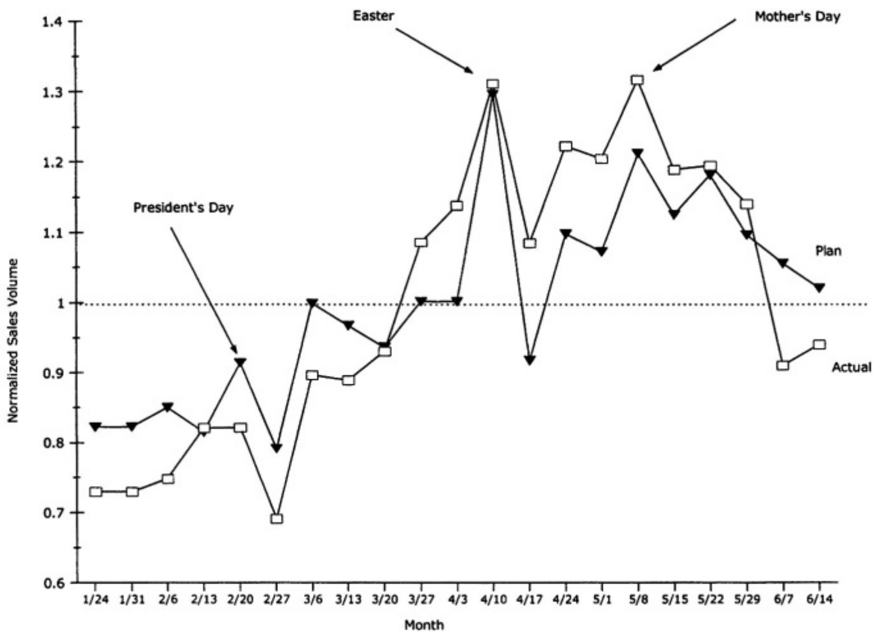

Figure 5.3. Sales volume for spring 1993: actual and planned.

pricing at an apparel retailer. The study provides an example of how such models can be applied and gives an indication of their potential impact.

The firm studied by Heching et al. [247] was a women's specialty apparel retailer with approximately 50 stores in the United States. The firm sold primarily its own private-label products and generally stocked items once at the beginning of the season. It then used markdowns to clear slow-selling merchandize.

The data set included the majority of the firm's sales over the spring 1993 season, spanning 184 styles in 25 groups (a collection of related styles). Weekly sales were obtained for each style sold during this period. Of all the styles in the data set, the firm took markdowns on 60 (the *markdown styles*). The remaining 124 styles had no price changes. While representing only one-third of all styles, markdown styles accounted for 42% of gross sales revenue.

There were strong seasonalities in sales due to major holidays and traditional shopping seasons as shown in Figure 5.3. Total weekly sales ranged from roughly 70% of average in slow weeks to 130% of average in the strongest weeks. The data also indicated that demand was indeed price-sensitive. After adjusting for seasonalities, the conditional probability of a sales increase, given a markdown, was 85%, while the unconditional probability of an increase was only 38%.

Sales of nearly all styles also tended to decline over time. Figure 5.4(i) plots weekly sales for one style that maintained the same price over its entire 12-week selling season. The weekly sales figures have been adjusted to eliminate any seasonality that can be attributed to traditional shopping seasons. Figure 5.4(ii) plots weekly sales for one of the markdown styles. A 28.6% markdown was implemented in week 6. The graph indicates a decline in sales over the weeks prior to the price change, as well as a decline in sales after the markdown price is implemented. Explanations for this declining-sales phenomenon include saturation of the customer base, loss of customers to competitors, a decline in the perceived value of an item as the selling season progresses and depletion of inventories of individual stock-keeping units.

The following demand model was used to model these features,

$$d(t,p) = w_t(a + bp)e^{-\alpha(t-t_0)},$$

where w_t is a seasonality factor, α is an age factor, and a and b are demand-function parameters. The seasonality factor was estimated from aggregate chainwide data. The age factor was estimated at the group level, while the demand function coefficients a and b were estimated using regression at the individual-style level. While there were significant errors in the prediction of individual weekly sales using this simple model, the average error in total revenues at the style level was only 1.2%; the error in the total revenue of all 60 markdown styles was only 0.53%.

The model was then used to estimate the effects of changes in the firm's markdown policy on the 60 markdown styles. The firm's markdown policy was compared with the markdowns recommended by a RM model that combined a simple online forecasting method with a deterministic dynamic-pricing model. Each week the demand function was reestimated, and an optimal price was computed based on this demand estimate. The new price was implemented if it was at least 20% lower than the initial price (a minimum markdown of 20%). The results are shown in Table 5.5. Note that model-based policy marks down only 33 of the 60 styles and that its average markdown week is much earlier than the firm's, though the average markdown is approximately the same. The estimated increase in revenue is 4.8%. This gain is due to (1) a better selection of which styles to markdown and (2) taking earlier markdowns on the styles that were marked down.

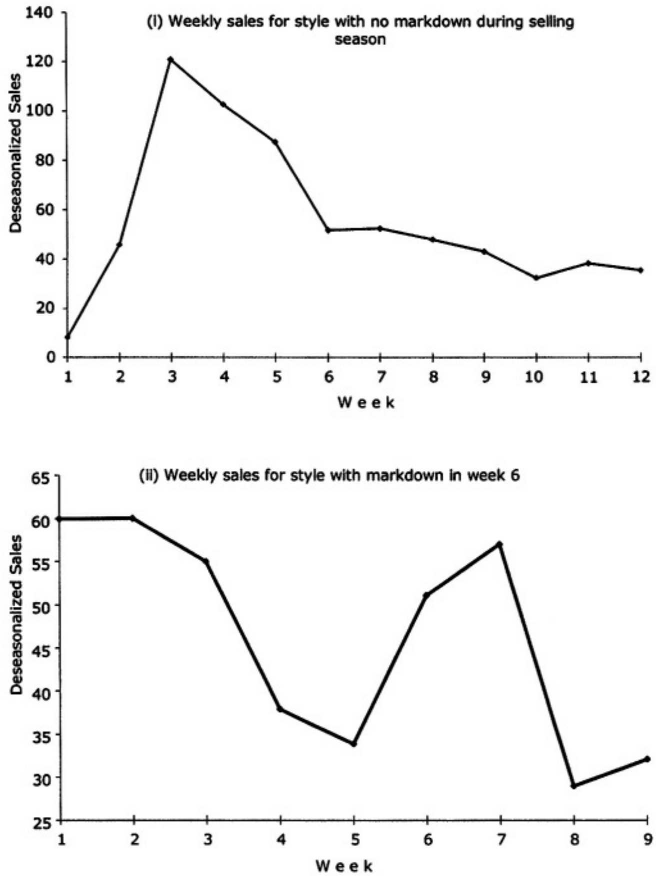

Figure 5.4. Effect of markdowns on two sample styles (sales adjusted for seasonality).

Table 5.5. Results of different markdown policies on 60 markdown styles.

	Model-Based Policy	*Firm's Policy*
Number of markdowns	33	60
Average markdown	25.3%	25.8%
Average markdown week	4.3	8.6
Revenue increase	4.8%	—

5.2.2 Stochastic Models

Next, we look at the case where the price-sensitive demand is stochastic. We separate the case of continuous-demand models from the case of Bernoulli (discrete Poisson) demand, though qualitatively the two cases are similar. We assume that the stochastic regularity Assumption 7.6 (namely, that the demand has bounded variance) holds throughout.

5.2.2.1 Continuous Demand

Here we assume demand in each period is a continuous random variable $D(t,p,\xi_t)$ of the form discussed in Section 7.3.4 with expectation $d(t,p) = E[D(t,p,\xi_t)]$. Capacity is continuous as well. Also, we assume initially that prices in every period have no constraint other than being nonnegative.

As in the deterministic case, we assume that the demand function $d(t,p)$ has an inverse $p(t,d)$. As a result, there is a one-to-one correspondence between prices p and mean demand d in each period, so we can express the random demand as a function of d. That is, $D(t,d,\xi_t)$ is the demand in period t where $E[D(t,d,\xi_t)] = d$. In this way, we can view the mean demand d as our decision variable. We require the following convexity assumption on the random demand:

ASSUMPTION 5.1 *For all t, the random demand $D(t,d,\xi_t)$ is convex and increasing in d on the set $\{d : d \geq 0\}$ for every value ξ_t. That is,*

$$D(t, \alpha d_1 + (1-\alpha)d_2, \xi_t) \leq \alpha D(t,d_1,\xi_t) + (1-\alpha)D(t,d_2,\xi_t),$$

for all $d_1 \geq 0, d_2 \geq 0$ and for all $0 \leq \alpha \leq 1$.

This simply says the demand function is convex in d for each realization of the random-noise term ξ_t. (Such a random function is called strongly stochastically convex; see Appendix B.) Note that both the additive- and the multiplicative-demand models satisfy this convexity assumption, as do combinations of the two models.

We also define the following truncated expected revenue function:

$$r^+(t,d,x) = p(t,d)E\left[\min\{D(t,d,\xi_t),x\}\right]. \tag{5.10}$$

This is interpreted as follows. Given a remaining capacity x and a price $p(t,d)$ in period t, then $r^+(t,d,x)$ is the expected revenue received, since what we sell is the minimum of the demand $D(t,d,\xi_t)$ and the capacity available x. We make the following additional assumption:

ASSUMPTION 5.2 For *all t and for every value ξ_t, both the inverse-demand function $p(t,d)$ and the random revenue $p(t,d)D(t,d,\xi_t)$ are concave in d on the set $\Omega_d(t)$.*

While somewhat restrictive, one can show that this assumption holds for both the additive- and the multiplicative-demand models provided the inverse-demand function $p(t,d)$ and revenue function $r(t,d)$ are concave, which is true, for example, for the linear- and log-linear-demand functions.

The optimization problem can then be formulated as follows:

$$\begin{aligned} V_t(x) &= \max_{d\geq 0} E\left[p(t,d)E\left[\min\{D(t,d,\xi_t),x\}\right] + V_{t+1}(x - D(t,d,\xi_t))\right] \\ &= \max_{d\geq 0} \left\{r^+(t,d,x) + G_{t+1}(x,d)\right\}, \end{aligned} \tag{5.11}$$

with boundary conditions are $V_{T+1}(x) = 0$ for all x and $V_t(0) = 0$ for all t, where we define

$$G_{t+1}(x,d) \equiv E[V_{t+1}(x - D(t,d,\xi_t))].$$

This function is like the value function, in that it gives the expected revenue to go in stage $t+1$ as a function of certain state variables—in this case, the current inventory x and the demand rate decision d. The difference is that it replaces the future inventory state x_{t+1} in the value function $V_{t+1}(x_{t+1})$ by the two variables that determine x_{t+1}—namely, x and d.

The following proposition characterizes the properties of the functions $V_t(x)$ and $G_t(x,d)$:

PROPOSITION 5.1 *If Assumptions 5.1 and 5.2 hold, then for all t,*
(i) $G_t(x,d)$ is jointly concave in x and d,
(ii) $V_t(x)$ is concave in x, and
(iii) $\frac{\partial}{\partial d}G_t(x,d)$ is increasing in x and decreasing in d.

This proposition is proved in Appendix 5.A and has important consequences for the optimal pricing policy. First, under Assumption 5.2 $r^+(t,d,x)$ is concave in d (it is the minimum of two concave functions), and from Proposition 5.1(i) we know that G_{t+1} is concave in d as well. Therefore, a necessary and sufficient condition for an optimal d^* is obtained by differentiating the term inside the maximization in the dynamic program (5.11) and setting the result to zero, which yields

$$\frac{\partial}{\partial d}r^+(t,d,x) = -\frac{\partial}{\partial d}G_{t+1}(x,d^*).$$

By Proposition 5.1(iii), the right-hand side above is decreasing in x, and since $\frac{\partial}{\partial d}r^+(t,d,x)$ is decreasing in d, this means that higher inventory levels x imply a higher optimal sales rate d^*—and consequently a lower optimal price p^*—in any period t. That higher inventories lead to lower optimal prices is certainly intuitive.

5.2.2.2 Bernoulli Demand

If the random demand is Bernoulli (discrete Poisson), then a different analysis is required. Here we assume there is only one customer per period and the customer in period t has a willingness to pay v_t; that is, a random variable with distribution $F(t, v) = P(v_t \leq v)$. Therefore, if the firm offers a price of p in period t, it will sell exactly one unit if $v_t > p$ (with probability $1 - F(t,p)$). Letting $d(t,p) = 1 - F(t,p)$ denote the (average) demand rate, we can define an inverse-demand function, $p(t,d) = F_t^{-1}(1 - d(t))$ and revenue-rate function, $r(t,d) = dp(t,d)$, as before. The inventory and demand in this case are both assumed to be discrete.

Letting $V_t(x)$ denote the optimal expected revenue to go, the problem can be formulated in terms of demand rates $d(t)$ using the Bellman equation:

$$\begin{aligned} V_t(x) &= \max_{d \geq 0} \{d(p(t,d) + V_{t+1}(x-1)) + (1-d)V_{t+1}(x)\} \\ &= \max_{d \geq 0} \{r(t,d) - d\Delta V_{t+1}(x)\} + V_{t+1}(x) \end{aligned} \tag{5.12}$$

with boundary conditions $V_{T+1}(x) = 0$ for all x and $V_t(0) = 0$ for all t, where $\Delta V_t(x) = V_t(x) - V_t(x-1)$ is the expected marginal value of capacity. Under the monotonicity Assumption 7.2 and assuming an interior solution, necessary and sufficient conditions for the optimal rate d^* are

$$J(t, d^*) = \Delta V_{t+1}(x), \tag{5.13}$$

which again, as in the deterministic case of equation (5.2), has the interpretation that we set the marginal revenue equal to the marginal opportunity cost in every period t. One can show

PROPOSITION 5.2 *If Assumption 7.2 holds, then the expected marginal value of capacity, $\Delta V_t(x)$, of the dynamic program (5.12) is decreasing in t and x—that is, $\forall x, t$*
(i) $\Delta V_{t+1}(x) \leq \Delta V_t(x)$ and (ii) $\Delta V_t(x+1) \leq \Delta V_t(x)$.

Again, this monotonicity has intuitive implications for the optimal price. Consider, for simplicity, the case where the marginal revenue is not time dependent, so $J(t,d) = \frac{\partial r(t,d)}{\partial d} = J(d)$. Note that (5.13) and Assumption 7.2 (that $J(d)$ is decreasing) together imply that higher marginal values correspond to lower optimal-demand rates—and hence higher optimal prices. Thus, Proposition 5.2(i) above says that with more time remaining, the marginal value of capacity increases and there-

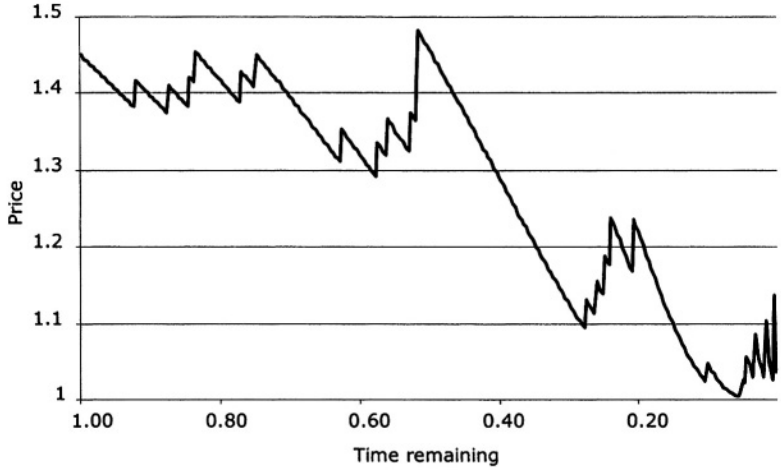

Figure 5.5. An example of the optimal price path in the stochastic case (25 units; exponential demand: $a = -1.1$, $b = 1$).

fore the optimal price increases as well.[2] Conversely, if time elapses without any sales taking place, the optimal price will fall. Proposition 5.2(ii) says the opposite is true of capacity; the more capacity remaining at any given point in time, the lower the optimal price. A numerical example illustrates this behavior:

Example 5.4 Consider a problem with $T = 333$ time-periods, an initial inventory of $C = 25$ units, and a time-homogeneous, exponential-demand function $d(t,p) = ae^{-bp}$ in each period t with parameters $a = 1.1$ and $b = 1$. A sample of the optimal-price path is shown in Figure 5.5. The time axis is normalized to one and represents the fraction of total time remaining. The points at which the price jumps correspond to sales; each sale results in a step increase in price. As time elapses without any sales taking place, prices decline. This is exactly the behavior implied by Proposition 5.2.

5.2.2.3 Comparing the Deterministic and Stochastic Models

One fact that is useful theoretically and computationally is that the deterministic model (5.1) provides an upper bound on the expected revenue from the stochastic model (5.11). This can be shown in a variety of ways. For example, by relaxing the capacity constraint in the stochastic

[2]Note this behavior does not necessarily hold if the marginal revenue varies with time, since in such cases whether the condition $J(t,d^*) = \Delta V_{t+1}(x)$ results in d^* rising or falling over time for a fixed x depends on how both $J(t,d)$ and $\Delta V_{t+1}(x)$ vary with time.

problem with a multiplier $\pi \geq 0$, we can form the relaxed problem

$$\max_{d(t)} E\left[\sum_{t=1}^{T} D(t, d(t), \xi_t) p(t, d(t)) + \pi(C - \sum_{t=1}^{T} D_t(d(t)))\right], \qquad (5.14)$$

where the random variable $D_t(d(t)) = 1$ if there is an arrival in period t using the control $d(t)$, and $D_t(d(t)) = 0$ otherwise. Note an optimal policy for the original stochastic problem (5.11) satisfies $x(0) \geq \sum_{t=1}^{T} D_t(d(t))$ (a.s.); therefore, since $\pi \geq 0$ if we evaluate the objective function of (5.14) for such an optimal policy, it will give an upper bound on the optimal expected revenue for the original problem (5.11). Hence, maximizing (5.14) over all $d(t)$ certainly provides an upper bound as well. However, this function is separable in time t, so we can choose the control in each period to maximize $E[D_t(d(t))(p(t, d(t)) - \pi)]$ at each time t. Since $E[D_t(d(t))] = d(t)$, (5.14) is equivalent to maximizing

$$\sum_{t=1}^{T} [r(t, d(t)) - \pi d(t)] + \pi C.$$

This is solved as in the deterministic case by setting the marginal revenue $J(t, d) = \pi$ in each period t. Since this upper bound is valid for any $\pi \geq 0$, we can take $\pi = \pi^*$, the optimal dual price in the deterministic problem. This results in $d(t) = d^*(t)$, the optimal solution of the deterministic problem. Moreover, by the complementary slackness condition (5.3), the optimal dual price satisfies $\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0$, so the bound (5.14) becomes

$$\sum_{t=1}^{T} d^*(t) p(t, d^*(t)),$$

which is exactly the optimal deterministic revenue. Hence, the optimal deterministic revenue is an upper bound on the optimal expected stochastic revenue.

It's also possible to show that the solution produced by the deterministic dynamic-pricing problem is a reasonably good heuristic for the stochastic-pricing problem. Numerically, it performs well, and theoretically it can be shown to be asymptotically optimal for problems with large demand volumes (such as a large number of time-periods) and large initial inventories (see Gallego and van Ryzin [198]). Such properties provide support for using deterministic models as an approximation. The following example illustrates the deterministic approximation:

Example 5.5 Consider a variation of Example 5.4, where we have $T = 333$ periods and a time-homogenous exponential demand function $d(t, p) = ae^{-bp}$ with parameters

$a = 1$ and $b = 2.718$. Since the demand function is not time varying, the optimal deterministic price is constant. We denote this p^{DET}. The unconstrained revenue-maximizing price is $p^0 = 1$. Starting inventories C range from 1 to 20. Table 5.6 shows the prices and resulting revenues for this problem. As Table 5.6 shows, the

Table 5.6. **Example performance of the deterministic-price heuristic.**

C	p^{DET}	$V_T(C)$	$\frac{V_T^{DET}(C)}{V_T(C)}$
1	3.30	2.40	0.871
2	2.61	4.11	0.926
3	2.20	5.43	0.945
4	1.92	6.47	0.954
5	1.69	7.30	0.956
6	1.51	7.96	0.956
7	1.35	8.49	0.952
8	1.22	8.89	0.946
9	1.11	9.22	0.937
10	1.00	9.46	0.925
11	1.00	9.64	0.951
12	1.00	9.77	0.970
13	1.00	9.85	0.982
14	1.00	9.91	0.990
15	1.00	9.95	0.995
16	1.00	9.97	0.997
17	1.00	9.99	0.999
18	1.00	9.99	0.999
19	1.00	10.00	1.000
20	1.00	10.00	1.000

relative performance of the deterministic heuristic is poorest at $C = 1$ (13% below the optimal revenue) and $C = 10$ (7.5% below the optimal revenue) but otherwise performs reasonably well, especially in the very unconstrained case of initial inventory approaching 20. Note that $Td(p^0) = 10$, so $C = 10$ is the boundary between the constrained and unconstrained regions of the deterministic problem (the constrained region is where the multiplier $\pi^* > 0$ and the stock-clearing price $\bar{p}$ is used; the unconstrained region is where $\pi^* = 0$ and the revenue-maximizing price p^0 is used).

Intuitively, the deterministic prices perform well because they capture the correct "first-order" effect. That is, they maximize revenue subject to the constraint that the mean demand is within the capacity constraint. The stochastic policy does this as well but also adjusts prices dynamically to respond to fluctuations about the mean demand. In addition, the stochastic policy has a tendency to price higher earlier in the sales process (Figure 5.5), which reflects the *option value* of keeping initial prices high in the event realized demand is stronger than average. These two "second-order" adjustments result in the improvement in revenue exhibited in Table 5.6.

Hence, there are really two separate benefits to dynamic pricing. The first is simply to exploit the time-varying price sensitivity of customers; if the demand function $d(t,p)$ varies with t, then even the optimal deterministic price will vary with t due to the optimality condition (5.2). But in addition if demand is stochastic, dynamic pricing helps compensate for random fluctuations in demand and the option value of holding rather than selling units. This is seen in Figure 5.5, where the optimal stochastic prices vary despite the fact that the optimal deterministic prices for this example are constant. In general, both factors will be present in practical problems, but it is useful to distinguish the different forces at work in each case.

5.2.2.4 Prices Constrained to a Discrete Set

Just as in the deterministic case, it may be desirable in practice to constrain the prices to a finite set, $p(t) \in \Omega_p$, where $\Omega_p = \{p_1, \ldots, p_k\}$. Equivalently, the sales rate d are constrained to a discrete set, $d(t) \in \Omega_d(t)$, where as before, $\Omega_d(t) = \{d_1(t), \ldots, d_k(t)\}$, $d_i(t) = d(t, p_i)$ denotes the sales rate at time t when using the price p_i, and $r_i(t) = p_i d_i(t)$ denotes the corresponding revenue rate. For simplicity, we consider only the Bernoulli demand case here (see Section 7.3.4.3).

Computationally, using discrete prices is not a difficult change and in fact reduces the complexity of the dynamic program (5.12) because the search at each stage is now reduced to a finite set of prices. As in the deterministic case, the finite set of prices can be further reduced to only those prices defining the maximum concave envelope (the efficient prices) by using the efficiency criteria (5.7). The reasoning is identical to the deterministic case; inefficient prices produce less expected revenue and have a higher probability of consuming capacity than done by efficient prices (or mixtures of efficient prices) and therefore are never an optimal choice.

Theoretically, the analysis of the discrete problem can again be relaxed to put it in a form similar to the unconstrained price case by allowing the firm to randomize over the discrete set of prices. As in the deterministic case, define new variables $\alpha_i(t)$ that represent convex weights, and in each period replace the variable $d(t)$ in (5.12) with the convex combination

$$d(t) = \sum_{i=1}^{k} \alpha_i(t) d_i(t)$$

and replace the constraint $d(t) \in \Omega_d(t)$ with the constraint

$$\boldsymbol{\alpha}(t) \in W \equiv \{\boldsymbol{\alpha} : \sum_{i=1}^{k} \alpha_i = 1, \boldsymbol{\alpha} \geq 0\}.$$

The dynamic program then becomes

$$V_t(x) = \max_{\boldsymbol{\alpha}(t) \in W} \left\{ \sum_{i=1}^{k} \left[r_i(t) - d_i(t) \Delta V_{t+1}(x) \right] \alpha_i(t) \right\} + V_{t+1}(x)$$

with the usual boundary conditions $V_{T+1}(x) = 0,\ \forall x$ and $V_t(0) = 0$ for all t. In the stochastic case, a fractional solution $\boldsymbol{\alpha}(t)$ can be directly interpreted as a randomization of the prices in $\Omega_p = \{p_1, \ldots, p_k\}$. Also, one can eliminate inefficient prices using the maximum concave envelope as in the deterministic case.

We can put this in a form similar to (5.12) by noting there is a correspondence between the optimal choice of $\boldsymbol{\alpha}(t)$ and the optimal value of $d(t) = \sum_{i=1}^{k} \alpha_i(t) d_i(t)$, since for any fixed sales rate d, the optimal $\boldsymbol{\alpha}(t)$ that achieves this sales rate must maximize the expected revenue—that is, it solves the linear program

$$\begin{aligned} \hat{r}(t,d) = \quad \max \quad & \sum_{i=1}^{k} \alpha_i r_i(t) \\ \text{s.t.} \quad & \sum_{i=1}^{k} \alpha_i d_i(t) = d \\ & \sum_{i=1}^{k} \alpha_i d_i(t) = 1 \\ & \boldsymbol{\alpha} \geq 0. \end{aligned}$$

The resulting $\hat{r}(t,d)$ in fact will define the maximum concave envelope of the fixed set of prices. Hence, the optimization problem can be formulated as

$$V_t(x) = \max_{d} \{ \hat{r}(t,d) - d \Delta V_{t+1}(x) \} + V_{t+1}(x), \tag{5.15}$$

which has exactly the same form as (5.12) except that the maximum concave envelope function $\hat{r}(t,d)$, though continuous and concave, is no longer differentiable. (Like all objective functions of a maximization linear program, $\hat{r}(t,d)$ is a concave and piecewise linear function of the right-hand side d.) Thus, the optimality condition (5.13) must be replaced by

$$\Delta V_{t+1}(x) \in \partial r(t, d^*(t)), \tag{5.16}$$

where $\partial r(t, d^*(t))$ denotes the set of subgradients (the subdifferential) of $r(t, d^*(t))$ at the value $d^*(t)$. (See Appendix C for a definition and discussion of subgradients and nondifferentiable optimization.)

Practically speaking, the above condition implies that the optimal d will most often be at a corner point of the function $\hat{r}(t, d)$ (or there will be multiple optimal solutions along an interval containing two adjacent corner points) and we can always find one of the fixed prices that is optimal without randomizing. However, by formulating the problem this way, we preserve the concavity of $\hat{r}(t, d)$. Therefore, the structure of the optimal solution to (5.15) is the same as that of (5.12), and Proposition 5.1 continues to hold for this case.

5.3 Single-Product Dynamic Pricing with Replenishment

We next consider situations in which inventory can be replenished at a cost in each period, as in many production and supply-chain-management contexts. In such cases, both pricing and inventory decisions need to be made; pricing decisions are used to control demand, while replenishment decisions are used to control supply. The central problem is to optimally coordinate these demand and supply decisions.

As in the finite-supply case, we first look at deterministic models of this problem and then examine stochastic models.

5.3.1 Deterministic Models

We assume a single good with an end-of-period inventory, denoted $x(t)$, that can be replenished over time. There is a per-unit holding cost h_t for inventory in period t and a unit cost for replenishment c_t. We let $y(t)$ denote the amount ordered in period t. As in the finite-supply case, we can formulate the problem in terms of the sales rate $d(t)$, in which case we let $r(t, d(t))$ and $J(t, d)$ denote, respectively, the revenue rate and marginal revenue as before. Again, we assume that these functions satisfy the regularity conditions in Assumptions 7.1, 7.2, and 7.3 unless otherwise specified.

5.3.1.1 Unconstrained Capacity

We first consider the case where is no capacity constraint on the amount ordered in each period. The problem can be formulated as finding a set of rates $d^*(t)$ and reorder quantities $y^*(t)$ that solve

$$\max \sum_{t=1}^{T} r(t, d(t)) - h_t x(t) - c_t y(t) \tag{5.17}$$

$$\text{s.t.} \quad x(t) = x(t-1) - d(t) + y(t), \quad t = 1, \ldots, T$$
$$d(t), x(t), y(t) \geq 0, \quad t = 1, \ldots, T,$$

where we assume the initial inventory $x(0) = 0$ for simplicity.

The problem as stated above is not difficult to solve. Indeed, for $s \leq t$, define the coefficients

$$\gamma_{st} = c_s + \sum_{k=s}^{t-1} h_k,$$

and note that α_{st} is the cost of satisfying demand in period t with supply from period s. Let

$$\gamma^*(t) = \min_{s \leq t}\{\gamma_{st}\},$$

denote the lowest cost for supplying period t, and let $s^*(t)$ denote an index that achieves the minimum on the right-hand side above.

The optimal sales rate in any period t, $d^*(t)$, is then determined by equating the marginal revenue to this lowest marginal cost,

$$J(t, d^*(t)) = \gamma^*(t), \quad t = 1, \ldots, T.$$

And the optimal quantity to order in period s is simply determined by adding up the sales rates from later periods t whose lowest-cost supply is from period s,

$$y^*(s) = \sum_{t:s^*(t)=s} d^*(t), \quad s = 1, \ldots, T.$$

An interesting observation for this problem is that even if the demand functions are time-invariant ($r(t, d) = r(d)$ for all t), the optimal price can still vary over time due to changes in the cost of supply. In other words, because the optimality conditions equate marginal revenue to marginal cost, $J(d^*(t)) = \gamma^*(t)$, changes in the costs $\gamma^*(t)$ over time will lead to time-varying prices, even though the marginal revenue function is time-invariant.

5.3.1.2 Capacity Constraints on Ordering

The problem becomes somewhat more complex when there are capacity constraints on the order quantities of the form

$$y(t) \leq b_t, \quad t = 1, \ldots, T.$$

Such constraints, for example, could be due to limited production, transportation, or handling capacity. While (5.17) can be solved as a nonlinear program with these added capacity constraints, there is a simpler approach. If one discretizes the sales quantities, we can solve the problem using a greedy algorithm under the assumption that the marginal revenue in each period is decreasing. (See Chann, Simchi-Levi, and Swann [105] for a proof.)

The greedy algorithm proceeds as follows. For a fixed vector of rates $\mathbf{d} = (d(1), \ldots, d(T))$, define

$$f(\mathbf{d}) = \sum_{t=1}^{T} r(t, d(t)) - g(\mathbf{d}), \tag{5.18}$$

where $g(\mathbf{d})$ is the minimum cost for meeting the sales rates **d**, defined by fixing **d** and solving the following optimization problem in the variables $x(t), y(t), t = 1, \ldots, T$:

$$\begin{aligned} g(\mathbf{d}) = \quad \min \quad & \sum_{t=1}^{T} h_t x(t) + c_t y(t) \\ \text{s.t.} \quad & x(t) = x(t-1) - d(t) + y(t), \quad t = 1, \ldots, T \\ & y(t) \leq b_t, \quad t = 1, \ldots, T \\ & x(t), y(t) \geq 0, \quad t = 1, \ldots, T. \end{aligned} \tag{5.19}$$

Thus, $f(\mathbf{d})$ is the optimal profit given the demand rates **d**. Computing $f(\cdot)$ is efficient because the minimization problem to determine $g(\cdot)$, (5.19), is simply a minimum-cost network-flow problem.

For notational convenience, let $\mathbf{e}_t$ denote the t^{th} unit vector (the vector with a 1 in the t^{th} component and a zero in all other components), and let δ denote the discretization increment (all components of the vector d are assumed to be integral multiples of δ).

The greedy algorithm is as follows.

STEP 0 (Initialize): Initialize solution

$$\mathbf{d} = (d(1), \ldots, d(T)) = (0, \ldots, 0).$$

Calculate $f(\mathbf{d})$ using (5.18).

STEP 1 (Compute marginal values): FOR $t = 1, \ldots, T$, DO:
Compute $f(\mathbf{d} + \delta \mathbf{e}_t)$ from (5.18).

STEP 2 (Find largest marginal increase): Chose the index t^* for which the marginal gain $f(\mathbf{d} + \delta \mathbf{e}_t) - f(\mathbf{d})$ is largest.
IF $f(\mathbf{d} + \delta \mathbf{e}_{t^*}) - f(\mathbf{d}) \leq 0$, STOP (optimal solution found);
ELSE, update **d:**

$$\mathbf{d} \leftarrow \mathbf{d} + \delta \mathbf{e}_t$$

and GOTO STEP 1.

In words, at each stage the algorithm simply adds an increment δ of demand to the period t that yields the largest net gain $f(\mathbf{d}+\delta\mathbf{e}_t)-f(\mathbf{d})$ and stops when no period produces a positive net gain. Biller et al. [67] report a test of this model and algorithm on data from the automobile industry.

5.3.2 Stochastic Models

A stochastic version of the dynamic-pricing problem with replenishment can also be formulated as follows: As in Section 5.2.2, let $x(t)$ denote the inventory at the end of period t and T be the number of periods in the horizon. (We consider an infinite-horizon, stationary version of the problem in Section 5.3.2.2.) Because demand is random, it is possible that demand in a period can exceed the available inventory. In such cases, we assume that the firm can back-order demand, and this is represented by a negative inventory $x(t)$.

As before, we represent demand in each period as a random variable $D(t,p,\xi_t)$, of the form discussed in Section 7.3.4, with expectation $d(t,p)=E[D(t,p,\xi_t)]$, with a unique inverse $p(t,d)$. We assume that the quantities and demand are continuous. Also, we assume that prices in every period are unconstrained (with $p \geq 0$ the only requirement). Finally, we assume that the demand $D(t,d,\xi_t)$ satisfies the regularity condition in Assumption 7.6 and the convexity condition in Assumption 5.1. The random revenue in each period is $R(t,d,\xi_t)=p(t,d)D(t,d,\xi_t)$.

The inventory after ordering is denoted $y(t)$, and hence the quantity ordered is $y(t)-x(t)$. For notational convenience, we use $y(t)$ as the quantity-decision variable. We assume that we cannot dispose of items, so $y(t) \geq x(t)$.

There is a per-unit ordering cost c_t in period t and a convex cost $h_t(x)$ on the ending inventory x in period t. This cost typically will penalize both positive inventories (due to capital costs, storage costs, and so on), and negative inventories (due to lost goodwill or penalties for late delivery). For example, a function of the form

$$h(x)=ax^{+}+bx^{-}$$

is commonly used, where $x^{+}=\max\{x,0\}$, $x^{-}=\max\{-x,0\}$, a is the cost of holding a unit, and b is the penalty cost for back-ordering a unit.

5.3.2.1 Finite-Horizon Problem

In the multi-period case, the optimization problem can then be formulated as follows:

$$V_t(x) \;=\; \max_{y\geq x, d\geq 0} E\left[R(t,d,\xi_t)-c_t(y-x)-h_t(y-D(t,d,\xi_t))\right.$$

$$+V_{t+1}(y - D(t, d, \xi_t))]$$
$$= \max_{y \geq x, d \geq 0} \{r(t, d) - c_t(y - x) + G_{t+1}(y, d)\}, \quad (5.20)$$

where we define

$$G_{t+1}(y, d) \equiv E\left[V_{t+1}(y - D(t, d, \xi_t)) - h_t(y - D(t, d, \xi_t))\right].$$

Using arguments that are essentially the same as those in Proposition 5.1, one can show the following:

PROPOSITION 5.3 *(i)* $G_t(y, d)$ *is jointly concave in* y *and* d.
(ii) $V_t(x)$ *is concave in* x.
(iii) $\frac{\partial}{\partial d} G_t(y, d)$ *is increasing in* y.
(iv) $\frac{\partial}{\partial y} G_t(y, d)$ *is increasing in* d.

Proposition (5.3) (iii) and (iv) imply that G_t is a *supermodular function.* (See Appendix C for a definition of the supermodularity property.) These properties allow us to characterize the optimal pricing and ordering policy.

Specifically, let $y^0(t)$ and $d^0(t)$ denote the values that maximize (5.20) without the constraint $y \geq x$; that is, they solve

$$\max_{d \geq 0, y} \{r(t, d) - c_t(y - x) + G_{t+1}(y, d)\}.$$

Further, for simplicity assume an interior optimal solution for d and y so that, by joint concavity of G_t, the necessary and sufficient conditions for $y^0(t)$ and $d^0(t)$ are then

$$J(t, d^0(t)) = -\frac{\partial}{\partial d} G_{t+1}(y^0(t), d^0(t))$$
$$c_t = \frac{\partial}{\partial y} G_{t+1}(y^0(t), d^0(t)).$$

(If there are two or more sets of values satisfying these conditions, take the pair $(y^0(t), -d^0(t))$ that is lexicographically the largest.)

It follows, then, that if $x \leq y^0(t)$, the optimal policy in period t is to order up to $y^0(t)$ and set the demand rate at $d^0(t)$ (that is, $y^* = y^0(t)$ and $d^* = d^0(t)$), since the unconstrained optimal solution $(d^0(t), y^0(t))$ is feasible. However, if $x > y^0(t)$, then one can show that it is optimal to order nothing (for example, set $y^* = x$) and choose a demand rate d^* that is higher than $d^0(t)$. Equivalently, set the price lower than $p(t, d^0(t))$. Moreover, the higher the inventory x, the higher the optimal rate d^* (equivalently, the lower the optimal price $p(t, d^*)$).[3] The resulting policy

[3]To see this, we can argue informally as follows. Suppose that the optimal y^* and d^* satisfy $y^* > x > y^0(t)$ and $d^* < d^0(t)$. Then since the constraint $y \geq x$, is not binding, these values

is called a *base-stock, posted-price policy.* If inventory is less than the *base-stock* level $y^0(t)$, then order up to this level, and price at the *posted price* $p(t, d^0(t))$. If inventory exceeds the optimal base-stock level $y^0(t)$, then order nothing, and discount the price below the posted price, with the discount being larger the more the inventory exceeds the optimal base-stock level.

5.3.2.2 Infinite-Horizon, Stationary Problem

One can also extend this same analysis to an infinite-horizon setting. We assume that all the parameters of the problem are time-invariant and profits are discounted by a factor $0 < \beta < 1$ in each period.[4]

The value function in this case is also time-homogenous. The formulation is

$$V(x) = \max_{y \geq x, d \geq 0} \{r(d) - c(y-x) + \beta G(y,d)\},$$

where

$$G(y,d) \equiv E\left[V(y - D(d,\xi)) - h(y - D(d,\xi))\right].$$

In this infinite-horizon case, one can show that a time-invariant, base-stock, posted-price policy is optimal. That is, there exist values y^0 and d^0 such that if $x \leq y^0$, it is optimal to order up to y^0 and price at $p^0 = p(d^0)$. If $x > y^0$, we order nothing, and the optimal demand rate d^* is greater than d^0 and increasing in x. Note that in this infinite-horizon case, once we reach a point where $x < y^0$, then in all remaining periods we simply price at the posted price p^0 and order up to y^0. In other words, we use dynamic pricing only to clear inventory that is higher than the optimal base stock y^0. However, since such high inventory levels are only transient, in the long run, the policy ends up using a constant price.

---

must satisfy the first-order condition,

$$\begin{aligned} J(t,d^*) &= -\frac{\partial}{\partial d} G_{t+1}(y^*, d^*) \\ c_t &= \frac{\partial}{\partial y} G_{t+1}(y^*, d^*). \end{aligned}$$

But this contradicts the fact the $(y^0(t), -d^0(t))$ are the lexicographically largest pair of values satisfying the first-order conditions. Therefore, we must have $y^* = x$ and $d^* \geq d^0(t)$. Since $y^* = x$, the fact that d^* is increasing in x now follows from the fact that $J(t,d^*) = -\frac{\partial}{\partial d} G_{t+1}(x, d^*)$, that $J(t,d)$ is decreasing in d, and that $-\frac{\partial}{\partial d} G_{t+1}(x,d)$ is decreasing in y. See Federgruen and Heching [183] for a complete proof of these properties.

[4] Similar results hold for the case of the average profit criteria by considering the discounted problem with $\beta \to 1$. See Federgruen and Heching [183].

5.3.2.3 Fixed Costs

Another variation of the problem is to include a fixed cost for ordering. The finite-horizon version of this problem was studied by Chen and Simchi-Levi [113]. In this case, the cost function becomes

$$c_t(x) = \begin{cases} K_t + c_t x & \text{if } x > 0 \\ 0 & x = 0 \end{cases}$$

This results in a significantly more complex value function. However, one can show, in certain cases, that properties of the optimal policy are similar to those of classical inventory theory. For example, when the demand function has additive uncertainty, then the optimal ordering policy is of the (s_t, S_t) form, wherein we order only in period t if the inventory $x(t)$ drops below s_t, and in this case we order enough to restore the inventory to the target level S_t (order an amount $S_t - x$). However, this property does not hold for other stochastic-demand functions.[5]

Moreover, the optimal state-dependent price is quite a bit more complex. For example, Chen and Simchi-Levi [113] show that, as a function of the current inventory, the optimal price may not be decreasing in the inventory level between ordering epochs. This is because while there is an incentive to decrease price to reduce inventory, there is also an incentive to increase price to delay reordering and postpone incurring the fixed-ordering cost.

5.4 Multiproduct, Multiresource Pricing

Multiproduct, multiresource—or network—versions of dynamic pricing problems arise in many applications. Two fundamental factors typically link the pricing decisions for multiple products. First, demand for products may be correlated. For example, when products are substitutes or complements, the price charged for one product effects the demand for other related products. Then, a firm jointly managing the pricing of a family of such products must consider these cross-elasticity effects when determining its optimal pricing policy. Second, products may be linked by joint capacity constraints. For example, two products may require the same resource, which is available in limited supply. Even if there are no cross-elasticity effects between the two products, the pricing decision for one product will need to account for the joint effect on demand for the other product that uses the limited resource.

As in the case of capacity controls, most problems in real life are multiproduct problems, either because of cross-elasticity effects or because

[5] A somewhat more complex variant of this (s, S) policy does hold more generally, however; see Chen and Simchi-Levi [113].

of joint capacity constraints, or both. For example, a grocery store that is pricing brands in a food category—say, salty snacks—needs to consider the cross-elasticity effects of its pricing decision for all products in the category. An increase in the price of a packet of potato chips will not just cause a drop in demand for potato chips but will likely also increase the demand for corn chips. At the same time, these products may occupy the same limited shelf space, so stocking more of one product may require stocking less (or none) of other products.

We can model such situations using multiproduct-demand functions and joint capacity constraints on resources. However, like the network problems of capacity control, such formulations quickly become difficult to analyze and solve, which is the reason that many commercial applications of dynamic-pricing models make the simplifying assumption of unrelated products and independent demands and solve a collection of single-product models as an approximation.

Yet in cases where cross-elasticity or resource-constraint effects are strong—for example, when products are only slightly differentiated, customers are very price-sensitive, or joint capacity constraints are tight—then ignoring multiproduct effects can be severely suboptimal. In such cases, we must solve a pricing problem incorporating these effects—or at least approximating them in some fashion. In this section, we look as such multiproduct, multiresource models and methods.

5.4.1 Deterministic Models Without Replenishment

Under a deterministic-demand assumption, it is relatively straightforward to formulate a multiproduct, multiresource version of dynamic pricing similar to those described in Section 5.2. There are n products, indexed by j, and m resources, indexed by i. There is a horizon of T periods, with each period indexed by t. As in Section 7.3.2, let $\mathbf{d} = (d_1, \ldots, d_n)$ denote the demand rate for the n products and $\mathbf{p}(t, \mathbf{d})$ denote the inverse-demand function in period t. We further assume that the revenue-rate function $r(t, \mathbf{d})$ satisfies the regularity conditions of Assumption 7.4.

Product j uses a quantity a_{ij} of resource i. The matrix $\mathbf{A} = [a_{ij}]$ therefore describes the *bill of materials* for all n products. We assume there are limited capacities $\mathbf{C} = (C_1, \ldots, C_m)$ of the m resources.

The dynamic-pricing problem can then be formulated as finding a sequence of demand vectors $d^*(t)$ that maximizes the firm's total revenue subject to the capacity constraints C:

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} r(t, \mathbf{d}(t)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} \mathbf{A}\mathbf{d}(t) \leq \mathbf{C} \\ & \mathbf{d}(t) \geq 0, \quad t = 1, \ldots, T. \end{aligned} \tag{5.21}$$

By Assumption 7.4, $r(t, \mathbf{d})$ is concave in $\mathbf{d}$, and therefore, the following Kuhn-Tucker conditions are necessary and sufficient for characterizing an optimal solution $\mathbf{d}^*(t)$ to (5.21):

$$J(t, \mathbf{d}^*(t)) = \mathbf{A}^\top \boldsymbol{\pi}^* \tag{5.22a}$$

$$\boldsymbol{\pi}^{*\top}(\mathbf{C} - \sum_{t=1}^{T} \mathbf{A}\mathbf{d}(t)) = 0 \tag{5.22b}$$

$$\boldsymbol{\pi}^* \geq 0 , \tag{5.22c}$$

where $J(t, \mathbf{d}) = \nabla_d r(t, \mathbf{d})$ is the marginal-value vector and $\boldsymbol{\pi}^*$ is the optimal dual price on the joint-capacity constraints, having the usual interpretation as the vector of marginal opportunity costs (marginal values) for the m resources. Condition (5.22a) says that at the optimal sales rate, the marginal revenue for each product j should equal the marginal opportunity cost of the resources used by product j, or $\boldsymbol{\pi}^{*\top}\mathbf{A}_j$. Condition (5.22b) says that the marginal opportunity cost of resource i can be positive only if the corresponding capacity constraint for resource i is binding. Finally, (5.22c) requires that the marginal opportunity costs be nonnegative.

The nonlinear program (5.21) is relatively easy to solve numerically, since the objective function is concave and the constraints are linear. (See Bertsekas [58, 59] for specific techniques.)

Example 5.6 Consider the six-node airline network shown in Figure 5.6. Nodes 2 and 3 are "hub" nodes. (Leg seat capacities are as indicated in the figure.) For a given path j on the network, the revenue function is time homogeneous and log-linear

$$d_j(p_j) = a_j e^{-\epsilon_j (p_j/\bar{p}_j - 1)},$$

where $\bar{p}_j$ is interpreted as a reference price for itinerary j, a_j is the demand rate at the reference price, and ϵ_j is the magnitude of the elasticity of demand at the reference price. Demand-function parameters for all O-D pairs are shown in Table 5.7 along with the path (itinerary) used by each O-D pair.

Because the demand functions are time-homogeneous, optimal prices are constant over time. The optimal O-D prices and demand are shown in the last two columns in Table 5.7. The solution gives a total revenue of \$661, 200 across all O-D pairs.

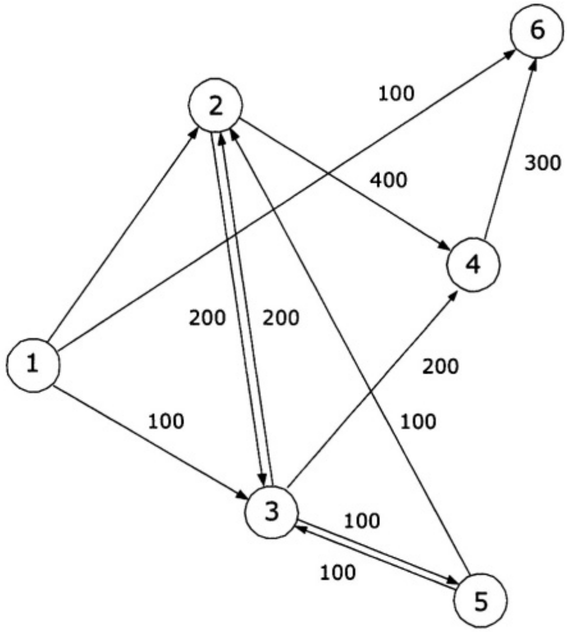

Figure 5.6. A six-node, two-hub airline network.

5.4.2 Deterministic Models with Replenishment

We can formulate deterministic multiproduct models with replenishment, analogous to those in Section 5.3 as follows:

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} r(t, \mathbf{d}(t)) - \mathbf{h}_t^\top \mathbf{x}(t) - \mathbf{c}_t^\top \mathbf{y}(t) \qquad (5.23) \\ \text{s.t.} \quad & \mathbf{x}(t) = \mathbf{x}(t-1) - \mathbf{A}\mathbf{d}(t) + \mathbf{y}(t), \quad t = 1, \ldots, T \\ & \mathbf{y}(t) \leq \mathbf{b}_t, \quad t = 1, \ldots, T \\ & \mathbf{d}(t) \geq 0, \quad t = 1, \ldots, T \\ & \mathbf{x}(t), \mathbf{y}(t) \geq 0, \quad t = 1, \ldots, T, \end{aligned}$$

where $\mathbf{x}(t)$ is an m-vector of inventory levels at the end of period t, $\mathbf{y}(t)$ is an m-vector of order quantities in period t, $\mathbf{h}_t$ is a vector of holding costs, $\mathbf{c}_t$ is a vector of ordering costs, and $\mathbf{b}_t$ is a vector of capacity constraints on the order quantities.

The introduction of an inventory-state variable makes this a more difficult problem to solve. However, in certain specialized cases the greedy allocation algorithm of the type described in Section 5.3.1.2 can be used to solve it exactly. (See Swann [497] for details.) This greedy algorithm can also be used as a heuristic in more general cases.

Table 5.7. Demand-function parameters, itineraries, and optimal solution for Example 5.6.

Market		*Demand Function*				*Optimal Solution*	
O	D	a_j	ϵ_j	$\bar{p}_j$	*Path*	d_j^*	p_j^*
1	2	300	1.0	220	1–2	135	$396.62
1	3	300	1.2	220	1–3	67	$495.86
1	4	300	2.0	400	1–2–4	165	$520.11
1	5	300	1.0	250	1–3–5	33	$752.04
1	6	300	0.8	200	1–6	100	$525.58
2	3	300	1.0	230	2–3	168	$364.28
2	4	300	0.9	200	2–4	143	$365.74
2	5	300	2.0	200	2–3–5	32	$423.79
2	6	300	1.0	200	2–4–6	92	$436.80
3	2	300	1.0	200	3–2	200	$281.76
3	4	300	2.0	230	3–4	131	$325.30
3	5	300	2.0	120	3–5	35	$249.51
3	6	300	2.0	150	3–4–6	14	$378.60
4	6	300	1.0	150	4–6	162	$243.30
5	2	300	1.0	200	5–2	100	$420.39
5	3	300	2.0	150	5–3	47	$289.90
5	4	300	1.0	160	5–3–4	21	$585.20
5	6	300	1.0	230	5–3–4–6	32	$748.50

5.4.3 Stochastic Models

Stochastic multiproduct pricing problems, like stochastic multiproduct capacity-allocation problems, are quite difficult to solve exactly. While in principle they can be formulated as dynamic programs, the size of the state space is often prohibitively large. Therefore, approximations offer the only practical hope to solve such problems.

One natural approach for a stochastic multiproduct problem is to approximate it by its deterministic equivalent problem, which as we've seen in Section 5.2.2.3 are reasonably easy to solve. As in the case of the single-product problem discussed in Section 5.2.2.3, one can indeed show that deterministic solutions are asymptotically optimal (in the same fluid scaling of the problem) in certain cases. That is, suppose the revenue in period t, $R(t, \mathbf{d}, \boldsymbol{\xi}_t)$, is random and we consider a deterministic problem that replaces this random demand by its mean, $r(t, \mathbf{d}) = E[R(t, \mathbf{d}, \boldsymbol{\xi}_t)]$. Then the optimal deterministic price trajectory from the resulting deterministic problem, when applied as an open-loop control for the stochastic problem, produces an expected revenue that is provably close to the optimal stochastic expected revenue.

For example, Gallego and van Ryzin [199] show that for a continuous time version of the multiproduct pricing problem of Section 5.4.1

with Poisson uncertainty, the solution to the equivalent deterministic problem is asymptotically optimal for the stochastic problem as the capacities and time horizon are scaled up proportionally. The arguments and formal definition of the scaling are similar to the asymptotic analysis of network capacity control problems presented in Section 3.6.2 and 3.6 and are omitted. However, the result does provide some intuition into the connection between these two problems.

5.4.4 Action-Space Reductions

One simplification that is useful for multiproduct dynamic pricing problems is to express the problem in terms of resource-consumption rates rather than the demand rates **d**. This yields an equivalent formulation with often a greatly reduced dimensionality that can be much easier to solve. The approach is due to Maglaras and Meissner [354].

To illustrate the main idea, consider the case of the deterministic model (5.21) where there is only $m = 1$ resource but $n > 1$ products. For example, this could be a situation similar to the traditional single-resource problem of Chapter 2 but one in which we control the demand for each product j, d_j, by adjusting its price p_j. The deterministic problem (5.21) in this case is then

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} r(t, \mathbf{d}(t)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} \sum_{j=1}^{n} d_j(t) \leq C \\ & \mathbf{d}(t) \geq 0, \quad t = 1, \ldots, T. \end{aligned} \tag{5.24}$$

To reduce the dimensionality of this problem, we express the problem in terms of the aggregate-demand rate rather than the individual demand rates **d**. To this end, define the aggregate-demand rate

$$\hat{d} = \sum_{j=1}^{n} d_j,$$

and for a given $\hat{d}$ define the maximized revenue-rate function by

$$\begin{aligned} \hat{r}(t, \hat{d}) = \quad \max \quad & r(t, \mathbf{d}) \\ \text{s.t.} \quad & \sum_{j=1}^{n} d_j = \hat{d} \\ & \mathbf{d} \geq 0. \end{aligned} \tag{5.25}$$

That is, $\hat{r}(t, \hat{d})$ is the instantaneous maximum revenue rate given that the total demand rate (equivalently, the resource *consumption rate*) is constrained to be $\hat{d}$. It is easy to show that if $r(t, \mathbf{d})$ is jointly concave in $\mathbf{d}$, then $\hat{r}(t, \hat{d})$ will be concave in $\hat{d}$.

Using these new variables, we can then formulate (5.24) as

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} \hat{r}(t, \hat{d}(t)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} \hat{d}(t) \leq C \\ & \hat{d}(t) \geq 0 \quad t = 1, \ldots, T. \end{aligned} \tag{5.26}$$

Note that this is now a problem that is equivalent to a single-product pricing problem of the same form as (5.1) with a scalar demand rate $\hat{d}$ and revenue-ratefunctions $\hat{r}(t, \hat{d})$. Once we solve for the optimal demand rates $\hat{d}^*(t)$, we can then convert these into optimal vectors of demand rates $\mathbf{d}^*(t)$ by inserting $\hat{d}^*(t)$ into the optimization problem (5.25). Thus, the solution proceeds in two steps: first solve (5.26) to determine the optimal aggregate sales rate, and then solve (5.25) at each time t to disaggregate this optimal aggregate rate into a optimal vector of sales rates (equivalently prices) for each product.

This same action-space-reduction approach also works for stochastic versions of this problem, of the types examined in Section 5.2.2. To illustrate, consider the dynamic program (5.11) for the continuous, additive-uncertainty-demand model, but now suppose there are n products. The n-**product** version of (5.11) yields the dynamic program

$$V_t(x) = \max_{\mathbf{d} \geq 0} E \left[R(t, \mathbf{d}, \boldsymbol{\xi}(t)) + V_{t+1}(x - \sum_{j=1}^{n} D_j(t, d_j, \xi_j(t))) \right], \tag{5.27}$$

where $D_j(t, d_j, \xi_j(t)) = d_j + \xi_j(t)$ is the random demand for product j.

To reduce the action space, we again define the aggregate-demand rate $\hat{d} = \sum_{j=1}^{n} d_j$ and a maximized expected revenue rate using (5.25), where now $r(t, \mathbf{d}) = E[R(t, \mathbf{d}, \boldsymbol{\xi}_t)]$. Also, let

$$\hat{\xi}(t) = \sum_{j=1}^{n} \xi_j(t)$$

denote the aggregate noise term and

$$\hat{D}(t, \hat{d}, \hat{\xi}(t)) = \hat{d} + \hat{\xi}(t)$$

denote the aggregate (random) demand. When these transformed variables are substituted into the dynamic program (5.27), it reduces to the following equivalent single-product formulation

$$V_t(x) \quad = \quad \max_{\hat{d}\geq 0} \left\{ \hat{r}(t,\hat{d}) + G_{t+1}(x,\hat{d}) \right\},$$

where

$$G_{t+1}(x,\hat{d}) \equiv E[V_{t+1}(x - \hat{D}(t,\hat{d},\hat{\xi}(t)))].$$

This has the same form as the single-product DP (5.11). Thus, the single-resource, n-product dynamic-pricing problem is really no more difficult to solve than the single-product problem.

This action-space reduction idea also extends to the general multi-product ($m > 1$), multiresource problem (5.21) *as* well. In this case, one can show that the problem can be reduced to one with only m demand rates (one for each resource) rather than the original n rates (one for each product). Namely, let $\hat{\mathbf{d}} = (\hat{d}_1, \ldots, \hat{d}_m)$ define the maximized revenue rate at each time t

$$\begin{aligned} \hat{r}(t,\hat{d}) = \quad & \max \quad r(t,\mathbf{d}) \\ & \text{s.t.} \quad \mathbf{A}\mathbf{d} = \hat{\mathbf{d}} \\ & \qquad \mathbf{d} \geq 0. \end{aligned}$$

This maximized revenue-rate function and the new demand-rate variables $\hat{\mathbf{d}}$ are then used to reformulate the general problem (5.21) as

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} \hat{r}(t,\hat{\mathbf{d}}(t)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} \hat{\mathbf{d}}(t) \leq \mathbf{C} \\ & \hat{\mathbf{d}}(t) \geq 0 \quad t = 1,\ldots,T. \end{aligned}$$

What these reductions show, in essence, is that the complexity of the multiproduct, multiresource dynamic-pricing problem is caused not by the number of products n but by the number of resources m, since ultimately m determines the dimensionality of both the state and action spaces.

5.5 Finite-Population Models and Price Skimming

We next consider what effect a finite-population assumption has on an optimal dynamic-pricing policy.[6] Recall that a finite-population model assumes that we sample customers without replacement from a finite number of potential customers. Thus, the history of demand (how many customers have purchased, how much they paid, and so on) affects the distribution of both the number and valuations of the remaining customers.

Because the finite-population assumption is more complex, we focus on deterministic models of this situation. However, we consider both a myopic and strategic customer version of the problem.

5.5.1 Myopic Customers

Recall that a myopic customer is assumed to purchase the first time the current price $p(t)$ drops below his valuation v. Combined with the finite-population assumption, this behavior can be exploited by the firm to achieve *price skimming*—a version of classical second-degree price discrimination.

Assume for simplicity that there is a finite population size N and that customers in this population have valuations v that are uniformly distributed on the interval $[0, \overline{v}]$. As an approximation, we assume that sales can occur in fractions, so the population can be regarded as continuous. The important point to note is that the fraction of customers who purchased until time t leave the population of customers for the remaining sale period.

As a result of the myopic-customer assumption, if the firm offers a price p, $N(1 - \overline{v}/p)$, customers will buy. And by the finite-population assumption, there will then be $N\overline{v}/p$ remaining customers, with valuations uniformly distributed on the interval $[0, p]$.

Now, consider a firm that sells a fixed capacity C of a product to this population over T time-periods. The firm is free to set different prices in each period. What is the optimal pricing strategy?

First, it is not hard to see that the optimal prices are decreasing over time, since (by the myopic-customer assumption) the only customers left at time t are those with values less than the minimum price offered in periods $1, \ldots, t-1$. (See Section 8.3 for further motivations for a

---

[6]Section 8.3.4 covers the economic aspects of a durable-goods monopolist under a finite-population, strategic-customer assumption. Here we concentrate on more operational results of dynamic pricing.

firm setting a decreasing schedule of prices.) Hence, the firm will sell nothing if it posts a price in period t that is higher than the minimum price offered in the past. This observation, applied inductively, shows that the optimal prices must decline over time. Moreover, note that if $p(t) \leq p(t-1)$ for all t, the revenue generated in period t is given by

$$p(t)\frac{N}{\overline{v}}(p(t-1) - p(t)),$$

where we define $p(0) = \overline{v}$. This is because $\frac{N}{\overline{v}}(p(t-1) - p(t))$ is the number of customers with valuations greater than $p(t)$ but less than the lowest previous price $p(t-1)$.

To see the effect the decreasing price schedule has on the optimal pricing policy, assume for simplicity that $C > N$, so the capacity constraint is never binding. In this case, the firm must solve

$$\begin{aligned} \max \quad & \sum_{t=1}^{T} \frac{N}{\overline{v}} p(t)(p(t-1) - p(t)) && (5.28\text{a}) \\ \text{s.t.} \quad & p(t) \leq p(t-1), \quad t = 1, \ldots, T && (5.28\text{b}) \\ & p(0) = \overline{v}, && (5.28\text{c}) \\ & p(t) \geq 0. && (5.28\text{d}) \end{aligned}$$

Note that the objective function is jointly concave in $p(t), t = 1, \ldots, T$. It is not hard to see that the constraints (5.28b) are redundant, since the objective function (5.28a) will penalize the use of a price $p(t) > p(t-1)$. Therefore, ignoring constraints (5.28b) and defining $p(T+1) = 0$, the first-order conditions imply the optimal unconstrained solution must satisfy

$$p(t) = \frac{p(t-1) - p(t+1)}{2}, \quad t = 1, \ldots, T.$$

One can easily verify that the solution

$$p^*(t) = \overline{v}\left(1 - \frac{t}{T+1}\right) \qquad (5.29)$$

satisfies these first-order conditions. Since the optimization problem (5.28a–d) is strictly concave and (5.29) satisfies the inequality constraints $p(t) \leq p(t-1)$ for all t, it is in fact the unique optimal solution for (5.28a–d). This solution is illustrated in Figure 5.7(i).

The optimal pricing strategy effectively exploits the myopic behavior of customers to segment them into T+1 groups based on their valuations, and then price discriminates based on this segmentation. Specifically, as shown in Figure 5.7, segment t consists of those customer whose

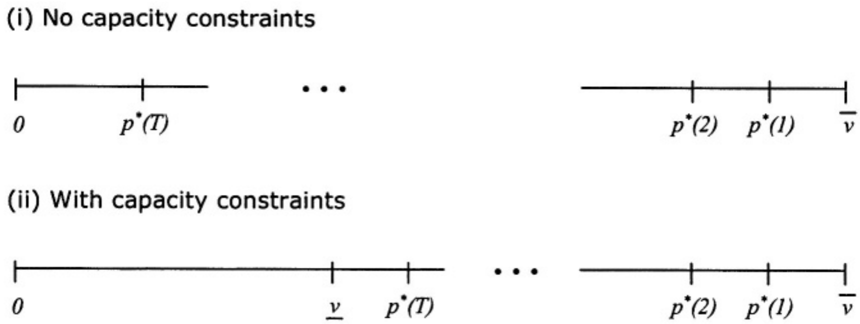

Figure 5.7. **Optimal price-skimming solution for myopic customers: (i) no capacity constraints, (ii) with capacity constraints.**

valuations are in the range $[p^*(t), p^*(t-1)]$, and these segments pay a declining price $p^*(t)$ given by (5.29). Segment $T+1$ has values in the range $[0, \overline{v}/(T+1)]$ and is not served at all.

There are several interesting observations about this solution. First, note we can write the optimal price in period t as

$$p^*(t) = \frac{p^*(t-1)}{2} + \frac{\overline{v}}{2}\left(1 - \frac{t}{T+1}\right).$$

The first term on the right, $p^*(t-1)/2$, is simply the single-period revenue-maximizing price, which follows from the fact that the remaining customers in period t have values uniformly distributed on $[0, p^*(t-1)]$. Therefore, the optimal price in period t is higher than the single-period revenue-maximizing price for period t (except in the last period $t = T$, where they are equal). Intuitively, this occurs because there is an additional benefit to the firm of raising its price in period t in the multiperiod setting; namely, it will have more customers to sell to in the future.

Second, note the price changes over time not because the distribution of valuations changes over time—as in the infinite-population model of demand—but because the firm seeks to price discriminate among the finite population of customers. For example, in an equivalent infinite-population model (essentially, the model of Section 5.2.1 with a linear-demand function), the distribution of values of customers is unaffected by past demand, and hence the distribution would still be uniform over $[0, \overline{v}]$ in each period. In this case, the optimal price to charge in each period would be a constant $\overline{v}/2$ rather than the declining price (5.29). Therefore, a finite population of customers creates an incentive to offer dynamically decreasing prices to achieve price discrimination, an incentive that is not present in infinite-population models.

Finally, note that if the number of periods T increases, the firm's revenues increase because one can show (after some algebra) that the optimal total revenue for T periods is

$$\sum_{t=1}^{T} p^*(t)\frac{N}{\overline{v}}(p^*(t) - p^*(t-1)) = \frac{N\overline{v}}{2}\left(\frac{T}{T+1}\right).$$

Indeed, as T tends to infinity, the firm achieves perfect price discrimination and captures the entire consumer surplus $N\overline{v}/2 = \int_0^{\overline{v}}(N/\overline{v})dv$; each customer ends up paying a price arbitrarily close to his valuation. In particular, a continuous-time model of this problem can achieve perfect price discrimination because the firm can continuously lower prices from $\overline{v}$ down to zero over the interval $[0, T]$. A number $dp(N/\overline{v})$ of customers with values $[p, p + dp]$ will buy when the price is p, so the firm achieves a revenue of $\int_0^{\overline{v}} p(N/\overline{v})dp = N\overline{v}/2$, which is the entire consumer surplus.

The introduction of a binding capacity constraint does little to change this basic story. Indeed, the solution (5.29) will not be feasible if $C < NT/(T+1)$. However, in this case, one can show that the optimal price is simply modified so that only those customers above a lower limit $\underline{v}$ are segmented, where

$$N(\overline{v} - \underline{v})\left(\frac{T}{T+1}\right) = C.$$

The optimal price in this case becomes

$$p^*(t) = \underline{v} + (\overline{v} - \underline{v})\left(1 - \frac{t}{T+1}\right),$$

and customers with valuations less than $\underline{v} + (\overline{v} - \underline{v})/(T+1)$ are not served. This solution is illustrated in Figure 5.7(ii).

5.5.2 Strategic Customers

One might question why customers would behave myopically when faced with a price-skimming strategy. Indeed, knowing that prices will decline over time, rational customers could do better (increase their net utility) by deviating from myopic behavior and delaying purchase until the price is much lower than their valuation. Such behavior is quite plausible and is a valid criticism of the myopic-customer model, but it complicates the analysis of the firm's optimal-pricing policy considerably. Most significantly, it turns the pricing problem into a game between the firm and its customers, in which we must analyze the equilibrium using game-theoretic tools.

Strategic customer behavior is, in fact, a central feature of the theory of optimal mechanism design discussed in Chapter 6 on auctions.

Both auction and list-price mechanisms are analyzed in Chapter 6, and we provide much of the analysis and insight about pricing with strategic customer behavior there (see also Section 8.3.4). Here we focus on the more limited topic of the effect of strategic customers on the price-skimming strategy alone. Further, to keep things simple, throughout this section we consider only the case where the firm has no capacity constraint ($C > N$).

To proceed, one first has to make assumptions about whether the firm can credibly commit to a schedule of prices over time or whether the firm must follow a subgame-perfect equilibrium-pricing strategy. (See Appendix F for a discussion of subgame perfection.) In our case, requiring a subgame-perfect equilibrium means that the strategy for the firm at each time t has to be an equilibrium for the residual revenue-maximization game over the horizon $t, t+1, \ldots, T$, given whatever state the firm and customers were in period t.

For example, if the firm can commit to a price schedule, then a rational customer will simply look at the schedule of prices and (assuming no discounting of utility) decide to purchase in the period with the lowest price, and only customers with valuations above this lowest price will decide to purchase. So effectively, it is only the lowest price among the T periods that matters to customers. Given this fact (and ignoring capacity constraints), the firm will then set this minimum price as the single-period revenue-maximizing price, which, in the case where customer valuations are uniformly distributed on $[0, \overline{v}]$, is just $\overline{v}/2$. The firm will then set arbitrary but higher prices in the other periods. Which period the firm chooses for the minimum price doesn't matter unless revenues are discounted, in which case the firm would prefer collecting revenues sooner rather than later and would choose period 1. The total revenue the firm receives is then $N\overline{v}^2/4$, which is just the product of the price $\overline{v}/2$ and the number of customers willing to pay that price, $N\overline{v}/2$. One can formalize this reasoning and show that this is indeed the equilibrium strategy in the case where the firm has to commit to a price schedule.

Note that the fact that customers are rational has eliminated the ability of the firm to price discriminate; the firm is forced to offer a single uniform price to all customers. Moreover, the firm's revenue is strictly worse under this model. This is to be expected, the firm ought to do worse when customers are "smarter."

However, the single-period strategy outlined above is not always subgame-perfect. To see why, suppose this lowest price $\overline{v}/2$ occurs in period 1. Then in period 2, there will be a population of customers with values less than $\overline{v}/2$ who have not purchased. If the firm has any remaining supply after period t, it would rather sell the remaining stock at some positive price than let it go unsold. Thus, it has an incentive to lower the price in period 2 to capture some of the remaining customers. However, rational customers realize the firm faces this temptation after period 1 and, anticipating the price drop, do not purchase in period 1, so offering the lowest price in period 1 cannot be a subgame-perfect equilibrium.

Besanko and Winston [47] analyze the subgame-perfect pricing strategy. The equilibrium is for the firm to lower prices over time, similar to the price-skimming strategy of Section 5.5.1. In the case where revenues are not discounted, this equilibrium results in the firm setting a declining sequence of prices, where the price in the last period T is simply the single-period optimal price $\overline{v}/2$; all customers buy only in the last period. This case is essentially equivalent to the case where the firm can commit to a schedule of prices, with the exception that the firm is forced to offer the lowest price only in the last period.

The situation is somewhat more interesting if revenues and customer utility are discounted at the same rate. In this case, the subgame-perfect equilibrium has customers with high values buying in the early periods and those with lower values buying in later periods, again, as in the price skimming case of Section 5.5.1. However, unlike the price-skimming case, the equilibrium price in each period is *lower* than the single-period revenue-maximizing price for the customers remaining in that period. In particular, in period 1 the equilibrium price is less than $\overline{v}/2$, and the equilibrium price declines in subsequent periods. Thus, the firm is strictly worse off than when it can commit to a price schedule. This is because when the firm can commit to its price schedule, it can force all customers to purchase in period 1 by simply offering very high prices in periods $t > 1$ while setting a price of exactly $\overline{v}/2$ in period 1. All customers will then buy in period 1 at a price of $\overline{v}/2$.

Besanko and Winston [47] show that with strategic customers, the firm is always better off with fewer periods; that is, the firm's equilibrium revenue is decreasing in the number of periods. This is because the inability of the firm to commit to prices in later periods hurts it, and the more periods, the more often the firm falls victim to the temptation to lower prices. That is, it discounts early and often. This is to be contrasted with the case of myopic customers, where the firm's revenues are increasing in the number of periods. Thus, although the strategy looks like price skimming, rational customers create a qualitatively different situation for the firm than do myopic customers.

5.6 Promotions Optimization

In this section we look at normative models for retail and trade promotions. We first discuss promotions in general and how they differ from the sorts of dynamic-pricing problems considered thus far. We then look at two specific models of promotion optimization.

5.6.1 An Overview of Promotions

As mentioned in this chapter's introduction, promotions are short-run, temporary price changes that are frequently applied to replenishable and consumable goods (such as CPG products). Promotions are run either by the manufacturer (trade promotions) or by retailers (retail promotions or consumer promotions). Manufacturer may also give a discount directly to the end customer in the form of coupons and rebates. While manufacturers are interested in increasing sales or profits for their brand, retailers are interested in overall sales or profits for an entire product category.

A promotion generally increases sales to both the retailer and manufacturer, but there are a variety of factors at work behind the increase. Customers may increase their consumption of the product due to two fundamental effects: higher household inventories lead to fewer stockouts and therefore an increase in consumption; and higher inventories give customers greater flexibility in consuming the product because they don't have to worry about replacing the inventory at higher prices. For instance, Wansink and Deshpande [553] and Chandon and Wansink [104] show that larger household inventory causes faster usage rates if product-usage occasions are flexible (snack foods), products need refrigeration, or products occupy a prominent place in the pantry (for empirical evidence of this based on scanner data, see Ailawadi and Neslin [5]). Some other reasons promotions cause an increase in demand include customers switching from nondiscounted brands to the discounted brands and customers (or retailers for trade deals) stockpiling to take advantage of the low price (forward buying).

Not surprisingly a dominating factor behind the demand increase is the type of product. For example, products such as yogurt and potato chips tend to see an increase because of increased consumption, while for products such as tomato ketchup, diapers, and toilet paper, the sales increase is primarily because of brand switching or stockpiling.

Promotions, in the framework of RM, can be thought of as either (1) a manufacturer using price to dispose excess inventory, (2) a manufacturer trying to gain market share to induce customers to try out its products, (3) retailers experimenting with price to find optimal price points, (4) separating price-sensitive customers, who are willing to use coupons or who wait for deals, (5) retailers trying to increase store traffic, as customers once inside the store are likely to purchase other, nonpromotional items, or (6) a tactic for store brands or small firms to compete against the large advertising budgets of the established brands.

5.6.1.1 Types of Promotions

As mentioned, the main dichotomy in promotions is between retail promotions and trade promotions. Many promotional events are in fact closely coordinated between the manufacturer and the retailer. For instance, if the retailer runs an advertised promotion, the manufacturer may agree to bear a share of the advertising cost, or the trade deal may involve running an in-store display supplied by the manufacturer.

Retail promotions can be advertised or unadvertised (in-store promotions), often coordinated with temporary in-store displays. The price promotion part may take the form of a simple percentage off, coupons, or a "multibuy" (discount for multiple items packaged together), or an extra free (such as three for the price of two; or 15% more free). The latter two types are usually manufacturer-driven, as packaging may have to be changed.

Trade promotions traditionally are in the form of *off-invoice* as a percentage off the amount ordered during the promotion period. Surprisingly, many off-invoice promotions do not require the retailers pass the discount on to the customer, so they may just purchase more during the promotion period and sell it at regular price. The manufacturer would simply see a drop in orders once the promotion period is over.

More effective for the manufacturer is the use of mail-in coupons (direct discount to the final customer), or *scan-back deals,* in which the manufacturer reimburses retailers a certain amount for each unit sold, so the discount is on units sold to end customers rather than on units purchased by the retailer. Scan-back deals eliminate forward buying by the retailer and aligns the retailer's objectives with the manufacturer's.

5.6.1.2 Empirical Findings

The promotions literature is rich in empirical work—based mostly on scanner POS data—that analyzes the effects of promotions on sales and profits in different categories. The common trends that emerge from this research, summarized by Blattberg, Briesch, and Fox [74] as *empirical*

Table 5.8. Empirical generalizations on promotions.[a]

	Finding	*Supporting Literature*
1.	Temporary retail price reductions substantially increase sales.	Woodside and Waddle [580] Moriarty [390] Blattberg and Wisniewski [77]
2.	Higher-market-share brands are less deal elastic.	Bolton [84] Bemmaor and Mouchoux [45] Vilcassim and Jain [535]
3.	The frequency of the deals changes the consumer reference price.	Lattin and Bucklin [329] Kalwani et al. [281] Kalwani and Yim [282] Mayhew and Winer [367]
4.	The greater the frequency of sales, the lower the height of the deal spike.	Bolton [84] Raju [434]
5.	Cross-promotional effects are asymmetric, and promoting higher-quality brands impacts weaker brands disproportionately.	Blattberg and Wisniewski [77, 78], Krishnamurthi and Raj [315, 315] Cooper [128], Walters [549]
6.	Retailers pass-through less than 100% of trade deals.	Chevalier and Curhan [114], Curhan and Kopp [138], Walters [549], Blattberg and Neslin [76]
7.	Display and feature advertising have strong effects on item sales.	Woodside and Waddle [580], Blattberg and Wisniewski [77], Kumar and Leone [316]
8.	Advertised promotions can result in increased store traffic.	Walters and Rinne [547], Kumar and Leone [316], Walters and MacKenzie [548], Grover and Srinivasan [225]
9.	Promotions affect sales in complementary and competitive categories.	Walters and Rinne [547], Walters and MacKenzie [546], Mulhern [396], Walters [549], Mulhern and Leone [395]

[a] *Source:* Blattberg, Briesch and Fox [74].

generalizations,[7] are valuable both for the practitioner as well as the academic researcher. Table 5.8 gives the main findings. In addition to the findings in Table 5.8, Blattberg, Briesch, and Fox [74] report some conflicting findings with respect to the following four questions:

[7]Blattberg, Briesch and Fox [74] define an empirical generalization as follows: (1) the topic being studied is well-defined; (2) there are at least three articles by at least three different authors in which empirical research has been conducted in the specific area, and (3) the empirical evidence is consistent.

- Does the majority of promotional volume come from switchers rather than from customers increasing their consumption or category volume growth? The most likely explanation for the variation in the findings here may be the differences in the nature of the products; one can well imagine promotions causing consumption increase for yogurt, but not say, for toilet paper or ketchup.
- Do promotional elasticities exceed long-run price elasticities? That is, because of the temporary nature of a promotion, does it cause a greater increase in demand than if the firm were to permanently lower its price?
- Is the trough after the deal due to customers' accelerating their purchases and stockpiling, creating a drop in the normal sales after a promotion? Somewhat surprisingly, there is no consensus whether this happens.
- Is there is a negative long-term effect to promotions? Are promotions detrimental to long-term brand equity? The findings have been mixed, with some studies discovering a long-term negative effect, and some finding both a positive and negative impact due to promotions.

5.6.2 Retailer Promotions

We next examine two normative models of promotion optimization. In the first model, due to Greenleaf [221], a monopolist retailer is assumed to maximize profits from promoting a particular brand. (A "brand" is a particular-size of a given product.) Customers are assume to have a reference price (see Appendix E), assumed to be an exponentially smoothed average of past prices, as follows:

$$\bar{p}(t) = \alpha\bar{p}(t-1) + (1-\alpha)p(t-1) + \xi_t, \tag{5.30}$$

where $0 \leq \alpha < 1$ is the smoothing parameter,[8] and ξ_t is a 0-centered random variable representing the error term.

Demand is assumed to be composed of two separable factors, a base demand $q(p(t))$ and a reference price factor $g(\bar{p}(t), p(t))$, as follows:

$$d(t) = q(p(t)) + g(\bar{p}(t), p(t)), \tag{5.31}$$

where

$$g(\bar{p}(t), p(t)) = \begin{cases} \delta(\bar{p}(t) - p(t)) & \text{if } \bar{p}(t) > p(t) \\ \gamma(\bar{p}(t) - p(t)) & \text{otherwise.} \end{cases} \tag{5.32}$$

[8] Based on scanner data, Greenleaf [221] finds $\alpha = 0.925$ for peanut butter, and Hardie, Johnson, and Fader [237] find $\alpha = 0.83$ for orange juice. It is also common in promotions models to assume *a priori* $\alpha = 0$—that is, the reference price is the previous period's price.

The parameters $\delta, \gamma > 0$ model customers' asymmetric price sensitivity to loss $(\bar{p}(t) \leq p(t))$ or gain $(\bar{p}(t) > p(t))$ perception: if customers value gains more than losses, $\delta > \gamma$, and if they are loss averse, $\delta \leq \gamma$. (Again, see Appendix E for a discussion of consumer-choice theory based on valuations of losses and gains.)

The reference-price dynamics given by (5.30) capture the effect of current promotions on future profits; frequent and deep promotions will reduce customers' reference price $\bar{p}(t)$ for the brand, and as a result they will start perceiving the normal price as a loss. So even though promotions generate short-run profits, it is in the retailer's long-run interest not to run promotions too frequently.

The retailer is assumed to maximize its discounted profits over an infinite planning horizon. The retailer's discount factor is β, $0 < \beta < 1$, and the marginal cost of production is c. This results in the following dynamic program:

$$\max_{p(t)} \sum_{t=0}^{\infty} \beta^t [(p(t) - c)d(t)], \tag{5.33}$$

with a state evolution equation given by (5.30).

Greenleaf [221], using simulations, shows that the optimal policy for the retailer obtained by solving (5.33) can be cyclical, oscillating between periods of high prices and periods of low prices. Kopalle, Rao, and Assunção [310], using analytical and numerical techniques, derive a number of interesting structural properties of (5.33). Specifically, they show that if the entire customer population is loss averse $(\delta \leq \gamma)$, a constant-price policy is optimal. On the other hand, if customers value gains more than losses $(\delta > \gamma)$, then a cyclical policy of hi-lo pricing is optimal for the retailer. In other words, the asymmetry in customer valuations for gains and losses can be sufficient motivation to run promotions. Moreover, they numerically show for this case that the difference between the high and low price increases as the gain coefficient δ increases for a fixed level of loss coefficient γ, and the high price increases as the memory parameter α in (5.30) decreases.

As we mentioned, retailers are more interested in category profits than in profits from promoting a particular brand. For a retailer managing n brands in a category, the objective therefore is to manage the n prices over time, $\mathbf{p}(t) = (p_1(t), \ldots, p_n(t))$. This requires solving the following optimization problem:

$$\max_{\mathbf{p}(t), t=1,2,\ldots} \sum_{t=1}^{\infty} \sum_{j=1}^{n} \beta^t [f_j(\mathbf{p}(t)) + (p_j(t) - c) g_j(\bar{p}_j(t), p_j(t))], \tag{5.34}$$

where $f_j(\mathbf{p}(t))$ is an aggregate profit function, dependent on the prices of all the brands but excluding reference price effects, which are captured by the second term, and $g_j(\bar{p}_j(t), p_j(t))$ is of the same form as (5.32). The state as before is the reference price, and the state equation is

$$\bar{\mathbf{p}}(t) = \alpha\bar{\mathbf{p}}(t-1) + (1-\alpha)\mathbf{p}(t).$$

Kopalle, Rao, and Assunção [310] analyze (5.34) and show once again that when $\delta_i > \gamma_i$ for all brands, a cyclical pricing policy of hi-lo pricing is optimal, and moreover, the cycles are *in phase*—that is, all the brands are priced high, or all the brands are discounted together. The reason is that hi-lo prices in phase minimize the cross-price effect, at the same time allowing the retailer to take advantage of the reference-price effect.

5.6.3 Trade-Promotion Models

As we discussed, a manufacturer offers rebates to retailers to promote its own brand. The cooperation might take many forms (such as joint advertising and store displays), and the contracts are varied (such as scan-backs and sale guarantees).

On the one hand, models for optimizing the manufacturer's promotions tend to be simpler than the retailer's problem, as the manufacturer is concerned with only one brand. But on the other hand, one has to model retailer pass-thru behavior. (Recall that pass-thru is the percentage of the discount the retailer passes on to the end consumer.) This requires modeling the vertical competition between manufacturer and retailer. In contrast to the previous section, however, one typically ignores reference-price effects because the discount is offered to the retailer rather than to the end customers.

The most widely used model for representing demand as a function of deal price and displays is the SCAN*PRO model of Section 9.6.4. Kopalle, Mela, and Marsh [309] analyze a Stackelberg game between a manufacturer and retailer, where the demand is given by the SCAN*PRO functional form. Silva-Russo, Bucklin, and Morrison [470] give a simpler mixed integer programming formulation (see also Tellis and Zufryden [507]), where the manufacturer assumes that retailers are passive, but they model retailers' pass-thru percentages. They report an implementation of the model at a large CPG manufacturer. The formulation does not by itself give insight into the optimal structure or policies for the manufacturer, but it is reasonably practical and captures the main concerns of the manufacturer in the formulation of its constraints.

5.7 Notes and Sources

The book by Nagle [400] provides a good general-management overview of pricing decisions. Elmaghraby and Keskinocak [177] provide a nice current survey on research in the area of dynamic pricing. As for the connection between pricing- and capacity-allocation decisions, see Walczak and Brumelle [543].

Smith and Achabal [480] study a continuous-time version of the problem with inventory-depletion effect as in Section 5.2.1.5. They also study the problem of selecting the optimal initial inventory and report summary results of tests of the model at several major retailers. Heching, Gallego, and van Ryzin [247] provide revenue estimates based on a regression test of this same type of deterministic model on data from an apparel retailer.

Gallego and van Ryzin [198] analyzed a continuous-time, time-homogeneous version of the stochastic model of Section 5.2.2, providing monotonicity properties of the optimal price, an exact solution in the exponential demand case, and proving the asymptotic optimality of the deterministic policy as in Section 5.2.2.3. Bitran and Mondschein [73] analyze a discrete-time model of the problem essentially the same as that presented in Section 5.2.2 and test in on apparel retail data. Zhao and Zheng [589] analyze the continuous-time model with a time-varying demand function and provide an alternative proof of monotonicity of the marginal values $\Delta V_t(x)$; they also provide results on the monotonicity of optimal prices over time. See also Kincaid and Darling [304] and Stadje [484]. Das Varmand and Vettas [145] analyze the problem of selling a finite supply over an infinite horizon with discounted revenues, where the discounting provides an incentive to sell items sooner rather than later and there is no hard deadline on the sales season.

Stochastic models with discrete price changes are analyzed in the continuous-time case in a series of papers by Feng and Gallego [185, 186] and Feng and Xiao [188, 189]. The problems differ in terms of whether there are two prices or more than two prices, whether the price changes are reversible or one-way changes. Feng and Gallego [186] extend the analysis also to the interesting case where demand is Markovian and may depend on the current inventory level—for example, as in the classical Bass model of new-product diffusion. The notion of the *maximum concave envelope* of prices is due to Feng and Xiao [188]. See also You [586] for a discrete-time analysis of the problem.

There is an extensive literature on production-pricing problems. Eliashberg and Steinberg [175] provide of review of joint pricing and production models. Single-period, convex-cost problems under demand uncertainty are analyzed by Karlin and Carr [293], Mills [386], and the early paper of Whitin [564]. The literature on single-period pricing under demand uncertainty (the price-dependent newsvendor problem) is surveyed by Petruzzi and Dada [417]. Multiperiod, convex cost models are analyzed by Hempenius [249], Thowsen [509], and Zabel [509].

Rajan, Rakesh, and Steinberg [433] analyze a deterministic model of dynamic pricing within an inventory replenishment cycle, where the motivation for dynamic pricing is the deterioration in the product as well as its declining market value with age (for example, pricing perishable foods).

The optimality of the greedy allocation algorithm for the deterministic production-pricing problem with capacity constraints was shown by Chann, Simchi-Levi, and Swann [105]. The optimality of the base-stock, posted-price policy discussed in Section 5.3.2 was proved by Federgruen and Heching [183]. The fixed-cost version of this problem was recently analyzed by Chen and Simchi-Levi [113].

Multiproduct, multiresource dynamic-pricing problems were analyzed in Gallego and van Ryzin [199], including bounds on the relationship between the stochastic and deterministic versions of the problem. The action-space-reduction approach is a recent result due to Maglaras and Meissner [354]. A related network pricing we have omitted is congestion pricing for communications service; see for example Pashalidis [413].

Stokey [490] analyzes a model of intertemporal price discrimination similar to that presented in Section 5.5.1. See also Kalish [279]. Stokey [491] analyzes a price-skimming model with rational customers under the assumption that the firm can commit to a price schedule. The material in Section 5.5.2 on the subgame-perfect pricing equilibrium for a firm faced with strategic customers is from Besanko and Winston [47].

The artificial-intelligence community also has recently become interested in dynamic pricing, using autonomous software agents. The approach is simulation based, with experiments using various strategies for the players. Although relevant, the approach is beyond the scope of this book, though the interested reader can refer to Morris, Ree, and Maes [393] and Morris and Maes [392].

The literature on promotions is rich in empirical work, which we have summarized, somewhat tersely, in Table 5.8. The material in Section 5.6.1.2 is entirely from Blattberg, Briesch, and Fox [74]. For more empirical generalization articles, see Bell, Chiang, and Padmanabhan [34] and Sethuraman and Srinivasan [458].

The standard reference on promotions is the book by Blattberg and Neslin [76]. There is a large body of work that tries to understand the interactions between the retailer and the manufacturer using game theory, which we do not have the opportunity to cover here—see Lal and

Villas-Boas [322, 323], Lal [325, 324], Rao, Arjuni, and Murthi [436], Gerstner and Hess [211] and Bell, Iyer, and Padmanabhan [35].

APPENDIX 5.A: Proof of Monotonicity Results

Proof of Proposition 5.1

Since there are multiple parts, we restate the proposition:

If Assumptions 5.1 and 5.2 hold, then for all t*:*

(i) $G_t(x,d)$ *is jointly concave in* x *and* d*.*

(ii) $V_t(x)$ *is concave in* x*.*

(iii) $\frac{\partial}{\partial d}G_t(x,d)$ *is increasing in* x *and decreasing in* d*.*

We first need a preliminary result:

LEMMA 5-5.A.1 *If Assumption 5.2 holds, then the truncated revenue function* $r^+(t,d,x)$ *is jointly concave in* x *and* d *on the for* $x \geq 0$ *and* $d \in \Omega_d(t)$*.*

Proof

By definition,

$$r^+(t,d,x) = p(t,d)E\left[\min\{D(t,d,\xi_t),x\}\right] = E\left[\min\{p(t,d)D(t,d,\xi_t),p(t,d)x\}\right].$$

By Assumption 5.2, the term $p(t,d)D(t,d,\xi_t)$ is concave in d, and $p(t,d)$ is concave as well. Also note if $x \geq 0$, then $p(t,d)x$ will also be jointly concave in x and d. This follows because the Hessian

$$\begin{vmatrix} xp''(t,d) & p'(t,d) \\ p'(t,d) & 0 \end{vmatrix}$$

is negative definite, since both $xp''(t,d) \leq 0$ and the determinant $-(p'(t,d))^2 \leq 0$. Therefore, $\min\{p(t,d)D(t,d,\xi_t),p(t,d)x\}$ is jointly concave because it is the minimum of two concave functions. Finally, taking expectations preserves concavity, hence $r^+(t,d,x)$ is jointly concave. *QED*

We are now ready to prove Proposition 5.1. Parts (i) and (ii) are related by induction. Indeed, we first show that if $G_{t+1}(x,d)$ is jointly concave in x,d and Assumptions 5.2 holds (so by Lemma 5-5.A.1 $r^+(t,d,x)$ is jointly concave in d and x), then $V_t(x)$ is concave in x. To do so, consider any two values nonnegative values x_1 and x_2 and any real α satisfying $0 \leq \alpha \leq 1$. For notational convenience define the convex combination $\bar{x} = \alpha x_1 + (1-\alpha)x_2$, and let d_i^* denote the value that maximizes $r(t,d) + G_{t+1}(x_i,d)$, for $i = 1,2$ and define $\bar{d} = \alpha d_1^* + (1-\alpha)d_2^*$. Then

$$\begin{aligned} V_t(\bar{x}) &= \max_{d\in\Omega_d(t)}\{r(t,d) + G_{t+1}(\bar{x},d)\} \\ &\geq r^+(t,\bar{d},\bar{x}) + G_{t+1}(\bar{x},\bar{d}) \\ &\geq \alpha(r^+(t,d_1^*,x_1) + G_{t+1}(x_1,d_1^*)) + (1-\alpha)(r^+(t,d_2^*,x_2) + G_{t+1}(x_2,d_2^*)) \\ &= \alpha V_t(x_1) + (1-\alpha)V_t(x_2), \end{aligned}$$

where the last inequality follows from the joint concavity of r^+ and G_{t+1}. So $V_t(x)$ is concave in x provided r is concave (Assumption 7.2), and G_{t+1} is jointly concave.

Likewise, we show that G_{t+1} is jointly concave if $V_{t+1}(x)$ is concave and Assumption 5.1 holds. To see this, consider any four nonnegative x_1, x_2, d_1, d_2 and any real

α satisfying $0 \le \alpha \le 1$. Define $\bar{x} = \alpha x_1 + (1-\alpha)x_2$ and $\bar{d} = \alpha d_1 + (1-\alpha)d_2$. Then

$$\begin{aligned} G_{t+1}(\bar{x}, \bar{d}) &= E[V_{t+1}(\bar{x} - D(t, \bar{d}, \xi_t)] \\ &\ge E[V_{t+1}(\bar{x} - (\alpha D(t, d_1, \xi_t) + (1-\alpha)D(t, d_2, \xi_t))] \\ &\ge E[\alpha V_{t+1}(x_1 - D(t, d_1, \xi_t)) + (1-\alpha)V_{t+1}(x_2 - D(t, d_2, \xi_t))] \\ &= \alpha G_{t+1}(x_1, d_1) + (1-\alpha)G_{t+1}(x_2, d_2), \end{aligned}$$

where the first inequality follows from Assumption 5.1, $V_{t+1}(x)$ is increasing in x, and the second inequality follows from the fact that $V_{t+1}(x)$ in concave.

Parts (i) and (ii) of Proposition 5.1 now follow from these two results using an induction argument and the fact that for all x, $V_0(x) = 0$, which is concave.

Finally, to show part (iii), note that the fact that $\frac{\partial}{\partial d} G_t(x, d)$ is decreasing in d follows from the concavity of G_t in d (part (i)). To show it is increasing in x, take a nonnegative α and d note that the difference

$$V_{t+1}(x - D(t, d+\alpha, \xi_t)) - V_{t+1}(x - D(t, d, \xi_t))$$

is oppositive since $D(t, d+\alpha, \xi_t) \ge D(t, d, \xi_t)$ by Assumption 5.1 and therefore is increasing in x by the concavity of $V_{t+1}(\cdot)$. Therefore, taking expectations above we then have that the difference

$$G_{t+1}(x, d+\alpha) - G_{t+1}(x, d)$$

is increasing in x as well. Since

$$\frac{\partial}{\partial d} G_t(x, d) = \lim_{\alpha \to 0} \frac{1}{\alpha}(G_{t+1}(x, d+\alpha) - G_{t+1}(x, d)),$$

it therefore follows that $\frac{\partial}{\partial d} G_t(x, d)$ is increasing in x. *QED*

Proof of Proposition 5.2

The proof here is essentially identical to Proposition 2-2.A.4 result for the discrete-choice single-resource model in Appendix 2.A.

We first show that $\Delta V_t(x)$ is decreasing in x. The proof is by induction on t. First, this is trivially true for $t = T+1$ by the boundary conditions $V_{T+1}(x) = 0$ for all x. Assume it is true for period $t+1$, and consider period t. Let d_i^* denote the optimal solution to (5.12) for inventory level $x+i$; that is, it is an optimal solution in the recursion

$$V_t(x) = \max_{d \in \Omega_d(t)} \{r(t,d) - d\Delta V_{t+1}(x+i)]\} + V_{t+1}(x)$$

and note that since $\Delta V_t(x+i) = V_t(x+i) - V_t(x+i-1)$, we can write

$$\begin{aligned} \Delta V_t(x+2) - \Delta V_t(x+1) &= \Delta V_{t+1}(x+2) - \Delta V_{t+1}(x+1) \\ &\quad +(r(t, d_2^*) - d_2^* \Delta V_{t+1}(x+2)) \\ &\quad -(r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1)) \\ &\quad -(r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1)) \\ &\quad +(r(t, d_0^*) - d_0^* \Delta V_{t+1}(x)) \end{aligned}$$

From the optimality of d_1^*, the following inequalities hold:

$$r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1) \ge r(t, d_2^*) - d_2^* \Delta V_{t+1}(x+1)$$

and

$$r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1) \geq r(t, d_0^*) - d_0^* \Delta V_{t+1}(x+1).$$

Substituting into (5.A.0) we obtain

$$\begin{aligned} \Delta V_t(x+2) - \Delta V_t(x+1) \quad \leq \quad & \Delta V_{t+1}(x+2) - \Delta V_{t+1}(x+1) \\ & +(r(t, d_2^*) - d_2^* \Delta V_{t+1}(x+2)) \\ & -(r(t, d_2^*) - d_2^* \Delta V_{t+1}(x+1)) \\ & -(r(t, d_0^*) - d_0^* \Delta V_{t+1}(x+1)) \\ & +(r(t, d_0^*) - d_0^* \Delta V_{t+1}(x)). \end{aligned}$$

Rearranging and canceling terms yields

$$\begin{aligned} \Delta V_t(x+2) - \Delta V_t(x+1) \quad \leq \quad & (1 - d_2^*)(\Delta V_{t+1}(x+2) - \Delta V_{t+1}(x+1)) \\ & + d_0^*(\Delta V_{t+1}(x+1) - \Delta V_{t+1}(x)). \end{aligned}$$

By induction, $\Delta V_{t+1}(x+2) - \Delta V_{t+1}(x+1) \leq 0$ and $\Delta V_{t+1}(x+1) - \Delta V_{t+1}(x) \leq 0$ and since d values are at most one (expected demand in a period is at most one in the discrete Poisson case), $1 - d_2^* \geq 0$ and $d_0^* \geq 0$. Therefore, $\Delta V_t(x+2) - \Delta V_t(x+1) \leq 0$. (Note the concavity of $r(\cdot)$ is not required for this part of the proof.)

To show monotonicity in t, using the same notation note that

$$\begin{aligned} & \Delta V_t(x+1) - \Delta V_{t+1}(x+1) & (5.A.1) \\ = \quad & (r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1)) - (r(t, d_0^*) - d_0^* \Delta V_{t+1}(x)) \\ \geq \quad & (r(t, d_1^*) - d_1^* \Delta V_{t+1}(x+1)) - (r(t, d_0^*) - d_0^* \Delta V_{t+1}(x+1)) \\ = \quad & (r(t, d_1^*) - r(t, d_0^*)) - \Delta V_{t+1}(x+1)(d_1^* - d_0^*), & (5.A.2) \end{aligned}$$

where the first inequality above follows by the fact that $\Delta V_{t+1}(x+1) \leq \Delta V_{t+1}(x)$. Now by the concavity of $r(t, d)$ we have that

$$(r(t, d_1^*) - r(t, d_0^*)) \leq \frac{\partial}{\partial d} r(t, d_1^*)(d_1^* - d_0^*).$$

But the first-order conditions imply $\frac{\partial}{\partial d} r(t, d_1^*) = \Delta V_{t+1}(x+1)$, so substituting above we have that

$$r(t, d_1^*) - r(t, d_0^*) \geq \Delta V_{t+1}(x+1)(d_1^* - d_0^*).$$

Substituting into (5.A.2) implies $\Delta V_t(x+1) - \Delta V_{t+1}(x+1) \geq 0$. *QED*