

Inventory, Discounts, and the Timing Effect

Hyun-Soo Ahn

Department of Operations and Management Science, Ross School of Business, University of Michigan,
Ann Arbor, Michigan 48109, hsahn@bus.umich.edu

Mehmet Gümüş

Desautels Faculty of Management, McGill University, Montreal, Quebec H3A 1G5, Canada,
mehmet.gumus@mcgill.ca

Philip Kaminsky

Department of Industrial Engineering and Operations Research, University of California, Berkeley,
Berkeley, California 94720, kaminsky@ieor.berkeley.edu

We introduce and analyze a model that explicitly considers the timing effect of intertemporal pricing—the concept, found in practice, that demand during a sale is increasing in the time since the last sale. We present structural results that characterize the interaction between the decision to hold a sale and the inventory-ordering decision. We show that the optimal inventory-ordering policy is a state-dependent base-stock policy; however, the optimal pricing policy can be quite complicated due to both the value and the cost of holding inventory and delaying sales. In our computational analysis, we find that compared to a fixed-price policy, we see an average gain in profit of almost 5% from optimally varying promotion and inventory decisions accounting for intertemporal demand, and we find that this potential profit gain increases as demand variability decreases. We also develop a heuristic based on deterministic pricing and find that it performs well relative to the optimal policy.

Key words: inventory; policies; pricing; uncertainty; stochastic

History: Received: September 17, 2007; accepted: August 29, 2008. Published online in *Articles in Advance* October 24, 2008.

1. Introduction

The success of dynamic pricing and revenue management in the air travel, hotel, and car rental industries has naturally led to the desire to extend these concepts to other industries, including those with nonperishable products and inventory replenishment. In the past, the applications of these tools have been primarily limited to markdowns and promotion pricing to eliminate excess inventory of seasonal products, or products with short life cycle. However, as data-processing technology and e-commerce have spread, dynamic pricing has become more accessible as a tool to help retailers better match supply with demand and increase operating profit. Many of the assumptions, models, and results of traditional revenue management settings have to be modified for these sometimes very different environments. In particular, in an environment in which inventory can be replenished and customers have multiple opportunities to buy the same product, effective ordering

and pricing strategies must account for intertemporal demand interaction.

In this context, intertemporal demand interaction refers to the sensitivity of current demand not only to current pricing, but also to past pricing decisions. A quick review of the advertising circulars in the Sunday paper suggests that retailers use price reductions for more than just eliminating excess inventory of out-of-season products—these retailers are instead attempting to benefit from intertemporal demand interactions in order to increase profits. This notion of intertemporal demand interaction is elegantly modeled and empirically validated by Pesendorfer (2002). Motivated by data he collects from supermarket chains, Pesendorfer models high-valuation customers who make purchases at the regular retail price when they enter the market, and low-valuation customers who may remain in the market to see if a sale is offered. Focusing initially on ketchup sales, he finds that the demand level is significantly higher

if previous prices were relatively high than if they were low, and that the level of demand during a sale increases in the time since the last sale in the store. Pesendorfer presents a deterministic model where a fixed number of low- and high-valuation customers enter the market each period, and shows that periodically offering sales is the optimal retailer's strategy. He then validates his findings with further empirical analysis. We call Pesendorfer's observation that demand increases in the depth of discount the *level effect*, and in Ahn et al. (2007) we build a deterministic model to explore the impact of the level effect on production and pricing policy in a capacitated system. We call Pesendorfer's observation that demand under the sale price increases in the time since the last sale the *timing effect*, and in this paper, we propose a stochastic demand model consistent with Pesendorfer's findings in order to increase our understanding of the impact of the timing effect on optimal inventory replenishment policy.

In this paper, we present a stylized model that assesses the impact of coordinating inventory ordering policy and the timing of sales in the presence of intertemporal demand effects. In particular, we explicitly capture the impact of the time since the last time a sale price was offered on the demand at a given sale price, and show that the presence of inventory has a nontrivial effect on dynamic pricing decisions.

Marketing researchers as well as economists have studied temporal price dispersion; however, most of this literature has focused on the effects of intertemporal demand interactions on pricing decisions while ignoring inventory considerations. Conlisk et al. (1984) and Sobel (1984, 1991) consider durable goods markets where low-valuation customers accumulate and wait for a sale. They show that firms engage in cyclic pricing behavior, employing a traditional "skimming" strategy where high-valuation customers are "skimmed off the top" with higher prices, and then low-valuation customers are charged lower prices later in each cycle. Others consider stockpiling behavior (Assuncao and Meyer 1993), change in customer's goodwill (Slade 1998, 1999), and information asymmetry and competition (Varian 1980) as key drivers of intertemporal price dispersion. In contrast to these papers, our paper explicitly considers the interplay of inventory and pricing decisions.

The revenue management literature accounts for the impact of inventory position on pricing decisions, but typically does not allow for inventory replenishment (e.g., Gallego and van Ryzin 1994 and Bitran and Mondschein 1997). Although there are a few recent papers that deal with more sophisticated consumer behavior (e.g., Aviv and Pazgal 2008, Elmaghraby et al. 2008, Su 2007, and Zhou et al. 2007), almost all of them—with the exception of Cachon and Swinney (2009), who allow initial replenishment before the selling season and a one-time quick-response replenishment during the selling season—assume no replenishment during the selling season, and thus do not consider inventory/production factors. In contrast, our paper explicitly considers the interplay of inventory replenishment and pricing decisions.

Another stream of operations management literature, including papers by Federgruen and Heching (1999), Chen and Simchi-Levi (2004a, b), Polatoglu and Sahin (2000), Chen et al. (2006), and Song et al. (2009), considers the coordination of pricing and inventory control (replenishment) with independent demand. Additionally, our paper is related to work on joint production-pricing problems; see review papers by Eliashberg and Steinberg (1993), Elmaghraby and Keskinocak (2003), and Chan et al. (2004). In almost all of these models, however, the demand in each period is assumed to be independent of price history.

In contrast, we explicitly model the relationship between current demand and price history. In this respect, our model is most closely related to two papers: Cheng and Sethi (1999) and Ahn et al. (2007). Cheng and Sethi consider a stochastic inventory model with Markov-modulated demand, where demand in each period is influenced by a state variable, which may in turn be influenced by the promotion decision in previous periods. However, in our model, the demand distribution is a function of the state variable and the price in the current period and in previous periods. This distinction allows us to capture the timing effect discussed above. Ahn et al. (2007) consider a deterministic capacitated pricing and production model in the presence of intertemporal demand interaction, and focus on determining the sequence of optimal prices in various scenarios. In this paper, we focus on the timing of sales and the

inventory policy in an environment where demand is stochastic, and is impacted by whether or not a sale is offered in the current period, and by the time since the last sale was offered.

2. Model Formulation

We consider a discrete-time, T -period finite-horizon, single-item stochastic demand model where at each period the retailer offers one of two possible prices, a (regular) *retail* price or a (discounted) *sale* price, both of which are exogenously determined. Let p^r and p^s ($p^r > p^s$) be the retail price and the sale price, respectively. At the start of each period ($t = 1, \dots, T$), the retailer makes inventory ordering (i.e., how much to order) and pricing (retail or sale) decisions after observing the starting inventory level and the number of periods since the last sale. Once decisions are made, demand is realized and holding or penalty cost is incurred.

Specifically, let x_t be the starting inventory level at the beginning of period t prior to inventory replenishment, k_t be the number of periods since the last sale before period t (where $k_t = 1$ if there was a sale in period $t - 1$), y_t be the inventory level in period t after inventory replenishment (we assume zero lead time) but before demand realization, and $p_t \in \{p^r, p^s\}$ be the price charged in period t . Our objective is to determine the optimal pricing and inventory policy that maximizes the expected (discounted) profit for a finite-horizon problem. We assume no fixed ordering cost and no maximum ordering capacity. For most of the paper, we assume lost sales, but we extend the results to the backorder case in the final section of the paper.

The explicit representation of intertemporal demand in a stochastic setting is fundamental to our model. As in Pesendorfer (2002), we conceptually divide potential customers into two groups: high-valuation consumers who will purchase independent of whether or not the product is on sale (that is, at the retail price or at the sale price), and low-valuation consumers who will purchase only at the sale price. In each period, customers from both groups enter the system (that is, their need for the good arises). When the retailer charges the retail price, only high-valuation consumers will attempt to buy, but some of the low-valuation consumers from the current and

previous periods will remain in the system, waiting to see if the sale price is charged in subsequent periods. On the other hand, when the retailer charges the sale price in period t , all of the customers who entered the market in period t (both high- and low-valued), as well as remaining low-valuation customers from previous periods, will buy the product (at the sale price). This conceptual demand formation process can be viewed as a stochastic extension of the model considered in Ahn et al. (2007), with price restricted to two possible levels.

Recall that we are interested in a *stochastic demand* setting that is consistent with the empirical findings of Pesendorfer (2002). Thus, our demand model needs to reflect the fact that demand at the sale price is likely to be higher than demand at the retail price, and the fact that demand at the sale price is likely to be higher when more periods have passed since the last sale than when fewer periods have passed. To do this, we assume that the demand distribution in each period t , $t = 1, \dots, T$, is determined by two factors: the price charged in period t , p_t , and the number of periods since the last sale, k_t .

To formally capture this, let $\xi_t(p_t, k_t)$ be a nonnegative and continuous random variable that represents the demand in period t when the retail price is p_t and the number of periods since the last sales is k_t , and let $\Phi_t(\xi | p_t, k_t)$ and $\phi_t(\xi | p_t, k_t)$ be its cumulative distribution function (c.d.f.) and probability density function (p.d.f.), respectively. Given p_t and k_t , $\xi_t(p_t, k_t)$ is defined as follows:

$$\xi_t(p_t, k_t) \sim \begin{cases} \xi^r, & \text{if } p_t = p^r; \\ \xi_k^s, & \text{if } p_t = p^s \text{ and } k_t = k. \end{cases} \quad (1)$$

Observe that when the retail price is charged, demand at the retail price, ξ^r , is independent of the past price trajectory and identically distributed with c.d.f. $\Phi^r(\xi)$, p.d.f. $\phi^r(\xi)$, and mean $\mu^r = E[\xi^r]$. On the other hand, when the sale price is charged after k , its resultant demand (ξ_k^s) is dependent on the time since the last sale and follows a distribution $\Phi_k^s(\xi)$ with density $\phi_k^s(\xi)$ and mean $\mu_k^s = E[\xi_k^s]$. This demand model enables us to capture Pesendorfer's notion of high-valuation customers who arrive at the system and buy immediately, regardless of the price and low-valuation customers who may accumulate until the

sale. To do this, we introduce the notion of stochastic order as follows: We say random variable X is stochastically smaller than random variable Y if $P(X \leq u) \geq P(Y \leq u)$ for all real u , and denote this by $X \leq_{ST} Y$. Then, the following assumption captures the aforementioned consumer behavior:

ASSUMPTION 1. $\xi^r \leq_{ST} \xi_1^s$ and ξ_k^s is stochastically increasing in k .

This assumption implies that in a stochastic sense, demand at the sale price is larger than demand at the retail price, and that demand at the sale price increases in time since the last sale.

It is also reasonable to assume that the marginal increase in demand at sale price is decreasing in time since the last sale in most practical settings, because one would not expect all low-valuation customers to wait indefinitely for a sale (and indeed, Pesendorfer made a similar observation in his empirical study). However, most of our results do not require this assumption—we formally introduce it when needed. Also, we note that one could propose an alternative model of intertemporal interaction in which the demand at the retail price is also affected by the frequency of sales. We do not do so, in order to keep the model tractable for the development of results and insights.

For our stochastic dynamic program, we use the inventory level at the start of period t and number of periods since the last sale, i.e., (x_t, k_t) , $t = 1, \dots, T$, as our state variables. In each period, the decision maker simultaneously sets price p_t to either the retail price p^r or the sale price p^s and raises the inventory level to y_t by ordering and receiving $y_t - x_t$ units of inventory before demand realization. Then, the inventory level and the number of periods since the last sale at the start of period $t + 1$ are as follows:

$$x_{t+1} = \max[y_t - \xi_t(p_t, k_t), 0] = [y_t - \xi_t(p_t, k_t)]^+,$$

$$k_{t+1} = \begin{cases} k_t + 1, & \text{if } p_t = p^r; \\ 1, & \text{if } p_t = p^s. \end{cases}$$

and

$$k_{t+1} = \begin{cases} k_t + 1, & \text{if } p_t = p^r; \\ 1, & \text{if } p_t = p^s. \end{cases}$$

We define $[u]^+ = \max[u, 0]$ and $[u]^- = \max[-u, 0]$ for any real number u , and observe that for a given

state (x_t, k_t) and retailer action (y_t, p_t) , the retailer realizes the following revenue and costs:

- $p_t \min[y_t, \xi_t(p_t, k_t)]$ is the revenue in period t .
- $c(y_t - x_t)$ is the cost of raising inventory from x_t to y_t in period t , where c is the unit ordering cost.
- $h(y_t - \xi_t(p_t, k_t)) = h^+[y_t - \xi_t(p_t, k_t)]^+ + h^-[y_t - \xi_t(p_t, k_t)]^-$ is the inventory and penalty cost at the end of period t , where h^+ is the per-unit holding cost and h^- is the per-unit penalty cost incurred upon stockout. Because we are already capturing the loss of revenue in our revenue term, h^- represents the additional loss of goodwill cost or any other lost-sales related cost not directly related to the current revenue. (When we consider the backorder case, the same cost function applies except that h^- represents the per-unit backorder cost.)

Let $V_t(x_t, k_t)$, $t = 1, \dots, T$, be the expected discounted profit-to-go function under the optimal policy starting from any admissible state (x_t, k_t) , and let $J_t(y_t, p_t; x_t, k_t)$ be the expected profit-to-go function if the retailer offers price p_t and raises the inventory level to y_t from state (x_t, k_t) in period t , and follows the optimal policy in subsequent periods. Then, the retailer's problem can be expressed as a stochastic dynamic program satisfying the following recursive relation:

$$V_t(x_t, k_t) = \max_{p_t \in \{p^r, p^s\}} \left\{ \max_{y_t \geq x_t} J_t(y_t, p_t; x_t, k_t) \right\} \quad \text{and}$$

$$J_t(y_t, p_t; x_t, k_t) = -c(y_t - x_t) + \int_0^\infty [p_t \min[y_t, \xi] - h^+[y_t - \xi]^+ - h^-[y_t - \xi]^- + \alpha V_{t+1}([y_t - \xi]^+, k_{t+1})] \phi_t(\xi | p_t, k_t) d\xi$$

where α , $0 < \alpha < 1$, is the discount factor. Let $V_{T+1}(x_{T+1}, k_{T+1}) = cx_{T+1}$ be the terminal value function that represents the salvage value of the inventory at the end of the planning horizon. This is a standard assumption in many inventory models (c.f., Porteus 2002), and will facilitate our analysis.

Following a standard approach in the stochastic inventory literature, we find it useful to work with transformed versions of our value functions. We let $\hat{J}_t(y, p; k) = J_t(y, p; x, k) - cx$ and $\hat{V}_t(x, k) = V_t(x, k) - cx$. Then,

$$\hat{V}_t(x, k) = \max_{p \in \{p^r, p^s\}} \left\{ \max_{y \geq x} \hat{J}_t(y, p; k) \right\} \quad (2)$$

where $\hat{J}_t(y, p; k)$ satisfies

$$\begin{aligned} \hat{J}_t(y, p; k) = & \int_0^\infty [p \min\{y, \xi\} - cy + \alpha c[y - \xi]^+ \\ & - h^+[y - \xi]^+ - h^-[y - \xi]^- \\ & + \alpha \hat{V}_{t+1}([y - \xi]^+, k_{t+1})] \phi_t(\xi | p, k) d\xi. \end{aligned}$$

In this transformed formulation, $\hat{J}_t(y, p; k)$ does not depend on x and the terminal condition becomes $\hat{V}_{T+1}(\cdot) = 0$. Note that with two decision variables (price and inventory level) and two state variables, the optimal policy can be quite complex and may depend on the state variables and time-to-go. In the next section, we analyze the form of these optimal decisions and characterize the structural relationships between the optimal inventory ordering/pricing decisions and the state variables. In §4, we explore a deterministic dynamic program that corresponds to the make-to-order case of our model. In §5, we use the insight from this exploration to develop an effective heuristic for our model and present the results of our computational study, and in §6, we present several extensions of our model and conclude.

3. Optimal Policy

The typical strategy employed to characterize the structure of the optimal policy in a periodic-review inventory model includes a proof that the value function is concave or quasi-concave (convex or quasi-convex under minimization) in inventory order-up-to level, a proof that this concavity (or quasi-concavity) is preserved under the dynamic programming operator, and an induction-based argument to show the desired structural result (see, e.g., Porteus 2002). If the value function in our model were concave, one could use the fact that concavity is preserved under expectation operator without any distributional assumption. Unfortunately, with joint maximization over price and inventory decisions, even the single-period profit function for our model is not concave (although it is quasi-concave) in order-up-to level, and proving that the expected profit-to-go function is quasi-concave (unimodal) for multiple periods is not trivial, because one needs to show that the property is preserved over induction and joint maximization over price and inventory decisions.

Thus, we are motivated to use a different approach—we restrict demand at retail and sale prices to distributions with strongly unimodal densities, the class of strongly unimodal distributions. Ibragimov (1956) introduced the class of *strongly unimodal* distributions and defined it as follows:

DEFINITION 1. A distribution of a random variable is said to be *strongly unimodal* if its convolution with any unimodal function is unimodal.

In fact, characterization of strongly unimodal distribution is convenient for a continuous random variable because the distribution is strongly unimodal if it has a log-concave density (Dharmadhikari and Joag-Dev 1988 and Fox et al. 2006). We use the fact that the unimodality is closed under integral convolution for strongly unimodal distributions to characterize the structure of the optimal inventory policy in our model by making the following technical assumption.

ASSUMPTION 2. *The distribution of demand at the retail price and the distributions of demand at the sale price (i.e., the distributions of ξ^r and ξ_k^s , $k = 1, 2, \dots$) are strongly unimodal.*

Although this assumption may sound technical and restrictive, a number of important distributions are indeed strongly unimodal, including, among others, normal, truncated normal, uniform, exponential, gamma with shape parameter $\alpha \geq 1$, and beta with $p \geq 1$ and $q \geq 1$ (Dharmadhikari and Joag-Dev 1988). For additional reference, results, and applications of strongly unimodal distribution, see Fox et al. (2006) and Dharmadhikari and Joag-Dev (1988).

With this assumption, we now show that the optimal inventory level for a given price is determined by a state-dependent base-stock policy:

THEOREM 1. *Suppose that Assumptions 1 and 2 hold. In each time period t , for each p and k , there exists an optimal base-stock level $s_t(p, k)$ such that if the starting inventory level $x_t < s_t(p, k)$, it is optimal to raise the inventory level to $s_t(p, k)$, and otherwise it is optimal to do nothing.*

PROOF. The proof is in the online appendix. \square

Theorem 1 characterizes the optimal order quantity for a given price when the starting inventory level is x and k periods have passed since the last sale. Note

that $s_t(p, k)$ is the optimal base-stock level in period t when the offered price is p and the number of periods since the last sale is k , and is given by

$$s_t(p, k) = \arg \max_{y \geq 0} \hat{J}_t(y, p; k).$$

In any period and state, the difference between the regular retail price base-stock level and the sale price base-stock level is driven by two factors: the difference between the retail and sale prices, and the number of periods since the last sale. The sale demand increases in the number of periods since the last sale, k , and all things being equal, this increases the sale price base-stock level as k increases. On the other hand, as the sale price decreases relative to the regular retail price, the sale price base-stock level tends to go down. Depending on which force is more significant, the sale price base-stock level can in general be larger or smaller than the regular retail price base-stock level.

However, to facilitate subsequent analysis, we make the following technical assumption that demand at the sale price (even when $k = 1$) is sufficiently high that the single-period order-up-to level at the sale price is higher than that at the retail price:

ASSUMPTION 3. *The order-up-to level of a single-period newsvendor problem with the retail and sale price demand distributions of our model is decreasing in price. That is,*

$$\Phi_1^{s(-1)} \left(\frac{p^s + h^- - c}{p^s + h^+ + h^- - \alpha c} \right) \geq \Phi_1^{r(-1)} \left(\frac{p^r + h^- - c}{p^r + h^+ + h^- - \alpha c} \right).$$

Indeed, for many realistic scenarios, it is reasonable to think that demand at the sale price will be considerably higher than demand at the retail price, which suggests that this assumption likely holds. In fact, Pesendorfer’s (2002) ketchup demand was on average seven times higher during sales.

Assumption 3, along with Assumptions 1 and 2, enables us to further characterize the structure of the optimal policy.

THEOREM 2. *Under Assumptions 1–3:*

(i) *When the retail price is offered, the base-stock level is independent of the number of periods since the last sale and the number of remaining periods in the planning horizon. That is, there exists a constant $s^r > 0$ such that $s_t(p^r, k) = s^r$ for all t and k .*

(ii) *The base-stock level under the sale price is always larger than the base-stock level under the retail price: $s_t(p^s, k) \geq s^r$ for all k and t .*

(iii) *For a given $k \geq 1$, $\hat{V}_t(x, k)$ is constant in $x \in [0, s^r]$ for all t .*

PROOF. The proof is in the online appendix. \square

Utilizing Theorems 1 and 2, we can begin to characterize the optimal pricing and inventory policy. For this purpose, in addition to the base-stock level under the retail price, s^r , and sale base-stock levels, $s_t(p^s, k)$, we introduce a third critical level $\tilde{s}_t(k)$ that represents the starting inventory level below which it is optimal to charge the retail price.

By Theorem 2.(iii), $\tilde{s}_t(k)$ must belong to one of the following three cases: (i) $\tilde{s}_t(k) = 0$, (ii) $\tilde{s}_t(k) \in (s^r, s_t(p^s, k))$, or (iii) $\tilde{s}_t(k) \geq s_t(p^s, k)$. The structure of the optimal joint pricing and inventory policy depends on $\tilde{s}_t(k)$.

THEOREM 3. *Suppose Assumptions 1–3 hold. If there have been k periods since the last sale, the optimal price and inventory policy in period t takes one of the following three forms depending on the initial inventory level and the values of $\tilde{s}_t(k)$, s^r , and $s_t(p^s, k)$:*

Case (i) $\tilde{s}_t(k) = 0$: *If $x \leq s_t(p^s, k)$, it is optimal to order up to $s_t(p^s, k)$ and offer the sale price p^s . Otherwise, it is optimal to not order and to charge a state-dependent price, $p_t^*(x, k) = \arg \max_{p \in \{p^r, p^s\}} \hat{J}_t(x, p; k)$ (see Figure 1).*

Case (ii) $s^r < \tilde{s}_t(k) < s_t(p^s, k)$: *If $x \in [0, s^r)$, it is optimal to order up to s^r and to sell at the retail price p^r . If $x \in [s^r, \tilde{s}_t(k))$, it is optimal not to order, and to sell at the retail price p^r . If $x \in [\tilde{s}_t(k), s_t(p^s, k)]$, it is optimal to order up to $s_t(p^s, k)$ and to sell at the sale price p^s . If $x > s_t(p^s, k)$, it is optimal to not order, and to follow a state-dependent pricing policy, $p_t^*(x, k) = \arg \max_{p \in \{p^r, p^s\}} \hat{J}_t(x, p; k)$ (see Figure 2).*

Case (iii) $\tilde{s}_t(k) \geq s_t(p^s, k)$: *If $x \leq s^r$, it is optimal to order up to s^r and sell at the retail price p^r . If $x \in (s^r, \tilde{s}_t(k))$,*

Figure 1 Optimal Inventory and Pricing Policy for Case (i) $\tilde{s}_t(k) = 0$





Figure 2 Optimal Inventory and Pricing Policy for Case (ii)
 $\tilde{s}_t(k) \in (s^r, s_t(p^s, k))$




it is optimal to not order and to sell at the retail price p^r . If $x_t \geq \tilde{s}_t(k)$, it is optimal to not order and to follow a state-dependent price $p_t^*(x, k) = \arg \max_{p \in \{p^r, p^s\}} \hat{J}_t(x, p; k)$ (see Figure 3).

PROOF. The proof is in the online appendix. \square

In general, when the starting inventory level is higher than the maximum of the appropriate base-stock levels (s^r and $s_t(p^s, k)$) and the threshold $\tilde{s}_t(k)$, the optimal pricing policy is state dependent. Therefore, the optimal price can switch between sale and retail in this region. In this region the optimal pricing policy can be quite complicated with respect to both state variables.

As an example, consider the relationship between the optimal pricing decision and the number of periods since the last sale, that is, the sensitivity of $p_t^*(x, k)$ to k for a fixed x . One might expect that if it is optimal to offer the sale price for a given starting inventory when the last sale was k periods ago, it would still be optimal to offer a sale if the last sale was $k + 1$ periods ago because the demand at the sale price when the last sale was $k + 1$ periods ago is stochastically larger than the demand at the sale price when the last sale was k periods ago. However, as the next example shows, the optimal price is not necessarily monotone in k .

Figure 3 Optimal Inventory and Pricing Policy for Case (iii)
 $\tilde{s}_t(k) \in [s_t(p^s, k), \infty)$



We consider a two-period example satisfying Assumptions 1–3, where retail and sale demands are drawn from uniform distributions:


EXAMPLE 1.

$$\xi_t(p, k) = \begin{cases} U(\mu_r, v_r), & \text{if } p = p^r; \\ U(\mu_k, v_k), & \text{if } p = p^s, \end{cases}$$

where $U(\mu, v)$ is a uniform random variable that has a support from $\mu - v$ to $\mu + v$. Let $\mu_r = 6$, $v_r = 3$, $\mu_k = 9 \sum_{i=1}^k \beta^{i-1}$, and $v_k = v_r \sqrt{\sum_{i=1}^k \beta^{2(i-1)}}$ with $\beta = 0.90$. Finally, $p^r = 30$, $p^s = 15$, $h^+ = 5$, $h^- = 0$, $c = 10$, $\alpha = 1$, and the planning horizon is $T = 2$ periods.

In Figure 4 we present the optimal first-period pricing policy for Example 1 as a function of initial inventory level for $k = 2, 3$, and 4. Observe that in the highlighted range, there are a range of initial inventory levels for which it is optimal to offer a sale when $k_1 = 2$, but not when $k_1 = 3$ or 4. For a given initial inventory level, the optimal price does not necessarily decrease in k . This counterintuitive behavior is a result of two opposing effects. On one hand, as k increases from 2, the demand in the second period if the sales price is not offered in the first period will increase, making delaying the sale more appealing. On the other hand, delaying the sale (for a given inventory level) will increase the holding cost in the first period if the starting inventory in the first period is sufficiently large. For this example, when the starting inventory is around 17, the increase in demand if k is allowed to grow to 4 in the second period from 3 in the first period (or from 4 to 5) dominates the extra holding cost, and therefore the retail price is optimal when $k = 3$ or 4. On the other hand,

Figure 4 Optimal Policy Is Nonmonotonic in k in the Highlighted Region



the increase in demand if k is allowed to grow to 3 in the second period from 2 in the first period does not fully compensate for the extra holding cost, so offering a sale is optimal.

Similarly, one might expect that $p_t^*(x, k)$ would decrease in starting inventory x because higher starting inventory levels lead to higher inventory holding costs, suggesting that inventory should be liquidated through a sale. However, we can create similar examples in which increasing inventory increases the value of waiting an additional period for higher demand under the sale price, and hence the price is nonmonotonic in starting inventory level. For instance, for certain examples we observe that the sale price is offered at moderate inventory levels while the retail price is offered at higher inventory levels. In these cases, moderate initial inventory implies that most of this inventory can be liquidated by offering a sale, so offering a sale becomes optimal. On the other hand, relatively high initial inventory implies that offering a sale is likely to liquidate only a portion of the existing inventory, so it may be optimal in this case to offer the retail price and delay the sale in order to increase sale demand level in subsequent periods.

We note that Assumption 3 plays a crucial role in the characterization of the policy in Theorems 2 and 3. If Assumption 3 does not hold, parts (ii) and (iii) of Theorem 2 do not hold, and, in turn, in Theorem 3 the same approach cannot be used to characterize the structure of the optimal policy when the starting inventory is between 0 and s^r .

We conjecture that a version of Theorem 3 could be extended to the infinite-horizon problem under some technical conditions. However, because we are only able to partially characterize the structure of value function as a function of initial inventory and the number of periods since the last sale, classical results (such as Iglehart 1963) do not directly apply. Although one might be able to use an approach similar to that of Chen and Simchi-Levi (2004b) and Feinberg and Lewis (2007), this would not be trivial and merits further research.

4. Make-to-Order System

As we demonstrated in the previous section, the optimal policy of our model is sometimes quite complicated, because pricing and inventory decisions are

influenced by intertemporal demand effects, the starting inventory level, and the time remaining in the planning horizon. In general, the optimal pricing policy may not be monotone either in starting inventory level or in the time since the last sale. To isolate the intertemporal demand effect on the pricing decision free of inventory considerations, we next consider a simplified version of the model, in which we assume that the retailer places an order after observing the demand realization—a “make-to-order” system. An immediate consequence of this restriction in our model is that we never pay inventory-related costs, and we satisfy all realizations of demand. Thus, the decision problem in each period can be recast as a *deterministic* dynamic program by replacing the random variable representing demand with its mean.

Observe that the only state variable in this make-to-order model is the time since the last discount was offered, k . The dynamic program is as follows:

$$V_t(k_t) = \max_{p_t \in \{p^r, p^s\}} [\pi(p_t, k_t) + \alpha V_{t+1}(k_{t+1})]$$

where $\pi(p^r, k_t) = E[(p^r - c)\xi^r] = (p^r - c)\mu^r$ and $\pi(p^s, k) = E[(p^s - c)\xi_k^s] = (p^s - c)\mu_k^s$ are the expected one-profit profits at the retail price and the sale price, respectively; $\alpha \leq 1$ is the discount factor; and $V_{T+1}(k) = 0$ for all k . Noting that the demand at the retail price is independent of the number of periods since last sale, we omit k in $\pi(p^r, k)$ and write it as $\pi(p^r)$.

We assume that the expected demand under the sale price, $E[\xi(p^s, k)] = \mu_k^s$, is strictly concave and increasing in k , that is, $\mu_{k+1}^s - \mu_k^s > \mu_{k+2}^s - \mu_{k+1}^s \geq 0$, so that the marginal increase of the expected demand at the sale price decreases as k increases. Under this assumption, we show that a threshold policy is indeed optimal:

THEOREM 4. *Let $p_t^*(k)$ be the optimal price at period t when k periods have passed since the last sale. Then,*

(i) $p_t^*(k)$ is nonincreasing in k . In other words, there exists a k_t^* such that it is optimal to charge the retail price, p^r , if $k \leq k_t^*$, and to offer the sale price p^s otherwise.

(ii) $V_t(k+1) - V_t(k) \leq \pi(p^s, k+1) - \pi(p^s, k)$ for all $k = 1, \dots, t$ and $t = 1, \dots, T$.

PROOF. We prove this result by induction in the online appendix. \square

For finite-horizon problems, finding the optimal pricing policy is quite straightforward because one simply needs to solve the corresponding shortest-path problem. Indeed, this motivates us to explore the effectiveness of this type of threshold policy as a heuristic for our original make-to-stock model. We investigate this in the computational section of the paper.

The threshold levels suggested by Theorem 4 are in general time dependent, so that there exists k_t^* for each period t such that if the number of periods since the last sale is less than k_t^* , it is optimal to offer the retail price and otherwise it is optimal to offer the sale price. To better understand the nature of these threshold levels, we next consider the infinite-horizon version of this problem in both the discounted and average profit cases.

For both cases, in §E of the online appendix, we prove that a time-invariant optimal threshold level k^* exists for each period t . In other words,

LEMMA 1. *There exists an optimal stationary policy for both discounted and average profit cases.*

Thus, there exists an optimal policy where the sale price is offered every k^* periods (i.e., a cyclic discount policy). In order to characterize the optimal cycle length, we write the expected profit associated with a policy in which the sale price is offered every k periods and find the optimal cycle length. Let $\Pi^\alpha(k)$ and $\Pi^A(k)$ be the discounted and average expected profit of a k -period cyclic policy with initial state $k_0 = 1$, respectively. After some algebraic manipulation, we get, for all $k \geq 1$,

$$\Pi^\alpha(k) = \frac{\pi(p^r)}{1-\alpha} + \alpha^{k-1} \frac{\pi(p^s, k) - \pi(p^r)}{1-\alpha^k} \quad (\text{Discounted profit}),$$

and

$$\Pi^A(k) = \pi(p^r) + \frac{\pi(p^s, k) - \pi(p^r)}{k} \quad (\text{Average profit}).$$

The derivations of both expressions above are found in §F of the online appendix. Observe that a periodic sale dominates selling at the retail price in every period if and only if there exists a k such that $\pi(p^s, k) - \pi(p^r) \geq 0$. In such a case, the optimal cycle

length is the one that maximizes the difference in profits:

$$k^{\alpha*} = \arg \max \frac{\alpha^{k-1}(\pi(p^s, k) - \pi(p^r))}{1 - \alpha^k} \quad (\text{Discounted profit}),$$

and

$$k^{A*} = \arg \max \frac{\pi(p^s, k) - \pi(p^r)}{k} \quad (\text{Average profit}).$$

Furthermore, because $\pi(p^s, k)$ is increasing in k , it suffices to consider $k \geq \underline{k} = \min\{k \mid \pi(p^s, k) - \pi(p^r) \geq 0\}$. The following two lemmas are useful for characterizing the optimal cycle length.

LEMMA 2. *In the infinite-horizon discounted profit problem, for $k \geq \underline{k}$,*

(i) *The k -period cyclic policy is better than the $k+1$ -period cyclic policy if and only if*

$$\frac{\pi(p^s, k+1) - \pi(p^s, k)}{\pi(p^s, k) - \pi(p^r)} \leq \frac{1/\alpha}{\sum_{i=1}^k \alpha^{i-1}} = \frac{1-\alpha}{\alpha} \frac{1}{1-\alpha^k}. \quad (3)$$

(ii) $\Pi^\alpha(k)$ is unimodal in k .

LEMMA 3. *In the infinite-horizon average profit problem, for $k \geq \underline{k}$,*

(i) *The k -period cyclic policy is better than a $k+1$ -period cyclic policy if and only if*

$$\frac{\pi(p^s, k+1) - \pi(p^s, k)}{\pi(p^s, k) - \pi(p^r)} \leq \frac{1}{k}. \quad (4)$$

(ii) $\Pi^A(k)$ is unimodal in k .

The proofs of these lemmas are in §G of the online appendix. Conditions (3) and (4) imply that if the marginal benefit of extending the cycle by one period becomes sufficiently small, then it is optimal to not extend the length of a cycle. In fact, both conditions describe precisely the minimum marginal gain required to optimally extend the cycle length by at least one period. Solving for the optimal cycle length is not hard because in both cases, the profit function is unimodal in k .

We employ Lemmas 2 and 3 to characterize the optimal pricing policy for the infinite-horizon problem.

THEOREM 5. *In the infinite-horizon discounted profit problem,*

(i) *It is optimal to sell at the retail price in every period if and only if $\pi(p^r) \geq \pi(p^s, k)$ for all $k \geq 1$.*

(ii) *Otherwise, it is optimal to use the k^* -period cyclic pricing policy where k^* is the smallest integer greater than \underline{k} satisfying condition (3).*

The results hold for the average profit case with condition (4) replacing condition (3).

PROOF. The proof is in §G of the online appendix. \square

We note that this result is similar in spirit to the results established in other papers examining intertemporal demand effects for both durable and nondurable goods, including the papers of Conlisk et al. (1984), Sobel (1991), Pesendorfer (2002), and Ahn et al. (2007). Although the settings and model details are different in each of these papers, they all conclude that some form of periodic sale is optimal for capturing the demand of low-valuation customers.

5. Computational Study

We use a computational study to develop managerial insights into the benefit of jointly making pricing and inventory decisions (either optimally or heuristically) in the presence of intertemporal demand interactions, as well as to explore the impact of various system characteristics on effective pricing and inventory-ordering policies. Before discussing our results, we detail the problem parameters that we use in our experiments.

5.1. Problem Parameters

In our computational study, we hold the retail price, sale price, unit cost, expected demand at the retail price, and discount factor constant with the following values: $p^r = 20$, $p^s = 16$, $c = 12$, $\mu^r = 10$, $\alpha = 0.9$. We consider the lost-sales version of our model (described in §2) with a 12-period planning horizon, $T = 12$, and vary the remaining parameters as follows in order to explore a variety of different settings satisfying Assumptions 1–3:

Cost parameters. We use stationary linear production cost and inventory holding cost functions $c_t(x) = cx$ and $h_t(x) = h^+x$ for all periods in the planning horizon. We do not consider shortage costs

explicitly because the effect of the lost sales is already captured through the lost revenue. We vary the holding cost as a fraction of unit production cost, and use h to denote this fraction; i.e., $h = h^+/c$. In this computational study, we consider the following parameter values for h : {0.05, 0.10, 0.15, 0.20}.

Demand parameters. We model demand with truncated normal random variables, where $N^+(\mu, \sigma)$ represents the positive part of normal random variable with mean μ and standard deviation σ . We define the demand in period t as follows:

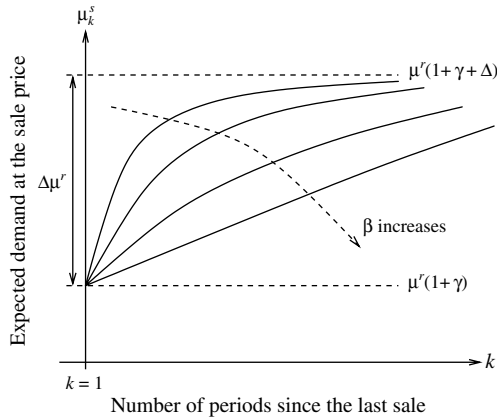
$$\xi_t(p, k) = \begin{cases} N^+(\mu^r, c_v \mu^r), & \text{if } p = p^r; \\ N^+(\mu_k^s, c_v \mu_k^s), & \text{if } p = p^s. \end{cases}$$

where c_v is the coefficient of variation (i.e., the proportion of standard deviation to the mean ($0 < c_v < 1$)), and μ^r and μ_k^s are related as follows:

$$\mu_k^s = \mu^r [1 + \gamma + \Delta(1 - \beta^{k-1})]$$

where $0 < \beta < 1$, $\gamma > 0$, and $\Delta > 0$. Observe that for a given coefficient of variation c_v , the mean, and therefore the standard deviation and distribution of the demand at the sale price, are completely characterized by the three parameters: β , γ , and Δ . The minimum increase of the mean demand at the sale price from the mean demand at the retail price is represented by γ , whereas the maximum increase of the mean demand at the sale price (i.e., the mean demand one could achieve by letting k approach infinity) is equal to $\gamma + \Delta$. Thus, Δ corresponds to the difference between the maximum and the minimum mean demands at the sale price. For example, if $\gamma = 1.5$ and $\Delta = 2$, then the mean demand at the sale price is 1.5 times higher than the mean demand at the retail price when $k = 1$, and although it increases as k increases, it will be no more than 3.5 times higher. Finally, β controls the marginal increase in the mean demand at the sale price as the time since the last sale increases, and captures the rate at which the demand at the sale price accumulates. For example, if $\beta = 0.50$, the marginal increase in the mean demand at the sale price decreases by half if the sale is delayed an additional period. The effect of these parameters on the mean demand at the sale price is illustrated in Figure 5.

Figure 5 Expected Demand at the Sale Price as a Function of k



Note. The figure illustrates the effect of changes in β , γ , and Δ .

In our experiments, we consider the following parameter values:

- c_v : {0.05, 0.35, 0.65, 0.95};
- γ : {0.5, 1.0, 1.5, 2.0};
- β : {0.2, 0.4, 0.6, 0.8};
- Δ : {0.5, 1.0, 1.5, 2.0}.

We compare the performance of the following three strategies for $4^5 = 1,024$ problem instances. For all instances, we fix the initial state (x_1, k_1) to $(0, 1)$: no starting inventory and no accumulation of low-valuation customers from previous periods.

To keep the computational experiment manageable, we vary several key parameters that are critical to understanding intertemporal demand interactions while keeping other parameter values fixed. Although one could extend the numerical study to explore changes in the optimal policy with respect to each parameter, we elect to focus on how intertemporal demand interaction (in particular, the demand distribution at the sale price) impacts inventory and pricing policy.

5.1.1. Optimal Policy. We solve the corresponding dynamic program, find the optimal policy, and calculate the expected profit under the optimal policy. Recall that the optimal policy consists of both inventory and pricing decisions, i.e., $(y_t^*(x, k), p_t^*(x, k))$. We define Π^* to be the optimal expected profit with the initial state, $(x_1, k_1) = (0, 1)$:

$$\Pi^* = V_1(0, 1).$$

5.1.2. Heuristic Policy. Recall that the optimal pricing policy for the make-to-stock model can be quite complicated. Thus, we are motivated to consider a simple threshold policy building on our make-to-order model results. In particular, we solve for the pricing decisions as if our model were make-to-order, and then use a myopic newsvendor solution for the ordering quantity.

HEURISTIC POLICY.

1. Let k_t^* be the threshold level found by solving the deterministic dynamic programming problem as if our instance were a make-to-order problem.

2. Define

$$s^r = (\Phi^r)^{-1}(f^r) \text{ and } s(p^s, k) = (\Phi_k^s)^{-1}(f^s),$$

where f^r and f^s are the critical-fractiles of single-period problem for the retail and sale prices, respectively:

$$f_r = \frac{p^r + h^- - c}{p^r + h^+ + h^- - \alpha c} \quad \text{and} \quad f_s = \frac{p^s + h^- - c}{p^s + h^+ + h^- - \alpha c}.$$

3. If $k < k_t^*$, then offer the retail price, p^r , and raise the inventory level to s^r .

4. If $k \geq k_t^*$, then offer the sale price, p^s , and raise the inventory level to $s(p^s, k)$.

Define $V_t^h(x, k)$ to be the expected profit-to-go under the heuristic from period t with a starting state (x, k) . We define Π^h to be the expected profit under the heuristic for 12 periods with initial starting state, $(x_1, k_1) = (0, 1)$:

$$\Pi^h = V_1^h(0, 1).$$

5.1.3. Constant Price Policy. Finally, we test a constant price policy, in which the price is fixed to be either the retail price, p^r , or the sale price, p^s , throughout the entire planning horizon, and inventory is raised to the appropriate myopic base-stock level described in the heuristic policy. We calculate the expected profit under each price and pick the one with the highest expected profit. Let Π^c be the expected profit under the constant price policy with initial starting state $(x_1, k_1) = (0, 1)$.

5.2. Value of Effective Policies

In this section, we discuss our analysis of the effectiveness of joint pricing and inventory management in the presence of intertemporal demand interactions. We start by defining two measures. First, we

define D^h to be a measure of the relative benefit of employing the optimal policy rather than the heuristic policy described above:

$$D^h = 100 \times \frac{\Pi^* - \Pi^h}{\Pi^*}.$$

Recall that the heuristic policy corresponds to sequentially making first the pricing decision (because in the heuristic, the pricing decision is made based only on the number of periods since the last sale decision) and then the inventory-ordering decision. Therefore, D^h represents relative increase in profit by simultaneously considering pricing and ordering decisions.

Similarly, we define D^c to be a measure of the relative benefit of employing the optimal policy rather than a constant price policy:

$$D^c = 100 \times \frac{\Pi^* - \Pi^c}{\Pi^*}.$$

In other words, D^c measures the relative increase in profit achieved by varying prices and inventory levels to account for intertemporal demand interactions as opposed to using a constant price and base-stock level.

We calculate the value of both measures described above (recall that these are expected values that we calculate numerically) for each set of parameter settings. In Table 1, we present the summary statistics of this analysis.

In this table, we see that, on average, the heuristic performs quite well (less than 1% from optimal), whereas the constant price policy is 5% worse than the optimal. In addition, the percentage optimality gap of the constant price policy can be as large as 27% for some instances. To explore the impact of problem parameters on performance, in Figure 6, we graph values of D^h and D^c (the left bar

is D^c and the right bar is D^h) for several values of a particular problem parameter, averaged over values of the other parameters. We make the following observations:

- In general, the benefit of explicitly varying pricing to account for intertemporal demand interactions (i.e., D^c) increases as the variability of demand decreases (i.e., c_v decreases), the rate at which demand at the sale price accumulates increases (i.e., β decreases), and the difference between the maximum and the minimum mean demands at the sale price increases (i.e., Δ increases). In other words, as the relative amount of demand that is impacted by intertemporal interactions increases, the value of changing prices to explicitly account for this interaction also increases. This is quite intuitive.

- D^c , the benefit of explicitly varying pricing to account for intertemporal demand interactions, is the largest when γ is around 1. To understand this, recall that γ represents the minimum increase in the mean demand at the sale price (when $k = 1$). When γ is sufficiently large, offering a sale price itself substantially increases the mean demand independent of the number of periods since the last sale. Therefore, a pricing strategy that offers a sale every period performs very well. When γ is sufficiently small, on the other hand, the time to accumulate enough demand to justify the sale price could be very long, so that the profit of the optimal policy may not be very different from that of a policy that offers the retail price every period. When γ is in between these extremes, periodic sales are intuitively most valuable.

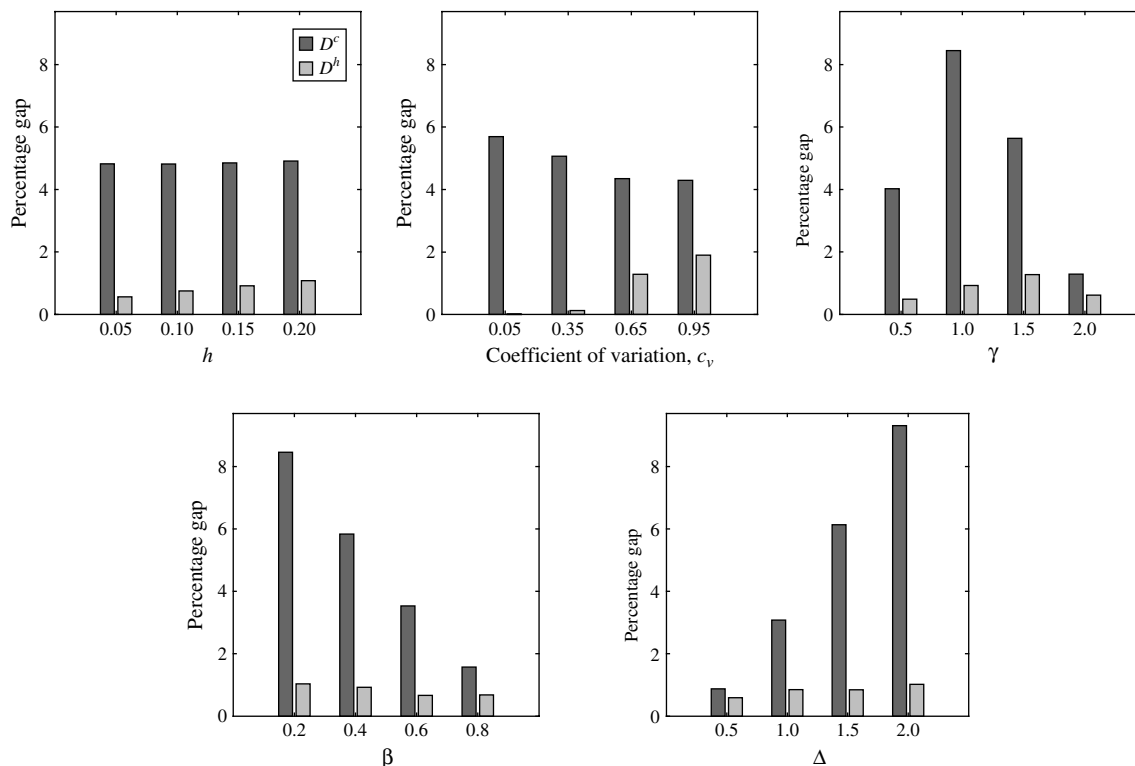
- The benefit of jointly rather than sequentially optimizing price and inventory (i.e., D^h) increases as holding cost and variability in demand increase. This is because the optimal policy in some sense strikes a balance between increased sales and increased inventory-related costs, but these inventory-related costs increase as holding cost or demand variability increase.

To further explore the impact of problem parameters, in Table 2 we explore the change in the two optimality gaps change with respect to specific Δ , β , and c_v values (averaging over other parameters) and observe that the heuristic policy that sequentially sets price followed by inventory level performs worse

Table 1 Summary Statistics of D^c and D^h

| | D^c | D^h |
|--------------------|-------|-------|
| Mean | 4.85 | 0.83 |
| Standard deviation | 5.79 | 1.29 |
| Minimum | 0.00 | 0.00 |
| 25th percentile | 0.00 | 0.00 |
| Median | 2.80 | 0.02 |
| 75th percentile | 8.00 | 1.35 |
| Maximum | 27.54 | 7.67 |

Figure 6 D^c and D^h for Different Values of Problem Parameters



Note. The left and right bars represent D^c and D^h , respectively.

than the constant price policy (that is, $D^c < D^h$) in the following cases:

- If demand is highly variable (c_v is large) and the difference between the maximum and minimum demand at the sale price is small, the heuristic policy on the average performs worse than the constant price policy. In this case, varying price and thus order quantity will lead to fluctuations in demand and inventory levels. Because the heuristic first sets pricing policy without considering inventory effect, it cannot account for inventory-related costs. As a result, the increase in the inventory-related costs and lost revenue, arising from fluctuations in inventory under the heuristic policy, outweighs the benefit of using periodic discounts. In other words, when demand is highly variable, unless you can carefully coordinate pricing and inventory, the benefit of varying price to take advantage of intertemporal demand effects is, at best, limited.

- When Δ decreases and β increases, the marginal benefit of taking advantage of intertemporal demand

interactions in terms of increased sales also decreases. Therefore, in this case also, the heuristic loses more by ignoring the relationship between pricing and inventory than it gains by manipulating prices.

5.3. Impact of Problem Parameters on Optimal Policy

In this portion of our study, we explore the impact of problem parameters on the optimal pricing and inventory policy. To do this, we examine the optimal inventory and pricing policy for the first period under a variety of scenarios, and illustrate the changes in the optimal policy as a function of starting inventory level in Figures 7 and 8.

In column (A) of Figure 7, we vary the number of periods since the last sale, k , while fixing other parameters. Observe that as k increases, the smallest inventory level at which it becomes optimal to offer a sale ($\tilde{s}_i(k)$, defined in §3) decreases and, at the same time, the base-stock level at the sale price (i.e., $s_i(p^s, k)$) increases. This is quite intuitive because

Table 2 Mean D^c and D^h for Various Levels of c_v , β , and Δ

| c_v | Δ | β | | | | | | | |
|-------|----------|---------|-------|-------|-------|-------|-------|-------|-------|
| | | 0.2 | | 0.4 | | 0.6 | | 0.8 | |
| | | D^c | D^h | D^c | D^h | D^c | D^h | D^c | D^h |
| 0.05 | 0.5 | 2.11 | 0.00 | 1.60 | 0.00 | 1.09 | 0.00 | 0.63 | 0.00 |
| | 1 | 6.70 | 0.00 | 4.48 | 0.00 | 2.83 | 0.00 | 1.51 | 0.00 |
| | 1.5 | 12.22 | 0.00 | 8.50 | 0.00 | 5.03 | 0.00 | 2.56 | 0.00 |
| | 2 | 18.01 | 0.00 | 12.39 | 0.00 | 7.70 | 0.00 | 3.66 | 0.00 |
| 0.35 | 0.5 | 1.42 | 0.01 | 0.96 | 0.05 | 0.64 | 0.00 | 0.29 | 0.16 |
| | 1 | 5.92 | 0.22 | 3.89 | 0.08 | 2.07 | 0.09 | 1.01 | 0.21 |
| | 1.5 | 11.51 | 0.26 | 7.63 | 0.06 | 4.70 | 0.06 | 1.94 | 0.06 |
| | 2 | 17.03 | 0.23 | 11.89 | 0.25 | 7.20 | 0.12 | 2.94 | 0.06 |
| 0.65 | 0.5 | 1.25 | 1.20 | 0.73 | 1.15 | 0.28 | 0.48 | 0.06 | 0.68 |
| | 1 | 4.74 | 1.59 | 3.20 | 1.30 | 1.88 | 1.42 | 0.61 | 0.92 |
| | 1.5 | 9.96 | 1.97 | 6.72 | 1.44 | 3.99 | 0.83 | 1.58 | 1.01 |
| | 2 | 14.82 | 1.88 | 10.59 | 1.97 | 6.38 | 1.37 | 2.78 | 1.33 |
| 0.95 | 0.5 | 1.35 | 1.77 | 0.89 | 1.80 | 0.50 | 1.01 | 0.20 | 1.12 |
| | 1 | 4.61 | 2.38 | 3.07 | 1.78 | 1.93 | 2.03 | 0.78 | 1.56 |
| | 1.5 | 9.49 | 2.61 | 6.60 | 2.26 | 3.90 | 1.21 | 1.78 | 1.73 |
| | 2 | 14.13 | 2.40 | 10.23 | 2.61 | 6.35 | 2.06 | 2.87 | 2.04 |

as k increases, the distribution of the demand at the sale price stochastically increases, so offering a sale becomes more appealing, and more inventory is needed to meet the demand at the sale price. Likewise, increasing γ (which represents the minimum increase of the mean demand at the sale price from the mean demand at the retail price) or increasing Δ (which corresponds to the increase in the maximum mean demand at the sale price) leads to the same qualitative behavior, as can be seen in columns (B) and (C) of Figure 7, respectively.

Column (A) of Figure 8 shows how the optimal policy changes in β , the rate at which the mean demand at the sale price increases in k . Observe that for a given k , the mean demand at the sale price increases as β decreases. Hence, as β increases, $\tilde{s}_i(k)$ increases, the base-stock level at the sale price (i.e., $s_i(p^s, k)$) decreases, and the region of starting inventory for which offering a sale is optimal decreases. Offering a sale therefore becomes less appealing. In Column (B) of Figure 8, we illustrate how the optimal policy changes when the coefficient of variation (c_v) of demand increases. Observe that as c_v increases, the standard deviation of demands at the retail and sale prices also increase. As a consequence, order-up-to levels increase and so do inventory-associated costs (both holding cost and lost revenue). Once

c_v becomes sufficiently large, these costs become significant enough so that offering a sale becomes less appealing and is only held to get rid of a high starting inventory, rather than to explicitly take advantage of intertemporal demand effects. In Column (C) of Figure 8, we observe that as the holding cost, h , increases, sales are held at lower inventory levels, and the base-stock levels at both the sale and retail prices are lower, as expected.

6. Extensions, Conclusions, and Future Research

6.1. Make-to-Stock Model with Backorder

In §3, we characterized the structure of the optimal policy for the lost-sales case. In this section, we extend the result to the case where unsatisfied demand is backlogged. We assume that backlogged buyers do not cancel and reorder even if the price goes down—this assumption is consistent with the models of Federgruen and Heching (1999) and Chen and Simchi-Levi (2004a, b). Hence, the inventory level x may be positive or negative. As in §2, h^+ is the per-unit holding cost, but h^- now represents the per-unit backorder cost. If the inventory level is raised to y_t from x_t , and the price is set to $p_t \in \{p^r, p^s\}$ when the state in period t is (x_t, k_t) , then the state in the next period is described as follows:


$$x_{t+1} = y_t - \xi_t(p_t, k_t), \quad t = 1, \dots, T, \quad \text{and}$$

$$k_{t+1} = \begin{cases} k_t + 1, & \text{if } p_t = p^r; \\ 1, & \text{if } p_t = p^s. \end{cases}$$

We assume that demand is fulfilled on a first-come-first-serve basis, so that no demand in the current period will be satisfied without first clearing backorders. Just as in the lost-sales case, we subtract the ordering cost cx in each period to transform the original model into a more tractable form for determining the optimal policy. As before, let $\hat{V}_t(x, k)$ be the optimal expected discounted revenue from period t and onwards, and $\hat{J}_t(y, p; k)$ be the expected discounted revenue of a policy that raises the inventory to y and sets the retail price p in period t and follows the optimal policy afterward. Then,

$$\hat{V}_t(x, k) = \max_{p \in \{p^r, p^s\}} \left\{ \max_{y \geq x} \hat{J}_t(y, p; k) \right\}$$

Figure 7 Optimal Policy in First Period ($t = 1$) as a Function of Initial Inventory to Illustrate Change in Optimal Policy in k , γ , and Δ , Respectively



Note. We hold the remaining problem parameters fixed: $c_v = 0.35$; $h = 0.05$; $\beta = 0.4$.

where $\hat{J}_t(y, p; k)$ satisfies

$$\begin{aligned} \hat{J}_t(y, p; k) &= \int_0^\infty [p\xi - cy + \alpha c(y - \xi) - h^+[y - \xi]^+ - h^-[y - \xi]^- \\ &\quad + \alpha \hat{V}_{t+1}(y - \xi, k_{t+1})] \phi_t(\xi | p, k) d\xi. \end{aligned}$$

Rather than Assumption 3, for this backorder version of our model we utilize the following technical assumption:

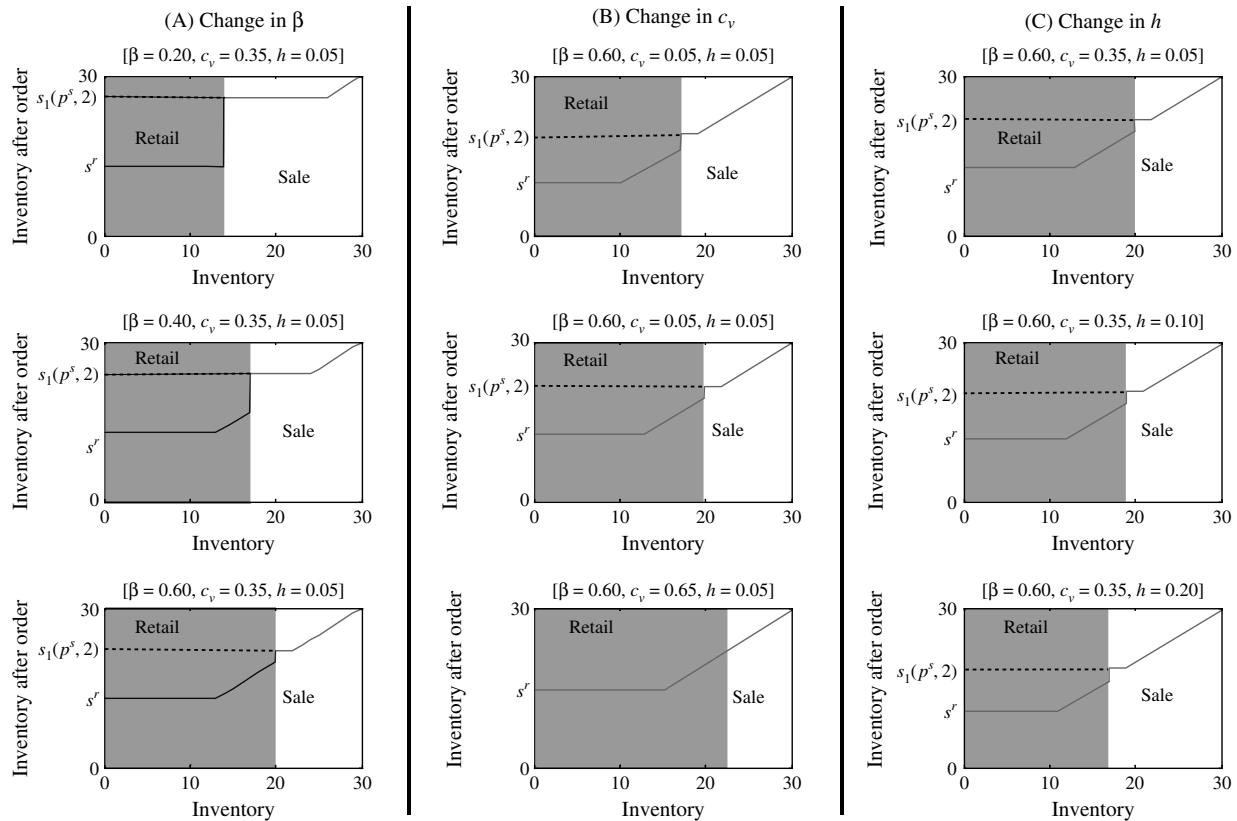
ASSUMPTION 4. $h^- > (1 - \alpha)c$.

Assumption 4 implies that satisfying the demand in the current period is better than delaying one period and satisfying it in the next period, and is frequently called the *nonspeculative* assumption in the inventory literature. Furthermore, the result is trivial when Assumption 4 does not hold: It is optimal to simply accrue all backorders until the end of horizon.

Utilizing Assumptions 1, 2, and 4, we can prove results analogous to Theorems 1 and 2 for this backorder case, and using these results, we are able to characterize the optimal inventory and pricing policy for the backorder case. As in the lost-sales case, $\tilde{s}_i(k)$ represents the starting inventory level below which it is optimal to charge the retail price. Replicating the analysis for the backorder case, it can be shown that $\tilde{s}_i(k)$ must belong to one of the following three cases: (i) $\tilde{s}_i(k) = -\infty$, (ii) $\tilde{s}_i(k) \in (s^r, s_t(p^s, k))$, or (iii) $\tilde{s}_i(k) \geq s_t(p^s, k)$.

THEOREM 6. *Suppose Assumptions 1, 2, and 4 hold. If there have been k periods since the last sale, the optimal pricing and inventory policy in period t takes one of the following three forms described in Theorem 3 with cases (i)–(iii) replacing cases (i)–(iii) in Theorem 3.*

The proof is similar to that of Theorem 3 and is therefore omitted.

Figure 8 Optimal Policy in First Period ($t = 1$) as a Function of Initial Inventory to Illustrate Change in Optimal Policy in β , c_v , and h , Respectively

Note. We hold the remaining problem parameters fixed: $k = 2$, $\gamma = 0.5$, $\Delta = 1.0$.

6.2. Conclusion

In this paper, we introduce and analyze a model that explicitly considers the timing effect of intertemporal pricing—the concept, found in practice, that the amount of demand during a sale is increasing in the time since the last sale. We present structural results that characterize the complex interaction between the decision to have a sale and inventory ordering decisions. Surprisingly, the sale decision is not necessarily monotonic in time since the last sale, and the sale/no sale decision is not necessarily monotonic in starting inventory level.

In our computational analysis, we find that compared to a fixed-price policy, we observe an average gain in profit of almost 5% from optimally making promotion and inventory decisions accounting for intertemporal demand, and develop a simple heuristic that achieves 70% to 80% of this gain. We find that this profit gain increases (to a maximum of about 25%) as the variability of demand

decreases, or the increase in demand resulting from delaying sales increases. From a policy perspective, we find that offering a sale becomes more appealing as the “gap” between demand at the sale price and demand at the regular price increases, but becomes less appealing as demand variation or holding cost increases. We believe that these insights contribute to the understanding of the impact of intertemporal demand effects on inventory and pricing and provide a rationale for the offering of periodic discounts.

Acknowledgments

This research was partially supported by NSF Grants DMI-0092854 and DMI-0200439. The authors thank the editor, associate editor, and three anonymous referees for their valuable and helpful insights.

References

- Ahn, H., M. Gümüş, P. Kaminsky. 2007. Pricing and manufacturing decisions when demand is a function of prices in multiple periods. *Oper. Res.* 55(6) 1039–1057.

- Assuncao, J. L., R. J. Meyer. 1993. The rational effect of price promotions on sales and consumption. *Management Sci.* **39**(5) 517–535.
- Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management.* **10**(3) 339–359.
- Bitran, G. R., S. V. Mondschein. 1997. An application of yield management to the hotel industry considering multiple day stays. *Oper. Res.* **43**(3) 427–443.
- Cachon, G. P., R. Swinney. 2009. Purchasing, pricing and quick response in the presence of strategic consumers. *Management Sci.* **55**(3) 497–511.
- Chan, L. M. A., Z. J. M. Shen, D. Simchi-Levi, J. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. S. David, S. D. Wu, Z. Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*. Springer Science-Business Media, New York, 335–392.
- Chen, F., S. Ray, Y. Song. 2006. Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales. *Naval Res. Logist.* **53**(2) 117–136.
- Chen, X., D. Simchi-Levi. 2004a. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* **52**(6) 887–896.
- Chen, X., D. Simchi-Levi. 2004b. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Math. Oper. Res.* **29**(3) 698–723.
- Cheng, F., S. P. Sethi. 1999. A periodic review inventory model with demand influenced by promotion decisions. *Management Sci.* **45**(11) 1510–1523.
- Conlisk, J., E. Gerstner, J. Sobel. 1984. Cyclic pricing by a durable goods monopolist. *Quart. J. Econom.* **99**(3) 489–505.
- Dharmadhikari, S., K. Joag-Dev. 1988. *Unimodality, Convexity, and Applications*. Academic Press Series in Probability and Mathematical Statistics. Academic Press, Boston.
- Eliashberg, J., R. Steinberg. 1993. Marketing-production joint decision-making. J. Eliashberg, G. L. Lilien, eds. *Handbooks in Operations Research and Management Science: Marketing*, Vol. 5. North Holland, Amsterdam, 827–880.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Management Sci.* **49**(10) 1287–1309.
- Elmaghraby, W., A. Gülcü, P. Keskinocak. 2008. Optimal pre-announced markdowns in the presence of rational customers with multi-unit demands. *Manufacturing Service Oper. Management.* **10**(1) 126–148.
- Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3) 454–475.
- Feinberg, E., M. E. Lewis. 2007. Optimality inequalities for average cost Markov decision processes and the stochastic cash balance problem. *Math. Oper. Res.* **32**(4) 769–783.
- Fox, E., R. Metters, J. Semple. 2006. Optimal inventory policy with two suppliers. *Oper. Res.* **54**(2) 389–393.
- Gallego, G., G. Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.
- Ibragimov, I. A. 1956. On the composition of unimodal distributions. *Theory Probab. Appl.* **1**(2) 255–260.
- Iglehart, D. L. 1963. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Sci.* **9**(2) 259–267.
- Lippman, S. A. 1975. On dynamic programming with unbounded rewards. *Management Sci.* **21**(11) 1225–1233.
- Pesendorfer, M. 2002. Retail sale: A study of pricing behavior in supermarkets. *J. Bus.* **75** 33–66.
- Polatoglu, H., I. Sahin. 2000. Optimal procurement policies under price-dependent demand. *Internat. J. Production Econom.* **65**(2) 141–171.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Slade, M. E. 1998. Optimal pricing with costly adjustment: Evidence from retail-grocery prices. *Rev. Econom. Stud.* **65**(1) 87–107.
- Slade, M. E. 1999. Sticky prices in a dynamic oligopoly: An investigation of (s, S) thresholds. *Internat. J. Indust. Organ.* **17**(4) 477–511.
- Sobel, J. 1984. The timing of sale. *Rev. Econom. Stud.* **51**(3) 353–368.
- Sobel, J. 1991. Durable goods monopoly with entry of new customers. *Econometrica* **59**(5) 1455–1485.
- Song, Y., S. Ray, T. Boyaci. 2009. Technical note: Optimal dynamic joint inventory-pricing control for multiplicative demand with fixed order costs and lost sales. *Oper. Res.* **57**(1) 245–250.
- Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Sci.* **53**(5) 726–741.
- Varian, H. R. 1980. A model of sale. *Amer. Econom. Rev.* **70** 651–659.
- Zhou, Y. P., M. Fan, M. Cho. 2007. On the threshold purchasing behavior of customers facing dynamically priced perishable products. Working paper, University of Washington, Seattle.