Available online at www.sciencedirect.com

SCIENCE DIRECT

Operations Research Letters 34 (2006) 323–332

Operations Research Letters

www.elsevier.com/locate/orl

Coordinating inventory control and pricing strategies: The continuous review model

Xin Chen[a,*], David Simchi-Levi[b]

[a]Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, USA
[b]The Engineering Systems Division and the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology (MIT), USA

Received 24 June 2004; accepted 30 April 2005
Available online 27 June 2005

Abstract

We analyze an infinite horizon, single product, continuous review model in which pricing and inventory decisions are made simultaneously and ordering cost includes a fixed cost. We show that there exists a stationary (s, S) inventory policy maximizing the expected discounted or expected average profit under general conditions.

Keywords: Inventory; Pricing; Stochastic; Continuous review; Fixed cost

1. Introduction

In recent years, many retailers have recognized the importance of coordinating replenishment strategies and pricing policies so as to boost the company bottom line. The academic community has also recognized the importance of models that coordinate pricing and inventory decisions, starting with the work of Whitin [12] who analyzed the celebrated newsvendor problem with price dependent demand. For a review, the reader is referred to Elmaghraby and Keskinocak [7] or Chan et al. [2].

To date, the literature has confined itself mainly to periodic review models. A general uncapacitated model is analyzed in Chen and Simchi-Levi [5,6] who considered a periodic review, single product model with stochastic demand with fixed ordering cost. For the finite horizon periodic review model, Chen and Simchi-Levi [5] proved that when the demand process is additive, i.e., the demand process has two components, a deterministic part which is a function of the price and an additive random perturbation, an $(s, S, \mathbf{p})$ policy is optimal. In such a policy 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. Price depends on the initial inventory level at the

* Corresponding author.
E-mail addresses: xinchen@uiuc.edu (X. Chen), dslevi@mit.edu (D. Simchi-Levi).

0167-6377/$ - see front matter
doi:10.1016/j.orl.2005.04.012

beginning of the period. Unfortunately, for general demand models, including multiplicative demand processes, i.e., the demand process has two components, a deterministic part which is a function of the price and a multiplicative random perturbation, an $(s, S, \mathbf{p})$ policy is not necessarily optimal. Specifically, Chen and Simchi-Levi introduce a new concept, symmetric k-convexity, and use it to prove that in this case an $(s, S, A, \mathbf{p})$ policy is optimal. In such a policy, the optimal inventory strategy at each time period is also characterized by two parameters (s_t, S_t) and a set $A_t \in [s_t, (s_t + S_t)/2]$, possibly empty depending on the problem instance. When the inventory level x_t at the beginning of period t is less than s_t or $x_t \in A_t$, an order of size $S_t - x_t$ is made. Otherwise, no order is placed. Price depends on the initial inventory level at the beginning of the period.

Surprisingly, by employing the symmetric k-convexity concept, Chen and Simchi-Levi [6] prove that in the infinite horizon periodic review model, a stationary $(s, S, \mathbf{p})$ policy is optimal for both additive demand and general demand processes specified in Chen and Simchi-Levi [6]. This is true for discounted as well as average profit criteria.

In this paper, we extend our approach for the infinite horizon periodic review model to a corresponding continuous review model. Here are some important features of the model analyzed in the current paper: demand arrives randomly at discrete time; pricing and replenishment decisions are made after serving the demand; the distribution of the inter-arrival time and the size of the demand depend on the selling price set at the previous decision epoch; the ordering cost includes both a fixed cost and a variable cost. The objective is to find an inventory policy and pricing strategy maximizing expected discounted or average profit over the entire planning horizon. For this model, we prove that a stationary $(s, S, \mathbf{p})$ policy is optimal for the infinite horizon continuous review problem under both the discounted and average profit criteria given rather general conditions.

Of course, one might think of transforming the continuous review model into a corresponding periodic review model through the classic Lippman's uniformization technique [9] and directly applying the results in Chen and Simchi-Levi [6] to the transformed model. However, it turns out that the transformed model is quite different from the periodic review model analyzed in [6], as we will see later on in this paper. Furthermore, the difference allows us to identify the structure of the optimal policy for the continuous review model under weaker conditions and to show some new properties that are not available for the periodic review model. In particular, no condition is imposed on the structure of the demand-price function and the distribution function of the inter-arrival time can be very general as compared with the infinite horizon periodic review model. Moreover, for the average profit model, under certain conditions, we prove that the *list price policy*, with respect to the inventory holding and shortage cost, is optimal, i.e., the higher the inventory holding and shortage cost, the lower the optimal selling price. Finally, we point out that the concept of symmetric k-convexity is employed to analyze the periodic review model in [6], while here this concept is not needed.

Our model and results generalize that of Feng and Chen [8]. Feng and Chen consider a model similar to ours, in which the *inter-arrival time* is assumed to be *exponential* and *demand size* is one unit. They focus on the long-run average profit and restrict prices to a discrete set. For this model, the authors show that inventory is managed based on an (s, S) policy and price is a function of the inventory level when a decision is made. Recently, Chao and Zhou [3] identify further structures of the optimal policies to facilitate the design of algorithms.

The paper is organized as follows. In Section 2, we review the main assumptions of our model. In Section 3, we prove some useful bounds on the reorder points and the order-up-to levels for the corresponding truncated finite horizon problems. We start in Section 4 by identifying properties of the best (s, S) inventory policy for both the discounted and average profit cases. These properties enable us to construct solutions for the optimality equations of the discounted and average profit problems in Section 5. In Sections 6 and 7, these equations are used to prove the optimality of a stationary $(s, S, \mathbf{p})$ policy for the infinite horizon problems with the discounted and average profit criteria, respectively. In Section 8, we illustrate the structure of the optimal price through a numerical example. Finally, in Section 9 we discuss extensions and provide concluding remarks.

2. The model

Consider a firm that has to make joint ordering and pricing decisions over an infinite time horizon with time independent demand process and costs functions. This problem is modeled as a semi-Markov decision process as follows. Customers arrive randomly and their inter-arrival time is independent, so is the demand size of each arrival. The firm fulfills the demand of each arrival from its inventory and unsatisfied demand is backlogged. After serving an arrival, replenishment and pricing decisions are made based on the current inventory level. Order lead time is assumed to be zero.

The ordering cost function includes both a fixed cost and a variable cost and is calculated as follows. Let u be the amount ordered, then the ordering cost function is

$$k\delta(u) + cu,$$

where

$$\delta(u) := \begin{cases} 1 & \text{if } u > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The selling price p is restricted to a compact set P. The inter-arrival time and the demand size of the next arrival depend on the price set at the latest arrival. Specifically, immediately after serving a customer, the firm decides on the ordering quantity and the price, p. The distribution of the next inter-arrival time is parameterized by p, which we denote by $F(t, p)$. Thus, the next arrival pays the selling price p and the demand size of this arrival is a random variable $d(p, \varepsilon)$, which is assumed to be a continuous function of p and a random perturbation ε independent of p. For technical reasons, we assume that $\Pr(d(p, \varepsilon) = 0)$ is bounded above by a positive constant less than 1 for any $p \in P$. Finally, let $R(p) = pE\{d(p, \varepsilon)\}$ be the expected revenue as a function of the selling price p.

Let x be the inventory level at time t. Since we allow backlogging, x may be positive or negative. In addition, let $h(x)$ be the rate at which the inventory holding and shortage cost is incurred.

Given a continuous time discount parameter γ with $\gamma \geqslant 0$, the objective is to decide on ordering and pricing policies so as to maximize the expected discounted profit ($\gamma > 0$) or the expected average profit ($\gamma = 0$) over the infinite horizon. In particular, let t_n be the epoch of the nth arrival. Let π be a policy representing a sequence of ordering and pricing decisions. For any natural number N and initial inventory level x, define

$$\begin{aligned} V^{\gamma}_{N,\pi}(x) = E_{\pi} \Bigg\{ & \mathrm{e}^{-\gamma t_N} c x_N \\ & + \sum_{n=0}^{N-1} \Bigg[-\mathrm{e}^{-\gamma t_n} (k\delta(y_n - x_n) \\ & + c(y_n - x_n)) + \mathrm{e}^{-\gamma t_{n+1}} p_n E\{d(p_n, \varepsilon)\} \\ & - h(y_n) \int_{t_n}^{t_{n+1}} \mathrm{e}^{-\gamma t}\, \mathrm{d}t \Bigg] \Bigg\}, \end{aligned} \tag{1}$$

where $x_0 = x$ is the inventory level at time $t_0 = 0$, x_n is the inventory level upon finishing serving the nth arrival, y_n is the inventory level after placing an order, p_n is the selling price set upon finishing serving the nth arrival and the subscript π for the expectation is used to emphasize the policy employed.

In the infinite horizon expected discounted profit problem ($\gamma > 0$), the objective is to determine an ordering and pricing policy π so as to maximize

$$\limsup_{N \to \infty} V^{\gamma}_{N,\pi}(x),$$

for any initial inventory level x. In the infinite horizon expected average profit problem, the objective is to find an ordering and pricing policy π maximizing

$$\limsup_{N \to \infty} \frac{1}{E_{\pi}\{t_N\}} V^{0}_{N,\pi}(x),$$

for any initial inventory level x. Note that for the semi-Markov decision processes, a different objective for the infinite horizon expected average profit problem is maximizing the lim sup of the expected average profit accumulated from time 0 to time t as t goes to ∞. However, under the assumptions specified in this paper, the two objectives are equivalent (see [10]).

Given a real value σ, consider the following dynamic programming problem. For $n = 1, 2, \ldots, N$,

$$\phi^{\gamma}_{n,\sigma}(x) = \max_{y \geqslant x, p \in P} -k\delta(y - x) + g^{\gamma}_{n,\sigma}(y, p), \tag{2}$$

with $\phi^{\gamma}_{0,\sigma}(x)=0$ for any x, where

$$g^{\gamma}_{n,\sigma}(y,p)=\tau^{\gamma}(p)(H^{\gamma}(y,p)-\sigma)+(1-\gamma\tau^{\gamma}(p))\times E\{\phi^{\gamma}_{n-1,\sigma}(y-d(p,\varepsilon))\},$$

$$H^{\gamma}(x,p)=\frac{1-\gamma\tau^{\gamma}(p)}{\tau^{\gamma}(p)}\hat{R}(p)-h^{\gamma}(x),$$

and

$$h^{\gamma}(x)=h(x)+\gamma cx, \quad \text{and}$$

$$\hat{R}(p)=(p-c)E\{d(p,\varepsilon)\}.$$

Furthermore,

$$\tau^{\gamma}(p)=E\left\{\int_0^{\xi} \mathrm{e}^{-\gamma t}\,\mathrm{d}t\right\},$$

and ξ is the inter-arrival time with distribution $F(t,p)$. Observe that $\tau^0(p)=E\{\xi\}$ and $1-\gamma\tau^{\gamma}(p)=E\{\mathrm{e}^{-\gamma\xi}\}>0$ for $p\in P$.

It is easy to verify that $\phi^{\gamma}_{n,0}(x)+cx$ is the maximum expected discounted profit accrued over a horizon starting at some epoch upon finishing serving an arrival with an initial inventory level x and ending at the epoch of the next nth arrival. Moreover, observe that if for any policy π and any initial inventory level x

$$\limsup_{N\to\infty}\frac{1}{E_{\pi}\{t_N\}}(V^0_{N,\pi}(x)-\sigma E_{\pi}\{t_N\})\leqslant 0,$$

then σ is an upper bound of the optimal value for the infinite horizon expected average profit problem. Thus, if $\limsup_{n\to\infty}((\phi^{\gamma}_{n,\sigma}(x)+cx)/E_{\pi}\{t_N\})\leqslant 0$ for any feasible policy π, then σ is an upper bound of the optimal value for the infinite horizon expected average profit problem.

Define

$$\underline{\tau}^{\gamma}=\inf_{p\in P}\tau^{\gamma}(p) \quad \text{and}$$

$$\bar{\tau}^{\gamma}=\sup_{p\in P}1-\gamma\tau^{\gamma}(p)=1-\gamma\underline{\tau}^{\gamma}.$$

Assumption 1. For any $\gamma\geqslant 0$, $\tau^{\gamma}(p)$ is continuous for $p\in P$ and $\underline{\tau}^{\gamma}>0$. As a consequence, $\bar{\tau}^{\gamma}=1-\gamma\underline{\tau}^{\gamma}<1$ for $\gamma>0$.

Thus, the assumption implies that the expected discounted inter-arrival time, $\tau^{\gamma}(p)$, is bounded below by a positive constant.

Observe that in (2), $H^{\gamma}(x,p)$, the profit rate function, is separable, i.e., it is the difference of a function with a single variable x and another function with a single variable p. However, this is not the case for the periodic review model analyzed in Chen and Simchi-Levi [6]. Notice that, in Chen and Simchi-Levi [6], an inventory holding and shortage cost is incurred at the end of each period and thus, the single period profit function cannot be separated. This difference is quite significant. Indeed, even though the analysis is similar, we are able to show for the continuous review model that the structural properties hold under weaker conditions and this allows us to identify some new properties that are not available for the periodic review model.

For technical reasons, we need the following assumption on the expected revenue and the inventory holding and shortage cost functions.

Assumption 2.

(a)

$$\bar{R}^{\gamma}=\sup_{p\in P}\frac{1-\gamma\tau^{\gamma}(p))}{\tau^{\gamma}(p)}\hat{R}(p)<\infty.$$

(b) h^{γ} is continuous and quasi-convex and

$$\lim_{|x|\to\infty}h^{\gamma}(x)=\infty.$$

(c)

$$\lim_{|x|\to\infty}\sup_{p\in P}(1-\gamma\tau^{\gamma}(p))R(p)-\tau^{\gamma}(p)h(x)-cx=-\infty.$$

Observe that we assume that at the end of the Nth arrival there is a unit salvage value c. Assumptions 1 and 2(c) allow us to show that the difference between the models with or without salvage values vanishes as $N\to\infty$. Hence our results apply to a model with zero salvage value.

3. Finite horizon model

In this section, we prove that the classical (s,S) inventory policy is optimal for problem (2). This approach is similar to that by Veinott [11].

Lemma 1. *Let x^γ be any minimum point of h^γ. We have*

(a) *For any $n \geqslant 1$, $g^\gamma_{n,\sigma}(x, p)$ is a non-decreasing function of x for any $x \leqslant x^\gamma$ and $p \in P$. As a consequence, $\phi^\gamma_{n,\sigma}(x)$ is non-decreasing for $x \leqslant x^\gamma$.*
(b) *For any $n \geqslant 1$, $\phi^\gamma_{n,\sigma}(x) \geqslant \phi^\gamma_{n,\sigma}(x') - k$ for any $x \leqslant x'$. Furthermore, $g^\gamma_{n,\sigma}(x, p) \geqslant g^\gamma_{n,\sigma}(x', p) - k$ for any $x^\gamma \leqslant x \leqslant x'$ and $p \in P$.*
(c) *For any $n \geqslant 1$, there exist $s^\gamma_{n,\sigma}$ and $S^\gamma_{n,\sigma}$ with $s^\gamma_{n,\sigma} \leqslant x^\gamma \leqslant S^\gamma_{n,\sigma}$ such that the $(s^\gamma_{n,\sigma}, S^\gamma_{n,\sigma})$ inventory policy is optimal for* (2).

Proof. We prove by induction. For $n=0$, $\phi^\gamma_{n,\sigma}(x)=0$ for any x. Assume $\phi^\gamma_{n-1,\sigma}(x)$ is non-decreasing for $x \leqslant x^\gamma$. This assumption, together with the fact that $\tau^\gamma(p)$, $1-\gamma\tau^\gamma(p) > 0$ for any $p \in P$ and $h^\gamma(y)$ is non-increasing for $y \leqslant x^\gamma$, implies that $g^\gamma_{n,\sigma}(x, p)$ is a non-decreasing function of x for any $x \leqslant x^\gamma$ and $p \in P$ and consequently $\phi^\gamma_{n,\sigma}(x)$ is non-decreasing for $x \leqslant x^\gamma$. Hence part (a) holds.

From (2), it is obvious that the first part of (b) holds for any n. The second part of (b) follows from the first part of (b) and the fact that $\tau^\gamma(p)$, $1-\gamma\tau^\gamma(p) > 0$ for any $p \in P$ and $h^\gamma(y)$ is non-decreasing for $y \geqslant x^\gamma$.

Part (c) is a direct consequence of parts (a) and (b). □

Let $\underline{s}^\gamma$ and $\bar{S}^\gamma$ be two real numbers such that

$$\underline{s}^\gamma \leqslant x^\gamma \leqslant \bar{S}^\gamma,\ h^\gamma(\underline{s}^\gamma) = h^\gamma(x^\gamma) + \frac{k}{\underline{\tau}^\gamma}, \quad \text{and}$$

$$h^\gamma(\bar{S}^\gamma) = h^\gamma(x^\gamma) + \left(\frac{1}{\underline{\tau}^\gamma} - \gamma\right)k.$$

Lemma 2. $\underline{s}^\gamma \leqslant s^\gamma_{n,\sigma} \leqslant x^\gamma \leqslant S^\gamma_{n,\sigma} \leqslant \bar{S}^\gamma$ *for any $n \geqslant 1$ and σ.*

Proof. For any fixed $p \in P$, we have by Lemma 1 part (a) and the quasi-convexity of h^γ that

$$g^\gamma_{n,\sigma}(x, p) - g^\gamma_{n,\sigma}(x^\gamma, p) \leqslant -k, \quad \text{for } x \leqslant \underline{s}^\gamma$$

and by Lemma 1 part (b) and the quasi-convexity of h^γ that

$$g^\gamma_{n,\sigma}(x, p) - g^\gamma_{n,\sigma}(x^\gamma, p) \leqslant 0, \quad \text{for } x \geqslant \bar{S}^\gamma.$$

The lemma follows from the above two inequalities and Lemma 1 part (c). □

We thus conclude that the optimal policy for problem (1) is obtained by solving problem (2) with $\sigma=0$. The inventory policy is an (s, S) policy while the price depends on the inventory level when making the decision. Furthermore, the parameters of the optimal policy are uniformly bounded, which is useful for our analysis later on. Finally, we point out that for the finite horizon periodic review model, an $(s, S, \mathbf{p})$ policy is not necessarily optimal, see Chen and Simchi-Levi [5].

4. Characterization of the best (s, S) policy

In this section, we characterize the properties of the best (s, S) policy, which allow us to construct solutions for the optimality equations in the next section.

Consider a stationary $(s, S, \mathbf{p})$ policy defined by two inventory levels s, S and a price $\mathbf{p}(x)$ which is a function of the inventory level x. Define $\tau(s, x, \mathbf{p})$ to be the length of a horizon starting at some epoch upon finishing serving an arrival with an initial inventory level x and ending at the earliest epoch with an inventory level no more than s. Therefore, we have $\tau(s, x, \mathbf{p})=0$ for $x \leqslant s$ and

$$\tau(s, x, \mathbf{p}) = \xi + \tau(s, x - d(\mathbf{p}(x), \varepsilon), \mathbf{p}) \quad \text{for } x > s,$$

where ξ is a random variable with distribution $F(t, \mathbf{p}(x))$.

Let $I^\gamma(s, x, \mathbf{p})$ be the expected γ-discounted profit incurred during a horizon that starts with an initial inventory level x achieved upon finishing serving an arrival epoch and ends at the earliest epoch when the inventory level drops to level s or below. Let $M^\gamma(s, x, \mathbf{p}) = E\{\tau(s, x, \mathbf{p})\}$ for $\gamma = 0$, and

$$M^\gamma(s, x, \mathbf{p}) = (1 - E\{\mathrm{e}^{-\gamma\tau(s,x,\mathbf{p})}\})/\gamma, \quad \text{for } \gamma > 0.$$

Observe that whenever $x \leqslant s$, we have $I^\gamma(s, x, \mathbf{p}) = 0$ and $M^\gamma(s, x, \mathbf{p}) = 0$. On the other hand when $x > s$ we have

$$I^\gamma(s, x, \mathbf{p}) = \tau^\gamma(\mathbf{p}(x))H^\gamma(x, \mathbf{p}(x)) + (1 - \gamma\tau^\gamma(\mathbf{p}(x))) \times E\{I^\gamma(s, x - d(\mathbf{p}(x), \varepsilon), \mathbf{p})\}, \tag{3}$$

and

$$M^\gamma(s, x, \mathbf{p}) = \tau^\gamma(\mathbf{p}(x)) + (1 - \gamma\tau^\gamma(\mathbf{p}(x))) \times E\{M^\gamma(s, x - d(\mathbf{p}(x), \varepsilon), \mathbf{p})\}. \tag{4}$$

The recursion (4) thus implies that $M^\gamma(s, x, \mathbf{p})$ is the expected γ-discounted time to drop from the initial inventory level x to or below s.

Define

$$c^\gamma(s, S, \mathbf{p}) = \frac{-k + I^\gamma(s, S, \mathbf{p})}{M^\gamma(s, S, \mathbf{p})}.$$

From the elementary renewal reward theory (see [10]), we know that for $\gamma = 0$, $c^\gamma(s, S, \mathbf{p})$ is exactly the long-run average profit. For $\gamma > 0$, the infinite horizon expected discounted profit starting with an initial inventory level x is

$$I^\gamma(s, x, \mathbf{p}) + E\{e^{-\gamma\tau(s,x,\mathbf{p})}\}c^\gamma(s, S, \mathbf{p})/\gamma,$$

which is equal to

$$c^\gamma(s, S, \mathbf{p})/\gamma + I^\gamma(s, x, \mathbf{p}) - c^\gamma(s, S, \mathbf{p})M^\gamma(s, x, \mathbf{p}). \tag{5}$$

We continue by assuming that the demand size is positive. Formally, this assumption says that for any realization of the random variables ε, $d(p, \varepsilon) \geqslant \eta > 0$ for some η and any $p \in P$. In addition, to avoid mathematical complications, we assume that the feasible set P of the selling prices has finite number of elements. Our results in this paper still hold when these assumptions are relaxed; see Chen [4] for details.

For any given (s, S), let $c^\gamma(s, S)$ be the optimal value of problem

$$\max_{\mathbf{p}:\mathbf{p}(x)\in P} c^\gamma(s, S, \mathbf{p}). \tag{6}$$

Define

$$\phi^\gamma(x, s, S, s') = \begin{cases} 0 & \text{for } x \leqslant s', \\ \max_{p\in P} g^\gamma(x, s, S, s', p) & \text{for } x > s', \end{cases} \tag{7}$$

where

$$g^\gamma(x, s, S, s', p) = \tau^\gamma(p)(H^\gamma(x, p) - c^\gamma(s, S)) + (1 - \gamma\tau^\gamma(p)) \times E\{\phi^\gamma(x - d(p, \varepsilon), s, S, s')\}.$$

Let $\phi^\gamma(x, s, S) = \phi^\gamma(x, s, S, s)$.

If for any given $x > s'$, $\mathbf{p}^*(x)$ is optimal for problem (7), one can show, from the definition of $c^\gamma(s, x, \mathbf{d})$, that for $x > s'$, we have

$$\phi^\gamma(x, s, S, s') = k + (c^\gamma(s', x, \mathbf{p}^*) - c^\gamma(s, S)) \times M^\gamma(s', x, \mathbf{p}^*). \tag{8}$$

Furthermore, in the case with $s' = s$, one can show that the optimal solution $\mathbf{p}^*(x)$, constructed recursively from (7), is optimal for problem (6) and in particular, $\phi^\gamma(S, s, S) = k$.

Let c^γ be the optimal value of problem

$$\max_{(s,S,\mathbf{p})} c^\gamma(s, S, \mathbf{p}) = \max_{(s,S)} c^\gamma(s, S). \tag{9}$$

Define

$$F^\gamma := \left\{(s, S) \,|\, Q^\gamma(S) \geqslant Q^\gamma(s) = c^\gamma(s, S) \geqslant \max_{x,p\in P} H^\gamma(x, p) - \frac{k}{\tau^\gamma(p)}\right\},$$

where $Q^\gamma(x)$ is the maximum profit rate for a given inventory level x, i.e.,

$$Q^\gamma(x) = \max_{p\in P} H^\gamma(x, p) = \bar{R}^\gamma - h^\gamma(x).$$

In the following proposition we show that it is sufficient to focus on the set F^γ in the search for the best (s, S) policy.

Proposition 1. $c^\gamma = \max_{(s,S)\in F^\gamma} c^\gamma(s, S)$.

A similar result holds for the corresponding infinite horizon, periodic review model in Chen and Simchi-Levi [6]; the proof of this proposition follows a somewhat similar logic and hence is omitted. For details, see Chen and Simchi-Levi [6] or Chen [4].

To provide some intuition, we point out that $Q^\gamma(x)$ is the maximum expected profit rate when we start with an inventory level x; $c^\gamma(s, S)$ can be viewed as the average discounted profit per unit time for a given (s, S) policy and its associated best price strategy. Thus, if $Q^\gamma(s) \neq c^\gamma(s, S)$, one can change the reorder point, s^γ, and improve the average discounted profit per unit time. If, on the other hand, $Q^\gamma(S) < c^\gamma(s, S)$, one can decrease S^γ and increase average discounted profit per unit time.

For any $(s, S) \in F^\gamma$, since $Q^\gamma(s) = c^\gamma(s, S)$, one can show that $\phi^\gamma(x, s, S)$ is continuous in x and

$$\phi^\gamma(x, s, S) = \begin{cases} 0 & \text{for } x \leqslant s, \\ \max_{p\in P} g^\gamma(x, s, S, s, p) & \text{for } x \geqslant s. \end{cases}$$

Furthermore, by following an approach similar to the proof of item (ii), one can show that $c^\gamma(s, S)$ is continuous and hence the set F^γ is compact.

In the following lemma, we characterize the best (s, S) inventory policy. This lemma is key to our analysis of the discounted and average profit problems.

Lemma 3. *There exists an optimal solution* $(s^\gamma, S^\gamma) \in F^\gamma$ *to problem* (9) *such that the functions* $\phi^\gamma(x) := \phi^\gamma(x, s^\gamma, S^\gamma)$ *and* $h^\gamma(x)$, *satisfy the following properties.*

(a) $\phi^\gamma(x) \leqslant k$ *for any* x *and* $\phi^\gamma(S^\gamma) = k$.
(b) $h^\gamma(s^\gamma) = \bar{R}^\gamma - c^\gamma$.
(c) $h^\gamma(x) \leqslant \bar{R}^\gamma - c^\gamma$ *for* $x \in [s^\gamma, S^\gamma]$.
(d) $\phi^\gamma(x) \geqslant 0$ *for any* $x \leqslant S^\gamma$.
(e) $s^\gamma \leqslant x^\gamma \leqslant S^\gamma$, *where* x^γ *is any minimum point of* $h^\gamma(x)$.

Proof. The existence of the best (s, S) inventory policy follows from the continuity of $c^\gamma(s, S)$ and the compactness of the set F^γ. From now on, we use $(s^\gamma, S^\gamma) \in F^\gamma$ to denote one of the best (s, S) policies.

Part (a) follows from (8) and the fact that (s^γ, S^γ) solves problem (9).

Parts (b) and (c) hold since $(s^\gamma, S^\gamma) \in F^\gamma$ and h^γ is quasi-convex.

Part (d) follows from part (c) and the recursive definition of ϕ^γ in (7).

We now prove part (e). First, it is not difficult to show that s^γ can be chosen as the smallest element in the set $\{x | h^\gamma(x) = \bar{R}^\gamma - c^\gamma\}$. Hence for any minimum point x^γ of h^γ, $s^\gamma \leqslant x^\gamma$.

It remains to show that $x^\gamma \leqslant S^\gamma$. Assume there exists a minimum point x^γ of h^γ, such that $s^\gamma \leqslant S^\gamma < x^\gamma$. We prove by induction that $\phi^\gamma(x)$ is non-decreasing for $x \leqslant x^\gamma$ and consequently we can choose S^γ to be x^γ. Assume that it is true for any x with $x \leqslant y$ for some $y \leqslant x^\gamma$. For any x and x' such that $s^\gamma \leqslant x \leqslant x' \leqslant \min\{y + \eta, x^\gamma\}$, we have that

$$\begin{aligned}\phi^\gamma(x) &= \max_{p \in P} \tau^\gamma(p)(H^\gamma(x, p) - c^\gamma) \\ &\quad + (1 - \gamma\tau^\gamma(p))E\{\phi^\gamma(x - d(p, \varepsilon))\} \\ &\leqslant \max_{p \in P} \tau^\gamma(p)(H^\gamma(x', p) - c^\gamma) \\ &\quad + (1 - \gamma\tau^\gamma(p))E\{\phi^\gamma(x' - d(p, \varepsilon))\} \\ &= \phi^\gamma(x'),\end{aligned}$$

where the inequality follows from the definition of function ϕ^γ, the quasi-convexity of h^γ, and the induction assumption. Therefore, $c^\gamma = c^\gamma(s^\gamma, S^\gamma) \leqslant c^\gamma(s^\gamma, x^\gamma)$, and hence $(s^\gamma, x^\gamma) \in F^\gamma$, which implies that (a)–(d) hold for the policy (s^γ, x^γ). Thus part (e) holds. □

Similar properties hold for the corresponding infinite horizon, periodic review model, see Chen and Simchi-Levi [6]. However, due to the separability of the profit rate function $H^\gamma(x, p)$, part (e) here is stronger than its counterpart for the infinite horizon periodic review model analyzed in [6]. Finally, we point out that Lemma 3 is essentially parallel to the characterization of the best (s, S) policy, given by Lemma 1 in Zheng [13], for the standard infinite horizon inventory control models.

5. Optimality equations

In this section, employing the properties of the best (s, S) policy, we construct solutions for the optimality equations of the infinite horizon models. In particular, Lemma 3 allows us to show that (ϕ^γ, c^γ) satisfies the equation

$$\begin{aligned}\phi^\gamma(x) = \max_{y \geqslant x} \Bigg\{ &\max_{p \in P} -k\delta(y - x) \\ &+ \tau^\gamma(p)(H^\gamma(y, p) - c^\gamma) \\ &+ (1 - \gamma\tau^\gamma(p))E\{\phi^\gamma(y - d(p, \varepsilon))\} \Bigg\}\end{aligned} \quad (10)$$

and that the (s^γ, S^γ) policy attains the first maximization in Eq. (10).

Notice that when $\gamma = 0$, (10) is the optimality equation for the average profit problem. On the other hand, when $\gamma > 0$, define

$$\hat{\phi}^\gamma(x) = \frac{c^\gamma}{\gamma} + \phi^\gamma(x).$$

Then (10) implies that

$$\begin{aligned}\hat{\phi}^\gamma(x) = \max_{y \geqslant x, p \in P} &-k\delta(y - x) + \tau^\gamma(p)H^\gamma(y, p) \\ &+ (1 - \gamma\tau^\gamma(p))E\{\hat{\phi}^\gamma(y - d(p, \varepsilon))\},\end{aligned}$$

which is the optimality equation for the γ-discounted profit problem with $\gamma > 0$.

Theorem 5.1. (ϕ^γ, c^γ) *satisfies Eq.* (10) *and the* (s^γ, S^γ) *policy attains the first maximization in Eq.* (10).

A similar result holds for the corresponding infinite horizon, periodic review model in Chen and Simchi-Levi [6]; the proof of this theorem follows a somewhat similar logic and hence is omitted. For details, see Chen and Simchi-Levi [6] or Chen [4].

It is appropriate to point out that due to the separable property of the profit rate function $H^\gamma(x,p)$, the concept of symmetric k-convexity, which is important for the analysis for the periodic review model in Chen and Simchi-Levi [6], is not needed here.

6. Discounted profit case

Consider the discounted profit case with a discount parameter $\gamma > 0$ and recall the definition of $\hat{\phi}^\gamma(x)$. Formula 5 tells us that $\hat{\phi}^\gamma(x)$ is the infinite horizon expected discounted profit for the stationary (s^γ, S^γ) policy associated with the best price strategy when starting with an initial inventory level x.

The following convergence result relates $\hat{\phi}^\gamma(x)$ to the maximum total expected discounted profit function over a horizon starting at some epoch with an initial inventory level x upon finishing serving an arrival and ending at the epoch of the next nth arrival, $\phi_n^\gamma(x)$.

Theorem 6.1. *For any $M \geqslant \max\{\bar{S}^\gamma, S^\gamma\}$ and any $n \geqslant 1$, we have that*

$$\max_{x\leqslant M} |\phi_n^\gamma(x) - \hat{\phi}^\gamma(x)| \leqslant (\bar{\tau}^\gamma)^{n-1} \max_{x\leqslant M} |\phi_1^\gamma(x) - \hat{\phi}^\gamma(x)|.$$

The proof of the theorem is based on induction and is similar to the analysis of the periodic review model in Chen and Simchi-Levi [6]. Thus, the proof is omitted and for more details, one may refer to Chen [4].

The theorem thus implies that the function $\phi_n^\gamma(x)$, converges to the infinite horizon expected discounted profit function, $\hat{\phi}^\gamma(x)$, associated with the stationary (s^γ, S^γ) policy and its corresponding best price strategy and as a consequence, this policy is optimal for the infinite horizon expected discounted profit problem.

7. Average profit case

In this section, we analyze the average profit case and hence assume that $\gamma = 0$. We show that the long-run average profit of the best (s, S) inventory policy associated with its best price strategy, c^0, is the limit of the maximum average profit per unit time over a horizon starting at some epoch upon finishing serving an arrival and ending at the epoch of the next nth arrival as $n \to \infty$.

Theorem 7.1. *c^0 is the optimal value for the infinite horizon expected average profit problem.*

Proof. First we prove by induction that for any given $M \geqslant \{\bar{S}^0, S^0\}$ and any $\sigma > c^0$,

$$\phi_{n,\sigma}^0(x) - \phi^0(x) \leqslant -n(\sigma - c^0)\underline{\tau}^0 + \max_{x\leqslant M}(-\phi^0(x)), \quad \text{for } x \leqslant M. \tag{11}$$

From (2) and (10), we have that for any $x \leqslant M$,

$$\begin{aligned}
&\phi_{n,\sigma}^0(x) - \phi^0(x)\\
&\quad= \max_{M\geqslant y\geqslant x, p\in P} -k\delta(y-x) + \tau^0(p)(H^0(y,p)\\
&\qquad - \sigma) + E\{\phi_{n-1,\sigma}^0(y - d(p,\varepsilon))\}\\
&\qquad - \max_{M\geqslant y\geqslant x, p\in P} -k\delta(y-x) + \tau^0(p)(H^0(y,p)\\
&\qquad - c^0) + E\{\phi^0(y - d(p,\varepsilon))\}\\
&\quad\leqslant \max_{M\geqslant y\geqslant x, p\in P} \{-k\delta(y-x) + \tau^0(p)(H^\gamma(y,p)\\
&\qquad - \sigma) + E\{\phi_{n-1,\sigma}^0(y - d(p,\varepsilon))\}\\
&\qquad - (-k\delta(y-x) + \tau^0(p)(H^0(y,p) - c^0)\\
&\qquad + E\{\phi^0(y - d(p,\varepsilon))\})\}\\
&\quad\leqslant -(\sigma - c^0)\inf_{p\in P} \tau^0(p) + \max_{M\geqslant y\geqslant x, p\in P}\\
&\qquad \times E\{\phi_{n-1,\sigma}^0(y - d(p,\varepsilon)) - \phi^0(y - d(p,\varepsilon))\}\\
&\quad\leqslant -(\sigma - c^0)\underline{\tau}^0 + \max_{x\leqslant M} \phi_{n-1,\sigma}^0(x) - \phi^0(x).
\end{aligned}$$

Thus, by induction, (11) is true. Hence for any $x \leqslant M$,

$$\frac{\phi_{n,\sigma}^0(x)}{n\underline{\tau}^0} \leqslant -(\sigma - c^0) + \frac{\phi^0(x) - \min_{x\leqslant M}\phi^0(x)}{n\underline{\tau}^0} \to -(\sigma - c^0) < 0.$$

By choosing M arbitrarily large, we have that for any x and $\sigma > c^0$,

$$\limsup_{n\to\infty} \frac{\phi^0_{n,\sigma}(x)}{n\underline{\tau}^0} \leqslant 0.$$

Therefore, c^0 is the optimal value for the infinite horizon expected average profit problem. $\square$

The above theorem implies that the stationary $(s^0, S^0, p^0_{(s^0,S^0)})$ policy is optimal for the infinite horizon expected average profit problem. In this optimal policy, the pricing strategy has a special structure when demand size is independent of price while the expected inter-arrival time is a strictly increasing function of the price. Indeed, as is shown in the following theorem, the higher the inventory holding and shortage cost, the smaller the selling price. Specifically,

Theorem 7.2. *Assume that $d(p,\varepsilon) = d(\varepsilon)$, $\tau^0(p)$ is strictly increasing. Then $h^0(x') < h^0(x)$ implies that $p^0(x') \geqslant p^0(x)$ for $x, x' \geqslant s^0$, where $p^0(x) := p^0_{(s^0,S^0)}(x)$ for simplicity.*

Proof. Recall that

$$\begin{aligned} p^0(x) \in \arg\max_{p\in P}\ & \tau^0(p)(H^0(x,p) - c^0) \\ & + E\{\phi^0(x - d(\varepsilon))\}. \end{aligned}$$

For any $s^0 \leqslant x, x' \leqslant S^0$, we have

$$\begin{aligned} &\tau^0(p^0(x))(H^0(x,p^0(x)) - c^0) + E\{\phi^0(x - d(\varepsilon))\} \\ &\quad \geqslant \tau^0(p^0(x'))(H^0(x,p^0(x')) - c^0) \\ &\qquad + E\{\phi^0(x - d(\varepsilon))\} \end{aligned}$$

and

$$\begin{aligned} &\tau^0(p^0(x'))(H^0(x',p^0(x')) - c^0) + E\{\phi^0(x' - d(\varepsilon))\} \\ &\quad \geqslant \tau^0(p^0(x))(H^0(x,p^0(x)) - c^0) \\ &\qquad + E\{\phi^0(x' - d(\varepsilon))\}. \end{aligned}$$

Adding the above two inequalities gives

$$(\tau^0(p^0(x')) - \tau^0(p^0(x)))(h^0(x') - h^0(x)) \leqslant 0.$$

Therefore $h^0(x') < h^0(x)$ implies that $p^0(x') \geqslant p^0(x)$. $\square$

Thus, in the special case analyzed by the theorem, a **list price** policy, based on inventory holding and shortage cost, is optimal.

8. A numerical example

In this section, we illustrate, through a numerical example, the structure of the optimal price for the infinite horizon model under average profit criterion. The readers are referred to Chen [4] for more numerical experiments which illustrate that the change of the performance of dynamic pricing strategies as system parameters changes.

In the example, we use $k=25$, $c=10$, and we assume Poisson arrival and the arrival rate is a linear function of the selling price: $\lambda = \frac{121}{3} - \frac{7}{3}p$, where price is a continuous variable which can take values from the interval 13–16. We also assume unit demand for each arrival. The inventory holding and shortage cost takes the following form: $h(x) = \max\{0, x\} + 5\max\{0, -x\}$. These data are based on our experience with a large industrial manufacturing company. Thus, we focus on integral inventory levels. Notice that as we point out in Section 9, the structural results of the optimal policy can be extended to the case with integral inventory levels. Fig. 1 clearly illustrates the structure of the optimal price characterized by Theorem 7.2.

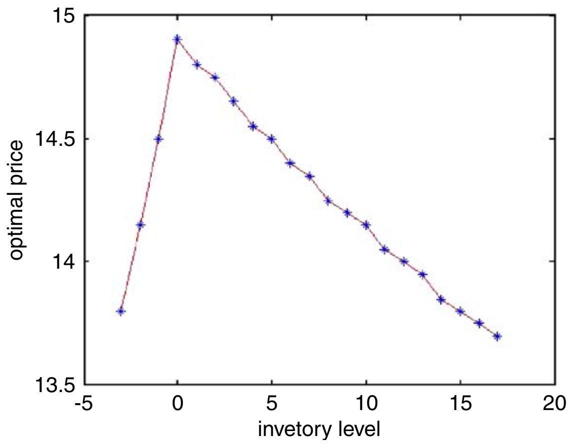

Fig. 1. Illustration of the structure of the optimal price.

9. Extensions and concluding remarks

In this section, we report on some important extensions of the model and results.

Interestingly, it is easy to show that the results can be extended to a model in which demand and inventory levels are restricted to discrete sets. In fact, Lemma 3 holds with part (b) replaced by the property that $h^{\gamma}(s^{\gamma}) \geqslant \bar{R}^{\gamma} - c^{\gamma} > h^{\gamma}(s^{\gamma}+1)$, and Theorems 5.1, 6.1, 7.1 and 7.2 follow directly from this fact.

It is also easy to show that our results can be extended to a model in which the demand size depends on the time that elapsed since the last customer arrival.

Thus, our model is more general than the model analyzed in [8] where the authors considered discrete price, unit demand and exponential inter-arrival time. In fact, Theorems 7.1 and 7.2 generalize their results to a model with general demand, and general inter-arrival time.

Finally, it is appropriate to point out an important limitation of our model, namely the assumption of zero lead time. Indeed, extending our results to constant lead time seems quite challenging, since for our model with lead time, the two decisions, the replenishment decision and the pricing decision, take effects at different times. Thus, the traditional approach in the standard stochastic inventory models, i.e., transforming the model with positive lead time to one with zero lead time, does not work here. Recall that in the case of the traditional continuous time inventory model, see Beckmann [1], an (s, S) policy is optimal under any constant lead time.

Acknowledgements

Research supported in part by the Singapore-MIT Alliance, ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DMI-9732795, DMI-0085683 and DMI-0245352. The authors would like to thank X. L. Chao and S. X. Zhou for providing their code which we modified for the purpose of our numerical study. We also thank the referee for insightful suggestions to improve the paper.

References

[1] M. Beckmann, An inventory model for arbitrary interval and quantity distributions of demand, Manage. Sci. 8 (1) (1961) 35–57.
[2] L.M.A. Chan, Z.J. Max Shen, D. Simchi-Levi, J. Swann, Coordination of pricing and inventory decisions: a survey and classification, in: D. Simchi-Levi, S.D. Wu, Z.J. Max Shen (Eds.), Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era, Kluwer, Boston, 2004 (Chapter 9).
[3] X.L. Chao, S.X. Zhou, Joint inventory and pricing strategy for a stochastic continuous-review system, Working paper, North Carolina State University, 2004.
[4] X. Chen, Coordinating inventory control and pricing strategies, Ph.D. Thesis, Massachusetts Institute of Technology, 2003.
[5] X. Chen, D. Simchi-Levi, Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: the finite horizon case, Oper. Res. 52 (6) (2004), pp 887–896.
[6] X. Chen, D. Simchi-Levi, Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: the infinite horizon case, Math. Oper. Res. 29 (3) (2004), pp. 698–723.
[7] W. Elmaghraby, P. Keskinocak, Dynamic pricing in the presence of inventory considerations: research overview, current practices, and future directions, Manage. Sci. 49 (2003) 1287–1309.
[8] Y. Feng, F. Chen, Joint pricing and inventory control with setup costs and demand uncertainty, The Chinese University of Hong Kong, 2002.
[9] S. Lippman, Applying a new device in the optimization of exponential queueing systems, Oper. Res. 23 (1975) 687–710.
[10] S. Ross, Applied Probability Models with Optimization Applications, Holden-Day, San Francisco, 1970.
[11] A.F. Veinott, On the optimality of (s, S) inventory policies: new conditions and a new proof, SIAM J. Appl. Math. 14 (5) (1966) 1067–1083.
[12] T.M. Whitin, Inventory control and price theory, Manage. Sci. 2 (1955) 61–80.
[13] Y.S. Zheng, A simple proof for optimality of (s, S) policies in infinite-horizon inventory systems, J. Appl. Prob. 28 (1991) 802–810.