OPERATIONS RESEARCH
Vol. 59, No. 5, September–October 2011, pp. 1144–1158
ISSN 0030-364X | EISSN 1526-5463 | 11 | 5905 | 1144

http://dx.doi.org/10.1287/opre.1110.0946

Integration of Inventory and Pricing Decisions with Costly Price Adjustments

Xin Chen
Industrial Enterprise and Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801,
xinchen@illinois.edu

Sean X. Zhou, Youhua (Frank) Chen
Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong
{zhoux@se.cuhk.edu.hk, ychen@se.cuhk.edu.hk}

Motivated by the widespread adoption of dynamic pricing in industry and the empirical evidence of costly price adjustments, in this paper we consider a periodic-review inventory model with price adjustment costs that consist of both fixed and variable components. In each period, demand is stochastic and price-dependent. The firm needs to coordinate the pricing and inventory replenishment decisions in each period to maximize its total discounted profit over a finite planning horizon. We develop the general model and characterize the optimal policies for two special scenarios, namely, *a model with inventory carryover and no fixed price-change costs* and *a model with fixed price-change costs and no inventory carryover*. Finally, we propose an intuitive heuristic policy to tackle the general system whose optimal policy is expected to be very complicated. Our numerical studies show that this heuristic policy performs well.

Subject classifications: inventory system; pricing; price adjustment costs; fixed costs; base-stock.
Area of review: Revenue Management.
History: Received March 2008; revisions received March 2009, December 2009, July 2010; accepted January 2011.

1. Introduction

The past decade has witnessed the proliferation of dynamic pricing practice in various industries (see McGill and van Ryzin 1999, Elmaghraby and Keskinocok 2003). Facilitated by sophisticated information technologies such as enterprise resource planning (ERP) systems and electronic tags, price changes are becoming easier. However, these changes are not costless. Indeed, two major types of price adjustment costs are identified in the economics literature: managerial costs and physical costs (or "menu costs" in the economics literature).

Managerial costs are directly related to "the time and attention required of managers to gather the relevant information and to make and implement decisions" (Bergen et al. 2003). Typically, these costs arise within a firm as a result of information gathering, decision making, and communications. Physical costs are incurred for retailers such as Best Buy and Target through the labor costs that result from manually changing thousands of shelf prices within their stores. Other physical costs that firms such as 3M, Ericsson, and Bed Bath and Beyond experience are the costs associated with producing, printing, and distributing their price books or catalogs.

Many empirical studies have shown that both managerial costs and physical costs are significant in retailing and other industries (Rotemberg 1982b, Levy et al. 1997, Slade 1998, Aguirregabiria 1999, Bergen et al. 2003, Zbaracki et al. 2004, Kano 2006). In particular, Levy et al. (1997) and Slade (1998) find through empirical studies that physical costs play a crucial role in the price-setting behavior of retail supermarkets. In a relatively recent study, Bergen et al. (2003) estimate the managerial costs incurred by firms when they change prices to be more than six times the magnitude of the physical costs. Indeed, as Bergen et al. (2003) further assert, the physical costs of changing prices have been reduced because of advances in information technology, whereas the managerial costs of doing so might actually have increased due to the added complexity of dealing with on-line and in-store pricing, the added data from customers buying through web sites, and the additional knowledge and systems required to understand the e-business.

In addition to these empirical findings, a number of quantitative models for analyzing optimal pricing strategies have also been developed in the economics literature. With a fixed price adjustment cost and stochastic inflation, Sheshinski and Weiss (1977) show that (s, S)-type pricing policies are optimal for a continuous-time model with deterministic price-dependent demand. Caplin and Spulber (1987) consider multiple firms that adopt an (s, S) pricing policy and show that the aggregate price of individual firms changes with money supply in the market over time and thus aggregate price stickiness disappears. See Sheshinski and Weiss (1993) and references therein for further discussion on costly price adjustment and (s, S) pricing policies.

Even though the economics literature illustrates that price adjustment costs do exist and play a crucial role in shaping firms' pricing strategies (e.g., Sheshinski and Weiss 1993), the literature on inventory and pricing coordination by the operations management community has largely ignored them. The purpose of this paper is to fill this gap by incorporating price adjustment costs into integrated inventory and pricing models and investigating the structural properties of optimal inventory and pricing policies.

More specifically, we consider a firm that manages a single product, periodic-review finite-horizon inventory system with stochastic price-dependent demand and costly price adjustments. Similar to Federgruen and Heching (1999), at the beginning of each period, an ordering decision is made and a linear ordering cost is incurred. The selling price of the period is also determined. The demand in each period is random and depends on the selling price set at the beginning of the period. However, different from Federgruen and Heching (1999) and other papers on inventory and pricing coordination, a price-adjustment cost is incurred when the price in one period is set to be different from that in the previous period. This might involve a fixed-cost component that is independent of the magnitude of the price changes and a variable-cost component that depends on their magnitude. The objective of the firm is to coordinate the pricing and inventory replenishment decisions in each period to maximize its expected total discounted profit over a finite planning horizon.

It turns out that the structure of the optimal policy in our model allowing for the full generality is rather complicated. Thus, we consider two special cases of our model: *a model with inventory carryover and no fixed price-change costs* and *a model with fixed price-change costs and no inventory carryover*. For the first model with a general convex price adjustment cost, we prove that a state-dependent base-stock inventory and list-price policy is optimal. In this policy, the base-stock level is nonincreasing in the price of the previous period. If the initial inventory level is below the base-stock level, then the list price is charged; otherwise, a discount is offered. Both the list price and the discount depend on the price of the previous period. When the variable cost is piece-wise linear (with two linear pieces) and convex, we can further prove that the optimal price follows a two-sided threshold-type policy, in which the threshold levels are nonincreasing in the initial inventory level and nondecreasing in the price of the previous period. For the second model, we employ Scarf's (1960) concepts of K-concavity and the symmetric K-concavity proposed in Chen and Simchi-Levi (2004a) to show that the optimal pricing strategy is the one-sided (s, S) policy if the price change is uni-directional and the two-sided (s, S, A) policy if it is bi-directional. For the first model, we assume that unsatisfied demand is backlogged, whereas for the second, we assume that it is filled by an emergency order or alternatively, by an order at the beginning of the next period. Under additional conditions, all our structural results hold for models with lost sales.

Finally, for the general problem with fixed plus linear variable price-adjustment costs, we develop an intuitive heuristic policy—a base-stock inventory and two-sided (s, S) pricing policy to manage inventory and set selling prices. We provide an approach to computing the policy parameters. Compared with the optimal policy, which is difficult to compute, our heuristic policy is amenable to practical implementation once the control parameters have been calculated. In addition, the numerical study demonstrates that the heuristic is quite effective.

Our paper falls within the growing research stream on inventory and pricing coordination that started with Whitin (1955). For a review of this literature, readers are referred to Elmaghraby and Keskinocak (2003), Chan et al. (2004), and Chen and Simchi-Levi (2011). Significant progress has been made recently in analyzing integrated inventory and pricing models with a fixed ordering cost and stochastic demand for both backlog (see Chao and Zhou 2006; Chen and Simchi-Levi 2004a, b, 2006; Huh and Janakiraman 2008) and lost sales (see Chen et al. 2005, Huh and Janakiraman 2008, Song et al. 2008) cases.

The previous studies, however, predominantly assume that the price adjustment is costless. To the best of our knowledge, there are only two papers that are closely related to ours: Aguirregabiria (1999) and Celik et al. (2009). Motivated by the phenomenon in practice of large periods without nominal price changes and short periods with low prices, Aguirregabiria (1999) proposes an inventory and pricing model that incorporates both fixed ordering costs and fixed price-adjustment costs but focuses more on empirical studies. Celik et al. (2009) focus on a continuous-time revenue management problem with costly price changes. They characterize the optimal pricing polices for the case of ample inventory and develop several heuristics based on a corresponding fluid model. However, their model does not take into account the inventory replenishment decision. It is also appropriate to mention here that Netessine's (2006) paper, which formulates and analyzes a deterministic model to optimize the timing of price changes, also recognizes the importance of the impact of these costs on inventory and pricing decisions.

The rest of this paper is organized as follows. In §2 we present the general model, introduce the notation, and define a class of general inventory and pricing policies. In §3 we consider the case without fixed costs. In §4 we analyze the model without inventory carryover. In §5 we provide the conditions under which the results of the backlog case can be extended to the lost sales case. We develop a heuristic policy for the general system in §6 and conduct a numerical study. We also extend the heuristic to the case with inventory dependent price adjustment costs. Finally, in §7 we conclude the paper with a discussion of possible extensions.

2. The General Model Setting

Consider a firm that manages a single-product, periodic-review inventory system with price-dependent demand in an N-period planning horizon. In each period, the firm needs to make both inventory replenishment and pricing decisions so as to maximize the total expected discounted profit over the planning horizon.

The time sequence of events is as follows. At the beginning of each period, the firm first reviews its inventory level and makes a replenishment decision. The order the firm places is received immediately (leadtime 0) and incurs a variable cost of c_n per unit. Then, based on the inventory level after replenishment and the selling price in the previous period, the firm decides whether to adjust the price and, if so, by how much. During this period, the random demand is realized, and all the revenues and costs are calculated at the end of the period.

Let p_n be the selling price of period n. For tractability, we assume that demand in period n, denoted by $D_n(p_n, \epsilon_n)$, depends linearly on a stochastic component, ϵ_n, and the selling price, p_n. Specifically,

$$D_n(p_n, \epsilon_n) = \beta_n - \alpha_n p_n,$$

where α_n and β_n are two positive random variables, and $\epsilon_n = (\alpha_n, \beta_n)$. For simplicity, ϵ_n are assumed to be independent of each other across time. The linear demand assumption is made to ensure that the one-period expected profit (excluding possible fixed cost components) is jointly concave in inventory and price. It should be noted that linear demand functions are commonly used in the literature and in practice (e.g., Petruzzi and Dada 1999, Simchi-Levi et al. 2005).

Let $\underline{p}$ and $\bar{p}$ be the lower and upper bounds of selling price p_n. We assume $D_n(p_n, \epsilon_n)$ to be nonnegative for any $p_n \in [\underline{p}, \bar{p}]$ and any realization of ϵ_n. (Therefore, if β_n is drawn from a positively valued range $[\underline{\beta}, \bar{\beta}]$, then α_n will take values in $[0, \underline{\beta}/\bar{p}]$, which ensures that $D_n(p, \epsilon_n)$ is nonnegative.) Let $\bar{D}_n(p_n) = \mathsf{E}[D_n(p_n, \epsilon_n)]$. It is clear that the expected revenue of period n, given as

$$R_n(p_n) = \mathsf{E}[p_n D_n(p_n, \epsilon_n)] = p_n[\mathsf{E}(\beta_n) - \mathsf{E}(\alpha_n) p_n],$$

is concave in p_n.

If the price-adjustment cost involves only a fixed component that is independent of the price-adjustment magnitude, the results in this paper hold for a more general demand function $D_n(p_n, \epsilon_n) = \beta_n - \alpha_n f(p_n)$, where $f(\cdot) \in [f(\underline{p}), f(\bar{p})]$ is a time-independent strictly increasing function, and the expected revenue $f^{-1}(d)[\mathsf{E}(\beta_n) - \mathsf{E}(\alpha_n) d]$ ($f^{-1}(\cdot)$: $[f(\underline{p}), f(\bar{p})] \to [\underline{p}, \bar{p}]$ is the inverse function of $f(\cdot)$) is concave in d. In this case, we can use the inventory level and d as the primal decision variables in our analysis rather than the inventory level and price.

Different from the majority of papers on inventory and pricing coordination, we assume that a price change from one period to the next is costly. Define $\delta(A) = 1$ if $A \neq 0$; otherwise, $\delta(A) = 0$. The cost of a price adjustment with magnitude Δ in period n is denoted by $C_n^p(\Delta)$, which is given as

$$C_n^p(\Delta) = K\delta(\Delta) + U_n(\Delta), \tag{1}$$

where the fixed cost K represents the menu cost or physical cost associated with a price change, and the variable cost $U_n(\cdot)$ represents the managerial or customer cost depending on the magnitude of the price adjustment. In the economics literature, it is commonly assumed that the variable cost $U_n(\cdot)$ is convex and increases with the size of the price change because the decision and internal communication costs are higher for larger price changes (see Zbaracki et al. 2004 for direct evidence from industrial markets and detailed analysis). Several forms of $U_n(\cdot)$ have been used in the economics literature, including piecewise linear functions $U_n(\Delta) = u_n|\Delta|$ (Tsiddon 1991, Slade 1998) and quadratic functions $U_n(\Delta) = u_n\Delta^2$ (Rotemberg 1982a, b; Roberts 1992).

These studies have been inconclusive about the relative magnitude of K and $U_n(\cdot)$, with some suggesting that the menu cost is small and insignificant (see McCallum 1986, Konieczny 1993), and others finding it to be large (see Levy et al. 1997, Bergen et al. 2003). Thus, several different forms of price adjustment costs have been used. For instance, Slade (1998) considers a model with both a fixed cost and piecewise linear variable costs and empirically shows that both exist. Rotemberg (1982a) considers a model with a quadratic variable cost alone and states that this cost primarily accounts for the implicit cost that results from customers' unfavorable reaction to large price changes. Rotemberg (1982b) further finds evidence of a significant variable price-adjustment cost from aggregate U.S. economic data, whereas Aguirregabiria (1999) and Kano (2006) focus only on fixed costs. A more general price-adjustment cost is proposed by Celik et al. (2009), who consider a menu cost K that may depend on the inventory level in addition to a convex variable price adjustment cost. Such a cost is reasonable in settings in which the price tag law requires a price tag for each item. We provide a brief discussion of the difficulty of including it in our model and extend our heuristic to this more general setting in §6.

Depending on whether inventory can be carried over between periods, we analyze two different models. The first is a durable product model in which inventory can be carried over from one period to another. In each period n, the ending inventory incurs a unit holding cost h_n, and any excess demand is assumed to be backlogged and to incur a unit shortage cost b_n. The second model is a perishable product model that does not allow inventory carryover. Any leftover inventory at the end of each period n is salvaged at r_n per unit. In addition, any unsatisfied demand is filled by an emergency order at cost b_n per unit. In the next two sections, we analyze these two models and characterize the optimal pricing and inventory replenishment policies.

At the beginning of each period, given starting inventory level x and the selling price of the previous period, p_{n-1}, let $v_n(x, p_{n-1})$ denote the total discounted optimal profit from period n to N minus $c_n x$ and γ be the discount factor, $0 \leqslant \gamma \leqslant 1$. The dynamic programming formulation of this problem is

$$v_n(x, p_{n-1}) = \max_{y \geqslant x,\, \underline{p} \leqslant p \leqslant \bar{p}} \{-C_n^p(p - p_{n-1}) + L_n(y, p)\}, \quad (2)$$

where

$$L_n(y, p) = R_n(p) + G_n(y, p) + \gamma c_{n+1} \mathsf{E}[y - D_n(p, \epsilon_n)] + \gamma \mathsf{E}[v_{n+1}(y - D_n(p, \epsilon_n), p)], \quad (3)$$

in which if excess demand is backlogged and inventory carryover is allowed, then

$$G_n(y, p) = -c_n y - h_n \mathsf{E}[(y - D_n(p, \epsilon_n))^+] - b_n \mathsf{E}[(D_n(p, \epsilon_n) - y)^+],$$

which is the modified one-period inventory cost function (note that $x^+ = \max(0, x)$, and for $n = 1$, $c_1 x$ is also included); if inventory carryover is not allowed (or excess demand is lost), $G_n(y, p)$ and $v_{n+1}(\cdot, \cdot)$ will be defined in §4 (or §5).

In the definition of $L_n(y, p)$, the first term is the expected revenue; the second is the ordering, inventory holding, and shortage costs; and the last two terms are the optimal expected discounted profit from period $n+1$ to the end of the planning horizon. Note that $G_n(\cdot, \cdot)$ is jointly concave because the demand function is linear, which is critical for our subsequent analysis. For ease of exposition, we assume that the initial price p_0 is given and $v_{N+1}(x, \cdot) = 0$ without loss of generality.

To end this section, we define the following class of inventory and pricing policies. We show that most of the policies we are about to discuss belong to this class.

Definition 1 (Base-Stock Inventory and Two-Sided (s, S) Pricing Policy). An inventory replenishment and pricing policy (y_n^*, p_n^*) is called a base-stock inventory and two-sided (s, S) pricing policy if, given the selling price p_{n-1} of period $n-1$ and the starting inventory level x of period n,

- inventory replenishment follows a base-stock policy with level $\bar{y}(p_{n-1})$, i.e., $y_n^* = \max\{\bar{y}(p_{n-1}), x\}$; and
- selling price p_n^* is determined by two pairs of parameters $(s(x), S(x))$ and $(z(x), Z(x))$ with $s(x) \leqslant S(x) \leqslant Z(x) \leqslant z(x)$ according to the following:

$$p_n^* = \begin{cases} S(x), & \text{if } p_{n-1} \leqslant s(x), \\ Z(x) & \text{if } p_{n-1} \geqslant z(x), \\ p_{n-1} & \text{otherwise.} \end{cases} \quad (4)$$

Note that all the parameters and sets are period-dependent.

For brevity, we denote this class of policies by $\mathfrak{P}$, which includes several simpler and commonly seen policies as special cases. When $s(x) = S(x)$ and $z(x) = Z(x)$, then it becomes a *two-sided threshold policy*. Finally, when only uni-directional price adjustment is allowed, the pricing policy is reduced to a *one-sided policy*.

3. Model with Inventory Carryover

In this section, we analyze our first model formulated as (2), in which inventory can be carried over from one period to another. When inventory can be carried over and fixed price-adjustment costs are incurred, problem (2) becomes very complicated. Indeed, concepts such as K-convexity and symmetric K-convexity, which are effective for analyzing stochastic inventory models and integrated inventory and pricing models with economies of scale, are unlikely to be applicable to a two-dimensional dynamic program (2) with a fixed price-adjustment cost, and thus a totally new convexity related concept might be needed. In the rest of the section, we therefore focus on the model in which a price adjustment incurs only variable costs, i.e., $K = 0$. As Rotemberg (1982b) observes, in some settings variable costs might be significantly more important than fixed costs. We provide a heuristic in §6 for the model with a positive fixed price-adjustment cost.

3.1. Convex Variable Price-Adjustment Costs

In this section, we assume the variable price-adjustment cost $U_n(\Delta)$ is convex and characterize the optimal pricing and inventory ordering policy.

The following two lemmas present some structural properties of the value functions, which will help us to characterize the optimal policies.

Lemma 1. (a) *$v_n(x, p_{n-1})$ is nonincreasing in x.*
(b) *$v_n(x, p_{n-1})$ is jointly concave in x and p_{n-1}.*

Proof. For part (a), note that the objective function in (2) is independent of x. The result is immediate because we are dealing with a maximization problem, and the feasible domain of y becomes smaller when x increases.

We prove part (b) by induction. Because $v_{N+1}(\cdot, \cdot) = 0$, the result holds for $N+1$. Suppose it is true for some $n+1$, $n \leqslant N$. For n, observe that the first two terms in (3) are concave in p and p_{n-1}, $G_n(y, p)$ is jointly concave in y and p, and $R_n(p) - \gamma c_{n+1} D_n(p) = (p - \gamma c_{n+1}) D_n(p)$ is concave in p as $D_n(p)$ is linear and decreasing in p. The concavity of the last term in $L_n(y, p)$ follows from the inductive assumption and the linearity of $D_n(p, \epsilon_n)$. Thus, the objective function in the maximization problem (2) is jointly concave in y, p, and p_{n-1}. This immediately implies that $v_n(x, p_{n-1})$ is jointly concave in x and p_{n-1} because concavity is preserved under maximization (see Proposition 2.2.15 of Simchi-Levi et al. 2005). □

Our next result indicates the submodularity of $v_n(x, p_{n-1})$ and $L_n(y, p)$.

LEMMA 2. *$v_n(x, p_{n-1})$ is submodular in x and p_{n-1} and $L_n(y, p)$ is submodular in y and p.*

PROOF. To facilitate the analysis, we define $w_n(x, \tilde{p}) = v_n(x, -\tilde{p})$ and $J_n(y, \tilde{p}) = L_n(y, -\tilde{p})$. Then,

$$w_n(x, \tilde{p}_{n-1}) = \max_{y \geqslant x,\, -\underline{p} \geqslant \tilde{p} \geqslant -\bar{p}} \{-U_n(\tilde{p}_{n-1} - \tilde{p}) + \tilde{L}_n(y, \tilde{p})\} \qquad (5)$$

with $w_{N+1}(\cdot, \cdot) = 0$, and

$$J_n(y, \tilde{p}) = R_n(-\tilde{p}) + \gamma c_{n+1} \mathsf{E}[y - D_n(-\tilde{p}, \epsilon_n)] + G_n(y, -\tilde{p}) + \gamma \mathsf{E}[w_{n+1}(y - D_n(-\tilde{p}, \epsilon_n), \tilde{p})]. \qquad (6)$$

To prove the lemma, it is sufficient to show that $w_n(x, \tilde{p}_{n-1})$ is supermodular in x and $\tilde{p}_{n-1}$ and $J_n(y, \tilde{p})$ is supermodular in y and $\tilde{p}$, which we prove in the following.

We prove that $w_n(x, \tilde{p}_{n-1})$ is supermodular in x and $\tilde{p}_{n-1}$ by induction. Note that the supermodularity of $J_n(y, \tilde{p})$ will be proved simultaneously. First, $w_{N+1}(x, \tilde{p}_N) = 0$ is obviously supermodular. Assume that $w_{n+1}(x, \tilde{p}_n)$ is supermodular. We now prove that $w_n(x, \tilde{p}_{n-1})$ is supermodular. Note that the function $-U_n(\tilde{p}_{n-1} - \tilde{p})$ is a composite of a one-dimensional concave function with $\tilde{p}_{n-1} - \tilde{p}$ and thus is supermodular in $\tilde{p}$ and $\tilde{p}_{n-1}$ (see Simchi-Levi et al. 2005, Theorem 2.3.6). The same argument can be applied to verify the supermodularity of $G_n(y, -\tilde{p})$ as $D_n(-\tilde{p}, \epsilon_n) = \beta_n + \alpha_n \tilde{p}$. Meanwhile, $R_n(-\tilde{p}) + \gamma c_{n+1} \mathsf{E}[y - D_n(-\tilde{p}, \epsilon_n)]$ is a separable function and thus is clearly supermodular in y and $\tilde{p}$.

It remains for us to prove that $w_{n+1}(y - D_n(-\tilde{p}, \epsilon_n), \tilde{p})$ is supermodular for a given sample path of ϵ_n. For any two pairs $(y, \tilde{p})$ and $(y', \tilde{p}')$ with $y > y'$ and $\tilde{p} < \tilde{p}'$, we have

$$\begin{aligned} &w_{n+1}(y - D_n(-\tilde{p}', \epsilon_n), \tilde{p}') - w_{n+1}(y' - D_n(-\tilde{p}', \epsilon_n), \tilde{p}') \\ &\quad \geqslant w_{n+1}(y - D_n(-\tilde{p}', \epsilon_n), \tilde{p}) - w_{n+1}(y' - D_n(-\tilde{p}', \epsilon_n), \tilde{p}) \\ &\quad \geqslant w_{n+1}(y - D_n(-\tilde{p}, \epsilon_n), \tilde{p}) - w_{n+1}(y' - D_n(-\tilde{p}, \epsilon_n), \tilde{p}), \end{aligned}$$

where the first inequality follows directly from the assumption of the supermodularity of $w_{n+1}(y, \tilde{p})$, and the second inequality follows from the concavity of $w_{n+1}(y, \tilde{p})$ in its first component. Thus, $w_{n+1}(y - D_n(-\tilde{p}, \epsilon_n), \tilde{p})$, and hence $\mathsf{E}[w_{n+1}(y - D_n(-\tilde{p}, \epsilon_n), \tilde{p})]$, is supermodular in y and $\tilde{p}$, which implies that $J_n(y, \tilde{p})$ is also supermodular in y and $\tilde{p}$. Therefore, the objective function in the maximization problem (5) is supermodular in y, $\tilde{p}$, and $\tilde{p}_{n-1}$. In addition, it is easy to see that the feasible set is a lattice. Hence, $w_n(x, \tilde{p}_{n-1})$ is supermodular, as the maximization of a supermodular function over a lattice remains still supermodular (see Theorem 2.7.6 of Topkis 1998, or Proposition 2.3.5 of Simchi-Levi et al. 2005). □

The submodularity of $v_n(x, p_{n-1})$ is quite intuitive and implies that the marginal optimal profit with respect to inventory level x decreases as the previous price p_{n-1} increases. The intuition behind this is that when a firm has an additional unit of inventory, it tends to set a lower price; with a higher starting price, the possible cost of price adjustment is also higher, thus causing a decrease in profit.

We now proceed to show the optimal policy. Let

$$\hat{P}_n(p_{n-1}, y) = \max\left\{ \arg\max_{\underline{p} \geqslant p \geqslant \bar{p}} \{-U_n(p_{n-1} - p) + L_n(y, p)\} \right\},$$

$$(p_n(p_{n-1}), y_n(p_{n-1})) = \max\left\{ \arg\max_{y,\, \underline{p} \geqslant p \geqslant \bar{p}} \{-U_n(p_{n-1} - p) + L_n(y, p)\} \right\}.$$ [1]

The following theorem characterizes the optimal pricing and inventory replenishment policy.

THEOREM 1. *For a general convex price-adjustment cost, if a bi-directional price change is allowed, then*

(a) *$y_n(p_{n-1})$ is nonincreasing and $p_n(p_{n-1})$ is nondecreasing in p_{n-1}, and $\hat{P}_n(p_{n-1}, x)$ is nonincreasing in x but nondecreasing p_{n-1}. Furthermore, $\partial \hat{P}_n(p_{n-1}, x)/\partial p_{n-1} \leqslant 1$.*

(b) *The optimal inventory policy $y_n^* = \min\{y_n(p_{n-1}), x\}$, and the optimal pricing policy $p_n^* = \min\{\hat{P}_n(p_{n-1}, x), p_n(p_{n-1})\}$.*

PROOF. From the definition of $J_n(y, \tilde{p})$ in the proof of Lemma 2, $-U_n(\tilde{p}_{n-1} - \tilde{p}) + J_n(y, \tilde{p})$ is supermodular in $(\tilde{p}, y, \tilde{p}_{n-1})$. Thus, as

$$(-p_n(\tilde{p}_{n-1}), y_n(\tilde{p}_{n-1})) = \max\left\{ \arg\max_{-\underline{p} \geqslant \tilde{p} \geqslant -\bar{p},\, y} \{-U_n(\tilde{p}_{n-1} - \tilde{p}) + J_n(y, \tilde{p})\} \right\},$$

the first half of part (a) follows from Theorem 2.8.1 in Topkis (1998), which concerns the monotonicity of the maximizer of a supermodular function. Applying a similar idea can show $\hat{P}_n(p_{n-1}, x)$ is nonincreasing in x but nondecreasing p_{n-1}.

To prove the second half of part (a), observe that the function $-U_n(\hat{p}) + L_n(x, p_{n-1} - \hat{p})$ is supermodular in $(\hat{p}, p_{n-1})$ for a given x from the concavity of $L_n(x, \cdot)$ and is supermodular in $(\hat{p}, x)$ for a given p_{n-1} from Lemma 2. Thus, $p_{n-1} - \hat{P}_n(p_{n-1}, x)$ is increasing in p_{n-1}, which implies that $\partial \hat{P}_n(p_{n-1}, x)/\partial p_{n-1} \leqslant 1$.

From the joint concavity of function $-U_n(p_{n-1} - p) + L_n(y, p)$, we can conclude that the optimal solution to problem (2) is given by $(p_n(p_{n-1}), y_n(p_{n-1}))$ if $x \leqslant y_n(p_{n-1})$ and $(x, \hat{P}_n(p_{n-1}, x))$ if $x > y_n(p_{n-1})$. □

The foregoing theorem implies that when the starting inventory level x is less than $y_n(p_{n-1})$, the firm should order up to $y_n(p_{n-1})$ and set the price at $p_n(p_{n-1})$; otherwise, it should not order and should sell the product at $\hat{P}_n(p_{n-1}, x)$. Note that when the variable price-adjustment cost takes a general convex form, the optimal policy is fully state-dependent. In particular, when the starting inventory level is higher than $y_n(p_{n-1})$, the optimal price that the firm should set depends on both state variables. Thus, the structural result in this theorem is only a partial characterization of the optimal policy.

3.2. Piecewise Linear Price-Adjustment Costs

When the variable price-adjustment cost is piecewise linear, i.e., $U_n(p-p_{n-1}) = u_n|p-p_{n-1}|$, we are able to derive sharper results and characterize the optimal policy in greater detail. To characterize the optimal policy with this form of price adjustment cost, we define several functions and parameters as follows. Let

$$g_n(\xi, x, p) = -\xi p + L_n(x, p),$$

where ξ is a real number. Define

$$P_n(\xi, x) = \arg\max_{\underline{p}\leqslant p\leqslant \bar{p}}\{g_n(\xi, x, p)\},^{2}$$

$$Y_n(\xi) = \max\left\{\arg\max_x\{g_n(\xi, x, P_n(\xi, x))\}\right\},$$

and

$$y_n(p) = \max\left\{\arg\max_y\{L_n(y, p)\}\right\}.$$

Note that if $p \geqslant p_{n-1}$ and $\xi = -u_n$ or if $p < p_{n-1}$ and $\xi = u_n$, then $\xi p_{n-1} + g_n(\xi, x, p)$ is the bracketed function in (2).

Let $p_n^-(x) = P_n(u_n, x)$, $p_n^+(x) = P_n(-u_n, x)$, $Y_n^+ = Y_n(u_n)$, $Y_n^- = Y_n(-u_n)$, $P_n^- = p_n^-(Y_n^+)$, and $P_n^+ = p_n^+(Y_n^-)$. It is straightforward to verify that the foregoing functions and quantities are well defined. Also note that $L_n(x, p)$ is jointly concave due to Lemma 1.

Lemma 3. (a) *$P_n(\xi, x)$ is nonincreasing in ξ and x. Thus, $p_n^-(x)$ and $p_n^+(x)$ are nonincreasing in x and $p_n^-(x) \leqslant p_n^+(x)$ for any x.*

(b) *$Y_n^+ \geqslant Y_n^-$, $P_n^- \leqslant P_n^+$.*

(c) *$y_n(p)$ is nonincreasing in p, $y_n(P_n^-) = Y_n^+$ and $y_n(P_n^+) = Y_n^-$.*

(d) *$p \leqslant p_n^+(y_n(p))$ if $p \leqslant P_n^+$ and $p \geqslant p_n^+(y_n(p))$ if $p \geqslant P_n^+$.*

(e) *$p \leqslant p_n^-(y_n(p))$ if $p \leqslant P_n^-$ and $p \geqslant p_n^-(y_n(p))$ if $p \geqslant P_n^-$.*

Proof. Define for any ξ, x, and $\tilde{p}$,

$$f(\xi, x, \tilde{p}) = \xi\tilde{p} + J_n(x, \tilde{p}),$$

where $J_n(x, \tilde{p})$ is given in (6). Because $\xi\tilde{p}$ is supermodular in ξ and $\tilde{p}$ and $J_n(x, \tilde{p})$ is supermodular in x and $\tilde{p}$ from Lemma 2, $f(\xi, x, \tilde{p})$ is supermodular in ξ, x, and $\tilde{p}$.

Observe that

$$-P_n(\xi, x) = \arg\max_{-\underline{p}\geqslant\tilde{p}\geqslant-\bar{p}}\{f(\xi, x, \tilde{p})\},$$

$$(Y_n(\xi), -P_n(\xi, Y_n(\xi))) = \max\left\{\arg\max_{x,\,-\underline{p}\geqslant\tilde{p}\geqslant-\bar{p}}\{f(\xi, x, \tilde{p})\}\right\}.^{3}$$

From the monotonicity of a maximizer of a supermodular function (Theorem 2.8.1 of Topkis 1998, or Theorem 2.3.7 of Simchi-Levi et al. 2005), $P_n(\xi, x)$ is nonincreasing in ξ and x. As a result, $p_n^-(x)$ and $p_n^+(x)$ are nonincreasing in x and $p_n^-(x) \leqslant p_n^+(x)$ for any x. In addition, $(Y_n(\xi), -P_n(\xi, Y_n(\xi)))$ is nondecreasing in ξ, and hence $Y_n^+ \geqslant Y_n^-$ and $P_n^- \leqslant P_n^+$. Thus, parts (a) and (b) hold.

We now prove part (c). Observe that

$$y_n(p) = \max\left\{\arg\max_y f(\xi, y, -p)\right\}$$

for any ξ. Hence, $y_n(p)$ is nonincreasing in p (again from Theorem 2.3.7 of Simchi-Levi et al. 2004). Because $y_n(p)$ is independent of ξ, we have $y_n(P_n^-) = Y_n^+$ and $y_n(P_n^+) = Y_n^-$. Thus, part (c) holds.

Now consider part (d). Let ξ be any element in the supergradient set of the concave function $L_n(y_n(p), p)$. We immediately have that $(y_n(p), p) \in \{\arg\max_{x,\,\underline{p}\leqslant p\leqslant\bar{p}} g_n(\xi, x, p)\}$. Because $L_n(y_n(p), p)$ is concave, if $p \leqslant P_n^+$, then we have $\xi \geqslant -u_n$; otherwise (i.e., if $p > P_n^+$), we have $\xi \leqslant -u_n$. In either case,

$$-p = \arg\max_{-\underline{p}\geqslant\tilde{p}\geqslant-\bar{p}} f(\xi, y_n(p), \tilde{p}).$$

Note that

$$-p_n^+(y_n(p)) = \arg\max_{-\underline{p}\geqslant\tilde{p}\geqslant-\bar{p}} f(-u_n, y_n(p), \tilde{p}).$$

Thus, from the supermodularity of $f(\xi, x, \tilde{p})$ in ξ and $\tilde{p}$, $p \leqslant p_n^+(y_n(p))$ for $p \leqslant P_n^+$ and $p \geqslant p_n^+(y_n(p))$ for $p \geqslant P_n^+$. Hence, part (d) holds.

Finally, part (e) can be proved similarly to part (d). □

A sketch of the functions $(y_n(p),\ p_n^+(x),\ p_n^-(x))$ and parameters $(Y_n^+, Y_n^-, P_n^+, P_n^-)$ is provided in Figure 2. The properties of these are important for characterizing the structure of the optimal policy, which we shall see in the following. We start with the optimal policies for the cases in which only single-directional price adjustments are allowed throughout the horizon. Note that in such cases, the corresponding dynamic recursion (2) is derived under the constraint $p_{n-1} \leqslant p \leqslant \bar{p}$ in the markup-only case while under $\underline{p} \leqslant p \leqslant p_{n-1}$ in the markdown-only case for each period n.

Theorem 2. *Consider a system with linear price-adjustment costs that allows only a single-directional price adjustment throughout the horizon. Given the starting inventory level x and the price of the previous period p_{n-1}:*

(a) *If only markup is allowed, the optimal inventory and pricing policy (y_n^1, p_n^1) belongs to $\mathfrak{P}$. Specifically, it follows a base-stock inventory and one-sided threshold pricing policy with parameters $\bar{y}(p_{n-1}) = \min\{Y_n^+, y_n(p_{n-1})\}$ and $S(x) = s(x) = \min\{P_n^-, p_n(x)\}$;*

(b) *If only markdown is allowed, the optimal inventory and pricing policy (y_n^2, p_n^2) belongs to $\mathfrak{P}$. Specifically, it follows a base-stock inventory and one-sided threshold pricing policy with parameters $\bar{y}(p_{n-1}) = \max\{Y_n^-, y_n(p_{n-1})\}$ and $Z(x) = z(x) = \min\{P_n^+, p_n^+(x)\}$.*

Figure 1. The optimal inventory and pricing policies for one-directional price changes.

(a) The optimal inventory and pricing policy with markup
(b) The optimal inventory and pricing policy with markdown

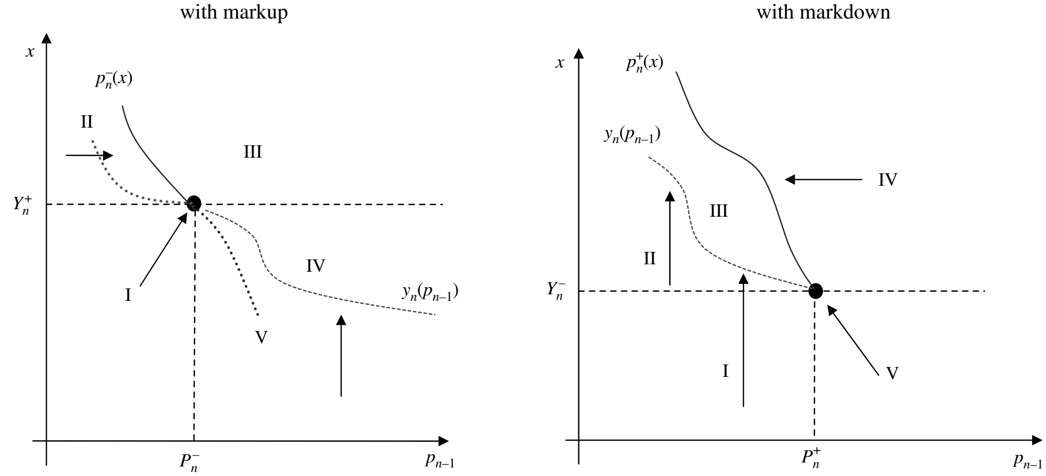

The optimal policy is illustrated in Figure 1. When only a markup is allowed (Figure 1(a)), if the starting state (x, p_{n-1}) of period n is in region I, then it is optimal to order up to Y_n^+ and to mark the price up to P_n^-. If the state is in region II, then it is optimal to order nothing and increase the price to $p_n^-(x)$. If the state is in regions III and IV, then it is optimal to do nothing. Finally, if the state lies in region V, then the price is kept unchanged, and it is optimal to order up to $y_n(p_{n-1})$. It is noteworthy that $p_n^-(x)$ and $y_n(p_{n-1})$ have the illustrated relationship and cross only once at (Y_n^+, P_n^-) because $y_n(P_n^-) = Y_n^+$, $p_n^-(Y_n^+) = P_n^-$ and because of part (e) of Lemma 3. When only a markdown is allowed (Figure 1(b)), if the initial state (x, p_{n-1}) lies in regions I and II, then the price should be kept unchanged, and it is optimal to order up to $y_n(p_{n-1})$. If the state is in region III, then nothing should be done. If the state is in region IV, then the firm should order nothing and mark the price down to $p_n^+(x)$. Finally, if the state is in region V, then it should order up to Y_n^- and mark the price down to P_n^+. Again, the visualized relationship between $p_n^+(x)$ and $y_n(p_{n-1})$ follows from $y_n(P_n^+) = Y_n^-$, $p_n^+(Y_n^-) = P_n^+$, and from part (d) of Lemma 3.

In general, when a bi-directional price change is allowed, with the properties presented in Lemma 3, the optimal policy is characterized in the following theorem. Its proof and the proof of Theorem 2 are provided in Appendix A in the electronic companion, which is available as part of the online version that can be found at http://or.journal.informs.org.

Theorem 3. *For a system with linear price-adjustment costs, given the starting inventory level x and the price of the previous period p_{n-1}, the optimal inventory and the pricing policy (y_n^*, p_n^*) of period n belongs to $\mathfrak{P}$. Specifically, it follows a base-stock inventory and two-sided threshold pricing policy with parameters* $\bar{y}(p_{n-1}) = \max\{\min(y_n(p_{n-1}), Y_n^+), Y_n^-\}$, $s(x) = S(x) = \min\{p_n^-(x), P_n^-\}$ *and* $z(x) = Z(x) = \min\{p_n^+(x), P_n^+\}$.

The optimal policy for the case with bi-directional price changes is illustrated in Figure 2. If the starting state $(x, p_{n-1}) \in \mathrm{I}$, then it is optimal to order up to Y_n^+ and to mark the price up to P_n^-. If the starting state $(x, p_{n-1}) \in \mathrm{II}$, then it is optimal to order nothing and to mark the price up to $p_n^-(x)$. If the starting state $(x, p_{n-1}) \in \mathrm{III} \cup \mathrm{V}$, then it is optimal to order nothing and to fix the price at p_{n-1}. If the starting state $(x, p_{n-1}) \in \mathrm{IV} \cup \mathrm{VI}$, then it is optimal to order nothing and to mark the price down to $p_n^+(x)$. If the starting state $(x, p_{n-1}) \in \mathrm{VIII} \cup \mathrm{IX}$, then it is optimal to order up to $y_n(p_{n-1})$ and to fix the price at p_{n-1}. Finally, if the starting state $(x, p_{n-1}) \in \mathrm{VII}$, then it is optimal to order up to Y_n^- and to mark the price down to P_n^+.

The optimal policy provided in Theorem 3 implies that the optimal price p_n^* is nondecreasing in p_{n-1} and nonincreasing in x (it also follows from the structural result for the case with a general convex price-adjustment cost). It should be pointed out that our structural results for the optimal policy extend those of Federgruen and Heching (1999), who illustrate in a corresponding integrated inventory and

Figure 2. The optimal inventory and pricing policy for bi-directional price changes.

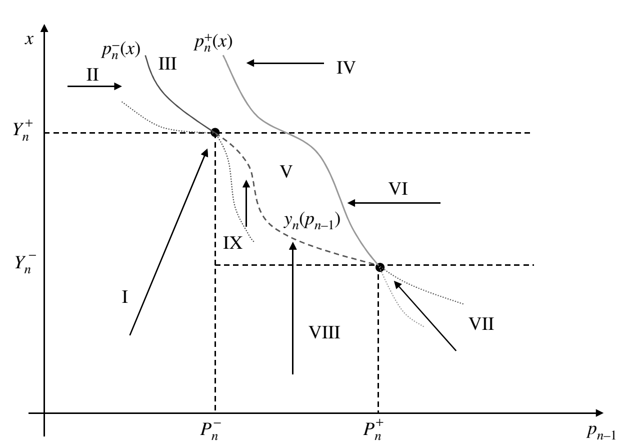

pricing model with costless price adjustments that a base-stock list price policy is optimal.

We end this section with the next result, which shows how the optimal policy parameters change with the unit price-adjustment cost u_n.

Proposition 1. *If the variable price adjustment cost u_n increases, then Y_n^-, P_n^-, and $p_n^-(x)$ decrease and Y_n^+, P_n^+, and $p_n^+(x)$ increase. However, $y_n(p_{n-1})$ is independent of u_n.*

Proof. The result follows directly from Lemma 3. □

We provide some intuition about the last result. Observe that for a given initial inventory level x in period n and selling price p_{n-1} in period $n-1$, the magnitude of the price change (if any) decreases if the variable price-adjustment cost u_n increases. Thus, the price raise-up-to levels P_n^- and $p_n^-(x)$ decrease, and the price decrease-down-to levels P_n^+ and $p_n^+(x)$ increase. At the same time, a lower price P_n^- implies greater demand and thus a higher order-up-to level Y_n^-, whereas a higher price P_n^+ implies less demand and thus a lower order-up-to level Y_n^+. It is clear that the optimal policy parameters after period n are independent of u_n. However, it is not clear how the optimal policy parameters before period n depend on u_n.

4. Model with No Inventory Carryover

In this section, we consider the second model in which no inventory can be carried over from one period to the next. In each period n, we assume that any unsatisfied demand must be fulfilled by an expedited order, which incurs an emergency ordering cost of b_n per unit. An alternative interpretation of emergency orders is that any unsatisfied demand is filled by an order in the following period, and a penalty is charged. It is worth noting that similar assumptions are made in previous studies of inventory models (e.g., Eeckhoudt et al. 1995, Simchi-Levi et al. 2005, Chen 2009). The unsold units at the end of each period n are salvaged with a per-unit value of r_n (if it is negatively valued, then the salvage value represents the disposal cost). The assumption that the unsold units of the product cannot be held in inventory for the next period, but instead yield a certain salvage revenue, is not as restrictive as it appears in that this assumption is appropriate to model settings that involve perishable products with a short shelf life.

To avoid trivialities, we assume that $b_n > c_n > r_n$. Although we allow the fixed price-change cost to be greater than zero (i.e., $K > 0$ in (1)) in the current model, for tractability, we require that $U_n(\Delta) = u_n|\Delta|$. In this setting, we are able to characterize the optimal pricing and ordering policy with the presence of both fixed and variable price-adjustment costs.

As we assume that all demands need to be satisfied and that the remaining inventory is salvaged at the end of each period, the starting inventory of every period is zero. Thus, it suffices to use the selling price of the last period as the sole state variable. The replenishment decision is thus simplified to the ordering quantity Q_n of each period n.

With a slight abuse of notation, we still let $v_n(p_{n-1})$ denote the total discounted optimal profit from periods n to N. The dynamic programming formulation of the problem is

$$v_n(p_{n-1}) = \max_{Q\geqslant 0,\, \underline{p}\leqslant p\leqslant \bar{p}} \{-C_n^p(p-p_{n-1}) + R_n(p) + G_n(Q,p) + \gamma v_{n+1}(p)\}, \tag{7}$$

where

$$G_n(Q,p) = -c_nQ - b_n\mathsf{E}[(D_n(p_n,\epsilon_n)-Q)^+] + r_n\mathsf{E}[(Q-D_n(p_n,\epsilon_n))^+]$$

models the ordering and inventory-related costs. Again, due to the linearity of the demand function, it is not difficult to see that $G_n(Q,p)$ is jointly concave in Q and p. Hence, we can first solve the optimal ordering quantity given the price p, $Q_n^*(p)$, which is the solution to

$$(b_n - c_n) - (b_n - r_n)\Pr(Q \geqslant D_n(p,\epsilon_n)) = 0. \tag{8}$$

The following result illustrates that the optimal ordering quantity in period n is nonincreasing in the selling price that is set for the period. Its proof is established by showing the submodularity of $G_n(Q,p)$, which follows a routine argument and thus is omitted here.

Lemma 4. *$Q_n^*(p)$ is nonincreasing in p.*

Substituting $Q_n^*(p)$ into (7), it is then simplified to

$$v_n(p_{n-1}) = \max_{\underline{p}\leqslant p\leqslant \bar{p}} \{-C_n^p(p-p_{n-1}) + L_n(p)\}, \tag{9}$$

where

$$L_n(p) = G_n(p) + R_n(p) + \gamma v_{n+1}(p),$$

and $G_n(p) := G_n(Q_n^*(p),p)$ is concave in p. Note that after optimizing the ordering quantity Q, which is myopic, the problem is reduced to one with a single decision, and the foregoing dynamic program recursion is almost identical to that which corresponds to another classical inventory model—the stochastic cash balance problem with symmetric ordering and return costs (see Appendix B for more details). Thus, the approach proposed in Chen (2003) and Chen and Simchi-Levi (2009) for analysis of the stochastic cash balance problem can be slightly modified to characterize the optimal pricing policy of the model developed here. In addition, the structure of our model's optimal pricing policy is essentially identical to that of the optimal policy of the stochastic cash balance problem.

The analysis in Chen (2003) and Chen and Simchi-Levi (2011) is built upon the following definition of symmetric K-concavity.

Definition 2. A real-valued function $f\colon [a, b] \to \Re$ is called symmetric K-concave for $K \geqslant 0$ if for any $x_0, x_1 \in [a, b]$ with $\lambda \in [0, 1]$,

$$f((1-\lambda)x_0 + \lambda x_1) \geqslant (1-\lambda) f(x_0) + \lambda f(x_1) - \max(\lambda, (1-\lambda))K. \quad (10)$$

(The definition in Chen and Simchi-Levi 2004a assumes that $a = -\infty$ and $b = \infty$. However, it is straightforward to extend the definition by restricting the function in a bounded interval $[a, b]$.) The concept of symmetric K-concavity is a generalization of Scarf's (1960) classical concept of K-concavity and is introduced by Chen and Simchi-Levi (2004a) when they analyze an integrated inventory and pricing model with fixed ordering cost and multiplicative demand. Interestingly, this concept has applications in the stochastic cash balance problem. We list the properties that are useful to our analysis in Appendix C.

Define

$$P_n^- = \min\Big\{ \operatorname*{argmax}_{p \in [\underline{p}, \bar{p}]} \{-u_n p + L_n(p)\} \Big\},$$

$$P_n^+ = \min\Big\{ \operatorname*{argmax}_{p \in [\underline{p}, \bar{p}]} \{u_n p + L_n(p)\} \Big\},$$

$$X_n^- = \{p \in [\underline{p}, P_n^-] \mid -u_n p + L_n(p) \geqslant -K - u_n P_n^- + L_n(P_n^-)\},$$

$$X_n^+ = \{p \in [P_n^+, \bar{p}] \mid u_n p + L_n(p) \geqslant -K + u_n P_n^+ + L_n(P_n^+)\}.$$

Let $s_n = \inf X_n^-$ if X_n^- is nonempty, and otherwise $s_n = \underline{p}$; $z_n = \inf X_n^+$ if X_n^+ is nonempty, and otherwise $z_n = \bar{p}$. It can easily be shown that $s_n \leqslant P_n^- \leqslant P_n^+ \leqslant z_n$ from their definitions.

If only a markup or markdown is allowed throughout the planning horizon, then the dynamic program recursion is almost identical to that which corresponds to the classical stochastic inventory model analyzed in Scarf (1960). In this case, we can apply the concept of K-concavity to prove that a one-sided (s, S) pricing policy is optimal. It should be pointed out that the policy parameters in the following result have different values than those in the case with bi-directional price changes, even though they share a common notation.

Proposition 2. (a) *If only markups are allowed in the entire planning horizon, then for period $n = 1, \ldots, N$, the $v_n(p)$ function is K-concave in $[\underline{p}, \bar{p}]$, and the optimal policy belongs to $\mathfrak{P}$. Specifically, the optimal inventory policy is base-stock with parameter $\bar{y}(p_{n-1}) = Q_n^*(p_n^*(p_{n-1}))$, whereas the optimal price p_n^* is a one-sided (s, S) pricing policy with $s(x) = s_n$ and $S(x) = P_n^-$.*

(b) *If only markdowns are allowed in the entire planning horizon, then for period $n = 1, \ldots, N$, the $w_n(\tilde{p}) = v_n(-\tilde{p})$ function is K-concave in $[-\bar{p}, -\underline{p}]$ and the optimal policy belongs to $\mathfrak{P}$. Specifically, the optimal inventory policy is base-stock with parameter $\bar{y}(p_{n-1}) = Q_n^*(p_n^*(p_{n-1}))$, whereas the optimal price p_n^* is a one-sided (s, S) pricing policy with $z(x) = z_n$ and $Z(x) = P_n^+$.*

Theorem 4. (a) *$v_n(p)$ and $L_n(p)$ are symmetric K-concave*; *and*

(b) *the optimal inventory policy is base-stock with parameter $\bar{y}(p_{n-1}) = Q_n^*(p_n^*(p_{n-1}))$, and the optimal price p_n^* is determined by*

$$p_n^* = \begin{cases} P_n^-, & \text{if } p_{n-1} \leqslant s_n, \\ \in \{p_{n-1}, P_n^-\} & \text{if } p_{n-1} \in (s_n, (s_n + P_n^-)/2), \\ p_{n-1} & \text{if } p_{n-1} \in ((s_n + P_n^-)/2, (P_n^+ + z_n)/2), \\ \in \{p_{n-1}, P_n^+\}, & \text{if } p_{n-1} \in ((P_n^+ + z_n)/2, z_n), \\ P_n^+ & \text{if } p_{n-1} \geqslant z_n. \end{cases} \quad (11)$$

Proof. Because our dynamic program recursion is almost identical to that for the stochastic cash balance problem with symmetric ordering and returning costs, we only sketch the proof of the result here; for a detailed proof, interested readers are referred to the corresponding recursion in Chen and Simchi-Levi (2009).

The theorem is proved by induction on n. By assuming $v_{n+1}(p)$ is symmetric K-concave, we can show that $L_n(p)$ and $v_n(p)$ are symmetric K-concave by Lemma 5 part (b) and Lemma 6 in Appendix B, respectively. The structure of the optimal policy at period n comes from the properties of the symmetric K-concavity of $L_n(p)$ and the concavity of $G_n(p)$ and $R_n(p)$ by using Lemma 5 part (d) for the markup direction and a symmetric argument for the markdown direction. □

The optimal pricing policy is visualized in Figure 3, in which the dashed lines specify different regions. When there is no inventory carryover, if p_{n-1} is in region I, then the price should always be marked up to P_n^-; if p_{n-1} is in region II and it is optimal to change the price, then it should be marked up to P_n^-; otherwise, it should be kept unchanged; if p_{n-1} is in region III, then the price should not be changed; if p_{n-1} is in region IV and it is optimal to change the price, then it should be marked down to P_n^+; otherwise, it should not be changed; and finally, if p_{n-1} is in region V, the price should be marked down to P_n^+. Note that the pricing policy is more complicated than the two-sided (s, S) policy defined in Definition 1 and can be regarded as a two-sided (s, S, A) policy (see more discussion on (s, S, A) policy in Chen and Simchi-Levi 2004), which includes regions $(s_n, (s_n + P_n^-)/2)$ and $((P_n^+ + z_n)/2, z_n)$ in which the optimal policy is ambiguous.

Sheshinski and Weiss (1977) also show the (s, S)-type pricing policy to be optimal for an infinite-horizon, deterministic, continuous-review problem with inflation and a

Figure 3. Optimal pricing policy with no inventory carryover.

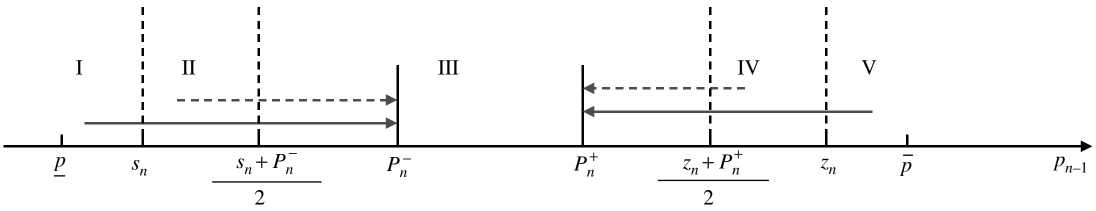

fixed price-adjustment cost. Nevertheless, their model setting is different from ours, and the main driver of the price change is inflation, i.e., inflation causes the selling price to drift continuously from the initial price level S to the level s at which point the price is marked up to S and this cycle then repeats itself. Nevertheless, it can be seen that the key driver of price changes in our setting is the nonstationarity of the system parameters and the market demand.

Remark. It is noteworthy that the assumption of stationary fixed costs can be relaxed. In particular, the preceding results hold even for time-dependent fixed costs, denoted by K_n, as long as $\gamma^{n-1}K_n$ is nonincreasing in n.

5. Lost Sales Model

In this section, we consider a lost sales model, in which unsatisfied demand in each period is lost. Different from the backlog case, we do not assume any specific form of demand function $D_n(p, \epsilon_n)$ here but only require it to be decreasing concave and twice differentiable in p and strictly increasing in ϵ_n (note that ϵ_n is a scalar in this case). For example, $D_n(p, \epsilon_n) = D_n(p) + \epsilon_n$ or $D_n(p, \epsilon_n) = D_n(p)\epsilon_n$, with $D_n(p)$ being decreasing concave. We start by discussing the case of inventory carryover.

The resulting dynamic program is as follows.

$$v_n(x, p_{n-1}) = \max_{y \geqslant x,\, \underline{p} \geqslant p \geqslant \bar{p}} \{-U_n(p_{n-1} - p) + L_n(y, p)\}, \quad (12)$$

with $v_{N+1}(\cdot, \cdot) = 0$, and

$$L_n(y,p) = (p - \gamma c_{n+1} + h_n)\mathsf{E}[\min\{D_n(p, \epsilon_n), y\}] + (\gamma c_{n+1} - h_n - c_n)y + \gamma \mathsf{E}[v_{n+1}((y - D_n(p, \epsilon_n))^+, p)]. \quad (13)$$

As in the backlogging model, we consider only variable price-adjustment costs here, and we do not consider any additional penalty costs of lost sales because it would become very difficult to find the conditions under which the value functions are concave. This setting is also adopted in other inventory-pricing models with lost sales (e.g., Song et al. 2008). We also require that $\underline{p} \geqslant \gamma c_{n+1} - h_n$, which is not too restrictive, because otherwise for those p such that $p < \gamma c_{n+1} - h_n$, the first term in $J_n(y, \tilde{p})$ (the expected net revenue) is negative.

It should also be noted that, different from the backlog case, the expected revenue in the lost sales case equals the expectation of the product of the selling price and the minimum of demand and the inventory level. This function is, in general, not jointly concave. Moreover, the requirement that the initial inventory at the beginning of any period must be nonnegative also makes the analysis more challenging. However, we are able to provide conditions under which the results of the backlog case can be extended to the lost sales case. The proofs of the results in this section are given in Appendix D.

Define

$$e_n(x, p) = -(p - \gamma c_{n+1} + h_n) \frac{[\Pr(D_n(p, \epsilon_n) > x)]'_p}{\Pr(D_n(p, \epsilon_n) > x)}. \quad (14)$$

This definition is similar to the elasticity of lost sales defined in Kocabiyikoglu and Popescu (2011) with the difference of an additional term $-\gamma c_{n+1} + h_n$ in the coefficient because we consider a multi-period setting. In the next two results, we provide conditions based on $e_n(x, p)$, under which the optimal policies we characterized for the backlog case continue to hold here.

Theorem 5. *If $e_n(x, p) \geqslant 1/2$ for all n, then*

(a) *$L_n(y, p)$ is jointly concave in y and p, $v_n(x, p_{n-1})$ is nonincreasing in x and jointly concave in x and p_{n-1}; and*

(b) *the optimal strategy (y_n^*, p_n^*) has the same structure as those presented in Theorem* 1.

For a single-period problem, Kocabiyikoglu and Popescu (2011) show that, under a similar condition (as noted, $e_n(x, p)$ here is defined slightly differently), the expected revenue $p\mathsf{E}[\min\{x, D_n(p, \epsilon_n)\}]$ is jointly concave. It happens that this condition leads to the joint concavity of our multi-period problem. Here, we give one example that satisfies the condition. If $D_n(p, \epsilon_n) = a_n - b_n p + \epsilon_n$ where ϵ_n is exponentially distributed with cdf $1 - e^{-\lambda x}$, $\lambda > 0$, then $e_n(x, p) \geqslant 1/2$ holds for all feasible x and p as long as $\underline{p} \geqslant 1/(2\lambda b_n) + \gamma c_{n+1} - h_n$. We would also like to point out that, other papers, such as Chen et al. (2005) and Song et al. (2008), provide conditions under which the single period profit, after optimizing the ordering quantity, is quasi-concave in the price. However, their approaches cannot be applied in our setting.

Under a stronger condition that $e_n(x, p) \geqslant 1$, we can further show that $L_n(y, p)$ is submodular, and hence the optimal policies when $U_n(\Delta)$ is piecewise linear have the same structure as those presented in the backlog case.

THEOREM 6. *If $e_n(x, p) \geqslant 1$ for all n and $U_n(\Delta) = u_n|\Delta|$, then*

(a) *$L_n(y, p)$ is jointly concave and submodular in y and p, and $v_n(x, p_{n-1})$ is jointly concave and submodular in x and p_{n-1}; and*

(b) *the optimal strategy (y_n^*, p_n^*) (markup only, markdown only bi-directional) has the same structure as those presented in Theorems* 2 *and* 3.

An example that satisfies $e_n(x, p) \geqslant 1$ can also be found similarly, as we did for the previous theorem.

We next consider a case in which there is no inventory carryover. Redefine

$$G_n(Q, p) = -c_n Q + r_n \mathsf{E}[(Q - D_n(p, \epsilon_n))^+],$$

where r_n still denotes the unit salvage value and $r_n \leqslant \underline{p}$. It is clear that the results derived in §4 continue to hold if

$$\bar{g}_n(p) = \max_Q \{p\mathsf{E}[\min\{D_n(p, \epsilon_n), Q\}] + G_n(Q, p)\} \tag{15}$$

is concave in p. Let $Q_n^*(p)$ be the optimal ordering quantity with a given price p. The next theorem provides the condition that (15) is concave.

THEOREM 7. *For the case in which no inventory is carried over between periods, replace $-\gamma c_{n+1} + h_n$ in* (14) *with $-r_n$. If the resulting $e_n(Q_n^*(p), p) \geqslant 1/2$ for all n, then*

(a) *$\bar{g}_n(p)$ is concave in p; and*

(b) *the optimal pricing policy for the lost sales case has the same structure as that presented in Theorem* 4.

Part (a) in Theorem 7 follows from the condition $e_n(Q_n^*(p), p) \geqslant 1/2$, which can be proved by similar arguments to those in the proof of Proposition 6 in Kocabiyikoglu and Popescu (2011), and thus we skip it here. Part (b) follows directly from part (a). When the demand function takes the form $D_n(p, \epsilon_n) = A_n(p)\epsilon_n + B_n(p)$, the condition holds if ϵ_n has a distribution with an increasing failure rate (IFR) and $(p - r_n)A_n'(p)$ and $(p - r_n)B_n'(p)$ are decreasing in p. See Kocabiyikoglu and Popescu (2011) for further discussion.

6. The General Model: Heuristic and Numerical Study

Based on the analysis and results presented in the previous sections, we can anticipate that the optimal policy of the general problem with both inventory carryover and a fixed price-adjustment cost will be state-dependent and might not have a simple structure. Recall that even for the model without fixed price-adjustment costs, the optimal policy is state-dependent. Therefore, the optimal policy might be too complicated to implement. To provide a control policy that is relatively easy to implement for a general system, in this section we propose an intuitive heuristic inspired by the analysis in the two preceding special cases. We focus only on $U_n(\Delta) = u_n|\Delta|$ and the backlog case, although the heuristic is also applicable to the lost sales case.

Due to the existence of fixed price-adjustment costs and based on the analysis in §§3 and 4, it is natural to construct a heuristic policy from $\mathfrak{P}$, i.e., the heuristic follows a base-stock inventory and two-sided (s, S) pricing policy. To make the policy easier to implement, we set $s(x) = s_n^h$, $S(x) = P_n^{h-}$, $z(x) = z_n^h$, $Z(x) = P_n^{h+}$ with state-independent parameters s_n^h, P_n^{h-}, z_n^h, and P_n^{h+}, which we later show how to compute. The parameter for the base-stock policy is $\bar{y}(p_{n-1}) = y_n^h(p_n^h(p_{n-1}))$, where $p_n^h(p_{n-1})$ is the selling price after the price-adjustment decision. We search $y_n^h(p_n^h(p_{n-1}))$ of the value function in a compact domain. We name this policy *Heuristic CZC*.

Because we are dealing with a finite-horizon, nonstationary problem, the policy control parameters described above need to be computed for each period recursively, starting from period N. Specifically, after the control parameters of period $n+1$ are computed, they are applied to derive the corresponding value function of period $n+1$ and, in turn, the control parameters of period n. Because the value function resulting from the heuristic policy holds no nice structural properties, the computation of the control parameters relies on complete enumerations over the feasible sets. For the price p in period n, there is a given feasible set $[\underline{p}, \bar{p}]$, whereas for the inventory order-up-to level, we need to impose an upper bound based on the range of possible realizations of demand. The detailed steps of the heuristic are presented in the following algorithm.

Heuristic CZC for the General System

- *Step* 1. Set $n = N$ and define $v_{N+1}^h(x, p) = v_{N+1}(x, p) = c_{N+1}x$. Let $l_N(y, p) = R_N(p) + G_N(y, p) + \gamma\mathsf{E}[v_{N+1}^h(y - D_N(p, \epsilon_N), p)]$.
- *Step* 2. Define

$$(Y_n^{h+}, P_n^{h-}) \in \Big\{\arg\max_{y,\, p\in[\underline{p},\bar{p}]} \{-u_n p + l_n(y, p)\}\Big\},$$

with

$$s_n^h = \min\{p \in [\underline{p}, P_n^{h-}] \mid -u_n p + l_n(Y_n^{h+}, p) \geqslant -K - u_n P_n^{h-} + l_n(Y_n^{h+}, P_n^{h-})\} \quad (\text{if empty, } s_n^h = \underline{p});$$

and

$$(Y_n^{h-}, P_n^{h+}) \in \Big\{\arg\max_{y,\, p\in[\underline{p},\bar{p}]} \{u_n p + l_n(y, p)\}\Big\},$$

with

$$z_n^h = \max\{p \in [P_n^{h+}, \bar{p}] \mid u_n p + l_n(Y_n^{h-}, p) \geqslant -K + u_n P_n^{h+} + l_n(Y_n^{h-}, P_n^{h+})\} \quad (\text{if empty, } z_n^h = \bar{p});$$

and $y_n^h(p_{n-1})$ is the global maximizer of

$$y_n^h(p_{n-1}) \in \Big\{\arg\max_y \{l_n(y, p)\}\Big\}.$$

• *Step* 3. Set the price of period n according to a two-sided (s, S) pricing policy with parameters (s_n^h, P_n^{h-}) and (z_n^h, P_n^{h+}), and then replenish the inventory following a state-dependent base-stock policy. Specifically, if $p_n = P_n^{h-}$, then $y_n^h = \max\{x, Y_n^{h+}\}$ and

$$v_n^h(x, p_{n-1}) = -K - u_n(P_n^{h-} - p_{n-1}) + l_n(\max\{x, Y_n^{h+}\}, P_n^{h-}) + c_n x;$$

or else if $p_n = P_n^{h+}$, then $y_n^h = \max\{x, Y_n^{h-}\}$ and

$$v_n^h(x, p_{n-1}) = -K + u_n(p_{n-1} - P_n^{h+}) + l_n(\max\{x, Y_n^{h-}\}, P_n^{h+}) + c_n x;$$

otherwise, if $p_n = p_{n-1}$, then $y_n^h = \max\{x, y_n^h(p_{n-1})\}$ and

$$v_n^h(x, p_{n-1}) = l_n(\max\{x, y_n^h(p_{n-1})\}, p_{n-1}) + c_n x.$$

• *Step* 4. If $n > 1$, then $n = n - 1$ and let $l_n(y, p) = R_n(p) + G_n(y, p) + \gamma \mathsf{E}[v_{n+1}^h(y - D_n(p, \epsilon_n), p)]$; go to Step 2. Otherwise, stop.

It is not difficult to see from the definition that $Y_n^{h+} = y_n^h(P_n^-)$, $Y_n^{h-} = y_n^h(P_n^+)$, and $s_n^h \leqslant z_n^h$. Because we have to derive the recursion for value function $v_n^h(x, p_{n-1})$ in the algorithm, the control parameters of the heuristic might not be easy to compute. However, we expect the computational effort to be less than that for the exact optimal policy, and the heuristic policy is also amenable to practical implementation once the control parameters are on hand.

To test the effectiveness of the heuristic compared to the optimal solutions, we conduct an extensive numerical study. We first consider instances with a short planning horizon, $N = 4$. Assume $D_n(p, \epsilon) = 80 - 2p + \epsilon$, in which ϵ is a negative binomial random variable ($P(\epsilon = i) = C_{i+r-1}^r(\tau)^r(1-\tau)^i$) with $\tau = 0.5$ and $r = 8$, so the expected value $\mathsf{E}[\epsilon] = 8$. The other basic parameters are, for $n = 1, \ldots, N$: $b_n = 10$, $h_n = 2$, $u_n = 3n - 2$, $K = 30$, $\alpha = 0.9$, $c_n = 6 - n^{0.5}$, $\underline{p} = 3$, and $\bar{p} = 35$.

Different instances are generated by alternating one of the basic system parameters. For each instance, we consider the possible combination of the initial inventory level $x \in [-20, 50]$ and price $p_0 \in [3, 35]$ to avoid the impact of the initial state on the performance of the heuristic. That is to say, for each set of parameters there are a total of 2,343 instances. We define the relative error of the heuristic as

$$\text{Error\%} = \max_{x, p_0} \left\{ \frac{v_1(x, p_0) - v_1^h(x, p_0)}{v_1(x, p_0)} \right\} \times 100\%,$$

where $v_1^h(x, p_0)$ is the resulting total discounted profit of Heuristic CZC given state (x, p_0). The optimal profit $v_1(x, p_0)$ is computed recursively through an exhaustive search over the feasible domain of price and ordering quantity (in this numerical study, because the largest expected one-period demand is 82 and the largest variance is 80, we set an upper bound of 200 for the ordering quantity).

We generate four groups of examples. By keeping $\mathsf{E}[\epsilon] = 8$, we generate Group 1 instances by considering the parameter $r \in [72, 32, 12, 8, 2]$. The resulting variances $V[\epsilon]$ are 80, 40, 20, 16, and 10. Group 2 is generated by alternating the unit ordering cost c_2 in period 2 from 2 to 10 with a step size of 2. For Group 3, we change the backlog cost b from 20 to 80 with a step size of 20. For the last group, Group 4, we change the unit price-adjustment cost u_1 of period 1 from 5 to 20 with a step size of 5. To see the impact of the fixed price-adjustment cost on the performance of the heuristic, we further test the preceding four groups of instances with a larger fixed cost of $K = 100$.

The results are reported in Table 1. We can see that the heuristic performs very well in these examples. Compared to the results of $K = 30$, we find that the performance of the heuristic deteriorates slightly as K becomes larger. However, it still performs quite well. The average computation time of the heuristic for each instance (using Matlab) is 0.76 seconds of CPU time, in contrast to 810.6 seconds of CPU time of the optimal policy on a PC with a 2.66-GHz CPU.

To further demonstrate the effectiveness of the heuristic, we consider a very simple myopic policy whose parameters are calculated based on the single-period profit and apply them to Heuristic CZC. Although computationally easier, we can see from Table 2 that its performance is considerably worse after comparing with Table 1.

Finally, to determine how well the heuristic performs with a longer planning horizon, we test the previous four groups of instances with $N = 12$ and $K = 100$. The average and maximum errors for the four groups are (0.00%, 0.00%), (0.15%, 0.36%), (0.92%, 2.29%), and (0.40%, 0.56%), respectively. These results show that the resulting profit of the heuristic is still quite close to the optimal when the planning horizon is long.

We also conduct numerical studies to compare Heuristic CZC with several other simple ones, and the results are reported in Appendix E.

Table 1. Performance of heuristic CZC.

	$K = 30$				$K = 100$			
Error %	GP 1 (r)	GP 2 (c_2)	GP 3 (b)	GP 4 (u_1)	GP 1(r)	GP 2(c_2)	GP 3(b)	GP 4(u_1)
Avg	0.25	0.27	0.07	0.55	0.32	0.70	0.80	1.01
Max	0.30	0.87	0.19	0.85	0.33	1.17	1.28	1.44

Table 2. Performance of myopic heuristic.

Error %	GP 1 (r)	GP 2 (c_2)	GP 3 (b)	GP 4 (u_1)	GP 1(r)	GP 2(c_2)	GP 3(b)	GP 4(u_1)
	$K=30$				$K=100$			
Avg	0.72	1.04	1.50	1.20	0.92	1.23	3.80	1.69
Max	1.96	2.89	1.53	2.34	6.31	6.98	5.90	6.58

Inventory-Dependent Price-Adjustment Cost

In the retail industry, the physical costs of a price adjustment might include inventory ticketing/relabeling costs, and hence they might depend on the existing inventory level (Celik et al. 2009). Let c_n^I be the price-adjustment cost per unit of inventory. It should be stressed that this cost would be incurred only when there are price changes between periods, and it is thus of the fixed-cost type. In addition, such an inventory dependent price-adjustment cost should be incurred only when there is positive inventory; for backlogged units, the customers pay the price of the last period, and thus there is no need to change the price tags when they become available. Also note that the labeling cost for the "fresh" inventory, brought by the order that has just been placed, is accounted for in the variable ordering cost. Therefore, when the starting inventory level of period n is x, the total price-adjustment cost is expressed as

$$C_n^p(\Delta, x) = (K + c_n^I x^+)\delta(\Delta) + U_n(\Delta).$$

With this more general cost function, we can anticipate that the problem will be even more complicated than the models analyzed in the previous sections. To tackle this complexity, we provide the following two heuristic policies. The first simply applies the heuristic we developed for the case with $c_n^I = 0$. The second slightly modifies the preceding heuristic to accommodate the new cost feature. As the fixed cost now depends on the inventory level, we redefine the parameters $s(x)$ and $z(x)$ as $s(x) = s_n^h(x)$ and $z(x) = z_n^h(x)$, where

$$s_n^h(x) = \min\{p \in [\underline{p}, P_n^{h-}] \mid -u_n p + l_n(Y_n^{h+}, p) \geqslant -K - c^I x^+ \\ - u_n P_n^{h-} + l_n(Y_n^{h+}, P_n^{h-})\} \quad (\text{if empty, } s_n^h(x) = \underline{p});$$

and

$$z_n^h(x) = \max\{p \in [P_n^{h+}, \bar{p}] \mid u_n p + l_n(Y_n^{h-}, p) \geqslant -K - c^I x^+ \\ + u_n P_n^{h+} + l_n(Y_n^{h-}, P_n^{h+})\} \quad (\text{if empty, } z_n^h(x) = \bar{p}).$$

It is clear that the pricing policy of the second heuristic becomes a state-dependent two-sided (s, S) pricing policy. Inventory replenishment still follows a state-dependent base-stock policy. It becomes evident that this policy is more difficult to compute and implement than the first one, because the thresholds that trigger price changes are now functions of the starting inventory level.

We next test the performance of these two heuristics. With $c_n^I \in \{1, 5\}$ for all n, we test the same four groups of numerical examples studied previously with $N = 4$. The results are tabulated in Tables 3 and 4, in which, Avg_1 (Avg_2) and Max_1 (Max_2) are the average and maximum errors for the first (second) heuristic. It can be observed from the results that the second heuristic outperforms the first, although the first is easier to implement. When c_n^I is small, both heuristics are quite effective. However, with a larger c_n^I, the performance of the first heuristic deteriorates quite significantly. This is intuitive, because it ignores the inventory-dependent price-adjustment costs. We therefore should not ignore such costs when designing the heuristic.

7. Conclusion

We investigate a multi-period inventory model with costly price adjustments in this paper. A firm needs to make its pricing and inventory ordering decisions simultaneously in each period to maximize total profits. Due to the complexity of a general system with both a fixed cost and inventory carryover, we characterize the optimal inventory replenishment and pricing strategy for two special cases of the general model: one with inventory carryover between periods, but no fixed costs for price adjustments; and the other with fixed costs for price adjustments, but no inventory carryover. For the general problem, we provide an intuitive heuristic policy and show with numerical examples that it is quite effective.

The results can also be extended to the case of infinite horizon. Consider an infinite horizon system with

Table 3. Performance comparison of two heuristics: $c_n^I = 1$.

Error %	GP 1 (r)	GP 2 (c_2)	GP 3 (b)	GP 4 (u_1)	GP 1(r)	GP 2(c_2)	GP 3(b)	GP 4(u_1)
	$K=30$				$K=100$			
Avg_1	2.01	1.94	1.84	2.25	1.64	1.89	2.38	2.79
Max_1	2.05	2.57	1.96	2.56	1.69	2.88	3.08	3.25
Avg_2	0.86	0.82	0.76	1.41	1.48	1.26	1.47	1.99
Max_2	1.16	1.39	0.92	1.91	1.69	1.86	1.74	2.22

Table 4. Performance comparison of two heuristics: $c_n^I = 5$.

	$K = 30$				$K = 100$			
Error %	GP 1 (r)	GP 2 (c_2)	GP 3 (b)	GP 4 (u_1)	GP 1(r)	GP 2(c_2)	GP 3(b)	GP 4(u_1)
Avg_1	9.66	9.70	9.69	9.59	10.47	10.71	10.93	11.25
Max_1	9.70	9.77	9.70	9.79	10.51	11.54	11.48	11.61
Avg_2	3.45	3.82	4.92	4.36	3.75	3.91	5.36	4.57
Max_2	3.62	4.54	5.58	5.09	3.97	4.19	6.20	5.51

stationary parameters. With several technical assumptions (see Federgruen and Heching 1999, Chen and Simchi-Levi 2004b), the optimal policy for the finite horizon case can be extended and is stationary. First consider the first model with markup only. If the starting state is in region I of Figure 1(a), then the price will be set at P^- in period one and never change afterward, and the firm just needs to maintain an order-up-to level Y^+ at the beginning of each period; if the starting state is in region II, then after a finite number of periods, the state will fall into region I and remain there forever; if the starting state is in region II, IV, or V, then the firm will never change the price and will replenish inventory only if necessary. Similar scenarios would occur in the markdown-only and bi-directional price change cases. For a model without inventory carryover, it can be seen that the firm would change the price at most once.

There are several other possible extensions to our models.

First, a possible extension is to study problems with capacity constraints of inventory replenishment. In this case, we expect that the optimal inventory policy would become a so-called modified base-stock-type policy, i.e., if ordering, then the firm orders up to the optimal base-stock level (orders the optimal quantity) if possible; otherwise, it orders to full capacity. The optimal pricing policy is expected to be more complicated in this case as the policy parameters are more likely to depend on the inventory level due to the capacity constraint.

Second, in this paper, we consider mainly symmetric price-adjustment costs (see §3.2), i.e., the same variable price-adjustment costs for markups and markdowns. However, the results in §3.2 hold even when markups and markdowns incur different variable price-adjustment costs. When there are asymmetric fixed price-adjustment costs, the results in §4 can also be extended by employing the approach and results for the stochastic cash balance problem with asymmetric costs in Chen and Simchi-Levi (2009).

Third, we assume that demand in a period depends only on the price in the current period. However, the marketing literature argues that demand might also depend on historical prices, i.e., the prices in previous periods. It would be interesting to incorporate these more general demand models into our problem and investigate the corresponding optimal policies. Finally, incorporating a fixed ordering cost into our model would be interesting and important. Analysis of such a model is challenging and deserves further exploration.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Endnotes

1. Observe that the objective function $-U_n(\tilde{p}_{n-1} - \tilde{p}) + J_n(y, \tilde{p})$ in problem (5) is supermodular. Thus the maximum is well defined. See, for instance, Theorem 2.3.7 of Simchi-Levi et al. (2005).
2. Because the expected revenue function is strictly concave in p, the set of optimal solutions here is a singleton for a given x.
3. The maximum is well defined because $f(\xi, x, \tilde{p})$ is supermodular (see Theorem 2.3.7 of Simchi-Levi et al. 2005).

Acknowledgments

The authors are grateful for the comments and suggestions received from the associate editor and two anonymous referees. They also thank Xiuli Chao, Jian Yang, and Yuhan Zhang for their helpful comments. The first author is partly supported by National Science Foundation grants CMMI-0653909, CMMI-0926845 ARRA, and CMMI-1030923; the second author is partly supported by Hong Kong GRF Fund CUHK #419408; and the third author is partly supported by Hong Kong GRF Fund CUHK #410907.

References

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

Bergen, M., S. Dutta, D. Levy, M. Ritson, M. Zbaracki. 2003. Shattering the myth of costless price changes. *Eur. Management J.* **21**(6) 663–669.

Caplin, A., D. F. Spulber. 1987. Menu costs and the neutrality of money. *Quart. J. Econom.* **CII**(4) 703–725.

Celik, S., A. Muharremoglu, S. Savin. 2009. Revenue management with costly price adjustments. *Oper. Res.* **57**(5) 1206–1219.

Chan, L. M. A., Z. J. Max Shen, D. Simchi-Levi, J. Swann. 2004. Coordination of pricing and inventory. D. Simchi-Levi, S. D. Wu, Z. J. Max Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Chapter 3. Kluwer, Boston.

Chao, X., S. X. Zhou. 2006. Joint inventory and pricing strategy for a stochastic continuous-review system. *IIE Trans.* **38**(5) 401–408.

Chen, X. 2003. Coordinating inventory control and pricing strategies. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Chen, X. 2009. Inventory centralization games with price-dependent demand and quantity discount. *Oper. Res.* **57**(6) 1394–1406.

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 with random demand and fixed ordering cost: The continuous review model. *Oper. Res. Lett.* **34**(3) 323–332.

Chen, X., D. Simchi-Levi. 2009. A new approach for the stochastic cash balance problem with fixed costs. *Probab. Engrg. Inform. Sci.* **23**(4) 545–562.

Chen, X., D. Simchi-Levi. 2011. Pricing and inventory management. O. Özer, R. Phillips, eds. *The Handbook of Pricing Management*. Forthcoming.

Chen, Y., S. Ray, Y. Song. 2005. Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales. *Naval Res. Logist.* **53**(2) 117–136.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. Risk averse (and prudent) newsboy. *Management Sci.* **41**(5) 786–794.

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.

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

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

Kano, K. 2006. Menu costs, strategic interactions, and retail price movement. Working paper, Queen's University, Kingston, Ontario, Canada.

Kocabiyikoglu, A., I. Popescu. 2011. The newsvendor with pricing: A stochastic elasticity perspective. *Oper. Res.* **59**(2) 301–312.

Konieczny, J. D. 1993. Variable price adjustment costs. *Econom. Inquiry* **31**(3) 488–498.

Levy, D., M. Bergen, S. Dutta, R. Venable. 1997. The magnitude of menu costs: Direct evidence from large U.S. supermarket chains. *Quart. J. Econom.* **112**(3) 791–825.

McCallum, B. T. 1986. On real and sticky-price theories of the business cycle. *J. Money, Credit, Banking* **18**(2) 397–414.

McGill, G., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.

Netessine, S. 2006. Dynamic pricing of inventory/capacity with infrequent price changes. *Eur. J. Oper. Res.* **174**(1) 553–580.

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

Roberts, J. 1992. Evidence on price adjustment costs in U.S. manufacturing industry. *Econom. Inquiry* **30**(3) 399–417.

Rotemberg, J. 1982a. Monopolistic price adjustment and aggregate output. *Rev. Econom. Stud.* **49**(4) 517–531.

Rotemberg, J. 1982b. Sticky prices in the United States. *J. Political Econom.* **90**(6) 1187–1211.

Scarf, H. 1960. The optimality of (s, S) policies for the dynamic inventory problem. *Proc. 1st Stanford Sympos. Math. Methods Soc. Sci.*, Stanford University Press, Stanford, CA.

Sheshinski, E., Y. Weiss. 1977. Inflation and costs of price adjustment. *Rev. Econom. Stud.* **44**(2) 287–303.

Sheshinski, E., Y. Weiss. 1993. *Optimal Pricing, Inflation, and the Cost of Price Adjustment*. MIT Press, Cambridge, MA.

Simchi-Levi, D., X. Chen, J. Bramel. 2005. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*, 2nd ed. Springer-Verlag, New York.

Slade, M. E. 1998. Optimal pricing with costly adjustment: Evidence from retail-grocery prices. *Rev. Econom. Stud.* **65**(1) 87–107.

Song, Y., S. Ray, T. Boyaci. 2008. Optimal dynamic joint inventory-pricing control for multiplicative demand with fixed order costs and lost sales. *Oper. Res.* **57**(1) 245–250.

Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.

Tsiddon, D. 1991. On the stubbornness of sticky prices. *Internat. Econom. Rev.* **32**(1) 69–75.

Whitin, T. M. 1955. Inventory control and price theory. *Management Sci.* **2**(1) 61–68.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, M. Bergen. 2004. Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Rev. Econom. Statist.* **86**(2) 512–533.