

Strategic Management of Distressed Inventory

Guillermo Gallego

Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027,
 gmg2@columbia.edu

Robert Phillips

Nomis Solutions, Inc., San Bruno, California 94066,
 robert.phillips@nomissolutions.com

Özge Şahin

Stephen Ross School of Business, University of Michigan, Ann Arbor, Michigan 48103,
 ozge@bus.umich.edu

It is well known that maximizing revenue from a fixed stock of perishable goods may require discounting prices rather than allowing unsold inventory to perish. This behavior is seen in industries ranging from fashion retail to tour packages and baked goods. A number of authors have addressed the markdown management problem in which a seller seeks to determine the optimal sequence of discounts to maximize the revenue from a fixed stock of perishable goods. However, merchants who consistently use markdown policies risk training customers to “wait for the sale.” We investigate models in which the decision to sell inventory at a discount will change the future expectations of customers and hence their buying behavior. We show that, in equilibrium, a single-price policy is optimal if all consumers are strategic and demand is known to the seller. Relaxing any of these conditions can lead to a situation in which a two-price markdown policy is optimal. We show using numerical simulation that if customers update their expectations of availability over time, then optimal sales limit policies can evolve in a complex fashion.

Key words: revenue management; consumer behavior; pricing

History: Received: May 2006; Accepted: September 2007, after 1 revision.

1. Introduction

A merchant holding a fixed inventory of a perishable good is better off (at least in the short run) discounting his stock than allowing it to perish. This is true whether the item being sold is physically perishable (as in bread), is a service that uses fixed capacity (such as an airline seat), or needs to be removed from the store by a certain “out date” (as in the case of fashion goods or discontinued consumer electronics). As a result, many sellers of constrained and perishable goods use a “markdown policy” whereby an item is initially sold at full price and then subjected to deeper and deeper discounts until inventory is completely sold or the out date arrives, whichever comes first. Under a markdown policy, the price of an item can only decrease over time. This is in contrast to a “promotions policy” in which an item may be subject to periodic promotions of limited duration (Warner and Barsky 1995). Industries in which markdown policies are common include fashion retailers, European tour operators, and consumer electronics, among others. In this paper we consider the problem of determining an optimal markdown policy when markdowns in a current season can influence customer expectations in future seasons.

A seller using a markdown policy faces the *markdown management* (or *markdown optimization*) problem of when to increase the discount (i.e., lower the price) on remaining stock and by how much. This problem has been addressed by a number of authors, including Feng and Gallego (1995, 2000), Bitran and Mondshein (1997), and Gallego and Van Ryzin (1994). Elmaghraby et al. (2004) consider the case in which customers may demand more than one unit of inventory. Heching et al. (2002) showed that optimizing the timing and magnitude of markdowns could significantly increase revenues at a women’s specialty apparel retail store. The benefits of markdown management are sufficiently substantial that a number of retailers, including ShopKo, JCPenney, Sears, and Circuit City, have purchased commercial software systems from vendors such as Profit-Logic (now owned by Oracle). Levy and Woo (1999) and Friend and Walker (2001) discuss some of the challenges encountered in implementing markdown management systems. Talluri and van Ryzin (2004) and Phillips (2005) provide overviews of common approaches to the markdown management problem.

A virtually universal assumption in the markdown management literature is that the goal of the seller is

to maximize revenue from his current stock of inventory. Markdown management typically occurs in the situation when inventory has already been purchased or capacity levels are fixed and purchase costs are therefore sunk so that maximizing revenue is equivalent to maximizing contribution. This ignores the fact that the markdown policy adopted by a seller will influence the future behavior of customers. A retailer who initiates a policy of drastic late-season discounts is likely to find that more and more of his customers will wait for the discounts. There is a widespread feeling both among North American retailers and European tour operators that customers are indeed increasingly reluctant to purchase at list price and instead “wait for the sale.” As one observer noted, “Unfortunately, the department stores have trained customers to shop during clearance sales or special promotions” (Fickes 2001). This feeling is supported by empirical evidence—the average price paid for all items purchased in department stores fell from 94% of list price in 1967 to 80% of list in 1997; by 1998 more than 72% of all fashion items were sold at a discount, compared with less than 20% in 1950 (Phillips 2005). An automated system has even been proposed to predict future online airline seat prices based on past seller behavior to counsel customers when to wait for a better deal (Etzioni et al. 2003, Cook 2004).

This paper focuses on how a seller’s markdown policy would change if he considered the influence of his current actions on the future behavior of customers. We consider a situation in which a retailer will interact with customers many times and customers learn from the retailer’s past actions. One example is fashion retailing, in which markdowns occur during well-established seasons. Levy and Woo (1999) describe a typical case in which a fashion item was discounted three times during a 26-week season: the first time by 20%, the second time by 50%, and the third time by 70%. Because stores often schedule markdowns to occur at the same times each year (and often at the same discount level), savvy customers quickly learn that they are likely to save if they are willing to wait. In a similar vein, European tour operators begin offering tours approximately 18 months prior to departure. At first, tours are offered at list price (the so-called “brochure price”), but over time, prices are reduced for tours that are not selling well. Customers know that they can realize discounts of 50% or more on the brochure price by waiting. In these markets, a seller who offers a late discount is not only influencing current sales, he is also training customers to wait. We investigate how this training effect might change the seller’s optimal markdown policy. Specifically, we address the question of whether or not there are situations in which a seller would be

better off allowing merchandise to perish than selling it at a discount.

The decision faced by a customer in a markdown market is whether to buy now or wait for a discount. This decision would incorporate the customer’s belief about whether or not the item in question will be discounted at some point in the future, when it will be discounted, and by how much, as well as the probability that it will still be available if she decides to wait. In other words, customers are strategic. The focus of this paper is on models in which customers update their expectations of a seller’s behavior over many different interactions. We show that this can result in optimal markdown policies that are quite different from the tactical solutions. Optimal pricing with strategic customers has been addressed by a number of authors, including Harris and Raviv (1981) and Besanko and Winston (1990) and more recently in a working paper by Cachon and Swinney (2007). However, our paper is, to our knowledge, the first to address the question of optimal markdown policies in the face of customers who learn from past behavior.

2. Outline of the Paper

In §3, we introduce a general two-period model in which a seller can offer all or some of his unsold inventory for sale in the second period at a discount. In §4, we consider a fluid version of the model in which the total number of consumers in the system is constant and the willingness-to-pay distribution is stationary over time. We show that for this model a “consistent expectations equilibrium” always exists in which the probability assigned by customers that discounted goods will be available in the second period is consistent with the amount of inventory that the seller actually makes available.

In §4 we start by considering the case in which a seller has the freedom to set prices and to determine how much inventory to sell in each period, all customers are strategic, and both the willingness-to-pay distribution and number of customers are known to the seller. In this case, we show that it is optimal for the seller to set a single price across both periods and not to limit second-period sales. We then show that this result requires each of the assumptions—if the population consists of a mixture of strategic and tactical customers, then a two-price markdown policy can be optimal. If prices are set exogenously, it can be optimal for the seller to limit second-period sales. If the seller is uncertain about the size or composition of demand, then a two-price markdown policy can be optimal.

In §5, we consider models in which total demand is a random variable and the seller wishes to maximize total revenue across a time horizon consisting of

many different interactions. We use numerical simulations to explore the type of behavior that can occur. A key insight is that the stability of demand and the optimal policy of the seller depends critically upon the speed with which customers update their expectations. If customers never update their expectations (i.e., they are “myopic”), then the policy of selling all unsold inventory in the second period at a discount is optimal. As customers begin to update expectations more rapidly, policies in which only some unsold inventory is released in the second period can become optimal. The character of the optimal policy depends upon the speed of updating, on the prices in the two periods, and on the underlying willingness-to-pay distribution.

In §6, we discuss the implication of our results as well as some opportunities for further research.

3. Model and Notation

A seller has a fixed stock C of a good that is available for sale over two periods. The item perishes at the end of the second period with no residual or salvage value to the seller. (The assumption of no salvage value is not restrictive. Our models could be modified to incorporate a nonzero salvage value for unsold inventory without changing the nature of the results.) There is a population of size D of potential customers. Each potential customer can be characterized by a vector (W_1, W_2) where $W_i \geq 0$ is willingness to pay in period $i = 1, 2$. We define $f(w_1, w_2)$ on \mathfrak{R}^2 as the density function over willingness to pay with $F(w_1, w_2)$ the corresponding cumulative distribution function (c.d.f.).

We consider the case in which $W_2 \leq W_1$ for all customers—that is, the situation in which purchasing the item in the second period is an inferior alternative. The assumption of nonincreasing willingness to pay and perishable inventory characterizes a wide range of so-called “markdown industries.” In a markdown industry, the value of a good declines over time. The reasons for the decline may be deteriorating quality (as in day-old bread), loss of use (as in a coat whose usefulness declines as winter progresses), fashion, or approaching obsolescence. Phillips (2005, Chapter 9) provides a broader discussion of these factors and their effects in different markdown industries.

At the beginning of the first period, the seller announces prices r_1 and r_2 . Customers know both W_1 and W_2 at the beginning of the first period. However, the seller only knows the density of the joint distribution $f(w_1, w_2)$. Because $W_1 \geq W_2$ for all customers, a revenue-maximizing seller will only consider prices such that $r_1 \geq r_2$. Each customer determines in which (if either) of the two periods she will seek to purchase. Assume that a risk-neutral customer has a prior probability, say p , at the beginning of period one that the

good will be available in the second period. The buying behavior of each customer can be characterized by the following:

$$W_1 \geq p(W_2 - r_2)^+ + r_1$$

→ seeks to purchase in period 1 (1)

$$W_2 \geq (W_1 - r_1)^+ / p + r_2$$

→ seeks to purchase in period 2 (2)

$$W_1 < r_1 \quad \text{and} \quad W_2 < r_2$$

→ does not seek to purchase. (3)

Note that the seller may sell all of his inventory in the first period. In this case none of the customers who decide to wait until period two will be able to buy. In fact, some of the period-one customers may also be unsatisfied. Another possibility is that all period-one customers are satisfied but the remaining inventory is not sufficient to satisfy all remaining period-two customers. When inventory is insufficient to satisfy demand, we assume that limited supply is rationed among potential customers in a way that is independent of the customers’ willingness to pay. Perhaps the most common practice is “first come, first served,” in which customers are served in order of arrival where order of arrival is independent of willingness to pay.

Denote the unconstrained demands experienced by the seller in periods one and two as D_1 and D_2 , respectively. In each period, the seller may choose to set a sales limit $0 \leq b_i \leq C$, $i = 1, 2$ on the amount of inventory that he will allow to be sold in that period. The expected total revenue (TR) over a single interaction, where E is the expected value operator, is given by

$$\begin{aligned} \text{TR}(b_1, b_2) = & r_1 E[\min(D_1, b_1)] \\ & + r_2 E \min[D_2, b_2, \max(C - D_1, C - b_1)]. \end{aligned}$$

The seller’s *tactical* problem is to determine the values of b_i —say \hat{b}_i , $i = 1, 2$ —that maximize $\text{TR}(b_1, b_2)$. Since inventory is perishable, it is always tactically optimal to have $\hat{b}_1 = \hat{b}_2 = C$. Note that this is due to the fact that $r_1 \geq r_2$. If $r_1 < r_2$, then by Littlewood’s rule, $\hat{b}_1 = (C - G_2^{-1}(1 - r_1/r_2))^+$, where G_2 is the c.d.f. on D_2 (Littlewood 1972).

We now extend the model to consider multiple interactions between customers and the seller. For fashion goods, each interaction would represent a season. For simplicity, we assume that each customer seeks to purchase at most a single unit. Let $D_i(n)$ denote the unconstrained demand in period i during interaction $n \geq 0$, $p_i(n)$ a customer’s subjective probability of availability during period i in interaction n , and $b_i(n)$ the amount of inventory that the seller will offer for sale at price r_i during interaction n . Let $x_i(n)$ denote sales and $y_i(n) = x_i(n)/D_i(n)$ the fraction of

unconstrained demand in period i that is satisfied in interaction n . That is, a customer who seeks to purchase in period i on interaction n will find the good available with probability $y_i(n)$. We consider models in which customers' posterior probability that a good will be available during a period is a function of the fraction of total demand that was satisfied during that period in past interactions:

$$p_i(n) = g[y_i(n-1), y_i(n-2), \dots, y_i(1)], \quad n \geq 2 \quad (4)$$

$$p_i(1) = p_i. \quad (5)$$

This is the basic framework that we will use to address two questions. First, under what conditions does it make sense for a seller to adopt a mark-down policy, that is, announce a lower price for second-period sales? Second, once prices have been announced, under what conditions does it make sense for a seller to set a binding second-period sales limit?

4. Fluid Model

We define a fluid model in which the number of customers in the system is constant, the willingness-to-pay distribution is stationary, and both the number of customers and the distribution are known to the seller. The elements of this model are

- $n = 1, 2, \dots =$ interaction number.
- $W_i \geq 0$ willingness to pay of a customer in period $i = 1, 2$.
- $D =$ total customer population assumed constant across interactions. For what follows we will, without loss of generality, set $D = 1$.
- $D_i(n) =$ demand in period i on interaction n .
- $b =$ sales limit applied in period 2 with $0 \leq b \leq 1$.
- $p(n) =$ common subjective probability of second-period availability.
- $p(1) = p_1 =$ initial probability, prior to the first interaction.
- $r_i =$ price in period $i = 1, 2$ with $r_1 > r_2 > c = 0$.

Consistent with the usual assumptions in markdown management, we assume that unit cost is zero. Thus revenue maximization is equivalent to profit maximization. Moreover, all inventory is available for sale at price r_1 in period one.

• $\bar{F}(x) =$ complementary cumulative distribution function of the willingness-to-pay distribution; that is, $\bar{F}(x) =$ the fraction of the population with willingness to pay greater than x . For this model, we assume that the willingness-to-pay distribution is continuous and constant across interactions and that the willingness to pay is the same for each customer in periods one and two.

Since $W_1 = W_2 = W$ and $r_1 > r_2$, then $W > r_1$ implies $W > r_2$, so Equation (1) holds if and only if $W \geq r(p) =$

$(r_1 - pr_2)/(1 - p)$, while Equation (2) holds if and only if $r_2 \leq W < r(p)$. As a result,

$$D_1(n) = \bar{F}(r(p(n))) \quad \text{and} \quad D_2(n) = \bar{F}(r_2) - D_1(n). \quad (6)$$

Notice that the total demand across both periods is equal to $D_1(n) + D_2(n) = \bar{F}(r_2)$.

The final element of the fluid model is to specify how $p(n)$ evolves from interaction to interaction and how this is affected by the sales limit b . We assume that on each interaction, customers observe the fraction of second-period demand that is satisfied, $y(n) = \min(1, b/D_2(n))$, and update their probabilities based on $y(n)$. Specifically, we consider updating formulas of the form $p(n+1) = Z[y(n), p(n)]$ such that $Z[y, p]$ is a continuous function of both variables and has the property

$$Z[y(n), p(n)] = p(n) \iff y(n) = p(n). \quad (7)$$

Property 7 states that the updated probability $p(n+1)$ equals the prior probability $p(n)$ if and only if the realized fraction $y(n)$ is equal to the prior. This requirement is nonrestrictive and is met by most common updating schemes such as exponential smoothing. For this family of updating mechanisms and any constant value of b , the following lemma guarantees the existence of a "consistent expectation" value of p such that $y(n) = p(n) = p(n+1) = p$.

LEMMA 1. *For any $Z[y, p]$ continuous in y and p and satisfying Property 7, and any $b(1) = b(2) = \dots = b \geq 0$, there exists a p such that $y(n) = p(n) = p \iff y(n+1) = p(n+1) = p$. If $D_2(n)$ is a strictly increasing function of $p(n)$, this p is unique.*

PROOF. Define $H[y(n), p(n)] = [y(n+1), p(n+1)]$. By the continuity of W and Z , H is a continuous function of $y(n)$ and $p(n)$ on the closed set $S = [0, 1] \times [0, 1]$, so the existence of a fixed point (y, p) such that $H[y, p] = (y, p) \in S$ follows from Brouwer's fixed-point theorem. The fact that $y = p$ is a consequence of Property 7. We prove uniqueness by contradiction. Let $D_2(n)$ be a strictly increasing function of $p(n)$ and assume that there are two distinct fixed points $p_1 > p_2$. If $p_1 < 1$, then $p_1 = b/D_2[p_1] < b/D_2[p_2] = p_2$, a contradiction. If $p_1 = 1$, then $D_2[p_1] \leq b < D_2[p_2]$, which is also a contradiction. \square

Note that at p , the prior probability that a customer will be able to purchase in the second period is equal to the fraction of customers that is actually served. We consider the problem in which the seller sets r_1, r_2 , and b to maximize equilibrium contribution. This may not be the optimal solution, starting at an arbitrary value of p_0 , due to the contribution of nonequilibrium states. (We consider the effects of such dynamics in §5.)

Because we are considering only the equilibrium strategy, for this section we drop the interaction index n and write $p(b)$ to highlight the dependence of p on b . Define $\text{TR}(b)$ as the steady-state revenue per interaction given $b \geq 0$. The problem facing the seller is then

$$\max_b \text{TR}(b) = r_1 \bar{F}(r(p(b))) + r_2 \min\{b, \bar{F}(r_2) - \bar{F}(r(p(b)))\}.$$

We can divide $\text{TR}(b)$ into two regions:

1. $b < \bar{F}(r_2) - \bar{F}(r(p(b)))$. Then, $\text{TR}(b) = r_1 \bar{F}(r(p(b))) + r_2 b$.

2. $b \geq \bar{F}(r_2) - \bar{F}(r(p(b)))$. Then $p(b) = 1$, so $r(p(b)) = \infty$, resulting in $\text{TR}(b) = r_2 \bar{F}(r_2)$.

In particular, if $b = 0$, then $p(b) = 0$ and $\text{TR}(0) = r_1 \bar{F}(r_1)$, whereas if $b \geq \bar{F}(r_2)$, then $p(b) = 1$ and $\text{TR}(\bar{F}(r_2)) = r_2 \bar{F}(r_2)$, so these two cases correspond to selling only in period one or selling only in period two. Since there is a one-to-one correspondence between b and p for b in the interesting region $b \in [0, \bar{F}(r_2) - \bar{F}(r(p(b)))]$, we can analyze the problem as a function of p only where $b(p) = p[\bar{F}(r_2) - \bar{F}(r(p))]$.

Letting $H(r) = r\bar{F}(r)$ and after some algebra, we have

$$\begin{aligned} \text{TR}(p) &= r_1 \bar{F}(r(p)) + pr_2(\bar{F}(r_2) - \bar{F}(r(p))) \\ &= (1-p)H(r(p)) + pH(r_2), \end{aligned}$$

so $\text{TR}(p)$ is a convex combination of $H(r(p))$ and $H(r_2)$.

Recall that without loss of generality we let the customer demand $D = 1$, and we assume that the capacity of the seller is C . Then the optimization problem faced by the seller is

$$\begin{aligned} \max_p \quad & \text{TR}(p) \\ \text{s.t.} \quad & (1-p)\bar{F}(r(p)) + p\bar{F}(r_2) \leq C \\ & p \in [0, 1]. \end{aligned}$$

PROPOSITION 1. *If all customers act strategically and $H(r)$ is unimodal with a maximizer $r^* \leq r_1$, then*

$$\begin{aligned} \text{TR}(p) &\leq \max(\text{TR}(0), \text{TR}(1)) \\ &= \max(H(r_1), H(r_2)) \leq H(r^*). \end{aligned}$$

PROOF. The proof will be by contradiction. Suppose first that $H(r_1) \geq H(r_2)$ and there exists an optimal $p \in (0, 1)$ such that

$$(1-p)H(r(p)) + pH(r_2) > H(r_1).$$

Then,

$$(1-p)H(r_1) + pH(r_2) \geq (1-p)H(r(p)) + pH(r_2) > H(r_1)$$

because $H(r_1) \geq H(r(p))$, because $r(p) > r_1 > r^*$ and H is unimodal, but then $H(r_2) > H(r_1)$, contradicting the hypothesis.

Suppose now that $H(r_2) \geq H(r_1)$ and there exists an optimal $p \in (0, 1)$ such that

$$(1-p)H(r(p)) + pH(r_2) > H(r_2),$$

but this implies that $H(r_1) \geq H(r(p)) > H(r_2)$, again contradicting the hypothesis. \square

Proposition 1 shows that in the absence of a capacity constraint, it is optimal to set $r_2 = r^*$, $r_1 \geq r_2$ and $p = 1$. If C is finite and $r(C) = \inf\{r > 0: \bar{F}(r) \leq C\}$, then the optimal value of r_2 is $\max(r^*, r(C))$. Notice that for any $r_1 > r_2 = \max(r^*, r(C))$, all customers will buy in period two. Thus, the same results can be achieved by setting $r_1 = r_2 = r^*$ and a two-price policy is not necessary to achieve optimality.

If the seller is free to select r_1 , r_2 , and b ; all customers act strategically; and the customer willingness-to-pay distribution is known and does not change between periods one and two, then a two-price mark-down strategy cannot increase revenue over the best single-price strategy $\max(r^*, r(C))$. In other words, at equilibrium, there is no opportunity for intertemporal segmentation when all customers are strategic and total demand and the willingness-to-pay distribution are both known by the seller. In the following sections, we consider cases in which one or more of these assumptions are relaxed.

The case is different if prices are set exogenously. This could be the case if a seller does not have the flexibility to change prices in response to short-term changes in demand. For example, Broadway theaters do not tend to adjust their prices dynamically in response to seasonal changes in demand (Leslie 2004). Demand may be greater on Thursday night than Wednesday night, but the list price and the same-day purchase discount price will be the same on both nights. Thus, the prices are likely to be nonoptimal for at least one of the two nights.

We consider the possibility that a seller might wish to pursue a *mixed strategy*—that is, offer all unsold second-period inventory for sale at a discount with some probability less than 1. As before, we assume sufficient capacity to serve all demand. Then at equilibrium, the probability assigned by a customer that he will be able to purchase at a discount in the second period will be the same as the probability that the seller is offering it for sale, which we will denote by p . We consider the situation in which the seller chooses p to maximize expected total revenue per interaction:

$$E \text{TR}(p) = r_1 \bar{F}(r(p)) + pr_2(\bar{F}(r_2) - \bar{F}(r(p))), \quad (8)$$

where $r(p) = (r_1 - pr_2)/(1-p)$, and as before, we assume without loss of generality (w.l.o.g.) that total demand $D = 1$.

PROPOSITION 2. If

$$\bar{F}'(r_1) > -\frac{r_2(\bar{F}(r_2) - \bar{F}(r_1))}{r_1(r_1 - r_2)} \quad \text{and}$$

$$r_1\bar{F}(r_1) \geq r_2\bar{F}(r_2),$$

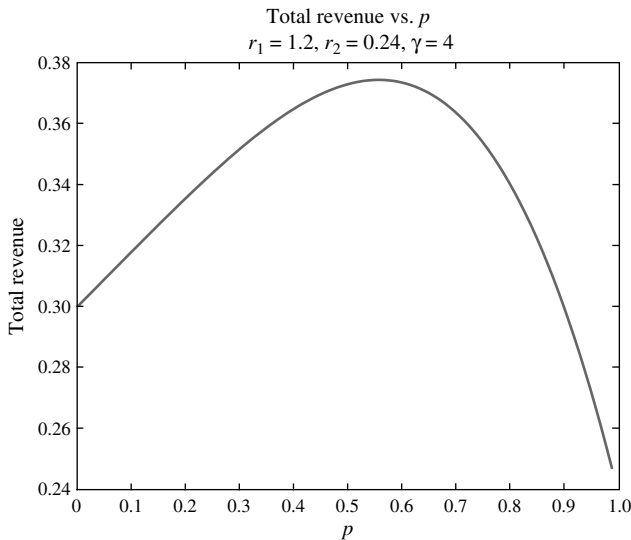
then there exists an optimal mixed strategy on the part of the seller with $p^* \in (0, 1)$.

PROOF. First note that $p = (r(p) - r_1)/(r(p) - r_2)$ is a continuous, strictly increasing function that maps $r(p) \in (r_1, \infty)$ to $p \in [0, 1)$. This means that any finite $r(p) > r_1$ maps one-to-one to a finite $p \in (0, 1)$. The first condition implies $d(ETR)/dr|_{p=0} > 0$. The second condition implies that $ETR(0) \geq ETR(1)$. These two conditions, along with the fact that $ETR(p) \leq r_1\bar{F}(r_2)$, imply the existence of an internal maximizer. \square

For example, assume that willingness to pay is distributed according to $\bar{F}(r) = \bar{F}(r_1)[1 - (r - r_1)^\gamma/(1 - r_2)^\gamma]$ for $\gamma > 1$ with $0 < (r_2/r_1)\bar{F}(r_2) \leq \bar{F}(r_1) < \bar{F}(r_2) < 1$. It is easy to show that this case satisfies the conditions of Proposition 2 and that a mixed strategy must therefore be optimal. Figure 1 illustrates the dependence of expected revenue on p for this distribution with $\gamma = 4$. In this case, the payoff from playing a mixed strategy over a pure strategy can be substantial. The best pure strategy is to set $p = 0$ with expected revenue $ETR(0) = 0.3$. Playing the best mixed strategy gives $ETR(0.6) = 0.375$, a 25% improvement.

We have shown that a single-price policy is optimal if all customers are strategic with a common willingness-to-pay distribution that is known to the seller and that does not change between periods. It is natural to speculate that the optimality of a

Figure 1 Expected Revenue as a Function of the Mixed Strategy Probability p for the Willingness-to-Pay Distribution
 $F(r) = \bar{F}(r_1)[1 - (r - r_1)^{0.4}/(1 - r_2)^{0.4}]$



single-price policy may be dependent on the fact that willingness to pay is the same in both periods; that is, $W_1 = W_2$ for all customers. Many goods that are marked down, such as fashion goods or day-old bread, are considered less desirable over time. However, we can show that a decline in customer willingness to pay such that $W_2 = W_1 - \theta$ with $\theta > 0$ in itself does not justify a two-price markdown policy.

Suppose that $W_2 = W_1 - \theta$ for all customers in the population. Given $p \in [0, 1]$ and $r_1 \geq r_2 + \theta$ define $r(p, \theta) = (r_1 - p(r_2 + \theta))/(1 - p)$. Then customers with $W_1 \geq r(p, \theta)$ will buy at r_1 , customers with $r(p, \theta) > W_1 \geq r_2$ will buy at r_2 , and customers with $W_1 < r_2$ will not buy. The corresponding revenue function is

$$TR(p) = (1 - p)H(r(p, \theta)) + p[H(r_2 + \theta) + \theta(\bar{F}(r(p, \theta)) - \bar{F}(r_2 + \theta))].$$

Notice that

$$TR(p) \leq (1 - p)H(r(p, \theta)) + pH(r_2 + \theta),$$

because $\bar{F}(r(p, \theta)) \leq \bar{F}(r_2 + \theta)$.

From this, it follows that $TR(p) \leq \max(H(r_1), H(r_2 + \theta)) \leq H(r^*)$. In conclusion, the revenues are bounded above by $TR(p) \leq H(r^*)$, and this upper bound can be achieved by selecting $r_1 = r^*$ and $p = 0$ with obvious modifications if C is binding.

4.1. Mixtures of Strategic and Myopic Customers

In the previous section we showed that, under reasonable assumptions, a single-price policy with no sales limit can maximize revenue when demand is known and all customers act strategically—even if they place a lower value on second-period consumption. Strategic customers can be contrasted with *myopic customers* who will purchase the product in period one if $W > r_1$. If $r_1 \geq W > r_2$, they will purchase in the second period. What if the seller faces a mixture of myopic and strategic customers? We show in this section that a two-price markdown policy can lead to significant gains in revenue compared with a single-price policy if even a small fraction of the customers act myopically and the rest of the customers act strategically.

Suppose that there are D customers and that $D\rho$ behave strategically and $D\bar{\rho}$ act myopically where $\bar{\rho} = 1 - \rho$, $\rho \in [0, 1]$. We assume that all customers have the same willingness-to-pay distribution in each of the two periods. Expected customer revenue for a particular choice of p , r_1 , and r_2 is

$$TR(p, r_1, r_2) = r_1D[\rho\bar{F}(r(p)) + (1 - \rho)\bar{F}(r_1)] + pr_2D[\rho(\bar{F}(r_2) - \bar{F}(r(p))) + (1 - \rho)(\bar{F}(r_2) - \bar{F}(r_1))].$$

Now $r_2 = r^*$ and $p = 1$ yields

$$TR(1, \rho) = DH(r^*) + D(1 - \rho)(r_1 - r_2)\bar{F}(r_1) \geq DH(r^*),$$

so two prices can be better than one!

We now consider the problem of determining the optimal values of p , r_1 , and r_2 when the population consists of a mixture of strategic and myopic customers.

PROPOSITION 3. If $\bar{F}(r_1)(\rho r_1 + \bar{\rho} r_2) < r_2 \bar{F}(r_2)$, then

$$TR(p, r_1, r_2) < TR(1, r_1, r_2).$$

PROOF. The condition implies

$$\rho H(r(p)) < \rho H(r_2) + \bar{\rho} r_2 [\bar{F}(r_2) - \bar{F}(r_1)]. \quad (9)$$

Because $H(r(p))$ is maximized at $p = 0$, a sufficient condition for the inequality (9) to hold is $\rho H(r_1) < \rho H(r_2) + \bar{\rho} r_2 [\bar{F}(r_2) - \bar{F}(r_1)]$, which simplifies to $\bar{F}(r_1)(\rho r_1 + \bar{\rho} r_2) < \bar{F}(r_2) r_2$. \square

Notice that $\bar{F}(r_1)(\rho r_1 + \bar{\rho} r_2) \leq H(r_1) \leq H(r^*)$, so the condition of Proposition 3 always holds if $r_2 = r^*$. This means that for a large choice of prices r_1 and r_2 , it is optimal to allow all the remaining inventory to be sold at r_2 ! In other words, it is optimal to set prices so that all of the strategic customers purchase in the first period. For $p = 1$, the optimization problem is given by

$$\begin{aligned} \max_{r_1, r_2} \quad & TR(1, r_1, r_2) = D[(r_1 - r_2)(\bar{\rho} \bar{F}(r_1)) + r_2 \bar{F}(r_2)] \\ \text{s.t.} \quad & D \bar{F}(r_2) \leq C. \end{aligned}$$

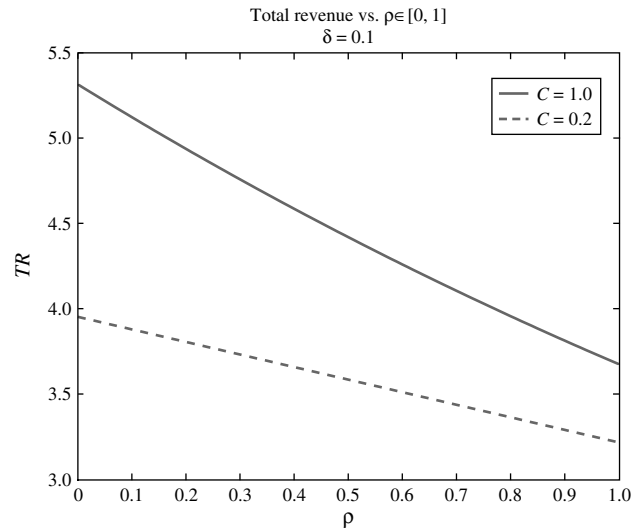
We can solve this problem explicitly for r_1 and r_2 for $p = 1$ in the case of an exponential willingness-to-pay distribution with $\bar{F}(r) = e^{-\delta r}$.

PROPOSITION 4. For the exponential willingness-to-pay distribution with parameter δ , the following policy is optimal:

$$\begin{aligned} r_1 &= r_2 + \frac{1}{\delta} \\ r_2 &= \max\left(\frac{1 - \bar{\rho} e^{-1}}{\delta}, r(C)\right) \\ p &= 1. \end{aligned}$$

The proof is straightforward and therefore omitted. The reader can verify that the solution for the exponential willingness to pay satisfies the conditions of Proposition 3. Figure 2 shows how revenues deteriorate as the proportion of strategic customers is increased. Notice that the decline is nearly linear and actually linear over the regions where capacity is tight. In this case the revenue decreases at rate $CE[W]/e$.

Figure 2 Revenue as a Function of $\rho \in [0, 1]$ Where ρ Is the Fraction of Strategic Customers in the Population and $1 - \rho$ Is the Fraction of Myopic Customers in the Case When Willingness to Pay Follows an Exponential Distribution with Parameter δ



4.2. Segmentation with Demand Uncertainty

In this section we will show that a markdown policy can be superior to a single-price policy when all customers are strategic if the seller is uncertain about the composition (or size) of his demand. Specifically, we assume that customer willingness to pay will follow one of two known distributions. In state i , the number of customers with willingness to pay greater than r is given by $d_i(r) = D \bar{F}_i(r)$. The seller announces prices r_1 and r_2 prior to knowing which state will hold. However, the seller knows the probability of state i , ρ_i for $i = 1, 2$ with $\rho_1 + \rho_2 = 1$. Once the seller announces prices r_1 and r_2 , nature chooses the state $i \in \{1, 2\}$. All customers are strategic. For what follows, we assume $d_1(r) > d_2(r)$ for all $r > 0$.

This could represent the situation in which a seller is uncertain whether a new fashion item is going to be a “hit” or a “miss” prior to ordering and setting prices. If it is a hit, demand will be high; if it is a miss, demand will be low. A similar but much simpler model was proposed by Lazear (1986) to explain the existence of markdown policies for fashion goods. Lazear’s model considered the case when all customers had the same willingness to pay for an item, but the seller was *ex ante* uncertain whether that willingness to pay was low or high. We extend the model to the case where customers follow a distribution of willingness to pay.

As before, a customer will purchase in period one if her willingness to pay exceeds $r(p) = (r_1 - pr_2)/(1 - p)$. Notice that r_1 , r_2 , and p completely determine $r(p)$. Define $\bar{r} = r(p)$ and note that

$$r_1 = pr_2 + (1 - p)\bar{r}. \quad (10)$$

In this section we will calculate the optimal policy in terms of \bar{r} , r_2 , and p and then derive the optimal value of r_1 using (10).

Fix \bar{r} , r_2 , and p and assume scenario i occurs. Then $d_i(r_2) - d_i(\bar{r})$ customers seek to purchase in period two but only $(C - d_i(\bar{r}))^+$ units are available. Let

$$\alpha_i = \min\left(1, \frac{(C - d_i(\bar{r}))^+}{d_i(r_2) - d_i(\bar{r})}\right).$$

Then α_i is the fraction of the second-period demand that will be satisfied if the seller does not limit sales. Assume that the seller does not limit sales. Then his expected revenue is given by

$$\sum_{i=1}^2 \rho_i [\alpha_i H_i(r_2) + (1 - \alpha_i) H_i(\bar{r})].$$

Because customers do not know which state holds prior to making their purchase choice, we must have

$$\sum_{i=1}^2 \rho_i \alpha_i = p.$$

We have shown that the best policy when the willingness-to-pay distribution is known is to set a first-period price to maximize expected revenue from first-period sales. Then in scenario i the seller would choose a price to maximize $H_i(r)$ subject to $d_i(r) \leq C$, resulting in optimal solution $\max(r_i^*, r_i(C))$ where r_i^* maximizes the revenue rate $H_i(r)$ and $r_i(C) = \inf\{r: d_i(r) \leq C\}$. Clearly, $r_1(C) > r_2(C)$, and it is reasonable to assume that $r_1^* > r_2^*$, so we will proceed under this assumption.

Now define $\bar{r}_1 = \max(r_1^*, r_1(C)) > \max(r_2^*, r_2(C)) = r_2$. We assume that $C \leq d_1(r_1^*)$, since otherwise there will be excess capacity in both scenarios. Under this additional assumption it follows that $\bar{r}_1 = r_1(C)$, so $d_1(\bar{r}_1) = C$ and $\alpha_1 = 0$. Because $d_2(\bar{r}_1) < d_1(\bar{r}_1) = C$, it follows that $\alpha_2 > 0$. Moreover, because $d_2(r_2) \leq C$, it follows that $\alpha_2 = 1$.

We have shown that if $d_1(r) > d_2(r)$, $r_1^* > r_2^*$ and $C \leq d_1(r_1^*)$, then it is optimal to set $r_2 = \max(r_2^*, r_2(C))$ and $\bar{r} = r_1(C)$, resulting in $\alpha_1 = 0$ and $\alpha_2 = 1$ with corresponding $p = \rho_2$ and revenues equal to

$$\rho_1 H_1(\bar{r}) + \rho_2 H_2(r_2) = (1 - p) H_1(\bar{r}) + p H_2(r_2)$$

and that this revenue is achieved by announcing prices r_2 and $r_1 = \rho_2 r_2 + (1 - \rho_2) \bar{r}$.

We will show briefly using examples that other forms of uncertainty such as random total demand and willingness-to-pay distribution with random parameter can also lead to a two-price markdown policy being optimal.

4.2.1. Random Total Demand. Suppose that the total demand D is a doubly stochastic Poisson process with random intensity Λ with any nonnegative distribution. In this case, the use of two prices can lead to increased revenue over a single-price policy. The proof of the optimality of two-price markdown policy is available from the authors. Suppose that Λ can take values 300 or 200 with probabilities $\rho_1 = 0.7$ and $\rho_2 = 0.3$, respectively. Assume that $C = 60$ and that willingness to pay is exponential with mean equal to 50. Setting $r_2 = 60.20$ and $r_1 = 78.67$ results in $p = \rho_2 = 0.3$ and total expected revenue of \$4,828.38 with $\alpha_1 = 0$ and $\alpha_2 = 1$. In contrast, the best the seller could do with a single price is \$4,345.48 with $r = 80.47 = p(C, 300)$, so using two prices increases revenues by 11%.

4.2.2. Willingness-to-Pay Distribution with Random Parameter. Assume that the willingness-to-pay distribution is exponential with random parameter Γ so that demand at price r is $\lambda e^{-r\Gamma}$, where λ is the expected total number of customers. For exponential willingness-to-pay distribution with parameter Γ that has nonnegative support, one can improve revenues by using a two-price markdown policy over a single-price policy. Here, $C = 60$, $\lambda = 250$, and $\Gamma = \gamma_1 = 1/70$ with probability 70% and $\Gamma = \gamma_2 = 1/50$ with probability 30%. The best single price is $p(C, \gamma_1) = 99.89$, resulting in revenue of \$5,211.77. In contrast, using price $r_2 = p(C, \gamma_2)$ and setting r_1 and $p = \alpha$ such that $r(p) = p(C, \gamma_1)$ is an optimal solution with two prices resulting in $r_2 = \$71.36$, $r_1 = \$91.33$, $\alpha = p = 0.30$ and revenues equal to \$5,480.27, representing a 5.15% improvement over the best single price.

5. Dynamic Models with Learning

In the previous section we considered the equilibrium situation in which the *ex ante* estimates of the probability of second-period availability are equal to the fraction of customers that will be served in the second period. We have shown that, in this situation, it is optimal for sellers to pursue a single-price policy if all customers are strategic and both total demand and the distribution of customer willingness to pay is known. If any one of these conditions is violated, a two-price markdown solution may become optimal.

Although equilibrium behavior is clearly important, it does not illuminate one of the most interesting features of markdown markets. It is in the interest of the seller to do everything possible to convince customers that little or no inventory will be available in the second period (“Buy now, supply is limited”). Yet when the second period arises, it is in the seller’s tactical interest to sell as much of his remaining inventory as he can. If customers form their expectations

of the seller's future behavior from his past behavior, this creates a tension between the seller's tactical interest in selling as much as he can now and his strategic desire not to "train" customers to expect high levels of future discount availability. In particular, it may be in the seller's interest to prevent the market from settling into equilibrium—if he can. In fact, we will show that, in a dynamic setting, complex nonequilibrium behavior is indeed possible.

In this section, we study multiperiod models in which aggregate demand D is a Poisson random variable with parameter λ . If all customers share the same value for p and the willingness to pay is drawn independently from the same distribution, then D_i , the demand in period $i = 1, 2$, is also Poisson with parameter λ_i given by

$$\lambda_1 = \lambda \Pr\{W_1 \geq w^*\} \quad (11)$$

$$\lambda_2 = \lambda \Pr\{W_1 < w^*, W_2 \geq r_2\}, \quad (12)$$

where $w^* = r_1 + p[W_2 - r_2]^+$. We consider the simple model in which $W_2 = W_1 - \rho$ for some $\rho \geq 0$; in other words, the difference between first-period and second-period willingness to pay is the same constant for every customer. In this case, we can specify the full dynamic system and Poisson demands $D_i(n)$ with parameters $\lambda_i(n)$, $i = 1, 2$:

$$p(0) = p_0 \quad (13)$$

$$p(n) = Z[y(n), p(n-1)] \quad (14)$$

$$\lambda_1 = \lambda P(W \geq r(p, \rho)) \quad (15)$$

$$\lambda_2 = \lambda P(r_2 \leq W \leq r(p, \rho)) \quad (16)$$

$$x_1(n) = \min[D_1(n), C] \quad (17)$$

$$x_2(n) = \min[D_2(n), C - x_1(n), b_2(n)] \quad (18)$$

$$y_i(n) = \frac{x_i(n)}{D_i(n)}, \quad i = 1, 2. \quad (19)$$

We consider the problem of setting optimal second-period sales limits for each of n successive sales interactions. This problem can be formulated as a dynamic program. However, the dynamic program using the simple updating rule $p(n) = \alpha y_2(n)/D_2(n) + (1 - \alpha)p(n-1)$ is computationally intractable. For this reason we use the updating approach

$$p(n+1) = \alpha \Pr\{D_2 \leq \min[C - D_1(n), b(n)]\} + (1 - \alpha)p(n). \quad (20)$$

That is, p is updated based on the probability that *all* second-period demand is met. This simplifies the computations considerably. α is a parameter that determines how rapidly customers update their expectations of inventory being available in the second period. If $\alpha = 0$, customers do not update at all

and $p(n) = p_0$ for any strategy used by the seller. If $\alpha = 1$, customers update their expectations immediately. Values of α between these extremes correspond to different updating speeds, with higher values representing more rapid updating.

We consider two types of policies. In a *dynamic sales limits* policy, the seller sets a new $b(n)$ in each interaction n . In a *static sales limits* policy, the seller sets a single $b = b(n)$ that applies during every interaction. We investigate the nature of the optimal policies in both cases.

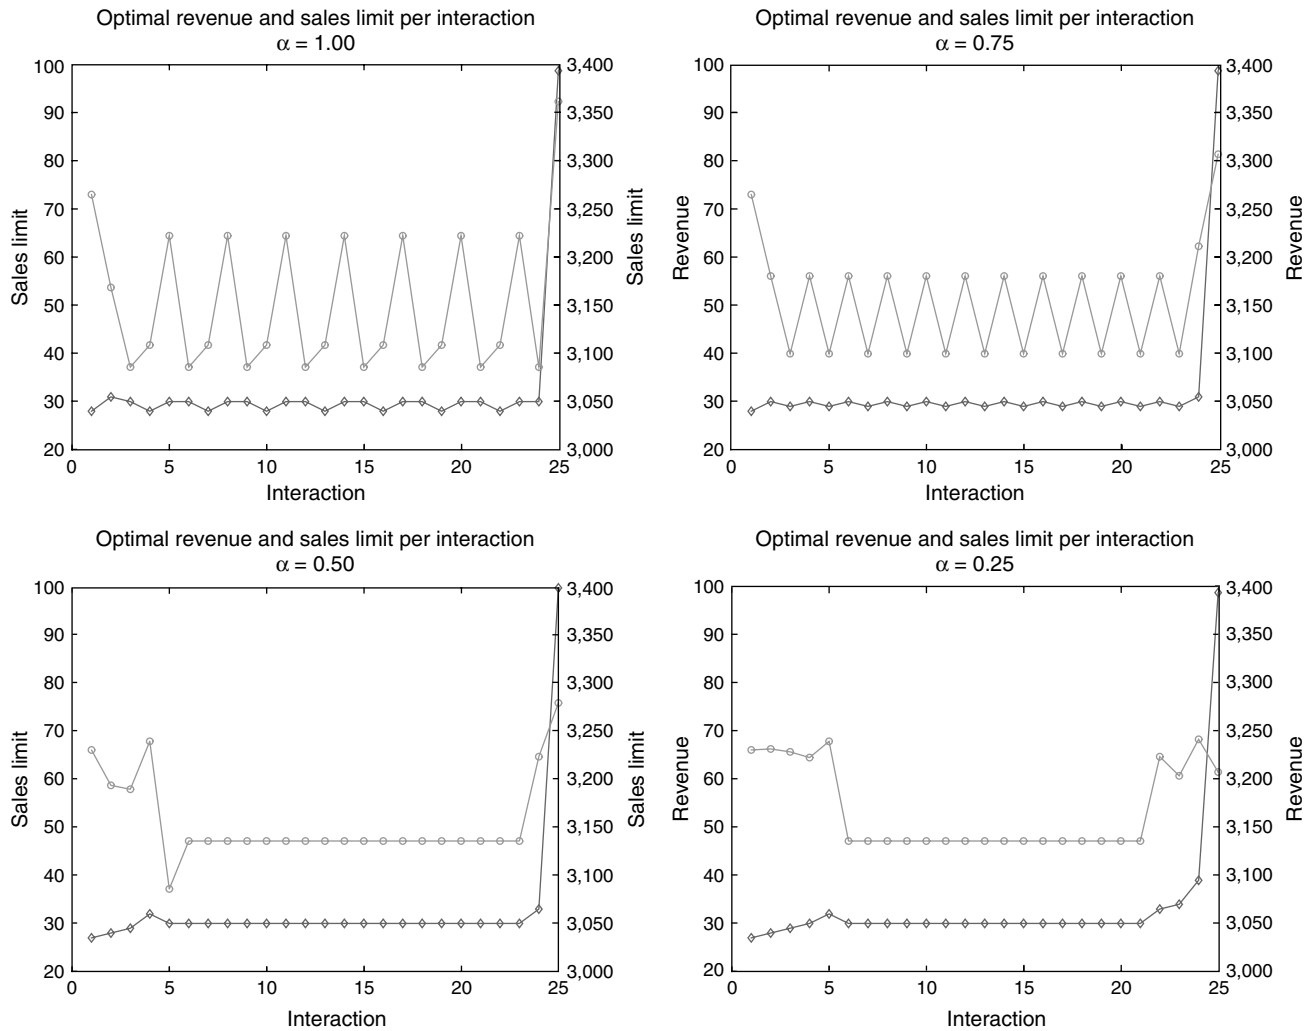
5.1. Dynamic Sales Limits

Let N be the total number of interactions. For each interaction, the decision variable is the second-period sales limit $b(n)$. We define $R(n, p(n))$ as the expected total revenue for periods $\{n, n+1, \dots, N\}$ given $p(n)$. Let β denote the discount factor between interactions. We can then formulate the problem as a dynamic program:

$$\begin{aligned} R(n, p(n)) &= \max_{0 \leq b(n) \leq C} \{r_1 E[\min(D_1(n), C)] \\ &\quad + r_2 E[\min(D_2(n), (C - \min(D_1(n), C)), b(n))] \\ &\quad + \beta E[R(n+1, p(n+1))]\} \\ p(n+1) &= \alpha \Pr\{D_2 \leq \min[C - D_1(n), b(n)]\} + (1 - \alpha)p(n) \\ R(N+1, p(N+1)) &= 0. \end{aligned}$$

Figure 3 illustrates the time tracks of optimal dynamic sales limits and revenues for four values of α when $N = 25$, $\lambda = 80$, $\beta = 1$, and $C = 100$. In each of the four panels in Figure 3, one track is optimal sales limit for each interaction, read from the left axis, and the other track is the corresponding revenue per interaction, read from the right axis. Note that in each case, the final-period optimal sales limit is $b(N) = C$. This is an end-of-horizon effect. However, in every period other than the last, it is optimal for the seller to set a sales limit < 100 . That is, it is optimal for the seller to restrict the amount of second-period inventory to be sold. The optimal limits converge to a three-point distribution with support at 28, 29, and 30 for $\alpha = 1$ and a two-point distribution with support at 29 and 30 for $\alpha = 0.75$. (Recall that α is a parameter that determines the speed of customer updating, with higher values of α reflecting more rapid updating.) For $\alpha = 0.25$ and $\alpha = 0.5$, the optimal sales limits converge to steady-state values. This suggests that if customers update quickly (large α), then the seller should use a strategy in which the sales limit is raised and lowered in alternate periods. For lower values of α , the seller should set a constant sales limit.

Figure 3 Optimal Revenue (◊) and Dynamic Sales Limits (◊) for Four Values of α with Poisson Arrivals



Note. For these runs, $N = 25$, $\lambda = 80$, and $C = 100$.

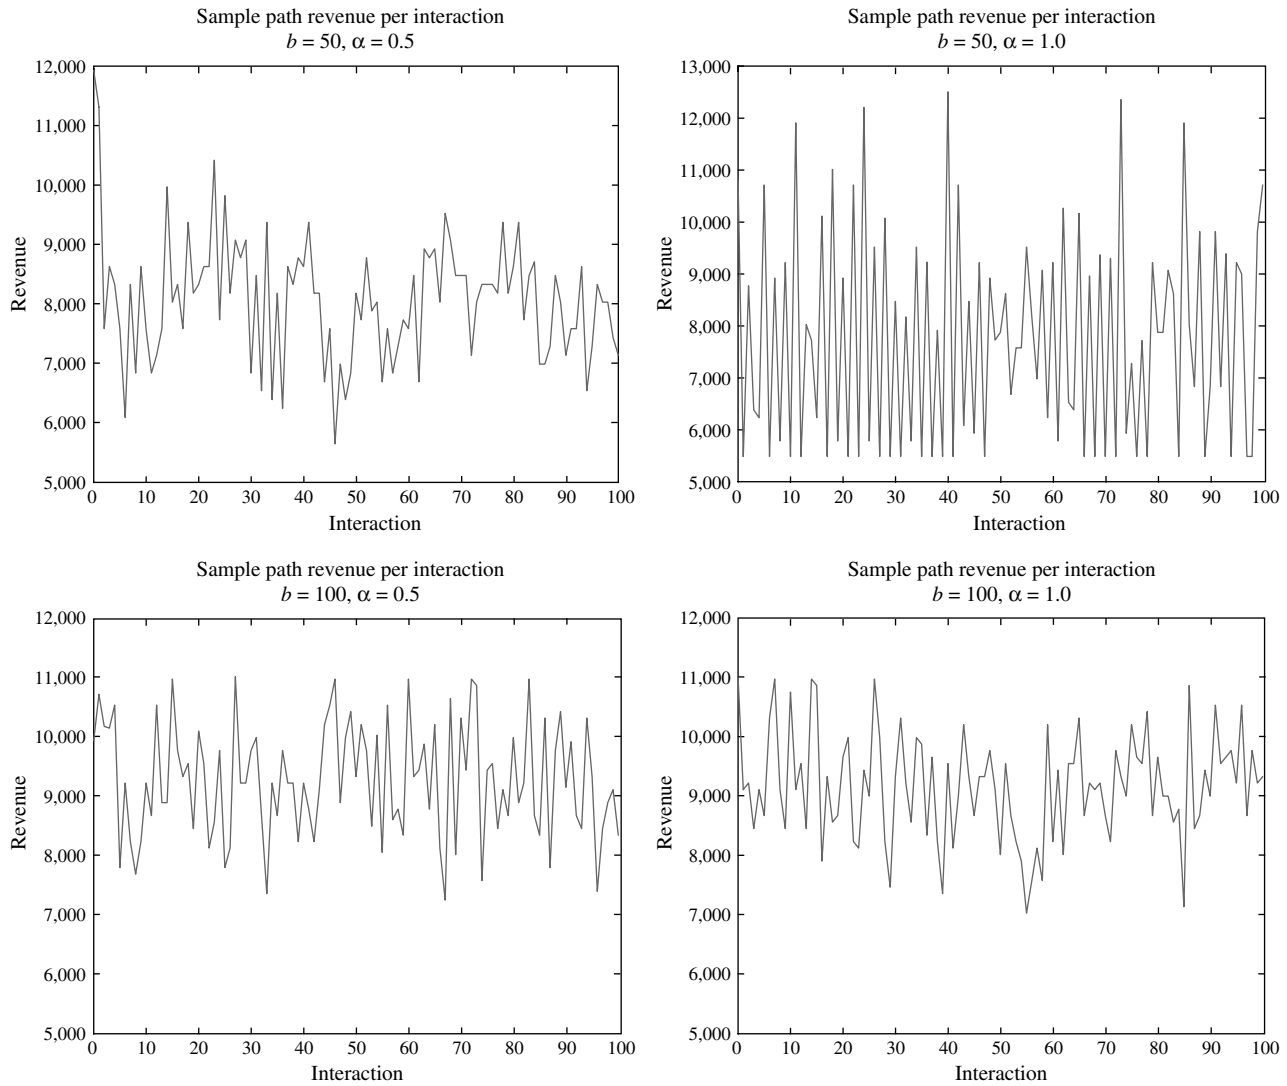
5.2. Static Sales Limits

The results from §5.1 suggest that a static sales limit can be optimal when the value of α is sufficiently low and the number of interactions is large. In this section, we compute the optimal *static* sales limit b . For these simulations, we assume a Poisson process with rates given by Equations (11) and (12). We consider exponential, triangular, and uniform willingness-to-pay distributions and set $W_2 = W_1 - \rho$ where $\rho \geq 0$. Inventory is fixed, $p(0) = 0$, and $p(n)$ is updated according to (20). The number of interactions is $N = 100$. To eliminate initialization effects, the results for the first 20 interactions are not shown.

Figure 4 shows some typical sample paths of revenue by interaction for different values of b and α in the case when willingness to pay follows an exponential distribution. These sample paths illustrate the fact that changes from interaction to interaction in customer behavior can lead to significantly greater variation in revenue as sales fluctuate beyond those arising from the underlying uncertainty in total demand. A

forecasting model that did not incorporate the dependence of future customer behavior on current and past sales limits would be unable to accurately forecast future demand. This could lead to a miscalculation of future sales limits and a corresponding sub-optimal revenue realization. The situation is reminiscent of the phenomenon described as “spiral down” by Homem-de-Mello et al. (2004) in the two-period revenue management model. Spiral down can occur when a revenue manager ignores the effects of changing the allocation for one booking class on bookings in other classes. Homem-de-Mello et al. show that spiral down can lead to significant loss of revenue. The sample paths in Figure 4 illustrate that sales variation among interactions tends to increase in α and decrease in b . This is intuitive—rapid updating by customers causes wider swings in behavior from interaction to interaction while p will tend to approach 1 as b increases, resulting in less variation from interaction to interaction. The dependence of demand variation on the sales limit suggests the

Figure 4 Some Typical Sample Paths for Different Values of α Where $\lambda = 250$, and Willingness to Pay Follows an Exponential Distribution with Mean 100, and the Prices Are $r_1 = 150$, $r_2 = 110$



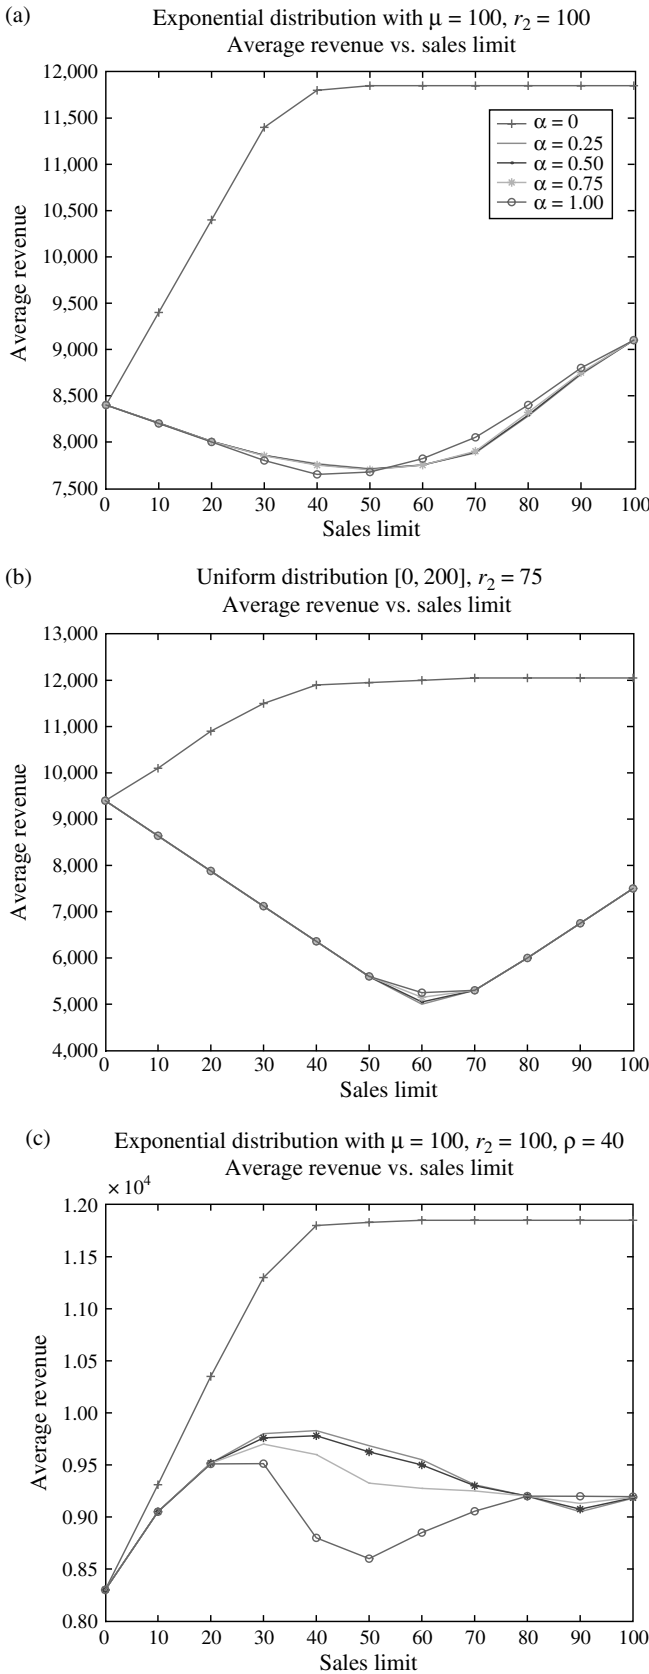
potential for a spiral down-like phenomenon in mark-down management—a planner who assumed that future demand was independent of past sales limits might make increasingly poor sales limit decisions.

Figure 5 shows expected revenue as a function of the sales limit for different willingness-to-pay distributions and different values of the updating parameter α . In these simulations, $N = 25$, $C = 100$, and $p_0 = 0$. For Figures 5(a) and 5(b) $\rho = 0$, and for Figure 5(c) $\rho = 40$. When customers do not update their expectations ($\alpha = 0$) and $p_0 = 0$, the optimal strategic value of b is the same as the optimal tactical value, namely, $b = C$. This situation is ideal for the seller—customers never anticipate the second-period sale so it is always optimal for the seller to allow all unsold inventory to be sold in the second period. Expected revenue drops radically as soon as customers begin to “wise up”—that is, $\alpha > 0$. In Figure 5(a), exponential distribution

with $\lambda = 250$ and common willingness to pay between the two periods ($\rho = 0$), the optimal strategic policy is still to allow all remaining inventory to be sold in the second period. However, expected revenue drops considerably from the nonanticipating case for all values of b . A similar drop is seen in the case of the uniform distribution (Figure 5(b)); however, in this case the optimal strategy for all values of α is for the seller never to allow any unsold inventory to be sold in the second period—that is, to set $b = 0$.

The case of the exponential distribution when $\rho = 40$ (Figure 5(c)) is more complex. Expected revenue still drops drastically as soon as customers begin to update. However, average revenue is now maximized at neither $b = 0$ nor $b = C$ but at an intermediate value that depends upon α . In fact, for $\alpha = 1$, average revenue has not only an intermediate maximum but also an intermediate minimum. This is intriguing

Figure 5 Expected Revenue as a Function of Sales Limit for Different Distributions and Values of α with $r_1 = 150$, $\lambda = 250$, $C = 100$



because it suggests that a seller who is trying to determine the optimal amount of inventory to withhold might miss the global maximum if he proceeds by a process of local adjustment. That is, if he starts by setting a sales limit of 70, which he gradually adjusts to increase revenue, he will end up ultimately setting $b = 100$ rather than at the global maximum near $b = 25$.

Also intriguing is the fact that $\alpha = 1$ results in higher revenues for all *positive* values of α tested over selected values of suboptimal sales limits, as can be seen in Figure 5. Notice, however, that this excludes $\alpha = 0$ and does not happen at optimal sales limits. The intuition for this is that if sales limits at r_2 are too high, then the seller may be better off if more customers wait to buy in period two, i.e., if more customers act strategically.

The results in Figure 5 are for a fixed r_2 . Define $R(b)$ as the expected revenue per interaction when b is the static sales limit. It is natural to speculate that $R(0) > R(C)$ when r_2 is small and $R(C) > R(0)$ when r_2 becomes sufficiently large. That is, when r_2 is small, it is better to withhold all second-period inventory, and when it is large, it is better for the seller to allow all second-period inventory to be sold. This is, in fact, the case. We note first that by the same method as Lemma 1 we can show that for any value of b , Equation (20) has a unique fixed point $p^*(b)$, utilizing the facts that $D = D_1 + D_2$ is independent of p and $\Pr\{\max[D - C, D_2(p) - b] \leq 0\}$ is a decreasing function of p . Moreover, $p^*(b)$ is increasing in b . This gives us what we need to prove the following:

PROPOSITION 5. *There exists an r_2^* such that $R(0) \geq R(C)$ if $r_2 \leq r_2^*$ and $R(0) < R(C)$ otherwise for $\alpha > 0$.*

PROOF. Let D_1^0 and D_1^C be the random variables denoting first-period demand when $b = 0$ and $b = C$, respectively. Note that $p^*(0) \rightarrow 0$ as $n \rightarrow \infty$ and thus,

$$R(0) = r_1 E \min(D_1^0, C),$$

where $D_1^0 \sim \text{Poisson}(\lambda \Pr(W > r_1))$. When $b = C$, (20) becomes

$$\begin{aligned} p(n+1) &= \alpha \Pr(D_2(n) \leq \min[C - D_1(n), C]) + (1 - \alpha)p(n) \\ &= \alpha \Pr(D_1(n) + D_2(n) \leq C) + (1 - \alpha)p(n), \end{aligned}$$

which goes to $\Pr(D \leq C)$ as $n \rightarrow \infty$ where

$$D \sim \text{Poisson}(\lambda \Pr(W \geq r_2)).$$

Therefore,

$$D_1^C \sim \text{Poisson}(\lambda \Pr(W > r(p)))$$

with $r(p) > r_1$, which implies $D_1^0 > D_1^C$. We also have

$$R(C) = r_1 E \min(D_1^C, C) + r_2 E \min[D_2^C, (C - D_1^C)^+].$$

$D_1^0 > D_1^C$ implies $r_1 E \min(D_1^0, C) > r_1 E \min(D_1^C, C)$; because $R(C)$ is a continuous function of r_2 , there exists r_2 such that $R(0) = R(C)$. \square

This suggests that the strategy of withholding all unsold inventory in the second period is superior to allowing it all to be sold when r_2 is small and vice versa. It is also clear that $b = 0$ is optimal when $r_2 = 0$ and $b = C$ is optimal when $r_2 = r_1$. Thus, as we increase r_2 from zero, it is initially optimal to withhold all unsold inventory. As r_2 increases, intermediate solutions are possible where it is optimal to set $b \in (0, C)$. As r_2 increases further, at some point it becomes optimal to allow all unsold inventory to be sold in the second period.

6. Summary and Discussion

We consider the situation in which a seller is offering a finite (although potentially very large) stock of perishable inventory over two periods, with inventory perishing at the end of the second period. We first consider the case in which

1. Sellers have the freedom to set prices in both periods;
2. All customers are “strategic”; that is, they choose in which period to purchase based on their expectation of prices;
3. Total customer demand is known to the seller; and
4. The willingness-to-pay distribution is also known to the seller.

Under these conditions, the optimal equilibrium policy is for the seller to set a single price and not restrict second-period sales. This can hold true even if customer willingness to pay is lower in the second period than in the first.

We have also shown that the four conditions listed above are all necessary for a single price to be optimal. If prices are set exogenously, then it can be optimal for the seller to limit second-period sales. If the seller is facing a population consisting of a mixture of strategic and myopic customers, a two-price markdown policy may be optimal. Finally, if the seller is uncertain about the magnitude of total demand or its composition, a two-price markdown policy may also be optimal. We note that both demand uncertainty and mixtures of strategic and myopic customers are likely to characterize most real-world markets and therefore help explain the widespread use of markdown policies by sellers in different markets.

We used numerical simulation to explore the case in which total demand is uncertain from interaction to interaction. In these cases, it is often optimal for the seller to limit second-period sales, even when all customers are strategic. The behavior of this optimal sales limit depends in part on the speed of customer updating. If customers update their expectations quickly,

then the optimal sales limit can oscillate between two values. If updating occurs more slowly, then the optimal sales limit tends to converge to a constant value. We also noted that faster customer updating results in higher variability in sales and revenue from interaction to interaction. This indicates that the dynamics of the system can increase variability significantly over that simply due to underlying demand uncertainty. A seller who assumed that demand was independent of his past policies might easily overestimate the variation in underlying customer demand.

There are several directions in which these models could be extended. For example, we have assumed that customers are risk neutral. We would expect in many cases customers would be risk averse. Risk neutrality implies that a customer with a willingness to pay of \$100 and facing a ticket price of \$80 would prefer to wait for the chance of a half-price ticket at TKTS if the probability of discount ticket availability in the second period was greater than 33%. This would not likely be the case for a parent who has promised to take her child to see *The Lion King*. Intuitively, a population of risk-averse customers would seem to make withholding second-period inventory more attractive since the corresponding reduction in probability of availability would make risk-averse customers even more likely to purchase in the first period.

References

- Besanko, D., W. L. Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* **36**(5) 555–567.
- Bitran, G. R., D. Mondsschein. 1997. Periodic pricing of seasonal products in retailing. *Management Sci.* **43**(1) 64–79.
- Cachon, G. P., R. Swinney. Purchasing, pricing, and quick response in the presence of strategic consumers. Working paper, University of Pennsylvania, Philadelphia.
- Cook, J. 2004. Venture capital: Airfare fluctuations launch startup that predicts prices. *Seattle Post-Intelligencer* (October 8), http://seattlepi.nwsource.com/venture/194294_vc08.html.
- Elmaghraby, W., A. Gülcü, P. Keskinocak. 2004. Optimal markdown mechanisms in the presence of rational customers with multi-unit demands. Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.
- Etzioni, O., R. Tuchinda, C. A. Knoblock, A. Yates. 2003. To buy or not to buy: Mining airfare data to minimize ticket purchase price. *Proc. 9th ACM SIGKDD Internat. Conf. Knowledge Discovery and Data Mining*, ACM Press, New York, 119–128.
- Feng, Y., G. Gallego. 1995. Optimal stopping times for end of season sales and optimal starting times for promotional fares. *Management Sci.* **41**(8) 1371–1391.
- Feng, Y., G. Gallego. 2000. Perishable asset revenue management with Markovian time dependent demand intensities. *Management Sci.* **46**(7) 941–956.
- Fickes, M. 2001. Hot topic: Doom to boom? *Retail Traffic* (April) 13, http://retailtrafficmag.com/mag/retail_hot_topic_doom.
- Friend, S., P. Walker. 2001. Welcome to the new world of merchandising. *Harvard Bus. Rev.* **79**(November) 133–141.
- Gallego, G., G. J. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.

- Harris, M., A. Raviv. 1981. A theory of monopoly pricing schemes with demand uncertainty. *Amer. Econom. Rev.* **71**(3) 347–365.
- Heching, A., G. Gallego, G. van Ryzin. 2002. A theoretical and empirical investigation of markdown pricing in fashion retailing. *J. Revenue Pricing Management* **1**(2) 139–160.
- Homem-de-Mello, T., W. L. Cooper, A. J. Kleywegt. 2004. Models of the spiral-down effect in revenue management. Working Paper 04-002, Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.
- Lazear, E. P. 1986. Retail pricing and clearance sales. *Amer. Econom. Rev.* **76**(1) 14–32.
- Leslie, P. 2004. Price discrimination in Broadway theater. *RAND J. Econom.* **35**(3) 520–541.
- Levy, M., J. Woo. 1999. Yield management in retail: The application of advanced mathematics to the retail pricing dilemma. *J. Professional Pricing* **8**(3, June) 23–26.
- Littlewood, K. 1972. Forecasting and control of passengers. *12th AGIFORS Sympos. Proc., Nathanya, Israel*, Airline Group of the International Federation of Operational Research Societies, New York, 95–128.
- Phillips, R. 2005. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA.
- Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer, New York.
- Warner, E. J., R. B. Barsky. 1995. The timing and magnitude of retail store markdowns: Evidence from weekends and holidays. *Quar. J. Econom.* **110**(2) 321–352.