Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost

Xin Chen
Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
xchen@mit.edu

Advisor: David Simchi-Levi
Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

1. Introduction

Traditional inventory models focus on effective replenishment strategies and typically assume that the commodity's price is exogenously determined. In recent years, however, a number of industries have used innovative pricing strategies to manage their inventory effectively. For example, techniques such as *revenue management* have been applied in the airlines, hotels, and rental car agencies—integrating price, inventory control, and quality of service, see Kimes (1989). In the retail industry, to name another example, dynamically pricing commodities can provide significant improvements in profitability, see Gallego and van Ryzin (1994).

These developments call for models that integrate inventory control and pricing strategies. These models are clearly important not only in the retail industry, where price-dependent demand plays an important role, but also in manufacturing environments in which production/distribution decisions can be complemented with pricing strategies to improve the firm's bottom line. In particular, the ability to dynamically adjust prices is critical for products facing seasonal demand and short product life cycle. In these situations, the so-called "dynamic pricing" strategies may have a huge impact on the performance of the firm.

Of course, many papers address the coordination of replenishment strategies and pricing policies, starting with the work of Whitin (1955) who analyzed the celebrated newsvendor problem with price-dependent demand. For a review, the reader is referred to Federgruen and Heching (1999).

To date, the literature has confined itself mainly to either (i) models with variable ordering costs but no fixed costs; (ii) models in which inventory cannot be carried over from one period to the next; or (iii) models in which replenishment decisions are made only at the beginning of the planning horizon, see Federgruen and Heching (1999). Recently, however, Chen and Simchi-Levi (2002a, b) analyzed a fairly general inventory/pricing model. Specifically, Chen and Simchi-Levi (2002a, b) consider a periodic review, single-product model with stochastic demand. Demands in different periods are independent of each other, and their distributions depend on the product price. Pricing and ordering decisions are made at the beginning of each period, and all shortages are backlogged. The ordering cost includes both a fixed cost and a variable cost proportional to the amount ordered. Inventory holding and shortage costs are convex functions of the inventory level carried over from one period to the next. The objective is to find an inventory policy and pricing strategy that maximizes expected profit over the entire planning horizon.

The model is similar to the model analyzed by Federgruen and Heching (1999), except that the latter assumes that the ordering cost is *proportional* to the amount ordered and thus does not include a fixed-cost component. In addition, the demand function is assumed to be a linear function of the price, see Lemma 1 in Federgruen and Heching (1999).

1523-4614/03/0501/0059$05.00
1526-5498 electronic ISSN

The paper by Thomas (1974) also considers a periodic review, finite-horizon model with a fixed ordering cost and stochastic, price-dependent demand. The paper postulates a simple policy, referred to by Thomas as (s, S, p), which can be described as follows. The inventory strategy is an (s, S) policy: If the inventory level at the beginning of period t is below the reorder point, s_t, an order is placed to raise the inventory level to the order-up-to level, S_t. Otherwise, no order is placed. The price, p, depends on the initial inventory level at the beginning of the period. Thomas provides a counterexample that shows that with a "few prices" (i.e., when price is restricted to a discrete set) this policy may fail to be optimal. Thomas goes on to say: "If all prices in an interval are under consideration, it is conjectured that a (s, S, p) policy is optimal under fairly general conditions" (p. 517).

In §2, we review the main assumptions of the model analyzed by Chen and Simchi-Levi (2002a, b). In §3 we characterize the optimal inventory and pricing policies for *additive demand functions*. We show that in this case *the policy proposed by Thomas* is indeed *optimal*. In §4 we analyze general demand functions that may be *nonadditive*. We demonstrate that in this case the profit-to-go function is not necessarily k-concave and an (s, S, p) policy is not necessarily optimal. We introduce the concept of symmetric k-convex functions and apply it to provide a characterization of the optimal policy. Finally, in §5 we extend the results obtained by Chen and Simchi-Levi (2002a) for the finite-horizon model to the infinite-horizon case under both discounted- and average-cost criteria.

2. The Model

Consider a firm that has to make production and pricing decisions over either a finite-time horizon with T periods or the infinite horizon. For convenience, in the finite-horizon case, we index periods from 1 to T, where 1 is the last period and T is the first period of the planning horizon. Notice that for the infinite-horizon case, we assume that all input parameters are stationary.

Demands in different periods are independent of each other. For each period t, $t = 1, 2 \dots$, let $d_t =$ demand in period t, $p_t =$ selling price in period t, and $\underline{p}_t, \bar{p}_t =$ lower and upper bounds on p_t, respectively. We concentrate on demand functions of the following forms:

Assumption 1. *For $t = 1, 2, \dots$, the demand function satisfies*

$$d_t = D_t(p_t, \epsilon_t) := \alpha_t D_t(p_t) + \beta_t, \tag{1}$$

where $\epsilon_t = (\alpha_t, \beta_t)$, and α_t, β_t are two random variables with $E\{\alpha_t\} = 1$ and $E\{\beta_t\} = 0$. The random perturbations, ϵ_t, are independent across time. Furthermore, $D_t(p_t)$ is a strictly decreasing function of p_t, and the expected revenue, $R_t(d) = dD_t^{-1}(d)$, is a concave function of expected demand d.

Observe that, by scaling and shifting, the assumptions $E\{\alpha_t\} = 1$ and $E\{\beta_t\} = 0$ can be made without loss of generality. A special case of this demand function is the *additive* demand function. In this case, the demand function is of the form $d_t = D_t(p) + \beta_t$. This implies that only β_t is a random variable, while $\alpha_t = 1$. Another special case of the demand function (1) is a model with *multiplicative* demand. In this case, the demand function is of the form $d_t = \alpha_t D_t(p)$, where α_t is a random variable. Finally observe that special cases of the function $D_t(p)$ include $D_t(p) = b_t - a_t p$ $(a_t > 0, b_t > 0)$ in the additive case and $D_t(p) = a_t p^{-b_t}$ $(a_t > 0, b_t > 1)$ in the multiplicative case; both are common in the economics literature (see Petruzzi and Dada 1999).

Let x_t be the inventory level at the beginning of period t, just before placing an order. Similarly, y_t is the inventory level at the beginning of period t after placing an order. The ordering cost function includes both a fixed cost and a variable cost and is calculated for every t, $t = 1, 2, \dots$, as

$$k_t \delta(y_t - x_t) + c_t(y_t - x_t),$$

where $\delta(0) = 0$ and $\delta(u) = 1$ for $u > 0$. Leadtime is assumed to be zero and hence an order placed at the beginning of period t arrives immediately before demand for the period is realized.

Unsatisfied demand is backlogged. Let x be the inventory level carried over from period t to the next period. Because we allow backlogging, x may be positive or negative. A cost $h_t(x)$ is incurred at the end of

period t that represents inventory holding cost when $x > 0$ and shortage cost if $x < 0$. As is customary in classic stochastic inventory problems, h_t is assumed to be convex.

The objective for the joint inventory/pricing problem is to decide on ordering and pricing policies so as to maximize total expected discounted profit over the entire planning horizon or the long-run average profit for the infinite-horizon case.

To find the optimal strategy for the joint inventory/pricing problem, let $v_t(x)$ be the maximum total expected discounted profit (discounted relative to period t) when t periods remain in the planning horizon and the inventory level at the beginning of period t is x for the finite-horizon problem. A natural dynamic program that can be applied to find the policy for the joint inventory/pricing problem is as follows. For $t = 1, 2, \ldots, T$,

$$v_t(x) = c_t x + \max_{y \geq x} -k_t \delta(y - x) + f_t(y, p_t(y)) \qquad (2)$$

with $v_0 = 0$, where γ is the given discount factor,

$$f_t(y, p) := -c_t y + E\{pD_t(p, \epsilon_t) - h_t(y - D_t(p, \epsilon_t)) + \gamma v_{t-1}(y - D_t(p, \epsilon_t))\},$$

and

$$p_t(y) \in \operatorname*{argmax}_{\bar{p}_t \geq p \geq \underline{p}_t} f_t(y, p). \qquad (3)$$

For technical reasons, we need the following assumption on the fixed ordering cost, as is commonly imposed in the traditional stochastic inventory problems.

Assumption 2. *In the finite-horizon model, the fixed ordering costs satisfy* $\gamma k_{t-1} \leq k_t$, *for* $t = 2, 3, \ldots$.

3. Finite-Horizon Model with Additive Demand

To characterize the optimal policy for the finite-horizon model with additive demand functions, Chen and Simchi-Levi (2002a) prove the following property.

Lemma 1. *Suppose there is a finite value* $p_t(y)$ *that maximizes* (3) *for any value of* y. *Then,* $y - D_t(p_t(y))$ *is a nondecreasing function of* y.

The lemma thus implies that the higher the inventory level at the beginning of time period t, y_t, the higher the expected inventory level at the end of period t, $y_t - D_t(p(y_t))$. Using this property, together with the definition of k-convex functions introduced by Scarf (1960), Chen and Simchi-Levi (2002a) prove,

Theorem 3.1. *For any* t, $t = T, T-1, \ldots, 1$, *we have*

(a) $f_t(y, p_t(y))$ *and* $v_t(x)$ *as k-concave.*

(b) *There exist* s_t *and* S_t *with* $s_t \leq S_t$ *such that it is optimal to order* $S_t - x_t$ *and set the selling price* $p_t = p_t(S_t)$ *when* $x_t < s_t$, *and not to order anything and set* $p_t = p_t(x_t)$ *when* $x_t \geq s_t$.

The theorem thus implies that the (s, S, p) policy introduced by Thomas (1974) is indeed optimal for additive demand processes. An interesting question is whether $p_t(y)$ is a nonincreasing function of y. Unfortunately, this property, which holds for the model with no fixed cost (see Federgruen and Heching (1999), does not hold for our model.

Proposition 1. *The optimal price,* $p_t(y)$, *is not necessarily a nonincreasing function of* y.

4. Finite-Horizon Model with General Demand

In this section, we focus on the finite-horizon model with general demand functions (1). Contrary to the additive demand case, Chen and Simchi-Levi (2002a) prove that

Lemma 2. *There exists an instance of Problem* (2) *with a multiplicative demand function and time-independent parameters such that the functions* $f_{T-1}(y, p_{T-1}(y))$ *and* $v_{T-1}(x)$ *are not k-concave.*

Lemma 3. *There exists an instance of Problem* (2) *with multiplicative demand functions where an* (s, S, p) *policy is not optimal.*

To overcome these difficulties, Chen and Simchi-Levi (2002a) propose a weaker definition of k-convexity, referred to as symmetric k-convexity:

Definition 4.1. A real-valued function f is called sym-k-convex for $k \geq 0$, if for any x_0, x_1, and $\lambda \in [0, 1]$,

$$f\big((1-\lambda)x_0 + \lambda x_1\big) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + \max\{\lambda, 1-\lambda\}k. \qquad (4)$$

A function f is called sym-k-concave if $-f$ is sym-k-convex.

Observe that k-convexity, and hence convexity, is a special case of sym-k-convexity. Interestingly, our analysis of sym-k-convex functions reveals that these functions have properties that are parallel to those of k-convex functions. Specifically, based on properties of these functions, Chen and Simchi-Levi (2002a) prove the following results.

THEOREM 4.1. *For any* t, $t = T, T-1, \ldots, 1$, *we have*

(a) $f_t(y, p_t(y))$ *and* $v_t(x)$ *as sym-k-concave.*

(b) *There exists* s_t *and* S_t *with* $s_t \le S_t$ *and a set* $A_t \subset [s_t, (s_t + S_t)/2]$, *such that it is optimal to order* $S_t - x_t$ *and set* $p_t = p_t(S_t)$ *when* $x_t < s_t$ *or when* $x_t \in A_t$ *and not to order anything and set* $p_t = p_t(x_t)$ *otherwise.*

Theorem 4.1 thus implies that an (s, S, A, p) policy is the optimal policy for Problem (2) under general demand processes. In such a policy, the optimal inventory strategy is characterized by two parameters s_t and S_t and a set $A_t \subset [s_t, (s_t + S_t)/2]$, possibly empty. When the inventory level, x_t, at the beginning of period t is less than s_t or if $x_t \in A_t$, an order of size $S_t - x_t$ is made. Otherwise, no order is placed. Thus, it is possible that an order will be placed when the inventory level $x_t \in [s_t, (s_t + S_t)/2]$, depending on the problem instance. On the other hand, if $x_t \ge (s_t + S_t)/2$, no order is placed. Price depends on the initial inventory level at the beginning of the period.

5. The Infinite-Horizon Case

In this section, we consider a model similar to the one analyzed in the previous sections except that in the infinite-horizon case all parameters are assumed to be *time independent*. Of course, it is tempting to try and extend the results of Theorem 4.1, which establishes the optimality of an (s, S, A, p) policy for the finite-horizon general demand model, to the infinite-horizon case. Surprisingly, the following theorem, proved in Chen and Simchi-Levi (2002b), shows that this intuition can be misleading.

THEOREM 5.1. *A stationary* (s, S, p) *policy is optimal for both the additive demand model and the general demand model under average- and discounted-cost criteria.*

Thus, the theorem suggests that in the infinite-horizon case, the optimal policy is an (s, S, p) policy, independent of whether demand is additive or not. Interestingly, our proof of the optimal policy for the general demand model is based on two key results: The first is that the long-run average (or discounted) profit function is symmetric k-concave, suggesting that a stationary (s, S, A, p) policy is optimal. Surprisingly, our second result shows that in the infinite-horizon case the set A is an empty set.

References

Chen, X., D. Simchi-Levi. 2002a. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. Working paper, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.

——, ——. 2002b. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. Working paper, Operations Researach Center, Massachusetts Institute of Technology, Cambridge, MA.

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

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.

Kimes, S. E. 1989. A tool for capacity-constrained service firms. *J. Oper. Management* **8**(4) 348–363.

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

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

Thomas, L. J. 1974. Price and production decisions with random demand. *Oper. Res.* **26** 513–518.

Whitin, T. M. 1955. Inventory control and price theory. *Management Sci.* **2** 61–80.