PRODUCTION AND OPERATIONS MANAGEMENT
Vol. 19, No. 3, May–June 2010, pp. 249–260
ISSN 1059-1478 | EISSN 1937-5956 | 10 | 1903 | 0249
DOI 10.3401/poms.1080.01099

On the Benefit of Inventory-Based Dynamic Pricing Strategies

Hong Chen
Sauder School of Business, University of British Columbia, Vancouver, British Columbia, Canada V6T 1Z2,
China Europe International Business School, Pudong, Shanghai 201206, China, hong.chen@sauder.ubc.ca

Owen Q. Wu
Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109, owenwu@bus.umich.edu

David D. Yao
Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027, yao@ieor.columbia.edu

We study the optimal pricing and replenishment decisions in an inventory system with a price-sensitive demand, focusing on the benefit of the inventory-based dynamic pricing strategy. We find that demand variability impacts the benefit of dynamic pricing not only through the magnitude of the variability but also through its functional form (e.g., whether it is additive, multiplicative, or others). We provide an approach to quantify the profit improvement of dynamic pricing over static pricing without having to solve the dynamic pricing problem. We also demonstrate that dynamic pricing is most effective when it is jointly optimized with inventory replenishment decisions, and that its advantage can be mostly realized by using one or two price changes over a replenishment cycle.

Key words: joint inventory and pricing decisions; demand variability; dynamic pricing; Brownian demand model
History: Received: May 2008; Accepted: July 2009, after 1 revision.

1. Introduction

Inventory-based pricing strategies have been practiced in many industry sectors. For example, Aguirregabiria (1999) investigated a supermarket chain and found that retail prices tend to decrease when procurement orders are placed and increase between two orders. In the personal computer industry, Byrnes (2003) made the following observation on Dell's pricing strategy: "While its competitor's prices were stable with periodic adjustments, Dell's prices varied significantly from week to week as the company modified its prices to push products where component inventory was building beyond prescribed levels." In the automotive industry, the transaction price is typically negotiable, and dealers are usually willing to lower the price if they have a high level of inventory. Zettelmeyer et al. (2006) empirically found that a dealership moving from a situation of inventory shortage to an average inventory level lowers transaction prices by about 1%, corresponding to 15% of the dealers' average profit margin per vehicle.

In these examples, prices vary with inventory levels, and so price fluctuations can be observed even when demand is stable. One reason for this inventory-based pricing is that a high inventory level provides the firm with an incentive to lower the price so as to stimulate demand and reduce inventory holding cost. This rationale has been formally studied in the literature under various settings, e.g., Aguirregabiria (1999), Federgruen and Heching (1999), and Chen and Simchi-Levi (2006).

Understanding the profit drivers of dynamic pricing is of significant value to industry practitioners and researchers alike. In the industry practices cited above, will a higher demand variability enhance or limit the advantage of dynamic pricing over fixed prices? How will this impact, if any, change with respect to other profit drivers of dynamic pricing, such as holding cost and revenue? Furthermore, when the demand and cost parameters change, adjusting price alone is a common practice, as opposed to a *joint* decision on price *and* inventory level (via adjusting replenishment quantities). This is often a reflection upon the lack of coordination between sales and operations units within a firm. But will the advantage of a joint pricing-replenishment optimization outweigh that of dynamic pricing per se, and will this advantage be impacted by demand variability as well?

Our study addresses these issues and aims at generating insights that can help managers understand when and how to use dynamic pricing as an effective profit enhancement tool. Some of our findings confirm existing practices. For instance, we find dynamic pricing to be most effective in settings where the finished-goods inventory is expensive to hold and expensive to replenish in small batches. This is consistent with the prevalent practice of dynamic pricing in personal computer and automobile industries. Besides the cost drivers, we have also identified other drivers to the profitability of dynamic pricing, such as marginal revenue and demand variability. Indeed, we have developed a formula to capture explicitly the composition effect of all these profit drivers *without* having to solve the dynamic pricing problem. In addition, our study has brought to bear some rather subtle new insights highlighted below:

(i) The impact of demand variability on the benefit of dynamic pricing is reflected not only in the magnitude of the variability but also, and often more significantly, through the functional form of the variability (e.g., additive, multiplicative, or others). This calls for more sophisticated approaches to measuring and modeling demand variability.
(ii) In a dynamic pricing strategy, the number of price changes used in a replenishment cycle has a strong diminishing marginal benefit. Using a *single* price change per cycle, one can already capture about 75% of the profit improvement, and this percentage is quite insensitive to all model parameters.
(iii) Inventory-based dynamic pricing is best implemented jointly with the right inventory replenishment decision. A suboptimal replenishment quantity, even if it is only 10% off the optimal level, can easily undo any benefit from dynamic pricing. In other words, dynamic pricing alone is less effective than a joint pricing-replenishment strategy. Furthermore, the advantage of the joint optimization over dynamic pricing alone remains largely intact as demand variability increases.

The joint pricing and inventory replenishment problems have been extensively studied in various settings. We refer the reader to four survey papers: Petruzzi and Dada (1999), Yano and Gilbert (2003), Elmaghraby and Keskinocak (2003), and Chan et al. (2004). Here we highlight some papers closely related to our study. Federgruen and Heching (1999) examined a periodic-review model with price-dependent demand and linear replenishment cost, without a setup cost. They show that a base-stock list-price policy is optimal for both average and discounted objectives. When a replenishment setup cost is included, the (s, S, p) policy has been proven optimal under various settings. Periodic review with backorder setting is considered by Chen and Simchi-Levi (2004a, b) and Feng and Chen (2004). Periodic review with lost sales setting is considered by Polatoglu and Sahin (2000) and Chen et al. (2006). Huh and Janakiraman (2008) provide an approach for proving and generalizing many of the early results for both backorder and lost sales settings. Continuous-review models are studied by Feng and Chen (2003) and Chen and Simchi-Levi (2006). Markov modulated demand models are analyzed by Yin and Rajaram (2007) and Gayon et al. (2009).

Quantifying the profit improvement of inventory-based dynamic pricing over static pricing is important as it can help a firm weigh the benefits against potential shortfalls in using dynamic pricing (e.g., customer dissatisfaction). Results along this line have been limited to numerical examples. Federgruen and Heching (1999) experimented with a multiplicative demand case in a periodic-review system, and found that the benefit of dynamic pricing increases as demand variability increases. They reported a maximum of 6.54% increase in profit compared with a fixed pricing strategy. With an order-setup cost, Feng and Chen (2004) showed that the profit improvement of dynamic pricing is limited, while the profit improvement as a percentage of static pricing profit could be large when the static pricing profit is low. For a periodic-review system with lost sales, Chen et al. (2006) found that the profit improvement of dynamic pricing increases in the fixed ordering cost. Yin and Rajaram (2007) found that the benefit of dynamic pricing increases in the fixed ordering cost and demand variability. Gayon and Dallery (2007) pointed out that dynamic pricing is potentially more beneficial when the replenishment process is partially controlled. In a setting with strategic consumers, Gallego et al. (2008) showed that static-pricing can be optimal under some situations. To the best of our knowledge, this is the first paper that develops theoretical bounds on the benefit of inventory-based dynamic pricing. These theoretical bounds reveal how various factors drive the profitability of dynamic pricing.

The key element in our study is the Brownian demand model, which leads to an explicit characterization of the various drivers to the profitability of dynamic pricing, and enables us to obtain the insights highlighted earlier. Many studies have employed Brownian motion to model cumulative demand and investigated inventory control problems. Bather (1966) considered optimizing an (s, S) policy for a continuous-review inventory system with Brownian demand. Puterman (1975) studied the optimal control policy in a production-inventory system. Constantinides and Richard (1978) proved the optimality of threshold policies for a class of diffusion control problems. Studies along this line include those by Vickson (1986), Sulem (1986), Chao (1992), and Beyer (1994). More recently, Bensoussan et al. (2005) proved the optimality of (s, S) policy when the demand process consists of a compound Poisson process and a Brownian motion; Benkherouf (2007) further extended the model to include general holding cost. All of the above studies employ the Brownian demand model and focus on inventory cost minimization. The joint pricing and replenishment strategy is the main subject of the current paper.

The paper is organized as follows. A formal description of our model is presented in section 2. The benefit of dynamic pricing is examined in section 3. Bounds and approximation for the benefit of dynamic pricing is developed in section 4. The importance of the joint optimization is emphasized in section 5. Concluding remarks are summarized in section 6.

2. Model Elements

2.1. Demand Model

The cumulative demand up to time t is denoted as $D(t)$, and modeled by a diffusion process:

$$D(t) = \int_0^t \lambda(p_u)du + \int_0^t \sigma(\lambda(p_u))dB(u), \quad t \geq 0, \quad (1)$$

where p_t is the price charged at time t, $\lambda(p_t)$ is the demand rate at time t, $\sigma(\lambda)$ measures the variability of the demand (or the error of demand forecast) when the demand rate is λ, and $B(t)$ denotes the standard Brownian motion.

ASSUMPTION 1. *The inverse demand function $p(\lambda)$ is strictly decreasing and twice continuously differentiable. The revenue rate $r(\lambda) = p(\lambda)\lambda$ is strictly concave in λ.*

Denote $\rho(\lambda) := \sigma(\lambda)^2/\lambda$, which is a relative measure of demand variability in our model. To motivate, consider the case of a renewal demand process with the inter-arrival time having a mean $1/\lambda$ and a standard deviation σ_0. We can approximate this renewal process by a Brownian motion: $\lambda t + \lambda^{3/2}\sigma_0 B(t)$ (see Chen and Yao 2001). Here, $\sigma(\lambda) = \lambda^{3/2}\sigma_0$; and hence, $\rho(\lambda) = \lambda^2\sigma_0^2$, which is the squared coefficient of variation of the inter-arrival time, a common measure for demand variability.

ASSUMPTION 2. *The relative demand variability $\rho(\lambda)$ is convex and twice continuously differentiable.*

A class of demand variability functions that satisfy the above assumption is $\sigma(\lambda) = \sigma\lambda^{\beta}$ and correspondingly $\rho(\lambda) = \sigma^2\lambda^{2\beta-1}$, where $\beta \leq \frac{1}{2}$ or $\beta \geq 1$. Three special cases are given below:

- The constant demand variability $\sigma(\lambda) = \sigma$: $\rho(\lambda) = \sigma^2/\lambda$;
- The linear demand variability $\sigma(\lambda) = \sigma\lambda$: $\rho(\lambda) = \sigma^2\lambda$;
- The square-root demand variability $\sigma(\lambda) = \sigma\sqrt{\lambda}$: $\rho(\lambda) = \sigma^2$.

The constant demand variability case is a continuous-time analog to the additive demand model in the periodic-review setting. When the demand variability is $\sigma(\lambda) = \sigma\lambda$, the demand model (1) can be written as $dD(t) = \lambda(p_t)(dt + \sigma dB(t))$, which is an analog to the multiplicative demand model in a discrete setting. The square-root demand variability function is motivated by matching mean and standard deviation with a Poisson process. (Note that a Poisson process with rate λ has mean λt and standard deviation $\sqrt{\lambda t}$.)

The functional form of the demand variability depends on the type of consumer and product. For example, the demand from price-sensitive consumers tends to be more variable when the average demand is higher, whereas the demand from price-insensitive consumers has a price-independent variability. Other exogenous factors, such as weather and economic conditions, will also scale the demand by a random factor. Depending on the magnitude of demand variability from different sources, the aggregate demand variability may exhibit additive, multiplicative, or more complicated forms. In this paper, we do not aim to find the most appropriate functional form of demand variability, which is an empirical issue by itself. We focus on analyzing inventory and pricing decisions under the general form of the demand variability. Wherever appropriate, we use the above three special functional forms as illustrative examples.

2.2. Pricing-Replenishment Policies and Long-Run Average Objective

We assume that the holding cost rate is linear in the quantity held, and let h denote the cost of holding one unit of inventory per unit of time. At any time, the firm can replenish its inventory with quantity S at a replenishment cost $c(S)$, which is an increasing function of S. Assume replenishment is instantaneous, i.e., zero lead time, which is appropriate when the lead time is insignificant relative to the length of the replenishment cycle. Thus, inventory is depleted by the demand stream in a continuous-time fashion and then replenished immediately when it drops to the reorder point. We assume that the demands must be satisfied immediately upon arrival. Consequently, the reorder point is zero, and the replenishment follows a simple order-up-to policy, with the order-up-to level denoted by S. A rigorous proof of the optimality of such a replenishment policy may follow an approach similar to that in Constantinides and Richard (1978).

We allow the pricing decisions to be dynamically adjusted within each replenishment cycle (a cycle is the duration between two consecutive replenishment epochs), but we do not adjust prices continuously over time. Instead, we explicitly consider the frequency of price changes. Let $N \geq 1$ be a given integer, and let $S = S_0 > S_1 > \cdots > S_{N-1} > S_N = 0$. Immediately after a replenishment at the beginning of a cycle, price p_1 is charged until the inventory drops to S_1; price p_2 is then charged until the inventory drops to S_2; ...; and finally when the inventory level drops to S_{N-1}, price p_N is charged until the inventory drops to $S_N = 0$, when another cycle begins. The same pricing strategy applies to all cycles. For simplicity, we set $S_n = S(N-n)/N$. That is, we divide the full inventory S into N equal segments, and price each segment with a different price as the inventory is depleted by the demand. In summary, the decision variables are $(S, p_1, \ldots, p_N)$, where N is fixed. We will study the effect of N in section 3.

Next, we derive the long-run average profit function. Without loss of generality, suppose at time zero the inventory is filled up to S. Let $T_0 = 0$, and T_n be the first time when inventory drops to S_n:

$$T_n := \inf\{\, t \ge 0 \;:\; D(t) = nS/N \,\}, \quad n = 1, 2, \ldots, N.$$

We refer to the duration when the price p_n is applied as period n. The length of period n is therefore $\tau_n := T_n - T_{n-1}$. Since T_n's are stopping times, by the strong Markov property of the Brownian motion, τ_n is just the time during which S/N units of demand has occurred under the price p_n:

$$\tau_n \overset{dist.}{=} \inf\{\, t \ge 0 \;:\; \lambda(p_n)t + \sigma(\lambda(p_n))B(t) = S/N \,\}. \quad (2)$$

Let $X(t)$ denote the inventory level at t. Clearly, $X(t) = S - D(t)$ in the first cycle $t \in [0, T_N)$. Since the replenishment and pricing policy is the same for all cycles, $X(t)$ is a regenerative process with the replenishment epochs being its regenerative points. Hence, to optimize the long-run average profit, it suffices to focus on the first cycle.

Applying integration by parts, and recognizing that $dX(t) = -dD(t)$, $X(T_{n-1}) = S_{N-1}$, and $X(T_n) = S_n$, we have

$$\begin{aligned}\int_{T_{n-1}}^{T_n} X(t)dt &= T_nS_n - T_{n-1}S_{n-1} - \int_{T_{n-1}}^{T_n} t dX(t) \\ &= T_nS_n - T_{n-1}S_{n-1} - \int_{T_{n-1}}^{T_n} T_{n-1}dX(t) \\ &\quad + \int_{T_{n-1}}^{T_n} (t - T_{n-1})dD(t) \\ &= \tau_nS_n + \int_{T_{n-1}}^{T_n} (t - T_{n-1})[\lambda(p_n)dt \\ &\quad + \sigma(\lambda(p_n))dB(t)].\end{aligned}$$

A simple change of variable yields

$$\int_{T_{n-1}}^{T_n} (t - T_{n-1})dt = \int_0^{\tau_n} u du = \frac{\tau_n^2}{2},$$

whereas

$$\begin{aligned}&\mathsf{E}\left[\int_{T_{n-1}}^{T_n} (t - T_{n-1})dB(t)\right] \\ &= \mathsf{E}\left[\int_{T_{n-1}}^{T_n} t dB(t)\right] - \mathsf{E}[T_{n-1}]\mathsf{E}[B(T_n) - B(T_{n-1})] \\ &= 0,\end{aligned}$$

which follows from the martingale property of $B(t)$ and the optional stopping theorem.

Let $v_n(S, p_n)$ denote the expected profit (sales revenue minus inventory holding cost) during period n. Then, making use of the above derivation, we have

$$\begin{aligned}v_n(S, p_n) &= \frac{p_nS}{N} - \mathsf{E}\left[\int_{T_{n-1}}^{T_n} hX(t)dt\right] \\ &= \frac{p_nS}{N} - h\mathsf{E}[\tau_n]S_n - \frac{1}{2}h\lambda(p_n)\mathsf{E}[\tau_n^2] \\ &= \frac{p_nS}{N} - \frac{hS^2(N - n + \frac{1}{2})}{N^2\lambda(p_n)} - \frac{h\rho(\lambda(p_n))S}{2N\lambda(p_n)}, \end{aligned} \quad (3)$$

where the last equation uses the moments of τ_n proved in Lemma 1 in Appendix A. Note in the above expression the first term is the sales revenue from period n, the second term is the inventory holding cost attributed to the deterministic part of the demand (i.e., the drift part of the Brownian motion), and the last term is the additional holding cost incurred by demand variability in period n.

For ease of analysis, we use the inverse of the demand rates $\{\mu_n = \lambda(p_n)^{-1}, n = 1, \ldots, N\}$ as decision variables and define $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)$. Then, from Lemma 1 in Appendix A, the expected length of a replenishment cycle is $\frac{S}{N}\sum_{n=1}^N \mu_n$, and the long-run average profit can be written as follows:

$$\begin{aligned}V(S, \boldsymbol{\mu}) &= \frac{\sum_{n=1}^N v_n(S, p_n) - c(S)}{\frac{S}{N}\sum_{n=1}^N \mu_n} \\ &= \frac{\sum_{n=1}^N \left[p(\mu_n^{-1}) - \frac{hS}{N}\left(N - n + \frac{1}{2}\right)\mu_n - \frac{1}{2}h\mu_n\rho(\mu_n^{-1}) - a(S)\right]}{\sum_{n=1}^N \mu_n},\end{aligned} \quad (4)$$

where $a(S) = c(S)/S$. In the above average objective, the term $\sum_{n=1}^N \frac{1}{2}h\mu_n\rho(\mu_n^{-1}) / \sum_{n=1}^N \mu_n$ represents the additional holding cost due to demand uncertainty. The firm's problem is

$$\max_{S, \boldsymbol{\mu}} V(S, \boldsymbol{\mu}). \quad (5)$$

3. Characterizing the Benefit of Dynamic Pricing

Understanding what enhances and what limits the benefit of dynamic pricing is crucial to its effective implementation. In this section, we numerically analyze the profit drivers of dynamic pricing strategies.

Our numerical studies reported below cover a wide range of scenarios. Specifically, we consider $c(S) = K + cS$, $\sigma(\lambda) = \sigma\lambda^{\beta}$, and the inverse demand function is $p(\lambda; a, b) = a - b\lambda$ or $p(\lambda; a, b) = -a\log(b\lambda)$, corresponding to the linear and exponential demand functions, respectively. Then, the optimization problem in (5) becomes

$$\max_{S,\boldsymbol{\mu}} V(S, \boldsymbol{\mu}) = \frac{\sum_{n=1}^{N}\left[p(\mu_n^{-1}; a, b) - \frac{hS}{N}(N - n + \frac{1}{2})\mu_n - \frac{1}{2}h\sigma^2\mu_n^{2-2\beta} - \frac{K}{S} - c\right]}{\sum_{n=1}^{N}\mu_n}.$$

Applying a change of variables, $\boldsymbol{\mu} = b\tilde{\boldsymbol{\mu}}$ and $S = K\tilde{S}$, we can rewrite the above problem as follows:

$$\max_{\tilde{S},\tilde{\boldsymbol{\mu}}} V(\tilde{S}, \tilde{\boldsymbol{\mu}}) = \frac{\sum_{n=1}^{N}\left[p(\tilde{\mu}_n^{-1}; a, 1) - \frac{Khb\tilde{S}}{N}\left(N - n + \frac{1}{2}\right)\tilde{\mu}_n - \frac{1}{2}h\sigma^2(b\,\tilde{\mu}_n)^{2-2\beta} - \tilde{S}^{-1} - c\right]}{b\sum_{n=1}^{N}\tilde{\mu}_n}. \tag{6}$$

The objective in (6) indicates that the optimal solution $(\tilde{S}^*, \tilde{\boldsymbol{\mu}}^*)$ depends only on six parameters: $(N, a, Khb, h\sigma^2 b^{2-2\beta}, \beta, c)$. Thus, to sweep across all possible scenarios and study when dynamic pricing is more profitable, we can choose to vary six independent parameters $(N, a, c, \sigma, \beta, h)$ while fixing (K, b). In the case of linear demand $p(\lambda) = a - b\lambda$, we can also fix c because parameters a and c appear in (6) in the form of $a - c$. In the following numerical analysis, the scale of certain parameters may appear to be too large (e.g., for holding cost h), but from (6) we can see that increasing the scale of certain parameters is equivalent to increasing the scale of others, which helps to sweep across a wide range of scenarios and improve the robustness of our results. In the rest of this section, we use V_N^* to denote the profit under the optimal N prices and replenishment level.

3.1. Impact of Demand Variability

We are particularly interested in the effect of (σ, β, h), since these are the parameters that determine the additional holding cost due to demand variability. The composite effect of these three parameters on the profit is shown in Figure 1. We have also tested these effects using various values of (N, a, c) and different demand functions (including $a = 50$, 75, 100 for the linear demand, $a = 50$, 75, 100 and $c = 5$, 10, 15 for the exponential demand, and $V_N^* - V_1^*$ with $N = 2, 3, \ldots, 7$). Figure 1 reports $V_8^* - V_1^*$ for a linear demand case. The following observations are robust across all settings.

First, demand variability always has a negative effect on V_1^*, the profit under a static-pricing strategy. Second, dynamic pricing is not necessarily more valuable with a higher demand variability. How the profit improvement $V_8^* - V_1^*$ changes with demand variability depends on the form of the demand variability function. Higher additive demand variability (Figure 1a) reduces the benefit of dynamic pricing, while higher multiplicative demand variability (Figure 1b) tends to enhance it. Demand variability has no effect on the improvement when $\sigma(\lambda) = \sigma\sqrt{\lambda}$ (Figure 1c). Numerically, we also found that the price spread (difference between the highest and the lowest prices) tends to decrease in the additive demand variability and increase in the multiplicative demand variability. Thus, higher profit improvement is associated with larger spreads among the optimal prices (or, more fluctuating prices).

Figure 1 Profit Improvement of Dynamic Pricing Over Static Pricing: $c(S) = 100 + 5S$, $p(\lambda) = 50 - \lambda$

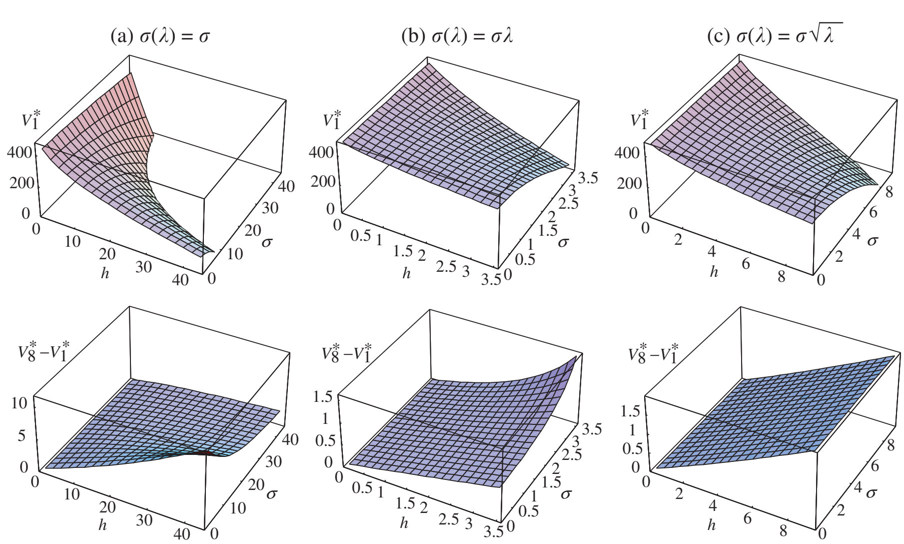

Figure 2 Effect of the Number of Price Changes: $c(S) = 100+5S$, $h = 20$, $\sigma(\lambda) = 20$, $p(\lambda) = 50 - \lambda$

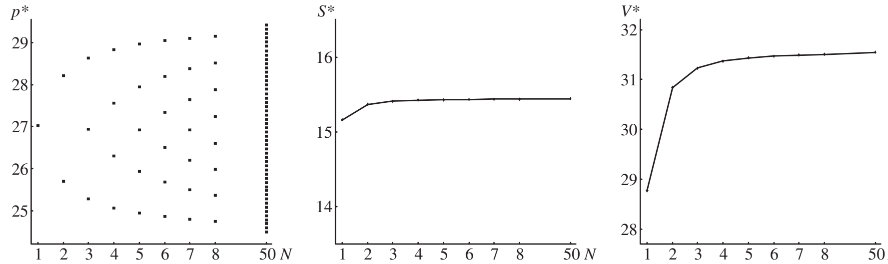

Refer to the discussion at the beginning of section 3 and note that this setting is equivalent to a class of settings with the same value of $(Khb, h\sigma^2 b^2, a - c)$, e.g., $c(S) = 20+50S$, $h = 5$, $\sigma = 2$, $p(\lambda) = 95 - 20\lambda$).

Another observation is that the *relative* profit improvement, $(V_8^* - V_1^*)/V_1^*$, increases in σ for all three forms of demand variability. This is clear for the cases $\sigma(\lambda) = \sigma\lambda$ and $\sigma\sqrt{\lambda}$, since $V_8^* - V_1^*$ is non-decreasing in σ, while V_1^* is decreasing in σ, as shown in Figure 1b and c. For $\sigma(\lambda) = \sigma$, when V_1^* approaches zero, the absolute improvement does not diminish. We can numerically verify that the relative profit improvement is indeed increasing in the additive demand variability and approaching infinity when V_1^* is close to zero. In a different model, Federgruen and Heching (1999) found that the percentage profit improvement increases in demand uncertainty. Our model yields similar results, but it is interesting to see that the same may not hold for the absolute profit improvement.

Our observation also concerns the impact of σ in connection with h. How these two parameters interact is ambiguous a priori. A higher h increases the need for inventory-based dynamic pricing, which may result in a higher impact of demand variability; but a higher h will likely lead to a lower replenishment level, which leaves less room for the dynamic pricing strategy to be effective. But our numerical results confirm that a higher h enhances the benefit of dynamic pricing and magnifies the impact of demand variability on the profit improvement, in both absolute and relative senses.

3.2. Single vs. Multiple Price Changes

We have observed in our numerical studies that the number of price changes has a strong diminishing marginal benefit. Figure 2 illustrates the optimal decisions and profit corresponding to different values of N. Here, using only two prices (i.e., a single price change within a replenishment cycle) already achieves about 75% of the profit improvement, and beyond $N = 8$, the marginal improvement is essentially nil. Thus, V_8^* is essentially the best profit that can be achieved via dynamic pricing.

This calls for further studies on the profit improvement associated with a single price change, i.e., a pricing strategy that is the least dynamic. We measure the effectiveness of a single price change by computing the ratio $(V_2^* - V_1^*)/(V_8^* - V_1^*)$ for thousands of problem instances with wide range of parameter values. We have tested a total of 7488 problem instances, and each instance actually represents a class of equivalent instances (with the same value of $(a, Khb, h\sigma^2 b^{2-2\beta}, \beta, c)$; refer to the discussion at the beginning of section 3). Table 1 reports the summary statistics of the effectiveness ratio. Remarkably, regardless of the model parameters, the two-price strategy *consistently* achieves about 75% of the profit improvement.

4. Bounds and Approximation for the Benefit of Dynamic Pricing

In this section, we derive bounds on the profit improvement of dynamic pricing strategy over static pricing. We further develop an explicit formula to directly quantify the benefit of dynamic pricing without having to solving the N-price problem. The bounds

Table 1 Effectiveness of the Two-Price Strategy: Statistics of $(V_2^* - V_1^*)/(V_8^* - V_1^*)$

Demand form	$\sigma(\lambda) = \sigma$	$\sigma(\lambda) = \sigma\lambda$	$\sigma(\lambda) = \sigma\sqrt{\lambda}$
Linear demand			
Mean (standard deviation)	0.759 (0.004)	0.761 (0.003)	0.760 (0.002)
Minimum / maximum	0.726 / 0.762	0.741 / 0.780	0.753 / 0.780
Number of instances	1200*	1200	1200
Exponential demand			
Mean (standard deviation)	0.759 (0.002)	0.761 (0.001)	0.760 (0.001)
Minimum / maximum	0.748 / 0.761	0.756 / 0.762	0.757 / 0.762
Number of instances	1296**	1296	1296
Overall			
Mean (standard deviation)		0.7600 (0.0025)	
Minimum / maximum		0.726 / 0.780	
Number of instances		7488	

*1200 instances are based on 20×20 different values of (h, σ) as in Figure 1 and $a = 50, 75, 100$.
**1296 instances are based on 12×12 different values of (h, σ) (same range as in Figure 1 but with a coarser mesh), $a = 50, 75, 100$, and $c = 5, 10, 15$.

and approximation formula reveal how various factors drive the profitability of dynamic pricing.

Intuitively, the benefit of dynamic pricing should be closely related to the degree of price fluctuation. Thus, we first investigate the range in which the optimal prices (or equivalently, optimal $\boldsymbol{\mu}$) may vary. Proposition 1 below gives this range for a fixed replenishment level S.

PROPOSITION 1. *For fixed S, let $\boldsymbol{\mu}^*$ be an optimal decision corresponding to the optimal N prices. Then for $1 \leq m \leq n \leq N$,*

$$\frac{S(n-m)}{N(G_1' + G_2'/h)} \leq \mu_n^* - \mu_m^* \leq \frac{S(n-m)}{N(G_1 + G_2/h)},$$

where

$$G_1 = \min_{\mu \in \mathscr{A}} \left\{ \frac{1}{2} \rho''(\mu^{-1})\, \mu^{-3} \right\}, G_1' = \max_{\mu \in \mathscr{A}} \left\{ \frac{1}{2} \rho''(\mu^{-1})\, \mu^{-3} \right\},$$

$$G_2 = \min_{\mu \in \mathscr{A}} \{ -r''(\mu^{-1})\, \mu^{-3} \}, \quad G_2' = \max_{\mu \in \mathscr{A}} \{ -r''(\mu^{-1})\, \mu^{-3} \}.$$

and $\mathscr{A}$ is any set that contains $[\mu_1^*, \mu_N^*]$.

Note that Assumptions 1 and 2 imply that

$$\frac{d^2 p(\mu^{-1})}{d\mu^2} = \mu^{-3} r''(\mu^{-1}) < 0 \quad \text{and}$$
$$\frac{d^2(\mu \rho(\mu^{-1}))}{d\mu^2} = \mu^{-3} \rho''(\mu^{-1}) \geq 0. \tag{7}$$

Hence, G_1, G_1', G_2, and G_2' are all positive, which implies that $\mu_n^* > \mu_m^*$ if $n > m$, i.e., the optimal prices are higher when the inventory level is lower.

In next proposition, we consider the profit difference between the optimal dynamic pricing and the optimal static pricing strategy, both being jointly optimized with the replenishment level.

PROPOSITION 2. *Let $(S_N^*, \boldsymbol{\mu}^*)$ be an optimal solution to the N-price dynamic pricing and replenishment problem (5) and let V_N^* be the corresponding optimal profit. Let (S^*, μ^*) be the optimal solution to the static (i.e., $N = 1$) pricing and replenishment problem. Then,*

$$V_N^* - V_1^* \leq \frac{h\, {S_N^*}^2 (1 - N^{-2})}{24 \bar{\mu}} \left[\frac{2}{G_1 + G_2/h} - \frac{G_1 + G_2/h}{(G_1' + G_2'/h)^2} \right],$$

$$V_N^* - V_1^* \geq \frac{h\, S^{*2} (1 - N^{-2})}{24 \mu^*} \left[\frac{2}{G'_1 + G'_2/h} - \frac{G_1' + G_2'/h}{(G_1 + G_2/h)^2} \right],$$

where $\bar{\mu} = \frac{1}{N} \sum_{n=1}^{N} \mu_n^*$, and

$$G_1 = \min_{\mu \in \mathscr{A}} \left\{ \frac{1}{2} \rho''(\mu^{-1})\, \mu^{-3} \right\}, \quad G_1' = \max_{\mu \in \mathscr{A}} \left\{ \frac{1}{2} \rho''(\mu^{-1})\, \mu^{-3} \right\},$$

$$G_2 = \min_{\mu \in \mathscr{A}} \{ -r''(\mu^{-1})\, \mu^{-3} \}, \quad G_2' = \max_{\mu \in \mathscr{A}} \{ -r''(\mu^{-1})\, \mu^{-3} \}.$$

and $\mathscr{A}$ is any set that contains $[\mu_1^*, \mu_N^*]$.

Motivated by the above upper bound and lower bound, we construct an approximation for the profit improvement. Let $G_1^0 = \frac{1}{2}\rho''(\mu^{*-1})\, \mu^{*-3}$ and $G_2^0 = -r''(\mu^{*-1})\, \mu^{*-3}$, and notice that

$$\frac{2}{G_1 + G_2/h} - \frac{G_1 + G_2/h}{(G_1' + G_2'/h)^2} > \frac{1}{G_1^0 + G_2^0/h}$$
$$> \frac{2}{G'_1 + G'_2/h} - \frac{G_1' + G_2'/h}{(G_1 + G_2/h)^2}.$$

In the lower bound for $V_N^* - V_1^*$, we replace the term in the brackets by $(G_1^0 + G_2^0/h)^{-1}$ and replace $(1 - N^{-2})$ by 1 when N is large. Then, the profit improvement of dynamic pricing can be approximated as follows:

$$V_N^* - V_1^* \approx \frac{h S^{*2}}{24 \mu^* (G_1^0 + G_2^0/h)}, \quad N \geq 8; \tag{8}$$

where $G_1^0 = \frac{1}{2}\rho''(\mu^{*-1})\, \mu^{*-3}$ and $G_2^0 = -r''(\mu^{*-1})\, \mu^{*-3}$ measure the convexity of the demand variability function and the concavity of the revenue function, respectively (refer to [7]). Note that the above approximation formula depends only on the static pricing strategy (S^*, μ^*) and allows us to estimate the

Table 2 Accuracy of the Approximation of Profit Improvement: Statistics of $\frac{hS^{*2}/(V_8^* - V_1^*)}{24\mu^*(G_1^0 + G_2^0/h)}$

Demand form	$\sigma(\lambda) = \sigma$	$\sigma(\lambda) = \sigma\lambda$	$\sigma(\lambda) = \sigma\sqrt{\lambda}$
Linear demand			
Mean (standard deviation)	1.003 (0.008)	1.011 (0.006)	1.008 (0.006)
Minimum/maximum	0.945/1.014	0.969/1.036	0.991/1.038
Number of instances	1200	1200	1200
Exponential demand			
Mean (standard deviation)	1.003 (0.007)	1.010 (0.003)	1.007 (0.004)
Minimum/maximum	0.971/1.012	0.999/1.015	0.996/1.014
Number of instances	1296	1296	1296
Overall			
Mean (standard deviation)		1.007 (0.007)	
Minimum/maximum		0.945/1.038	
Number of instances		7488	

All above statistics are based on the same problem instances as in Table 1.

benefit of dynamic pricing using the static pricing information.

To study the accuracy of the approximate formula in (8), we examine the ratio between the approximate value in (8) and the exact value of $V_8^* - V_1^*$ for wide range of parameter settings. A ratio that is close to 1 indicates high accuracy of this approximation. In Table 2 we list the statistics of this ratio, which are all quite close to 1. This indicates that the formula in (8) provides a good estimator for the profit improvement without having to solve the N-price replenishment problem.

When $N = 2$, multiplying the right hand side of (8) by 75% gives an approximation for the benefit from using only a single price change:

$$V_2^* - V_1^* \approx \frac{h\,S^{*2}}{32\mu^*\left(G_1^0 + G_2^0/h\right)}.$$

where $G_1^0 = \frac{1}{2}\rho''(\mu^{*-1})\,\mu^{*-3}$ and $G_2^0 = -r''(\mu^{*-1})\,\mu^{*-3}$.

Intuitively, the above formulae also suggest that for an inventory-based dynamic pricing strategy to yield a significant benefit, the products must be expensive to hold and expensive to replenish in small quantities, such as personal computers, home appliances and electronics, and automobiles. Not surprisingly, these are the kind of industry sectors where dynamic pricing is usually practiced. The formulae also suggest that both the deterministic ($r(\lambda)$) and stochastic ($\rho(\lambda)$) components of the demand must be taken into account when deciding whether dynamic pricing could lead to significant profit improvement. It is the concavity or convexity of these components that are driving the profitability of dynamic pricing. To provide intuition, we note that in the numerator of (5), the additional holding cost term $-\sum_{n=1}^{N}\frac{1}{2}h\mu_n\rho(\mu_n^{-1})$ and the revenue term $\sum_{n=1}^{N} p(\mu_n^{-1})$ are both concave in $\boldsymbol{\mu}$ (refer to [7]). The concavity of these two terms limits the extent to which the optimal prices vary, and thus limits the profit improvement of dynamic pricing. A measure of their concavity is G_1^0 and G_2^0. These are exactly the terms in the approximate formula.

5. Dynamic Pricing Alone vs. Joint Pricing Replenishment

It is commonly seen in practice that firms implement dynamic pricing without considering it jointly with the replenishment quantity. This is typically due to a lack of coordination between sales and operations units within an organization. Here, we study the benefit of a *joint* pricing-replenishment decision, as opposed to executing dynamic pricing alone.

Given a replenishment level S, let $V_N(S)$ denote the profit under the optimal N-price strategy. Consider the following two ratios:

$$\frac{V_8(S) - V_1(S)}{V_8^* - V_1(S)}, \quad \frac{V_8(S) - V_1(S)}{V_2^* - V_1(S)}.$$

The first ratio measures the proportion of the maximum improvement that can be achieved by implementing dynamic pricing alone. The second ratio compares dynamic pricing using multiple price changes against the joint pricing-replenishment strategy using only one price change. These ratios depend on how far the fixed variable S is away from the optimum S^*. We consider three scenarios: S is fixed at $(1 \pm 0.05)S^*$, at $(1 \pm 0.1)S^*$, and at $(1 \pm 0.15)S^*$. Then, for each scenario and each of the 7488 problem instances as in Tables 1 and 2, we compute the above two ratios. The results are summarized in Table 3, in terms of the medians of the ratios (we use median since the distribution is typically skewed).

The first ratio shows that when S deviates from the optimal S^*, dynamic pricing alone only achieves a fraction of the best possible improvement via joint optimization. The second ratio indicates that dynamic pricing alone generally performs no better than the joint optimization with only a single price change. In both cases, the performance of dynamic pricing alone deteriorates drastically when S deviates further from

Table 3 Median Profit Improvement Ratio

	$\sigma(\lambda) = \sigma$			$\sigma(\lambda) = \sigma\lambda$			$\sigma(\lambda) = \sigma\sqrt{\lambda}$		
Deviation of S	$\pm$ 5%	$\pm$ 10%	$\pm$ 15%	$\pm$ 5%	$\pm$ 10%	$\pm$ 15%	$\pm$ 5%	$\pm$ 10%	$\pm$ 15%
$\frac{V_8(S)-V_1(S)}{V_8^*-V_1(S)}$									
Linear demand	0.81	0.51	0.32	0.61	0.28	0.15	0.71	0.38	0.21
Exponential demand	0.82	0.53	0.33	0.68	0.34	0.19	0.77	0.46	0.27
$\frac{V_8(S)-V_1(S)}{V_2^*-V_1(S)}$									
Linear demand	1.00	0.59	0.35	0.72	0.30	0.16	0.85	0.42	0.22
Exponential demand	1.02	0.61	0.36	0.81	0.37	0.20	0.95	0.51	0.29

The medians are taken from the same 7488 problem instances as in Tables 1 and 2.

Figure 3 Effect of Demand Variability on the Median Profit Improvement Ratio: Linear Demand

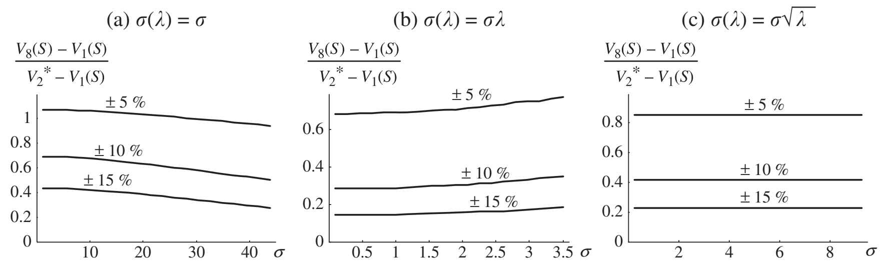

The medians are taken from the linear demand rate instances in Table 3.

the optimal level. The results strongly indicate the strength of a simple, two-price strategy along with a jointly optimized replenishment quantity.

Figure 3 further elaborates on Table 3 by showing the effect of demand variability on the above ratios. The figure for exponential demand is very similar and omitted. Notice that the ratios remain largely at the same level as demand variability changes, indicating the robustness of the findings in Table 3.

6. Concluding Remarks

In this paper, we have studied an inventory-based dynamic pricing strategy: the price is changed whenever the inventory drops by a fraction of the replenishment quantity. We address the drivers of the profit improvement of dynamic pricing. The benefit of dynamic pricing can be significant under certain conditions. First, higher holding cost, lower additive demand variability, higher multiplicative demand variability, and less concave revenue function all enhance the profit improvement of dynamic pricing. Second, dynamic pricing strategy should be implemented jointly with an optimal replenishment level; a suboptimal replenishment quantity can easily offset the benefit of dynamic pricing.

Through extensive numerical analysis, we found that 75% of the benefit of dynamic pricing can be captured by using only one price change. This implies that effective dynamic pricing does not necessarily use frequent price changes, especially when frequent price changes increase customer dissatisfaction. In this paper, we also provide a formula to directly estimate the benefit of dynamic pricing without the need to solve it.

Finally we make a remark on the timing of the price changes. This paper considers pricing based on an even division of the total replenishment quantity. In general, the segments are not necessarily equal, but our numerical experiments show that the optimal segments (optimized jointly with prices and replenishment level) are quite even and the corresponding profit is almost the same as using equal segments. Furthermore, in practical situations when the prices are chosen from a discrete set, the optimal prices may be equal across several segments, and the solution based on equal segments can effectively tell how many prices to use, when to change prices and what prices to change to, as evident from the following example.

Example. Let $\lambda(p) = 50 - p$, $c(S) = 100 + S$, $h = 1$, $\sigma(\lambda) = 10$. Suppose the prices have to be integer valued, and S has to be a multiple of 5. Using $N = 140$, we find the optimal policy is to order 70 units for each cycle, charge a price of \$25 until the inventory level drops to 67 units, charge \$26 until the inventory drops to 19, and charge \$27 until the inventory runs out. The policy uses only three prices (despite the choice of a large N value) and yields an average profit of \$528.745.

In other words, the equal partitioning of $[0, S]$ (into 140 segments) does not prevent us from finding the optimal partitioning (of three uneven segments). ■

Acknowledgments
Part of this research was undertaken while the first author was affiliated with the Cheung Kong Graduate School of Business, China, and the second author was affiliated with the University of British Columbia. The first author is supported in part by NSERC Grant; the third author is supported in part by HK/RGC Grants CUHK4179/05 and CUHK4182/07.

Appendix A

Lemma 1. *Let $\lambda_n = \lambda(p_n)$. For the stopping time τ_n defined in (2), we have*

$$\mathsf{E}[\tau_n] = \frac{S}{N\lambda_n} \quad \text{and}$$
$$\mathsf{E}[\tau_n^2] = \frac{\sigma(\lambda_n)^2 S}{N\lambda_n^3} + \frac{S^2}{N^2\lambda_n^2}.$$

Proof. (This lemma follows from well-known results regarding optional stopping applied to Brownian motion; refer to, e.g., Karlin and Taylor, 1975. The proof here is included for completeness.) Omit the subscript n, and let $x := S/N$. The definition of τ in (2) implies that $\lambda\tau + \sigma B(\tau) = x$. This, combined with $\mathsf{E}B(\tau) = 0$, which follows from applying optional stopping to $B(t)$, leads to $\mathsf{E}[\tau] = x/\lambda$. Next, we derive the second moment of τ. Since $\sigma B(\tau) = x - \lambda\tau$, we have

$$\sigma^2[B^2(\tau) - \tau] = (x - \lambda\tau)^2 - \sigma^2\tau.$$

Applying optional stopping to the martingale $B^2(t) - t$, we have

$$\mathsf{E}[(x - \lambda\tau)^2] = \sigma^2\mathsf{E}[\tau], \quad \text{or}$$
$$\lambda^2\,\mathsf{Var}(\tau) = \sigma^2\mathsf{E}[\tau].$$

This leads to

$$\mathsf{E}[\tau^2] = \frac{\sigma^2}{\lambda^2}\mathsf{E}[\tau] + \mathsf{E}^2[\tau] = \frac{\sigma^2 x}{\lambda^3} + \frac{x^2}{\lambda^2}. \qquad \square$$

Proof of Proposition 1. Based on (7), the numerator of the objective in (4) is strictly concave in $\boldsymbol{\mu}$. The ratio of a concave function over a positive linear function is known to be pseudo-concave (Mangasarian, 1970). Thus, the first-order conditions are sufficient. The first-order conditions for μ_m and μ_n give

$$\frac{\partial V(S,\boldsymbol{\mu}^*)}{\partial \mu_m} - \frac{\partial V(S,\boldsymbol{\mu}^*)}{\partial \mu_n} = \frac{\frac{dp}{d\mu}(\mu_m^{*\,-1}) - \frac{dp}{d\mu}(\mu_n^{*\,-1}) + \frac{hS}{N}(m-n) + h(f'(\mu_n^*) - f'(\mu_m^*))}{\sum_{n=1}^{N}\mu_n^*} = 0, \tag{A1}$$

where $f(\mu) := \frac{1}{2}\mu\rho(\mu^{-1})$. The relations in (7) imply that

$$G_1 \le f''(\mu) \le G_1', \quad \text{and}$$

$$G_2 \le -\frac{d^2p}{d\mu^2}(\mu^{-1}) \le G_2', \quad \forall \mu \in \mathscr{A}.$$

Hence, for $\mu_n^* \ge \mu_m^*$,

$$(\mu_n^* - \mu_m^*)G_1 \le f'(\mu_n^*) - f'(\mu_m^*) \le (\mu_n^* - \mu_m^*)G_1' \tag{A2}$$

$$(\mu_n^* - \mu_m^*)G_2 \le \frac{dp}{d\mu}(\mu_m^{*\,-1}) - \frac{dp}{d\mu}(\mu_n^{*\,-1}) \le (\mu_n^* - \mu_m^*)G_2', \tag{A3}$$

Combining (A1)–(A3) yields the desired inequalities in the proposition. $\square$

Proof of Proposition 2. The optimal single-price policy (S^*, μ^*) can achieve a profit no lower than the feasible single-price policy: $(S_N^*, \bar{\mu})$. Thus,

$$V_1^* \ge \frac{p(\bar{\mu}^{-1}) - \frac{hS_N^*}{2}\bar{\mu} - \frac{1}{2}h\bar{\mu}\rho(\bar{\mu}^{-1}) - a(S_N^*)}{\bar{\mu}}.$$
$$V_N^* - V_1^* \le \frac{\sum_{n=1}^{N}\left[p(\mu_n^{*\,-1}) - p(\bar{\mu}^{-1})\right] - \frac{hS_N^*}{N}\sum_{n=1}^{N}\left[\left(N - n + \frac{1}{2}\right)\mu_n^* - \frac{N}{2}\bar{\mu}\right] - \frac{h}{2}\sum_{n=1}^{N}\left[\mu_n^*\rho(\mu_n^{*\,-1}) - \bar{\mu}\rho(\bar{\mu}^{-1})\right]}{N\bar{\mu}}. \tag{A4}$$

Consider the first term in the numerator of (A4):

$$\begin{aligned}
p(\mu_n^{*\,-1}) - p(\bar{\mu}^{-1}) &= \int_{\bar{\mu}}^{\mu_n^*}\frac{dp}{d\mu}(\mu^{-1})d\mu \\
&= \int_{\bar{\mu}}^{\mu_n^*}\left[\frac{dp}{d\mu}(\bar{\mu}^{-1}) + \int_{\bar{\mu}}^{\mu}\frac{d^2p}{d\mu^2}(s^{-1})ds\right]d\mu \\
&\le \int_{\bar{\mu}}^{\mu_n^*}\left[\frac{dp}{d\mu}(\bar{\mu}^{-1}) - G_2(\mu - \bar{\mu})\right]d\mu \\
&= (\mu_n^* - \bar{\mu})\frac{dp}{d\mu}(\bar{\mu}^{-1}) - \frac{1}{2}G_2(\mu_n^* - \bar{\mu})^2,
\end{aligned}$$

where the inequality is due to $\frac{d^2p}{d\mu^2} \le -G_2$ (by definition of G_2 and (7)). Summing the above inequality over $n = 1, \ldots, N$ and dividing both sides by $N\bar{\mu}$, we have

$$\frac{\sum_{n=1}^{N}\left[p(\mu_n^{*\,-1}) - p(\bar{\mu}^{-1})\right]}{N\bar{\mu}} \le -\frac{G_2\sum_{n=1}^{N}(\mu_n^* - \bar{\mu})^2}{2N\bar{\mu}}$$

Note that $\frac{1}{N-1}\sum_{n=1}^{N}(\mu_n^* - \bar{\mu})^2$ is the variance of a random variable that is uniformly distributed over $\mu_1^*, \ldots, \mu_N^*$. That variance is minimal if $\mu_1^*, \ldots, \mu_N^*$ are minimally spaced out. From Proposition 1, the minimal distance between μ_n and μ_{n+1} is $\frac{S_N^*}{N(G_1' + G_2'/h)} \equiv \Delta$. Therefore, the above variance is bounded below by the variance of the random variable that is uniformly distributed over $\Delta, 2\Delta, \ldots, N\Delta$, which can be computed as $(N+1)N\Delta^2/12$. Thus,

$$\frac{\sum_{n=1}^{N}[p(\mu_n^{*-1})-p(\bar{\mu}^{-1})]}{N\bar{\mu}} \le -\frac{G_2}{2N\bar{\mu}}\frac{(N-1)(N+1)N\Delta^2}{12} = -\frac{G_2 {S_N^*}^2(1-N^{-2})}{24\bar{\mu}(G_1'+G_2'/h)^2}. \tag{A5}$$

The above lines of proof also apply if $p(\mu^{-1})$ is replaced by $-\frac{1}{2}\mu\rho(\mu^{-1})$ and G_2 is replaced by G_1. In parallel with (A5), we have

$$\frac{-\frac{h}{2}\sum_{n=1}^{N}[\mu_n^*\rho(\mu_n^{*-1})-\bar{\mu}\rho(\bar{\mu}^{-1})]}{N\bar{\mu}} \le -\frac{hG_1 {S_N^*}^2(1-N^{-2})}{24\bar{\mu}(G_1'+G_2'/h)^2}. \tag{A6}$$

Now we consider the second term in the numerator of (A4).

$$\frac{-\frac{hS_N^*}{N}\sum_{n=1}^{N}\left[\left(N-n+\frac{1}{2}\right)\mu_n^*-\frac{N}{2}\bar{\mu}\right]}{N\bar{\mu}} = \frac{hS_N^*\sum_{n=1}^{N}\left[\left(n-N-\frac{1}{2}\right)\mu_n^*+\frac{N}{2}\mu_n^*\right]}{N^2\bar{\mu}}$$
$$= \frac{hS_N^*[2\mu_1^*+4\mu_2^*+\cdots+2N\mu_N^*-(N+1)(\mu_1^*+\cdots+\mu_N^*)]}{2N^2\bar{\mu}}$$
$$= \frac{hS_N^*[(N-1)(\mu_N^*-\mu_1^*)+(N-3)(\mu_{N-1}^*-\mu_2^*)+\cdots]}{2N^2\bar{\mu}}$$

where in the last line, the series ends with $\mu_{N/2+1}^*-\mu_{N/2}^*$ if N is even, and ends with $2(\mu_{(N+3)/2}^*-\mu_{(N-1)/2}^*)$ if N is odd.

$$V_N^*-V_1^* \ge \frac{\sum_{n=1}^{N}\left[p(\tilde{\mu}_n^{-1})-p(\mu^{*-1})\right]-\frac{hS^*}{N}\sum_{n=1}^{N}\left[\left(N-n+\frac{1}{2}\right)\tilde{\mu}_n-\frac{N}{2}\mu^*\right]-\frac{h}{2}\sum_{n=1}^{N}\left[\tilde{\mu}_n\rho(\tilde{\mu}_n^{-1})-\mu^*\rho(\mu^{*-1})\right]}{N\mu^*}.$$

Applying Proposition 1 and the identity:

$$(N-1)^2+(N-3)^2+\ldots+\left(N+1-2\left\lfloor\frac{N}{2}\right\rfloor\right)^2=\frac{(N-1)N(N+1)}{6},$$

we have

$$\frac{-\frac{hS_N^*}{N}\sum_{n=1}^{N}\left[\left(N-n+\frac{1}{2}\right)\mu_n^*-\frac{N}{2}\bar{\mu}\right]}{\sum_{n=1}^{N}\mu_n^*} \le \frac{h{S_N^*}^2\left[(N-1)^2+(N-3)^2+\cdots+\left(N+1-2\left\lfloor\frac{N}{2}\right\rfloor\right)^2\right]}{2N^3\bar{\mu}(G_1+G_2/h)}$$
$$=\frac{h{S_N^*}^2(N-1)N(N+1)}{12N^3\bar{\mu}(G_1+G_2/h)} = \frac{h{S_N^*}^2(1-N^{-2})}{12\bar{\mu}(G_1+G_2/h)}. \tag{A7}$$

Combining the inequalities in (A4)–(A7), we have

$$V_N^*-V_1^* \le \frac{h{S_N^*}^2(1-N^{-2})}{24\bar{\mu}} \times \left[\frac{2}{G_1+G_2/h}-\frac{G_1+G_2/h}{(G_1'+G_2'/h)^2}\right].$$

Next we prove the lower bound for the profit improvement. Let (S^*,μ^*) be the optimal single price policy. Fixing $S=S^*$, we optimize over the N prices, and denote the optimal solution as $(S^*,\boldsymbol{\mu}^*)$. Then we construct a feasible policy $(S^*,\tilde{\boldsymbol{\mu}})$, where $\tilde{\mu}_n=\mu^*+\mu_n^*-\frac{1}{N}\sum_n\mu_n^*$. We assume that $\tilde{\mu}_n\in\mathscr{A}$. (If not, we could always scale $\tilde{\mu}_n$ and/or enlarge $\mathscr{A}$ such that $\tilde{\mu}_n\in\mathscr{A}$, and the resulting lower bound will be slightly different, but the insights remain the same.)

Since the feasible policy performs no better than the optimal one, we have

Then, following exactly the same logic of the proof for the upper bound, we have the following inequalities corresponding to (A5), (A6), and (A7), respectively:

$$\frac{\sum_{n=1}^{N}[p(\tilde{\mu}_n^{-1})-p(\mu^{*-1})]}{N\mu^*} \ge -\frac{G_2' S^{*2}(1-N^{-2})}{24\mu^*(G_1+G_2/h)^2}$$

$$\frac{-\dfrac{h}{2}\sum_{n=1}^{N}\left[\tilde{\mu}_n\,\rho(\tilde{\mu}_n{}^{-1})-\mu^*\rho(\mu^{*-1})\right]}{N\mu^*}$$
$$\geq -\frac{hG_1'\,S^{*2}(1-N^{-2})}{24\mu^*(G_1+G_2/h)^2}$$

$$\frac{-\dfrac{hS^*}{N}\sum_{n=1}^{N}\left[\left(N-n+\dfrac{1}{2}\right)\tilde{\mu}_n-\dfrac{N}{2}\mu^*\right]}{N\mu^*}$$
$$\geq \frac{h\,S^{*2}\left(1-N^{-2}\right)}{12\mu^*\left(G_1'+G_2'/h\right)}.$$

Combining the above inequalities gives the lower bound

$$V_N^*-V_1^*\geq \frac{h\,S^{*2}\left(1-N^{-2}\right)}{24\mu^*}\times\left[\frac{2}{G'_1+G'_2/h}-\frac{G_1'+G_2'/h}{(G_1+G_2/h)^2}\right].\qquad\square$$

References

Aguirregabiria, V. 1999. The dynamics of markups and inventories in retailing firms. *Rev. Econ. Stud.* **66**(2): 275–308.

Bather, J. A. 1966. A continuous time inventory model. *J. Appl. Probab.* **3**(2): 538–549.

Benkherouf, L. 2007. On a stochastic inventory model with a generalized holding costs. *Eur. J. Oper. Res.* **182**(2): 730–737.

Bensoussan, A., R. H. Liu, S. P. Sethi. 2005. Optimality of an (s, S) policy with compound Poisson and diffusion demands: A quasi-variational inequalities approach. *SIAM J. Control Optim.* **44**(5): 1650–1675.

Beyer, D. 1994. An inventory model with Wiener demand process and positive lead time. *Optimization* **29**(2): 181–193.

Byrnes, J. 2003 Dell manages profitability, not inventory. *Harvard Business School Working Knowledge*, June. Available at http://hbswk.hbs.edu/archive/3497.html (accessed July 31, 2009).

Chan, L. M. A., Z. J. Shen, D. Simchi-Levi, J. L. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. Simchi-Levi, D., D. Wu, Z. J. Shen (eds). *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*. Kluwer Academic Publishers, Boston, MA, 335–392.

Chao, H. 1992. The EOQ model with stochastic demand and discounting. *Eur. J. Oper. Res.* **59**: 434–443.

Chen, F. Y., S. Ray, Y. Song. 2006. Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales. *Nav. Res. Logist.* **53**(2): 117–136.

Chen, H., D. D. Yao. 2001. *Fundamentals of Queueing Networks: Performance, Asymptotics and Optimization*. Springer-Verlag, New York.

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.

Chen, X., D. Simchi-Levi. 2006. Coordinating inventory control and pricing strategies: The continuous review model. *Oper. Res. Lett.* **34**(3): 323–332.

Constantinides, G. M., S. F. Richard. 1978. Existence of optimal simple policies for discounted-cost inventory and cash management in continuous time. *Oper. Res.* **26**(4): 620–636.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Manage. Sci.* **49**(10): 1287–1309.

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3): 454–475.

Feng, Y., F. Y. Chen. 2003 Joint pricing and inventory control with setup costs and demand uncertainty. Working paper, Chinese University of Hong Kong.

Feng, Y., F. Y. Chen. 2004 Optimality and optimization of a joint pricing and inventory-control policy for a periodic-review system. Working paper, Chinese University of Hong Kong.

Gallego, G., R. Phillips, O. Şahin. 2008. Strategic management of distressed inventory. *Prod. Oper. Manag.* **17**(4): 402–415.

Gayon, J. P., Y. Dallery. 2007. Dynamic vs static pricing in a make-to-stock queue with partially controlled production. *OR Spectrum* **29**(2): 193–205.

Gayon, J. P., I. Talay-Değirmenci, F. Karaesmen, E. L. Örmeci. 2009. Optimal pricing and production policies of a make-to-stock system with fluctuating demand. *Probab. Eng. Inform. Sci.* **23**(2): 205–230.

Huh, W. T., G. Janakiraman. 2008. (s, S) optimality in joint inventory-pricing control: An alternate approach. *Oper. Res.* **56**(3): 783–790.

Karlin, S., H. Taylor. 1975. *A First Course in Stochastic Processes*. 2nd ed. Academic Press, New York.

Mangasarian, O. L. 1970. Convexity, pseudo-convexity and quasi-convexity of composite functions. *Cahiers du Centre d'Études de Recherche Opérationnelle* **12**: 114–122.

Petruzzi, N. C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* **47**(2): 183–194.

Polatoglu, H., I. Sahin. 2000. Optimal procurement policies under price-dependent demand. *Int. J. Prod. Econ.* **65**(2): 141–171.

Puterman, M. L. 1975. A diffusion process model for a storage system. Geisler, M. A. (ed.) *Logistics*. North-Holland Publishing Co., Amsterdam; American Elsevier Publishing Co., New York, 143–159.

Sulem, A. 1986. A solvable one-dimensional model of a diffusion inventory system. *Math. Oper. Res.* **11**(1): 125–133.

Vickson, R. G. 1986. A single product cycling problem under Brownian motion demand. *Manage. Sci.* **32**(10): 1336–1345.

Yano, C. A., S. M. Gilbert. 2003. Coordinated pricing and production/procurement decisions: A review. Chakravarty, A. K., J. Eliashberg (eds) *Managing Business Interfaces: Marketing, Engineering and Manufacturing Perspectives*. Kluwer Academic Publishers, Boston, MA, 65–104.

Yin, R., K. Rajaram. 2007. Joint pricing and inventory control with a Markovian demand model. *Eur. J. Oper. Res.* **182**(1): 113–126.

Zettelmeyer, F., F. S. Morton, J. Silva-Risso. 2006. Scarcity rents in car retailing: Evidence from inventory fluctuations at dealerships. NBER Working Paper No. 12177.