COMBINED PRICING AND INVENTORY CONTROL UNDER UNCERTAINTY

AWI FEDERGRUEN and ALIZA HECHING

Columbia University, New York, New York

(Received September 1996; revisions received June 1997, September 1997; accepted October 1997.)

This paper addresses the simultaneous determination of pricing and inventory replenishment strategies in the face of demand uncertainty. More specifically, we analyze the following single item, periodic review model. Demands in consecutive periods are independent, but their distributions depend on the item's price in accordance with general stochastic demand functions. The price charged in any given period can be specified dynamically as a function of the state of the system. A replenishment order may be placed at the beginning of some or all of the periods. Stockouts are fully backlogged. We address both finite and infinite horizon models, with the objective of maximizing total expected discounted profit or its time average value, assuming that prices can either be adjusted arbitrarily (upward or downward) or that they can only be decreased. We characterize the structure of an optimal combined pricing and inventory strategy for all of the above types of models. We also develop an efficient value iteration method to compute these optimal strategies. Finally, we report on an extensive numerical study that characterizes various qualitative properties of the optimal strategies and corresponding optimal profit values.

This paper addresses an important problem area in the interface between marketing and production/inventory planning—specifically, the simultaneous determination of pricing and inventory replenishment strategies in the face of demand uncertainty.

Recent developments in the area of yield and revenue management have demonstrated that major benefits can be derived by complementing a replenishment strategy with the *dynamic* adjustment of a commodity's price as a function of its prevailing inventory and the length of its remaining sales season. (See, e.g., Bitran and Mondschein 1993, 1995; Gallego and van Ryzin 1994, 1997; and Heching et al. 1999.) Conversely, a dynamic pricing strategy by itself is often insufficient to manage sales. For example, fashion items, with a short sales horizon relative to their long procurement lead times, and with correspondingly limited opportunities to adjust purchasing decisions, have traditionally been managed with a single purchase order delivered at the beginning of the season. More recently, one however observes attempts to mitigate the retailers' risk by the adoption of novel contractual arrangements between retailers and their suppliers. These arrangements, often referred to as backup arrangements, permit multiple deliveries during the season with the option of (partial) adjustments by the retailer after the first couple of weeks of the sales season. (See, e.g., Eppen and Iyer 1995, 1997; and Bassok 1994, 1995.)

More specifically, we analyze the following single item, periodic review model. Demands in consecutive periods are independent, but their distributions depend on the item's price in accordance with general stochastic demand functions. The price charged in any given period can be specified dynamically as a function of the state of the system. The company thus acts as a price setter or monopolist. Markets with perfect or limited competition can be analyzed only via much more complex game-theoretical models. A replenishment order may be placed at the beginning of some or all of the periods. Stockouts are fully backlogged. Ordering costs are proportional with order sizes, while inventory carrying and stockout costs all depend on the size of the end-of-the-period inventory level and shortfall, respectively, in accordance with given convex functions. Similarly, we assume that expected revenues in each period depend concavely on the item's price. This assumption is satisfied for many stochastic (in particular, linear) demand functions. We address both finite and infinite horizon models, with the objective of maximizing total expected discounted profit or its time average value, assuming that prices can either be adjusted arbitrarily (upward or downward) or that they can only be decreased.

We characterize the structure of an optimal combined pricing and inventory strategy for all of the above types of models. We also develop an efficient value iteration method to compute these optimal strategies. Finally, we report on an extensive numerical study which characterizes various qualitative properties of the optimal strategies and corresponding optimal profit values, e.g.:

(i) the benefits associated with a dynamic pricing strategy compared to a statically determined price,
(ii) the profit loss that occurs when prices are upwardly rigid, i.e., when they can only be reduced over time, and
(iii) the impact of demand uncertainty and price elasticities.

Pricing and replenishment strategies have traditionally been determined by entirely separate units of a company's organization, without proper mechanisms to coordinate

Subject classifications: Inventory/production: uncertainty, stochastic; operating characteristics; planning horizons. Marketing: pricing.
Area of review: Stochastic Models.

0030-364X/99/4703-0454 $05.00

these two planning areas. Current reengineering efforts, however, are geared towards the systematic elimination of organizational barriers between distinct functional areas within the same enterprise. This trend has fostered the need for planning models such as the ones treated in this paper, and corresponding decision support systems which cross traditional functional boundaries.

The same traditional dichotomy has been characteristic of the academic literature. There exists a plethora of literature on inventory planning, but it assumes, almost invariably, that the *demand* processes are *exogenously* determined, and therefore uncontrollable. In practice, a demand process can often be controlled by varying the price structure. The implication of an exogenous demand process, therefore, is that the price structure is exogenously determined as well. The more limited literature on pricing strategies, on the other hand, assumes by and large that the *supply* processes, i.e., the timing and size of purchases or production runs, are either entirely prespecified as exogenous input parameters or at best to be determined in a *static* manner.

More specifically, standard (single item) inventory models assume that the price to be charged, and hence the demand distribution pertaining to each period, is exogenously specified. Since expected revenues are constant under this assumption, these models focus on the minimization of expected operating costs. (See Porteus 1990 or Lee and Nahmias 1993 for recent surveys of this literature.) The literature on dynamic pricing strategies assumes by and large that one of the following situations prevails: (i) with the exception of an initial procurement at the beginning of the planning horizon, no subsequent replenishments can occur; (ii) no inventories can be carried from one period to the next, effectively decomposing the supply decisions on a period by period basis. As far as the former are concerned, we refer to Bitran and Mondschein (1993, 1995), Gallego and van Ryzin (1994 and 1997), Heching et al. (1999), and the references mentioned therein. The literature on the latter type of models focuses on adaptive learning regarding one or more of the parameters in the demand function (see Rothschild 1974, Grossman et al. 1977, McLennan 1984, Balvers and Cosimano 1990, and Braden and Oren 1994).

The need to integrate inventory control and pricing strategies was first propagated by Whitin (1955), in the embryonic days of inventory theory. Both Whitin (1955) and later Mills (1959, 1962) addressed the *single* period version of the model; here only a single price and supply quantity need to be determined. Subsequent work by Karlin and Carr (1962), Zabel (1970), Young (1978), Polatoglu (1991), Hempenius (1970), and Lau and Lau (1988) revisited the same single period model under alternative specifications of the (stochastic) demand function. Karlin and Carr (1962) also considered the infinite horizon version of the model; however, they did so under the assumption that a single constant price is to be specified at the beginning of the planning horizon.

The first treatments of dynamic combined pricing and inventory strategies (i.e., in a multiperiod setting) were undertaken under the assumption of deterministic demands. Thomas (1974) and Kunreuther and Schrage (1973) develop variants of the Wagner-Whitin (1958) dynamic lot sizing algorithm for settings where the demands can be controlled by selecting appropriate price levels. Rajan et al. (1992) analyze a continuous time version of the same model. See Eliashberg and Steinberg (1991) for a recent survey of integrated joint marketing-production decision models. Under demand uncertainty, the only existing results appear to be due to Zabel (1972) and Thowsen (1975); the former confined himself to a special class of stochastic demand functions where a price independent uniform or exponential distribution is added to or multiplied by a deterministic demand function. Thowsen (1975) extended Zabel's (1972) results for the case of an *additive* random term, to somewhat more general conditions that he admits "do not have any straightforward economic interpretation" and "will in some cases be difficult to verify." Thomas (1974) proposes a heuristic strategy for the multiperiod model. Amihud and Mendelson (1983) analyze the optimality equation that arises in the infinite horizon discounted profit model with bi-directional price changes, linear holding and backlogging costs, and additive error terms in the demand function. The objective of this paper is to demonstrate that price reactions to inventory changes are milder than what might be anticipated by the shape of the demand function. Li (1988) develops combined pricing and inventory strategies for a continuous-time model in which cumulative production and cumulative sales are both represented as (nonhomogeneous) Poisson processes with controllable intensities. The intensity of the demand process is controlled by varying the item's price.

The remainder of this paper is organized as follows. In §1 we introduce the basic model and its notation. In the remainder we systematically distinguish between the case where prices can be adjusted arbitrarily and settings where only markdowns are permitted. We refer to the former as the case of "bi-directional price changes" and to the latter as the "markdowns only" case. In §2 we characterize the optimal policy for a general finite planning horizon. Section 3 addresses the *infinite horizon discounted* profit model and characterizes its asymptotic behavior as the discount factor approaches 1. These results are used in §4 to characterize optimal policies for the long run average profit criterion. In §5 we discuss efficient methods to compute the optimal policies in the models addressed in §§ 2 through 4. Section 6 briefly covers a number of important extensions of the basic model. Specifically, we consider the impact of order leadtimes and upper limits on the maximum allowable price change or order size in any given period. In addition, we consider the case where stockouts are filled by emergency procurements at the end of the period in which they occur. Section 7 concludes our paper with an extensive numerical study evaluating scenarios based on actual sales data obtained from a major nationwide women's apparel retailer.

1. THE BASIC MODEL

In this section we specify the basic model and its notation. We consider a single item whose inventory and selling price are reviewed periodically. At the beginning of each period a simultaneous decision is made regarding the size of a new replenishment order (if any) as well as whether the price of the item is to be modified, and if so by what magnitude. We initially assume that replenishment orders become available instantaneously; see §6 for a treatment of positive lead times. Time-dependent upper limits may apply with respect to the order size and magnitude of price change in any given period. (Arrangements such as backup agreements for fashion items, where inventory replenishments can occur only at the beginning of the season as well as at a limited (prespecified) set of subsequent periods, can be modelled by setting the order size upper limits equal to zero in all periods excluding the prespecified replenishment periods.) We initially assume that no limits prevail with respect to replenishment order sizes or price changes; see §6 for a treatment of the model incorporating such limits. In case demand in a given period exceeds the available inventory, excess demand is (fully) backlogged. See §6 for a treatment of alternative assumptions, e.g., where stockouts are satisfied through emergency procurements at the end of the period in which they occur.

In models with a finite planning horizon, we index each period by the number of periods remaining until the end of the horizon.

Demands in consecutive periods are independent and nonnegative; demand in period t depends on the prevailing price according to a given general stochastic demand function:

$$D_t = d_t(p_t, \epsilon_t), \tag{1}$$

where

p_t = price charged in period t,
ϵ_t = random term with known distribution.

The set of feasible price levels is confined to the finite interval $[p_{\min}, p_{\max}]$ where

$p_{\min}$ = lowest possible unit price to be charged,
$p_{\max}$ = highest possible unit price to be charged.

An important special case of such stochastic demand functions arises when D_t is of the form:

$$D_t = \gamma_t(p)\epsilon_t + \delta_t(p), \tag{2}$$

with $\gamma(\cdot)$ and $\delta(\cdot)$ nonincreasing functions. (The cases of $\gamma_t(p) = 1$ and $\delta_t(p) = 0$ are often referred to as the *additive* and *multiplicative* model, respectively.) We assume that the demand function in each period t is nonincreasing and concave in the period's price and that expected demand is finite and strictly decreasing in the price:

Assumption 1. *For all* $t = 1, 2, \ldots$ *(i) the function* $d_t(p, \epsilon_t)$ *is nonincreasing and concave in* $p \in [p_{\min}, p_{\max}]$ *and (ii)* expected *demand* $\mathsf{E}d_t(p, \epsilon_t)$ *is finite and strictly decreasing in* p.

If the stochastic demand function is of the type given by (2), Assumption 1(i) is satisfied when $\gamma_t(p)$ and $\delta_t(p)$ are concave and nonincreasing functions of p, and at least one of these two functions is strictly decreasing in p. The latter holds, e.g., in the important special case where $\gamma_t(\cdot)$ or $\delta_t(\cdot)$ are linear (decreasing) functions of p, or when $\gamma_t(\cdot)$ or $\delta_t(\cdot)$ are power functions of the form, i.e., $\gamma_t(p)$ (or $\delta_t(p)$) $= c - kp^\nu$, with $\nu \geqslant 1$ for some positive constants $c, k > 0$.

Monotonicity of the demand functions is satisfied for all regular items; only special luxury items exhibiting the Veblen paradox are excluded. Thus, only the concavity assumption comes with some more significant loss of generality; it implies that the marginal *absolute* decrease in demand volume does not decrease as the price level is increased.

Rewards and costs in future periods are discounted with a discount factor $\alpha \leqslant 1$.

Let:

x_t = inventory level at the beginning of period t, *before* ordering,
y_t = inventory level at the beginning of period t, *after* ordering.

Two types of costs are incurred: end-of-the-period inventory carrying (and backlogging) costs and variable order costs. These are specified by:

$h_t(I)$ = inventory (or backlogging) cost incurred in a period whose ending inventory level equals I,
c_t = per unit purchase or production cost in period t.

Let

$$G_t(y, p) = \mathsf{E}h_t(y - D_t) = \mathsf{E}h_t(y - d_t(p, \epsilon_t)) \tag{3}$$

denote one-period expected inventory and backlogging costs for period t, $t = 1, 2, \ldots$, where the expectation here, as well as in the remainder of the paper, is taken over the distribution of the ϵ_t variables. We make the following assumptions regarding the functions G_t, their growth rate, and the finiteness of the moments of the demand distribution.

Assumption 2. $\lim_{y\to\infty} G_t(y, p) = \lim_{y\to-\infty}[c_t y + G_t(y, p)] = \lim_{y\to\infty}[(c_t - \alpha c_{t-1})y + G_t(y, p)] = \infty$ *for all* $p \in [p_{\min}, p_{\max}]$.

Assumption 3. $0 \leqslant G_t(y, p) = O(|y|^\rho)$ *for some integer* ρ.

Assumption 4. $\mathsf{E}[d_t(p, \epsilon_t)]^\rho < \infty$ *for all* $p \in [p_{\min}, p_{\max}]$.

Assumption 2 holds whenever the inventory (and backlogging) cost function h_t tends to infinity as the inventory level (or backlog size) increases to infinity; the latter applies to any reasonable inventory cost structure in which the loss associated with a stockout exceeds the unit's purchase price. Assumption 3 holds whenever the inventory cost functions $\{h_t\}$ are polynomially bounded, which is satisfied under all common cost structures. Finally, Assumption 4 often ensures that the G_t functions are well defined and finite. In addition, we shall assume h_t is convex and that the functions $G_t(y, p)$ are jointly convex in y and p:

Assumption 5. *For all $t = 1, \ldots, T$ h_t is convex and $G_t(y, p)$ is jointly convex.*

The following lemma shows that Assumption 5 is satisfied, for example, when the functions h_t are convex and the demand functions d_t are linear in p. (The former is satisfied under all common cost structures.)

Lemma 1. *Fix $t = 1, 2, \ldots$. Assume h_t is convex and the demand function d_t is linear in p. Then $G_t(y, p)$ is jointly convex in y and p.*

Proof. By the convexity of h_t and the linearity of the demand functions in p

$$
\begin{aligned}
h_t&\left(\frac{y_1 + y_2}{2} - d_t\left(\frac{p_1 + p_2}{2}, \epsilon_t\right)\right) \\
&= h_t\left(\tfrac{1}{2}[y_1 - d_t(p_1, \epsilon_t)] + \tfrac{1}{2}[y_2 - d_t(p_2, \epsilon_t)]\right) \\
&\leqslant \tfrac{1}{2} h_t(y_1 - d_t(p_1, \epsilon_t)) + \tfrac{1}{2} h_t(y_2 - d_t(p_2, \epsilon_t)),
\end{aligned}
$$

so that the functions $h_t(y, p, \epsilon_t)$ are jointly convex in (y, p). We conclude that the function $G_t(y, p) = \mathsf{E}_{\epsilon_t} h_t(y, p, \epsilon_t)$ is jointly convex in (y, p) as well. $\square$

Remark. The above representation of the one step expected inventory and backlogging cost functions $G_t(y, p)$, via (3), assumes implicitly that the functions $h_t(\cdot)$ are independent of the sales price p_t. Sometimes, the holding and backlogging costs associated with a given end-of-the-period inventory level may depend on the prevailing sales price, e.g., when $h_t(I) = h_t^+(p)I^+ + h_t^-(p)I^-$ with $h_t^+(p)$ and $h_t^-(p)$ given, say linear functions of p (i.e., $h_t^+(p) = a_t^+ + b_t^+ p$ and $h_t^-(p) = a_t^- + b_t^- p$). Such generalizations are easily incorporated as long as the resulting functions $G_t(\,\cdot\,,\,\cdot\,)$ continue to satisfy Assumptions 2, 3, and 5. Assumptions 2 and 3 invariably continue to hold; Assumption 5 is more restrictive but in the above example it continues to hold when b_t^+ and b_t^- are sufficiently small.

Finally, with respect to the timing of cash flows, we assume that revenues are received at the *end* of the period in which the sales occur. Further, all costs associated with a period must be paid at its *beginning*. Correspondingly, we assume that the price selected in any given period is always at least as large as the unit's replacement value, or variable order cost, in the next period, i.e.,

$$
p_t \geqslant c_{t-1}. \tag{4}
$$

(In most practical settings we have $c_t \leqslant p_{\min}$ for all $t = 1, 2, \ldots$ so that (4) is trivially satisfied.)

2. THE FINITE HORIZON PROBLEM

In this section we characterize the structure of a strategy maximizing expected discounted profit, under a given discount factor $\alpha < 1$. The planning horizon consists of T periods, numbered $T, T - 1, \ldots, 1$. For products sold over a specific sales season (e.g., fashion items or products with short life-cycles), T is naturally chosen to coincide with the length of the sales season. Other products, which are expected to be marketed over a long, indefinite length of time, require that T be chosen large enough to ensure that the computed optimal decisions pertaining to the first or an initial set of periods remain optimal under longer planning horizons.

Finite planning horizon models allow for arbitrary nonstationarities in the cost and revenue parameters as well as the demand functions. As mentioned above, we give separate treatment to the case where the price can be increased as well as decreased, and that where only markdowns are permitted.

2.1. Bi-directional Price Changes

If the price can be changed arbitrarily from period to period, the problem can be formulated as a Markov Decision Problem (MDP) with x_t as the state of the system at the beginning of period t. Thus, $S = \Re$ represents the state space. Let $v_t^*(x)$ denote maximum expected discounted profit for periods $1, 2, \ldots, t$ when starting period t in state x. The functions v_t^* satisfy $v_0^* \equiv 0$ and for $t = 1, 2, \ldots$

$$
v_t^*(x) = c_t x + \max_{\{y \geqslant x,\, \max(p_{\min}, c_{t-1}) \leqslant p \leqslant p_{\max}\}} J_t(y, p), \tag{5}
$$

where

$$
\begin{aligned}
J_t(y, p) = {}& \alpha p \mathsf{E} d_t(p, \epsilon_t) - c_t y - G_t(y, p) \\
&+ \alpha \mathsf{E} v_{t-1}^*(y - d_t(p, \epsilon_t)).
\end{aligned} \tag{6}
$$

We show that an optimal strategy employs a so-called base stock list price policy in each period. A base stock list price policy is characterized by a base stock level and list price combination, (y_t^*, p_t^*). If the inventory level is below the base stock level, it is increased to the base stock level and the list price is charged. If the inventory level is above the base stock level, then nothing is ordered, and a price discount is offered. In addition, the higher the excess in the initial inventory level, the larger the optimal discount offered. That is, the optimal price is a nonincreasing function of the initial inventory level, and no discounts are offered unless the product is overstocked. The term "base stock list price" policy was coined by Porteus (1990). It was first shown to be optimal in the special models considered by Thowsen (1975) and later by Young (1978).

First, let $V_t^*(x)$ denote the second term to the right of (5), i.e., $V_t^*(x) = v_t^*(x) - c_t x$ and set $V_0^* \doteq 0$. Rewriting (5) and (6) in terms of the functions J_t and V_t^* we obtain:

$$
V_t^*(x) = \max_{\{y \geqslant x,\, \max(p_{\min}, c_{t-1}) \leqslant p \leqslant p_{\max}\}} J_t(y, p), \tag{7}
$$

$$\begin{aligned} J_t(y, p) &= \alpha p \mathsf{E} d_t(p, \epsilon_t) - c_t y - G_t(y, p) \\ &\quad + \alpha c_{t-1}(y - d_t(p, \epsilon_t)) + \alpha \mathsf{E} V^*_{t-1}(y - d_t(p, \epsilon_t)) \\ &= \alpha(p - c_{t-1}) \mathsf{E} d_t(p, \epsilon_t) + (\alpha c_{t-1} - c_t) y \\ &\quad - G_t(y, p) + \alpha \mathsf{E} V^*_{t-1}(y - d_t(p, \epsilon_t)). \end{aligned} \tag{8}$$

Theorem 1. (a) *Fix $t = 1, \ldots, T$. The function $J_t(y, p)$ is jointly concave in y and p and the function $V^*_t(x)$ is concave and nonincreasing in x.*

(b) *Fix $t = 1, \ldots, T$. $J_t(y, p) = O(|y|^\rho)$ and $V^*_t(x) = O(|x|^\rho)$. $J_t(y, p)$ has a finite maximizer for all $t \geq 1$, denoted by (y^*_t, p^*_t). (In case of multiple maxima, select (y^*_t, p^*_t) to be the lexicographically largest.)*

(c) *If $x \leq y^*_t$ it is optimal to order up to the base stock level y^*_t and to charge the list price p^*_t; if $x > y^*_t$, it is optimal not to order.*

Proof. (a) By induction: Clearly, $J_1(\cdot, \cdot)$ is jointly concave: To verify joint concavity for the first term to the right of (8), fix a value for ϵ_t. Since the function $d_t(p, \epsilon_t)$ is concave in p (see Assumption 1), it follows that it possesses first and second order right and left derivatives. By straightforward calculus, one thus verifies that the function $(p - c_{t-1}) d_t(p, \epsilon_t)$ has nonpositive second order (left and right) derivatives for $p > c_{t-1}$, see (4), so that this function is concave as well. The same therefore applies to $\mathsf{E}_{\epsilon_t}(p - c_{t-1}) d_t(p, \epsilon_t)$. The second term to the right of (8) is linear in y while the third term is jointly concave in view of Assumption 5. Thus $V^*_1(x)$ is easily verified to be concave as well, and it is clearly non-increasing. Assume now that $J_{t-1}(\cdot, \cdot)$ is jointly concave for some $t = 2, \ldots, T - 1$ and that $V^*_{t-1}(\cdot)$ is also concave and nonincreasing. Then, $J_t(y, p)$ is jointly concave: joint concavity of the first three terms to the right of (8) is verified as for the case $t = 1$, above; to verify joint concavity of the last term in (8), note that for any given value of ϵ_t, $V^*_{t-1}(y - d_t(p, \epsilon_t))$ is jointly concave in y and p: For any pair of points (y_1, y_2) and (p_1, p_2), note by Assumption 1 that

$$d_t\left(\frac{p_1 + p_2}{2}, \epsilon_t\right) \geq \frac{1}{2} d_t(p_1, \epsilon_t) + \frac{1}{2} d_t(p_2, \epsilon_t).$$

Since the function $V^*_{t-1}(\cdot)$ is nonincreasing we have:

$$\begin{aligned} & V^*_{t-1}\left(\frac{y_1 + y_2}{2} - d_t\left(\frac{p_1 + p_2}{2}, \epsilon_t\right)\right) \\ &\quad \geq V^*_{t-1}\left(\frac{y_1 + y_2}{2} - \frac{1}{2} d_t(p_1, \epsilon_t) - \frac{1}{2} d_t(p_2, \epsilon_t)\right) \\ &\quad = V^*_{t-1}\left(\frac{1}{2}[y_1 - d_t(p_1, \epsilon_t)] + \frac{1}{2}[y_2 - d_t(p_2, \epsilon_t)]\right) \\ &\quad \geq \frac{1}{2} V^*_{t-1}(y_1 - d_t(p_1, \epsilon_t)) + \frac{1}{2} V^*_{t-1}(y_2 - d_t(p_2, \epsilon_t)) \end{aligned}$$

by the concavity of V^*_{t-1}. This implies that $\mathsf{E} V^*_{t-1}(y - d_t(p, \epsilon_t))$ is jointly concave in y and p as well. The concavity and monotonicity of V^*_t is immediate from (7).

(b) By induction: $J_1(y, p) = O(|y|^\rho)$ by Assumption 3. Also, $J_1(\cdot, \cdot)$ is jointly concave, and for all $p \in [p_{\min}, p_{\max}]$, $\lim_{|y| \to \infty} J_1(y, p) = -\infty$ by Assumption 2. This implies that J_1 has a finite maximizer. Assume now that for some $t = 2, \ldots, T$ $J_{t-1}(\cdot, \cdot) = O(|y|^\rho)$ and that J_{t-1} has a finite maximizer (y^*_{t-1}, p^*_{t-1}). It is easily verified that $V^*_{t-1}(x) = O(|x|^\rho)$ i.e. a constant $K > 0$ exists such that $V^*_{t-1}(x) \leq K(|x|^\rho + 1)$ for all x. (For $x \leq y^*_{t-1}$, the maximum in (7) is achieved in the point (y^*_{t-1}, p^*_{t-1}) while for $x > y^*_{t-1}$, $y = x$ achieves the maximum, see the proof of part (c).)

Thus

$$\begin{aligned} V^*_{t-1}(y - d_t(p, \epsilon_t)) &\leq K|y - d_t(p, \epsilon_t)|^\rho + K \\ &\leq K[|y| + d_t(p, \epsilon_t)]^\rho + K, \end{aligned} \tag{9}$$

and hence employing the Binomial expansion of the right-hand side of (9) and Assumption 4,

$$\begin{aligned} & \mathsf{E} V^*_{t-1}(y - d_t(p, \epsilon_t)) \\ &\quad \leq K \mathsf{E}\{[|y| + d_t(p, \epsilon_t)]^\rho + 1\} \\ &\quad \leq K + K \sum_{l=0}^{\rho} \binom{\rho}{l} |y|^\rho \max_{p_{\min} \leq p \leq p_{\max}} \mathsf{E} d_t^{\rho - l}(p, \epsilon_t). \end{aligned}$$

We conclude by Assumption 3 that $J_t(y, p_t) = O(|y|^\rho)$ as well. Since $J_{t-1}(\cdot, \cdot)$ has a finite maximizer, (y^*_{t-1}, p^*_{t-1}), we have for all y that $\mathsf{E} V^*_{t-1}(I(y, d_t(p, \epsilon_t))) \leq V^*_{t-1}(y^*_{t-1})$. It thus follows from (8), and Assumption 2, that $\lim_{|y| \to +\infty} J_t(y, p) = -\infty$ for all $p \in [p_{\min}, p_{\max}]$. Combined with the fact that J_t is jointly concave it follows that it has a finite maximizer as well, thus completing the induction step.

(c) Fix $t = 1, \ldots, T$. Since $J_t(\cdot, \cdot)$ is jointly concave, (y^*_t, p^*_t) is the optimal decision pair when $x \leq y^*_t$. Similarly, for $x > y^*_t$ it is optimal to choose $y = x$. (If for $x > y^*_t$, a decision pair (y, p') is chosen with $y > x$, then for the pair (x, p'') on the line connecting (y^*_t, p^*_t) with (y, p'), $J_t(x, p'') \geq J_t(y, p')$.) We conclude in particular that

$$y(x) \text{ is nondecreasing in } x. \quad \square \tag{10}$$

To prove that a base stock list price policy is optimal, it thus suffices to show that the optimal price to be selected in any given period is nonincreasing in the prevailing inventory level. In other words, under a higher starting inventory level, a price is selected that is no larger than the optimal price under a lower starting inventory, this to stimulate demand and promote a (larger) inventory reduction.

The proof of Theorem 1 shows that optimality of the ordering rule prescribed by the base stock list price policy depends on the joint concavity of the functions $J_t(y, p)$, $t = 1, \ldots, T$. Similarly, monotonicity of the optimal price level $p^*(x)$ in the prevailing inventory level, depends on a different property of the function J_t, namely, its *submodularity*. A function $f: \Re^2 \to \Re$ is said to be submodular (supermodular) if it has *antitone* (*isotone*) *differences*, i.e., for all $y_1 > y_2$ the difference functions $f(y_1, p) - f(y_2, p)$ are nonincreasing (nondecreasing) in p (see Topkis 1978 or Heyman and Sobel 1984).

Theorem 2. *Fix $t = 1, 2, \ldots, T$.*

(a) *The optimal price $p^*(x)$ is nonincreasing in x with $p^*(x) \leq p^*_t$, the list price for period t.*

(b) *A base stock list price policy with base stock* y_t^* *and list price* p_t^* *is optimal.*

Proof. (a) We first show that $J_t(y, p)$ is submodular. Since the sum of submodular functions is submodular, it suffices to establish submodularity for each term to the right of (8). The first and second terms are trivially submodular since they depend on only one of the two variables y, p. To show that $G_t(y, p)$ has isotone differences, fix ϵ_t and consider an arbitrary pair of inventory levels (y_1, y_2) and any pair of price levels (p_1, p_2) with $y_1 > y_2$ and $p_1 > p_2$. Let $\iota_1 = y_1 - d_t(p_1, \epsilon_t)$, $\iota_2 = y_1 - d_t(p_2, \epsilon_t)$, $\iota_3 = y_2 - d_t(p_1, \epsilon_t)$, and $\iota_4 = y_2 - d_t(p_2, \epsilon_t)$. By the monotonicity of the demand function d_t we have:

$$\iota_3 > \iota_4.$$

Thus, by the convexity of the h_t function we have:

$$\begin{aligned} h_t(\iota_1) - h_t(\iota_3) &= h_t(\iota_3 + (y_1 - y_2)) - h_t(\iota_3) \\ &\geqslant h_t(\iota_4 + y_1 - y_2) - h_t(\iota_4) \\ &= h_t(\iota_2) - h_2(\iota_4). \end{aligned}$$

We conclude that the function $h_t(y - d_t(p, \epsilon_t))$ has isotone differences in y and p, and the same supermodularity therefore applies to the function $G_t(y, p) = \mathsf{E}_{\epsilon_t} h_t(y - d_t(p, \epsilon_t))$. Finally, the submodularity proof for the last term in (8) is identical to that of $-G_t$ since V_{t-1}^* is concave.

The decision problem in period t can be viewed as consisting of two stages. In the first stage, the inventory level (after ordering) y is chosen and in the second stage the corresponding price p. The second stage problem thus has $S = \Re$ as its state space and $A = [\max(p_{\min}, c_{t-1}), p_{\max}]$ as the set of feasible (price) actions in each possible state $y \in S$. Since $J_t(y, p)$ is strictly concave in p in view of Assumption 1(ii), we have that the optimal price $p(y)$ is unique. Since $J_t(y, p)$ is submodular, it follows from Theorem 8-4 in Heyman and Sobel (1984) that the optimal price p is nonincreasing in the "state" y, and hence in x given (10).

(b) Immediate from part (a) and Theorem 1(c). □

Monotonicity of the price p as a function of the starting inventory x, extends the same monotonicity result obtained in Zabel (1972), Thowsen (1975), Young (1978), and Amihud and Mendelson (1983) for their special models treated.

2.2. The Model with Markdowns

When only markdowns are allowed, the state of the system at the beginning of period t is represented by the pair (x_t, p_{t+1}), with p_{t+1} the price in effect during the previous period. Let $v_t^*(x, p)$ denote the maximum expected discounted profit for periods $1, \ldots, t$ when starting in state (x, p). The functions v_t^* satisfy $v_0^* \equiv 0$ and

$$v_t^*(x, p) = c_t x + \max_{\{y \geqslant x,\ \max(p_{\min}, c_{t-1}) \leqslant p' \leqslant p\}} J_t(y, p'), \tag{11}$$

where $J_t(y, p') = \alpha p' \mathsf{E} d_t(p', \epsilon_t) - c_t y - G_t(y, p) + \alpha \mathsf{E} v_{t-1}^*(y - d_t(p', \epsilon_t), p')$.

As before, it is convenient to rewrite the recursion (11) in terms of the function $V_t^*(x, p) \equiv v_t^*(x, p) - c_t x$:

$$V_t^*(x, p) = \max_{\{y \geqslant x,\ \max(p_{\min}, c_{t-1}) \leqslant p' \leqslant p\}} J_t(y, p'), \tag{12}$$

where

$$\begin{aligned} J_t(y, p') = {} & \alpha(p' - c_{t-1}) \mathsf{E} d_t(p', \epsilon_t) \\ & + (\alpha c_{t-1} - c_t) y - G_t(y, p) \\ & + \alpha \mathsf{E} V_{t-1}^*(y - d_t(p', \epsilon_t), p'). \end{aligned} \tag{13}$$

In close analogy to the proofs of Theorems 1 and 2 we establish Theorem 3.

Theorem 3. *Fix* $t = 1, \ldots, T$.

(a) *The functions* $J_t(y, p')$ *and* $V_t^*(x, p)$ *are jointly concave in* (y, p') *and* (x, p) *respectively.* $V_t^*(x, p)$ *is nonincreasing in* x *and nondecreasing in* p.

(b) $J_t(y, p') = O(|y|^\rho)$ *and* $V_t^*(x, p) = O(|x|^\rho)$. *Also,* $J_t(y, p')$ *has a finite maximizer for all* $t \geqslant 1$, *denoted by* (y_t^*, p_t^*). *(In case of multiple maxima, select* (y_t^*, p_t^*) *to be the lexicographically largest.)*

(c) *Let* (y_1, p') *and* (y_2, p') *denote two pairs of inventory and price levels optimally selected in period* t *if period* t *starts in* (x_1, p_1) *and* (x_2, p_2), *respectively. If* $y_1 > y_2$ *then* $p_1 \leqslant p_2$.

Proof. See the Appendix.

We now characterize the structure of an optimal strategy in any given period $t = 1, \ldots, T$. Let $\hat{y}_t(p_{t+1})$ denote the (largest) maximizer of the function $J_t(\cdot, p_{t+1})$. (A maximizer exists since J_t is concave, see Theorem 3, and since $\lim_{|y| \to \infty} J_t(y, p_{t+1}) = -\infty$, see the proof of Theorem 1(b).) Note that $\hat{y}_t(p_{t+1}) = y_t^*$ if $p_{t+1} \geqslant p_t^*$ and that $\hat{y}_t(\cdot)$ is a nonincreasing function, see Theorem 1(c), a property which can be exploited in the computation of the optimal levels $\{\hat{y}_t(p): \max(p_{\min}, c_{t+1}) \leqslant p \leqslant p_{\max}\}$. We distinguish between two cases:

Case I: $p_{t+1} \geqslant p_t^*$: Apply the base stock list price rule described in §2.1.

Case II: $p_{t+1} < p_t^*$: If $x \leqslant \hat{y}_t(p_{t+1})$, choose the pair $(\hat{y}_t(p_{t+1}), p_{t+1})$. If $x > \hat{y}_t(p_{t+1})$, no order is placed, i.e., $y = x$ and the price is set at a value $p^*(x)$ that satisfies $p^*(x) \leqslant p_{t+1}$ and with $p^*(x)$ nonincreasing in x.

We continue to refer to the above rule as the base stock list price rule (with price-dependent order up to levels $\{\hat{y}_t(p_{t+1})\}$).

Theorem 4. *The base stock list price policy with order up to levels* $\{\hat{y}_t(p_{t+1})\}$ *is an optimal policy in period* t.

Proof. *Case I:* Consider the relaxation of (12) obtained by relaxing the constraint $p' \leqslant p_{t+1}$. Note that the base stock list price rule described in §2.1 is optimal for this relaxed problem and that for each inventory level x, a price $p' \leqslant p_t^* \leqslant p_{t+1}$ is chosen. In other words, the base stock list price rule is feasible and hence is optimal for the original problem (12) as well.

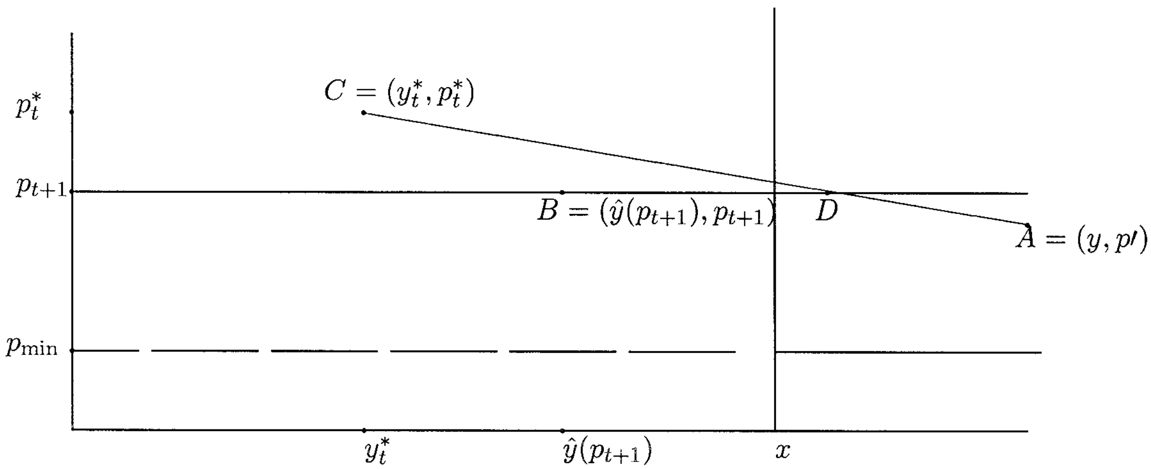

Figure 1. Proof of Theorem 3, Case II.

Case II, $x \leqslant \hat{y}(p_{t+1})$: By the definition of $\hat{y}(p_{t+1})$, the pair $(\hat{y}(p_{t+1}), p_{t+1})$ is best among all pairs in $\{(y, p_{t+1}): y \geqslant x\}$. Assume to the contrary that some pair $A = (y, p')$ with $y \geqslant x$ and $p' < p_{t+1}$ is strictly superior to $B = (\hat{y}(p_{t+1}), p_{t+1})$. Let $C = (y_t^*, p_t^*)$ and $D = (y^0, p_{t+1})$, the point of intersection of the line through A and C with the horizontal line $p' = p_{t+1}$. (Since $p_t^* > p_{t+1} > p'$ this point of intersection exists and $y^0 > 0$.) By the definition of $\hat{y}(p_{t+1})$ and the joint concavity of J_t, $J_t(B) \geqslant J_t(D) \geqslant J_t(A)$, which contradicts the strict superiority of (y, p').

Case II, $x > \hat{y}(p_{t+1})$: We first show that a pair (x, p') with $p' \leqslant p_{t+1}$ is optimal. (See Figure 1.) Note that the open rectangle bordered by the lines $p = p_{t+1}$, $p = p_{\min}$ and $y \geqslant x$ represents the region of feasible pairs and let A denote an optimal pair in this rectangle. Again, let $C = (y_t^*, p_t^*)$. Similar to the previous case, let D denote the point of intersection of the line through A and C with the boundary of the rectangle; by the concavity of J_t, $J_t(D) \geqslant J_t(A)$. By Theorem 1(c) we have that $y_t^* \leqslant \hat{y}(p_{t+1}) < x$. Therefore D lies on the vertical line $y = x$ or on the horizontal line $p = p_{t+1}$. In the first case, our claim is proven; in the latter, we have again by the concavity of $J_t(\cdot, p_{t+1})$ that $J_t(x, p_{t+1}) \geqslant J_t(D) \geqslant J_t(A)$, i.e., (x, p_{t+1}) is an optimal pair as well. It remains to be shown that the new price $p' \leqslant p_{t+1}$ is nonincreasing in $x \geqslant \hat{y}(p_{t+1})$. This can be established in close analogy to the proof of Theorem 2. □

3. THE INFINITE HORIZON DISCOUNTED PROBLEM

In this section, we consider an *infinite* planning horizon with stationary cost and revenue parameters as well as demand distributions. As discussed at the beginning of §2, this model is often suitable for basic goods with relatively long product life-cycles. In view of the stationarity of the model, we write $c_t = c$, $G_t = G$ and $d_t = d$ for all $t = 1, 2, \ldots$, while $\epsilon_1, \epsilon_2, \ldots$ are identically distributed as a random variable ϵ. In analyzing infinite horizon models, it is often useful to have one step expected net profits that are uniformly of the same sign. To achieve this, we subtract a constant $M = \max_{p_{\min} \leqslant p \leqslant p_{\max}} \alpha p \mathsf{E} d(p, \epsilon)$ uniformly from the one step expected profits. ($M < \infty$ since by Assumption 1 it is the maximum of a continuous function on a compact set.) We thus obtain shifted value functions $\hat{v}_t$ and $\hat{J}_t$ with

$$\hat{v}_t = v_t^* - \frac{M(1-\alpha^{t+1})}{1-\alpha} \quad \text{and} \quad \hat{J}_t = J_t - \frac{M(1-\alpha^{t+1})}{1-\alpha}.$$

3.1. Bi-directional Price Changes

When prices can fluctuate in both directions, the infinite horizon optimality equation (for the transformed model) is given by:

$$v^*(x) = cx + \max_{\{y \geqslant x,\ \max(p_{\min}, c) \leqslant p \leqslant p_{\max}\}} J(y, p), \tag{14}$$

where

$$J(y, p) = \alpha p \mathsf{E} d(p, \epsilon) - cy - M - G(y, p) + \alpha \mathsf{E} v(y - d(p, \epsilon)). \tag{15}$$

The following theorem describes the structure of an optimal policy in the infinite horizon model, and its relationship to that of the finite horizon models.

Theorem 5. *Assume prices can be changed in both directions.*

(a) $\hat{v} \doteq \lim_{t\to\infty} \hat{v}_t$, $v^* \doteq \lim_{t\to\infty} v_t^*$, $\hat{J} \doteq \lim_{t\to\infty} \hat{J}_t$ *and* $J^* \doteq \lim_{t\to\infty} J_t$ *all exist. Moreover,* $\hat{v} = v^* - M/(1-\alpha)$ *and* $\hat{J} = J^* - M/(1-\alpha)$ *and* $\hat{v}$ *and* v^* *equal the maximum infinite horizon discounted profit vector in the transformed and original models, respectively.*

(b) $\hat{v}$ *and* $\hat{J}$ (v^* *and* J^*) *satisfy the infinite horizon optimality equation* (15) *in the transformed* (*original*) *model.*

(c) $J^*(y, p)$ *and* $v^*(x)$ *are jointly concave in y and p and concave in x, respectively. Moreover,* $J^*(y, p)$ *has antitone differences and a finite maximizer* (y^*, p^*). *Finally,* $v^*(x) = O(|x|^{\rho+1})$.

(d) *A base stock list price policy with base stock level and list price combination* (y^*, p^*) *is optimal for the infinite horizon model.*

(e) *The sequence* $\{(y_t^*, p_t^*)\}$ *has at least one limit point* *and each such limit point* (y^*, p^*) *is a base stock/list price combination in an optimal base stock list price policy for the infinite horizon model. Moreover, there exist constants* $\underline{y} \leqslant 0 \leqslant \bar{y}$, *independent of* α, *such that for any optimal base stock/list price combination* (y^*, p^*), $y^* \in [\underline{y}, \bar{y}]$ *for* all $\alpha < 1$.

Proof. See the Appendix.

3.2. The Model with Markdowns

If only markdowns are permitted, the infinite horizon optimality Equation (14), for the transformed model, is replaced by:

$$v(x, p) = cx + \max_{\{x \leqslant y,\ \max(p_{\min}, c) \leqslant p' \leqslant p\}} J(y, p'), \tag{16}$$

where

$$J(y, p') = \alpha p' \mathsf{E} d(p', \epsilon) - cy - M - G(y, p') + \alpha \mathsf{E} v(y - d(p', \epsilon), p'). \tag{17}$$

The results of Theorem 5 easily extend to this model in which only markdowns are permitted. We thus obtain Theorem 6.

Theorem 6. *Assume that only markdowns are permitted. The results of Theorem 5 all hold for this model as well, replacing in part* (*c*) $v^*(x)$ *by* $v^*(x, p)$.

3.3. Asymptotic Behavior as the Discount Factor Approaches One

We conclude this section with a brief discussion of the behavior of the optimal policy and infinite horizon expected profit function v^*, as the discount factor α increases to one. To emphasize the dependency on α, we write (y^*_α, p^*_α) and v^*_α for (y^*, p^*) and v^*, respectively.

It is known that for Markov Decision Processes with *finite* state and action sets, a so-called Blackwell-optimal policy exists, i.e., the *same* policy is optimal for *all* sufficiently large discount factors α, and that the minimum cost function $v^*_\alpha = O((1-\alpha)^{-1})$ as $\alpha \to 1$. In our model, $v^*_\alpha = O((1-\alpha)^{-1})$ as $\alpha \to 1$ continues to apply. This follows from the inequalities:

$$-w^*(x) - \frac{\bar{g}}{(1-\alpha)} \leqslant v^*_\alpha \leqslant \frac{M}{(1-\alpha)}, \tag{18}$$

for a given function $w^*(x)$ and constant $\bar{g}$ that are independent of α. The upper bound in (18) follows from the fact that it represents the maximum present value of profits when all negative (cost) components are ignored. The lower bound follows from $v^*_\alpha \geqslant v^*_\alpha(x|p_{\min})$ (see the proof of Theorem 4(c)) and for this constant price model, $v^*_\alpha(\cdot|p_{\min}) \leqslant -w^*(x) - \bar{g}/(1-\alpha)$ (for appropriate choices of w^* and $\bar{g}$). (See, e.g., Aviv and Federgruen 1997.)

Let (y^*_α, p^*_α) denote the (lexicographically largest) maximizer of the function $J^*_\alpha(y, p)$. Theorem 5(d) establishes that a base stock list price policy with base stock/list price combination (y^*_α, p^*_α) is optimal for the infinite horizon model with discount factor $\alpha < 1$. We now show that the points $\{(y^*_\alpha, p^*_\alpha): 0 \leqslant \alpha < 1\}$ are contained in a closed rectangle.

Proposition 1. *There exist constants* $\underline{y}$ *and* $\bar{y}$, *independent of* α, *such that* $\underline{y} \leqslant y^*_\alpha \leqslant \bar{y}$ *for all* α *sufficiently close to one.*

Proof. Fix $\alpha < 1$. Since a base stock list price policy with base stock/list price combination (y^*_α, p^*_α) is optimal in the infinite horizon discounted model with discount factor α,

$$v^*_\alpha(y^*_\alpha) = \frac{[(\alpha p^*_\alpha - c)\mathsf{E} d(p^*_\alpha, \epsilon) - G(y^*_\alpha, p^*_\alpha)]}{(1-\alpha)}.$$

Thus, for α sufficiently close to one we have, in view of (18)

$$-\bar{g} \leqslant (\alpha p^*_\alpha - c)\mathsf{E} d(p^*_\alpha, \epsilon) - G(y^*_\alpha, p^*_\alpha),$$

or

$$G(y^*_\alpha, p^*_\alpha) \leqslant \lambda \doteq \bar{g} + \max_{\max(p_{\min}, c) \leqslant p \leqslant p_{\max}} (p - c)\mathsf{E} d(p, \epsilon) < \infty,$$

because $\mathsf{E} d(p, \epsilon)$ is concave and hence continuous in p by Assumption 1. Note that λ is *independent* of α, and by Assumption 2 there exist bounds $\underline{y}(\lambda)$ and $\bar{y}(\lambda)$ such that for all α close to one, $\underline{y}(\lambda) \leqslant y^*_\alpha \leqslant \bar{y}(\lambda)$. □

We also assume without loss of generality that

$$v^*_\alpha(y^*_\alpha) = \max_x v^*_\alpha(x) \geqslant 0, \tag{19}$$

for α sufficiently large, to preclude the trivial case where it is optimal to terminate the business for *any* starting condition.

In the next section (Theorem 7(a)) we show in fact that $\lim_{\alpha \to 1}(1-\alpha)v^*_\alpha = g^*$ where g^* denotes the long run average net profit. Moreover, although a Blackwell-optimal policy may fail to exist, Theorem 5(e) (and Theorem 6) shows that the optimal base stock level y^* is at least bounded in α.

4. THE AVERAGE PROFIT CRITERION

In this section we address the long-run average profit criterion. As with respect to the previously discussed performance criteria, we give separate treatment to (i) the model with bi-directional price changes and (ii) the model in which only markdowns are permitted. For the former, we show that a base stock list price policy continues to be optimal. Moreover, we show how this policy relates to policies that are optimal under the expected total discounted profit criterion. For the latter model (allowing only markdowns) we show that a policy of even simpler structure is optimal, i.e., a policy which adopts a constant price and employs a simple order-up-to rule. In other words, under the long-run average profit criterion, the markdown model reduces to a standard inventory model with a *fixed*, albeit controllable, price.

4.1. Bi-directional Price Changes

For the model in which prices can be increased as well as decreased at the beginning of each period, we establish the existence of a solution to the long-run average profit optimality equation. In addition, we show that a base stock list price policy achieves this optimality equation and is optimal. As in most other (applications of) Markov Decision Problems with infinite state spaces, this is most conveniently achieved when the state space is countable. We therefore discretize both the inventory level and price variables, and assume that the (demand) distribution of ϵ is discrete as well. The long-run average profit optimality equation is then given by:

$$h(x) + g = \max_{\{(y,p): y \geq x, p_{\min} \leq p \leq p_{\max}\}} \tag{20}$$

$$\cdot \{p\mathsf{E}d(p, \epsilon) - c(y - x) - G(y, p) + \mathsf{E}h(y - d(p, \epsilon))\}.$$

Also, let $g^*(x)$ denote maximum long run average profit when starting in state x.

We maintain Assumptions 1, 2, 3, 4, and 5 while requiring the finiteness of an additional moment of the demand distribution.

Assumption 4*. $\mathsf{E}[d^{\rho+1}] < \infty$ *for all* $p \in [p_{\min}, p_{\max}]$.

Theorem 7. *Assume Assumptions* 1, 2, 3, 4*, *and* 5 *hold*.

(a) *There exists a constant* g^* *such that* $g^*(x) = g^*$ *for all initial inventory levels* x. *Moreover there exists a sequence of discount factors* $\{\alpha_n\} \to 1$ *such that* $g^* = \lim_{n\to\infty}(1 - \alpha_n)v^*_{\alpha_n}(x)$ *for all inventory levels* x.

(b) *There exists a function* $h^*: \Re \to \Re$ *with* $h^*(x) \leq 0$ *and* $h^*(x) = O(|x|^{\rho+1})$ *such that* (h^*, g^*) *satisfies the optimality Equation* (20).

(c) *There exists a sequence of discount factors* $\{\alpha_n\} \to 1$ *such that the sequence of corresponding base stock list price policies converges to a limiting policy*, *say with base stock/list price combination* (y^*, p^*), *and this limiting policy is optimal for the long-run average profit criterion*. *Moreover*, *any policy that achieves the maximum in the optimality Equation* (20) *is optimal*.

(d) $\underline{y} \leq y^* \leq \bar{y}$ *with* $\underline{y}$ *and* $\bar{y}$ *defined in Proposition* 1.

Proof. Parts (a)–(c) of our theorem follow from the theorem in Sennott (1989) by establishing that Assumptions 1, 2, and 3* therein (here referred to as S-1, S-2, and S-3*, respectively) are satisfied; in particular, there exists a sequence of discount factors $\{\alpha_n\} \to 1$, such that:

$S-1$: $-\infty < v^*_{\alpha_n}(x) < \infty$ for all x and all α_n.

$S-2$: For some state y^0 let $h_\alpha(x) \doteq v^*_\alpha(x) - v^*_\alpha(y^0)$. $h_{\alpha_n}(x) \leq 0$ for all x and all $n = 1, 2, \ldots$.

$S-3^*$: There exists a nonnegative function $N(x) = O(|x|^{\rho+1})$ such that $h_{\alpha_n}(x) \geq -N(x)$ for all x and all n. Moreover, $\mathsf{E}N(y - d(p, \epsilon)) < \infty$ for all (y, p).

Sennott requires in addition that the action sets be finite. However, this assumption is made *only* to ensure that a sequence of discount factors $\{\alpha_n\} \to 1$ exists for which the corresponding sequence of optimal policies converges pointwise to a stationary policy. In our model the latter can be verified directly, in spite of the fact that the action sets are infinite. Indeed, it follows from Proposition 1 and the discreteness of the state space, that a base stock/list price combination (y^0, p^0) exists such that $(y^*_\alpha, p^*_\alpha) = (y^0, p^0)$ for a sequence of discount factors $\{\alpha_n\} \to 1$. Consider now the inventory level $x = y^0 + 1$, and recall that the prescribed optimal price $\underline{p}_\alpha(x)$ in this state satisfies $p_{\min} \leq \underline{p}_\alpha(x) \leq p^0$. Thus a subsequence of $\{\alpha_n\}$ can be constructed with a *common* value for $\underline{p}_\alpha(y^0 + 1)$. Similarly, one can construct a further subsequence of $\{\alpha_n\}$ with a common value for $\underline{p}_\alpha(y^0 + 2)$ as well as $\underline{p}_\alpha(y^0 + 1)$. Continuing via this diagonalization method, we construct the desired sequence of discount factors and hence limiting stationary policy.

Thus, with this choice of y^0 and the sequence $\{\alpha_n\}$, S-2 clearly applies. S-1 is established in Theorem 5. To verify S-3*, let $\underline{v}^*_\alpha(x)$ denote total expected discounted return of the policy with fixed price p^0, and order-up-to level y^0, when starting in state x. Let τ denote the (random) number of periods required until the inventory level is first increased to y^0, i.e., $\tau = \min\{t \geq 1: \sum_{i=1}^t d_i \geq x - y^0\}$ where $d_1, d_2, \ldots$ are independent random variables, all distributed like $d(p^0, \epsilon)$. Thus, τ denotes the number of renewals by time $[x - y^0]^+$ in the corresponding renewal process and hence $\mathsf{E}(\tau) = O(|x|)$ (see, e.g., Heyman and Sobel 1984, Equation 5–12). Thus,

$$\begin{aligned}
h_\alpha(x) &= v^*_\alpha(x) - v^*_\alpha(y^0) \\
&\geq \underline{v}^*_\alpha(x) - v^*_\alpha(y^0) \\
&= \mathsf{E}\left[\sum_{t=1}^{\tau} \alpha^t p^0 d(p^0, \epsilon)\right] \\
&\quad - \mathsf{E}\left[\sum_{t=1}^{\tau} \alpha^{t-1} G\left(x - \sum_{i=1}^{t-1} d_i, p^0\right)\right] \\
&\quad - c\mathsf{E}\left[y^0 - x + \sum_{i=1}^{\tau} d_i\right] + \mathsf{E}(\alpha^\tau - 1)v^*_\alpha(y^0) \\
&\geq -\mathsf{E}\left[\sum_{t=1}^{\tau} G\left(x - \sum_{i=1}^{t-1} d_i, p^0\right)\right] - c(y^0 - x) \\
&\quad - c\mathsf{E}(\tau)\mathsf{E}d(p^0, \epsilon) + \mathsf{E}(\tau)(\alpha - 1)v^*_\alpha(y^0) \\
&\geq -\mathsf{E}(\tau)\max\{G(x, p^0), G(y^0, p^0)\} \\
&\quad - c(y^0 - x) - \mathsf{E}(\tau)[c\mathsf{E}d(p^0, \epsilon) + M] \\
&\geq -K(|x|^{\rho+1} + 1) \\
&\doteq -N(x),
\end{aligned}$$

for an appropriate constant K. (The second inequality follows from $\alpha < 1$, $G \geq 0$, Wald's Lemma, and the inequalities $\alpha^\tau - 1 \geq \tau(\alpha - 1)$ for $\alpha < 1$, and $v^*_\alpha(y^0) \geq 0$, see (19). The third inequality follows from $y^0 \leq x - \sum_{i=1}^{\tau-1} d_i \leq x$ for all $t = 1, \ldots, \tau$ and the convexity of $G(\cdot, p^0)$, see Assumption 5, while $(\alpha - 1)v^*_\alpha(y^0) \geq -N$ follows from the definition of M or (18). Finally, the last inequality follows from $\mathsf{E}(\tau) = O(|x|)$ and Assumption 3.)

In addition,

$$\mathsf{E}N(y - d(p, \epsilon)) = K(\mathsf{E}|y| - d(p, \epsilon)^{\rho+1} + 1) \leqslant K(\mathsf{E}([|y| + d(p, \epsilon)]^{\rho+1})) < \infty,$$

by Assumption 4*, employing the binomial expansion of $[|y| + d(p, \epsilon)]^{\rho+1} + 1$. This establishes parts (a)–(c) of the theorem. (It follows from the proof of the Theorem in Sennott that $h^* = \lim_{n\to\infty} h_{\alpha_n} \leqslant 0$.) Moreover, since $y^* = \lim_{n\to\infty} y^*_{\alpha_n}$ for the above constructed sequence $\{\alpha_n\}$ and $\underline{y} \leqslant y^*_\alpha \leqslant \bar{y}$ for all α sufficiently close to one by Proposition 1, we obtain part (d). □

4.2. The Model with Markdowns

We now turn to the case where only price *reductions* are permitted and show that under the long-run average profit criterion, it is optimal to adopt a constant price and a simple order-up-to policy.

Theorem 8. *Assume Assumptions* 1, 2, 3, 4*, *and* 5 *hold and that the system is in state* (x, p). *Let* p^* *be a maximizer on* $[\max(p_{\min}, c), p_{\max}]$ *of the concave function* $\{(p - c)\mathsf{E}d(p, \epsilon) - \min_y G(y, p)\}$. *Under the long-run average profit criterion, it is optimal to* (*i*) *adopt the constant price* $p' = \min(p^*, p)$, *and* (*ii*) *follow a simple order-up-to policy with order-up-to level* $y^*(p')$.

Proof. For all $p \in [\max(p_{\min}, c), p_{\max}]$ let $(g^*|p)$ denote the long-run average profit for the model in which the price is kept constant at level p. Under our assumptions, $(g^*|p) < \infty$ for all $p \in [\max(p_{\min}, c), p_{\max}]$ (see Veinott 1966). Consider an arbitrary (possibly history dependent) policy π and let $\{p_t : t = 1, 2, \ldots\}$ denote the stochastic price process generated by this policy. Since $\{p_t\}$ is nonincreasing and the possible price range is finite, we have with probability one that p_t is constant after finitely many periods, after which point in time it is clearly optimal to adopt a simple order-up-to policy. This implies that the long-run average profit under policy π is given by a weighted average of the values $\{(g^*|\max(p_{\min}, c)), (g^*|\max(p_{\min}, c) + 1), \ldots, (g^*|p_{\max})\}$ and hence bounded from above by $\max_{p' \in [\max(p_{\min}, c), p_{\max}]}(g^*|p')$, a value that can be achieved by the policy that adopts a constant price p' achieving this maximum, and orders up to the corresponding order-up-to level. It remains to be shown that this maximizing constant price p' satisfies $p' = \min(p^*, p)$. Note that $(g^*|p') = (p' - c)\mathsf{E}d(p', \epsilon) - \min_y G(y, p')$ is a concave function of p' in view of Assumption 1 and Assumption 5. (Since G is jointly convex, $\min_y G(y, p')$ is convex in p'.) This implies that $(g^*|p')$ is nondecreasing for $p' \leqslant p^*$. □

5. COMPUTATIONAL METHODS

In this section we describe efficient methods to determine an optimal policy for each of the models discussed in §§ 2 through 4. The base stock list price policy that is optimal for the *finite* horizon models in §2 can clearly be computed by the recursions (7)–(8) and (12)–(13) for the model with bi-directional price changes and that with markdowns, respectively. By Theorems 4 and 6, the recursive schemes converge to the infinite horizon value functions under the total discounted profit criterion, and the sequences of optimal finite horizon policies converge to an optimal policy for the infinite horizon model as well.

This leaves us with the long-run average profit criterion. By Theorem 8, for the model in which only markdowns are permitted, the computational effort reduces to:

(i) Determining p^* by computing a maximizer (p^*, y^*) of the jointly concave function $[(p - c)\mathsf{E}d(p, \epsilon) - G(y, p)]$. For any starting price $p \geqslant p^*$, it is optimal to adopt the price p^* and in each period to order up to the level y^*.

(ii) For a starting price $p < p^*$, it is optimal to maintain the price p forever, and in each period to order up to a level $y^*(p)$ which minimizes the convex function $G(\cdot, p)$.

For the model with bi-directional price changes, we return to the recursive scheme (5)–(6), now with $\alpha = 1$. We now show that the sequence $\{v^*_t - tg^*\}$ converges pointwise to a function h^* such that (h^*, g^*) satisfies the optimality Equation (20). (By Theorem 7(c) any policy achieving the maximum in the optimality equation for this solution is optimal.) Note first from Theorem 7(d) that without loss of optimality the sets of feasible actions $\{A(x)\}$ may be restricted to sets $\{\hat{A}(x)\}$ as follows:

$$\hat{A}(x) = \{(y, p) : \max(x, \underline{y}) \leqslant y \leqslant \bar{y} \text{ and } \max(p_{\min}, c) \leqslant p \leqslant p_{\max}\}, \text{ if } x \leqslant \bar{y},$$

$$\hat{A}(x) = \{(x, p) : \max(p_{\min}, c) \leqslant p \leqslant p_{\max}\}, \text{ if } x > \bar{y}.$$

Assume therefore that the value iteration scheme (5)–(6) is implemented with these restricted action sets, i.e., by imposing an additional upper bound $y \leqslant \max(x, \bar{y})$ and modifying the lower bound $y \geqslant x$ to $y \geqslant \max(x, \underline{y})$ in the maximization problem (5). One easily verifies that these modified bounds do not affect the validity of the structural results in Theorem 8.

Theorem 9. *Assume Assumptions* 1, 2, 3, 4*, *and* 5 *hold. Let* $\{v^*_t\}$ *denote the sequence of value functions generated by* (5)–(6) *with the restricted action sets* $\{\hat{A}(x)\}$. *Then* $\{v^*_t - tg^*\}_{t=1}^{\infty}$ *converges to a function* h^* *such that* (h^*, g^*) *is a solution to the long-run average profit criterion* (20).

Proof. Theorem 1 in Aviv and Federgruen (1995) shows that once the existence of a solution (h^*, g^*) of the optimality Equation (20) has been established, convergence of $\{v^*_t - tg^*\}$ to such a solution can be guaranteed by the verification of a single additional condition regarding the growth rate of the function h. In particular, Theorem 7 establishes Assumption (A) in Aviv and Federgruen (1995) since no stationary policy has null-recurrent states. (Note that the inventory level after ordering is bounded from below by $\underline{y}$ and that the states $\{x : x \geqslant \bar{y}\}$ are all transient. Note also that the Markov chain induced by an optimizing base stock list price policy with base stock/list price combination (y^*, p^*) is aperiodic since any of the states $(y^* - d^0)$ with $\pi^0 = Pr[d(p^*, \epsilon) = d^0] > 0$, repeats itself after a single period with probability $\pi^0 > 0$.)

Recall from the proof of Theorem 7 that $N(x) = K|x|^{\rho+1} + K$ is a bounding function for the optimality equation h^*, i.e., $|h^*(x)| \leq N(x)$ for all x. Let $\{x_n\}_{n=1}^{\infty}$ represent the process of (start-of-period) inventory levels before ordering under an *arbitrary* policy. The additional condition to be verified is

$$(C): \mathsf{E}N(x_n|x_0 = x) = O(N(x)) \quad \text{for all } n \geq 1.$$

Note that $x_n = y_{n-1} - D_{n-1}$ with y_{n-1} and D_{n-1} independent of each other. Thus

$$\begin{aligned} N(x_n|x_0 = x) &= K|x_n|^{\rho+1} + K \\ &\leq K|y_{n-1} - D_{n-1}|^{\rho+1} + K \\ &\leq \max_{\underline{y} \leq y \leq \max(x, \bar{y})} K|y - D_{n-1}|^{\rho+1} + K \\ &= K[1 + \max\{|\underline{y} - D_{n-1}|^{\rho+1}, \\ &\qquad |\max(x, \bar{y}) - D_{n-1}|^{\rho+1}\}] \\ &\leq K\{1 + [|\underline{y}| + D_{n-1} + 1]^{\rho+1} \\ &\qquad + [\max(x, \bar{y}) + D_{n-1} + 1]^{\rho+1}\}, \end{aligned} \tag{21}$$

where the second equality follows from the function $|y - D_{n-1}|^{\rho+1}$ being convex in y and hence achieving its maximum in one of the extreme points of the interval $[\underline{y}, \max(x, \bar{y})]$. Condition (C) is verified by taking expectations over the distribution of D_{n-1} in (21), applying binomial expansions and invoking Assumption 4*. □

Instead of the function v_t^* that grows linearly with t, it is advisable to generate the normalized value-function w_t^* defined by $w_t^*(x) = v_t^*(x) - v_t^*(x^0)$ for some reference state x^0, e.g., $x^0 = y^*$. Note that the sequence $\{w_t^*\}$ can be generated from the recursion

$$\begin{aligned} w_t^*(x) = -cx + &\max_{\{y \geq x,\, \max(p_{\min}, c_{t-1}) \leq p \leq p_{\max}\}} \\ &\{\alpha p \mathsf{E} d_t(p, \epsilon_t) - c_t y - G_t(y, p) \\ &\quad + \alpha \mathsf{E} w_{t-1}^*(y - d_t(p, \epsilon_t))\} \\ &+ cx^0 - \max_{\{y \geq x^0,\, \max(p_{\min}, c_{t-1}) \leq p \leq p_{\max}\}} \\ &\{\alpha p \mathsf{E} d_t(p, \epsilon_t) - c_t y - G_t(y, p) \\ &\quad + \alpha \mathsf{E} w_{t-1}^*(y - d_t(p, \epsilon_t))\}. \end{aligned}$$

6. EXTENSIONS

In this section, we briefly discuss a number of extensions of our model.

First, we have assumed that stockouts are fully backlogged. In many settings stockouts are satisfied at the end of the very period in which they occur, through emergency orders or production runs. This variant of the model can be handled with minor adaptations. Assume, e.g., that inventory carrying costs in period t are given by a function $h_t^+(I)$ of the end-of-period inventory level I. Emergency purchases at the end of period t involve a cost of $\bar{c}_t \geq c_t$ per unit. Consider, e.g., the model with bi-directional price changes and define $v_t^*(x)$, $V_t^*(x)$ and $J_t(y, p)$ as before. The recursion (6) is modified to

$$\begin{aligned} J_t(y, p) = \alpha p \mathsf{E} d_t(p, \epsilon_t) &- c_t y - h_t^+([y - d_t(p, \epsilon_t)]^+) \\ &- \bar{c}_t \mathsf{E}[d_t(p, \epsilon_t) - y]^+ \\ &+ \alpha \mathsf{E} v_{t-1}^*([y - d_t(p, \epsilon_t)]^+), \end{aligned}$$

where $x^+ = \max(x, 0)$. Rewriting, as before, (22) in terms of $V_t^*(x) = v_t^*(x) - c_t(x)$, we obtain after some algebra:

$$\begin{aligned} J_t(y, p) = \alpha(p - c_{t-1}) \mathsf{E} d_t(p, \epsilon_t) &+ (\alpha c_{t-1} - c_t) y \\ &- G_t(y, p) + \alpha \mathsf{E} V_{t-1}^*([y - d_t(p, \epsilon_t)]^+), \end{aligned} \tag{23}$$

where

$$\begin{aligned} G_t(y, p) = \mathsf{E} h_t^+&([y - d_t(p, \epsilon_t)]^+) \\ &+ (\bar{c}_t + \alpha c_{t-1}) \mathsf{E}[d_t(p, \epsilon_t) - y]^+. \end{aligned} \tag{24}$$

Observe that the first three terms to the right of (23) are identical to those in (8), with $G_t(y, p)$ now defined as in (24). Thus, maintaining Assumptions 1–5, with $h_t(\cdot)$ replaced by $h_t^+(\cdot)$, all structural results in Theorems 1–8 continue to apply. (Note, e.g., that concavity of $J_t(y, p)$ and $V_{t-1}^*(\cdot)$ can be proven under these assumptions by the same induction proof; observe that if $V_{t-1}^*(\cdot)$ is concave and nonincreasing, for any given value ϵ_t, $V_{t-1}^*([y - d_t(p, \epsilon_t)]^+)$ is jointly concave in y and p as well, as the composition of a concave nonincreasing function with the jointly convex function $[y - d_t(p, \epsilon_t)]^+$. This implies that $\mathsf{E}_{\epsilon_t} V_t^*([y - d_t(p, \epsilon_t)]^+)$ is jointly concave in y and p as well.) Following the proof of Lemma 1, one verifies that joint convexity of the function $G_t(y, p)$ is guaranteed, for example, when the function $h_t^+(\cdot)$ is convex and the demand functions are linear. As in the model with full backlogging, $G_t(y, p)$ is easily verified to be supermodular when $h_t^+(\cdot)$ is convex (as assumed in Assumption 5). This permits us (as in Theorems 2 and 3) to show that the functions $J_t(y, p)$ are submodular for $t = 2, 3, \ldots$.

Thus, immediate clearance of stockouts via emergency orders results in a model with identical structural properties for its optimal policies as the basic model with full backlogging (given Assumptions 1–5 as specified above). The same cannot be said for the case where stockouts result in lost sales, since in this case the expected revenue term in (6) is given by $\alpha p \mathsf{E} \min(y, d_t(p, \epsilon_t))$ which fails to be (jointly) concave even when the demand functions are linear. Thus, the functions $J_t(\cdot, \cdot)$ may fail to be concave, Theorem 1 may fail to hold, and base stock list price policies may fail to be optimal. On the other hand, the existence of a stationary optimal policy in the discounted or long-run average profit *infinite* horizon model with stationary parameters can still be demonstrated, along with that of solutions to the optimality equations of these models. (The proofs are analogous to those of Theorems 5–8.)

We have also assumed that orders are received instantaneously. Often, a *positive* leadtime of $L \geq 1$ periods is incurred between the placement of an order and its receipt. (The leadtime L may be a deterministic constant, or a random variable in case the supplier's delivery times are subject to uncertainty.) In standard inventory models, with a *given* price strategy, positive leadtimes are easily incorporated, albeit only in the case of full backlogging. Here, the inventory level at the end of period $(t - L)$ may be expressed as

$$I_{t-L} = y_t - (D_t + D_{t-1} + \cdots + D_{t-L}),$$

with y_t reinterpreted as the inventory *position* at the beginning of period t (after ordering). Thus, in standard inventory models, I_{t-L} can only be controlled via y_t, enabling a dynamic programming formulation with x_t = inventory position at the beginning of period t (*before* ordering), as the single state variable, and charging expected holding and backlogging costs at the end of period $(t - L)$ to period t.

When the price may be varied in each period, each demand component $D_t, D_{t-1}, \ldots, D_{t-L}$ is controllable (along with y_t) via the prices charged in periods $t, t - 1, \ldots, t - L$. An exact formulation of the problem requires that the state of the system be described via an $(L + 1)$-dimensional vector consisting of the current inventory level and the sizes of the orders placed in the L preceding periods. (The same state representation is required when stockouts result in lost sales or under other types of inventory dynamics.) Such a formulation is intractable for all but the smallest values of L.

We therefore propose a *heuristic* treatment of the problem in which expected holding and backlogging costs at the end of period $(t - L)$ are charged to period t, according to one of the following two functions.

(I) $\hat{G}_t(y, p) = \mathsf{E}h_{t-L}(y - d_t(p, \epsilon_t) - d_{t-1}(p, \epsilon_{t-1}) - \cdots - d_{t-L}(p, \epsilon_{t-L}))$; i.e., assume that the price selected for period t is maintained over the next order leadtime of L periods. In close analogy to Lemma 1 it is easily verified that $G_t(\cdot, \cdot)$ is jointly convex if h_t is convex and the demand functions are linear (for example).

(II) $\underline{G}_t(y, p)$ = minimum expected holding and backlogging cost at the end of period $(t - L)$, if period t starts with an inventory position (after ordering) of y units, price p is chosen for period t, and an optimal price strategy is followed over the time interval $[t - 1, \ldots, t - L]$.

The complete function $\underline{G}_t(\cdot, \cdot)$ may be evaluated by the solution of an L-period horizon problem of the type discussed in §1, assuming that the order quantities in periods $t, t - 1, \ldots, t - L$ are fixed at the levels $(y - x_t), 0, \ldots, 0$ and that the holding and backlogging costs in periods $t, t - 1, \ldots, t - L + 1$ are zero. Also, choosing $\underline{G}_t(\cdot, \cdot)$ as the one-step expected cost function, results in an *upper bound* approximation for the entire problem. It follows from Theorem 1 that $\underline{G}_t(\cdot, \cdot)$ is again jointly convex provided Assumptions 1–5 continue to hold. Thus, under Assumptions 1–5 the structural results in Theorems 1–8 continue to apply to the (upper bound) approximation model.

We conclude that the models treated in §§ 2 through 4 can be solved in the case of *positive* leadtimes, via the same computational schemes and with the same structural results for the "optimal" policies, provided the one-step expected inventory/backlog cost function is chosen as $\hat{G}_t$ or $\underline{G}_t$.

Another complication often encountered but ignored by our basic model is the imposition of specific upper limits constraining order sizes and/or price fluctuations. Assume, e.g., that orders in period t are bounded by an amount b_t, the period's capacity.

All the results in §§ 2 through 4 continue to hold under these capacity restrictions. For the long-run average profit criterion, it is now necessary to impose an additional assumption to ensure that the system is stable, i.e., that the long-run capacity is sufficient to meet the demand. In the stationary model, this assumption reduces to $b > \mathsf{E}d(p_{\max}, \epsilon)$. Also, to ensure convergence of the value-iteration method in the *undiscounted* model (see Theorem 9), it is necessary to assume that $\mathsf{E}(d^{\rho+2}) < \infty$, a slight strengthening of Assumption 4 (4*). The proofs of Theorems 8 and 9 are considerably more involved, see Aviv and Federgruen (1997). Most importantly we obtain the same structural results and computational methods as in the basic model, with obvious adaptations for the capacity limits. For example, in the model with bi-directional price changes, under a base stock list price policy with parameters (y^*, p^*), one increases the inventory level to $y^0 = \min(y^*, x^0 + b)$, if before ordering the inventory level equals $x^0 \leqslant y^*$ units. Also, in this case, a price $p \geqslant p^*$ is chosen that maximizes the function $J(y^0, \cdot)$.

7. NUMERICAL STUDY

In this section we report on a numerical study conducted to assess the computational effort associated with the value-iteration methods, and more importantly, to attain qualitative insights into the structure of optimal policies and their sensitivity with respect to several parameters. Among the major questions investigated, we focus in particular on:

(i) the benefits of a dynamic pricing strategy compared to a fixed price strategy in settings with (a) continuous replenishment opportunities, (i.e., replenishment options at the beginning of each period) and (b) limited replenishment opportunities (e.g., one or two procurements during the entire selling season);

(ii) the benefits derived from bi-directional price changes, as opposed to those achievable when only markdowns are permitted;

(iii) the sensitivity of the optimal list price as a function of the initial inventory level; and

(iv) the sensitivity of the optimal base stock/list price combination with respect to the degree of variability and the seasonality patterns in the demands, as well as the price elasticities in the (stochastic) demand functions.

Our numerical study is based on data collected from a specialty retailer of high-end women's apparel. The retailer sells only its own private label, predominantly dresses, sportwear separates, and coordinated collections.

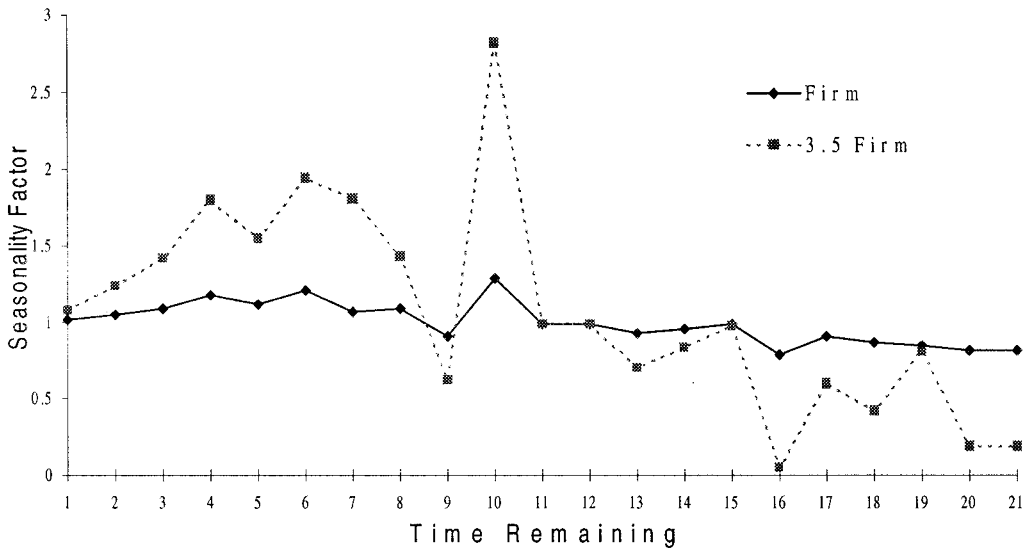

Figure 2. Seasonality factors.

To market its merchandise, the retailer utilizes a network of approximately 50 stores throughout the continental United States, all owned and managed by the parent company. For the purpose of this study, we consider only the aggregate sales process (throughout the United States), leaving the complications that arise due to geographic dispersion of these stores to a future publication.

Focusing on two of the company's fashion items, a dress and a skirt, we have generated for each item a total of 53 scenarios. For the basic pair of scenarios we have selected parameters to match actual sales, cost, and revenue data observed in the Spring 1993 selling season. The basic pair of scenarios has *stationary* data and *additive* stochastic demand functions (see (2)), i.e.,

$$D_t = \delta_t(p) + \epsilon_t, \tag{25}$$

where $\delta_t(p) = \delta(p)$ is assumed to be *linear*. (In an alternative set of scenarios we use the firm's actual 21 *seasonality factors*, $\{\gamma_t\}$, to gauge the impact of seasonalities; see Figure 2.) The choice of an additive model is justified by the fact that the standard deviation of weekly sales is independent of the price charged. For both items, the slope and intercept of the function $\delta(p) = a + bp$ were estimated as follows. We assume that the initial price p^0 (charged for most of the season), was selected by firm managers to maximize expected profits per week, and that the quantity purchased for the season equaled expected season-wide demand under this price. With Q this purchase quantity divided by the length of the season, these assumptions give rise to a pair of linear equations,

$$p^0 = \frac{1}{2}\left(\frac{-a}{b} + c\right) \quad \text{and} \quad a + bp^0 = Q,$$

from which the values of a and b can be computed.

In the base scenarios, the variables ϵ_t are independent and identically distributed, as the (discretized) truncation of a normal $\hat{\epsilon}$. The variables are truncated at $-\delta(p)$ to prevent negative demand realizations. The parameters μ and σ of the normal $\hat{\epsilon}$ are specified to ensure that $\mathsf{E}(\epsilon_t) = 0$ and $std(\epsilon_t)$, its standard deviation, equals $\delta(p) * c.v.$ with $c.v.$ a specified coefficient of variation. (Strictly speaking, the distribution of ϵ_t therefore depends on the price p, i.e., $\epsilon_t = \epsilon_t(p)$.) In our base scenarios we set $c.v. = 1$. (This order of magnitude of the coefficient of variation is often encountered for retail sales of fashion items.)

The variable manufacturing and distribution cost rates, $c_t = c$ represent actual dollar values experienced in the Spring 1993 sales season. Our base scenarios assume that bi-directional price changes are permitted, stockouts can be fully backlogged, and that holding and backlogging costs are proportional with the end-of-the-week inventory level or backlog, at rates h and π, respectively. (The cases of more limited pricing flexibility as well as emergency procurements to immediately fill excess demand are assessed in alternative sets of scenarios.) Because of the cost

Table I
Base Parameters for Dress and Skirt Items

Item	a	b	ϵ Dist'n	c.v.	Order Cost (c)	Holding Cost (h)	Penalty Cost (π)	Salvage (s)	Emergency Cost ($\bar{c}$)	Price Range	Max Inv.	Fixed Price
Dress	174	−3	Normal	1.0	22.15	0.22	21.78	17.72	221.50	25–44	400	40
Skirt	57	−1	Normal	1.0	14.05	0.17	16.83	11.24	140.50	15–44	400	37

Table II
Convergence Rates and Average Profit per Period: "(e/s)" = Emergency Shipments, "(m/d)" = Markdowns Only

	n^*	t^*	g^*
cv = 1.4	15	45	$ 901.85
cv = 1.2	10	30	$ 914.21
cv = 1.0	9	27	$ 925.54
cv = 0.75	6	18	$ 937.84
cv = 0.50	4	13	$ 947.76
cv = 0.25	3	8	$ 955.98
cv = 0.12	2	3	$ 960.10
$\pi/(\pi + h) = 0.99$	9	27	$ 925.52
$\pi/(\pi + h) = 0.97$	10	31	$ 931.65
$\pi/(\pi + h) = 0.95$	9	30	$ 934.80
$\pi/(\pi + h) = 0.90$	9	30	$ 939.52
b = −5	10	31	$ 975.64
b = −3	9	27	$ 925.52
b = −1	8	28	$1265.03
Normal	9	27	$ 925.52
Geometric	7	15	$ 910.99
Poisson	2	7	$ 959.37
cv = 1.4 (e/s)	13	36	$ 884.35
cv = 1.2 (e/s)	11	31	$ 900.92
cv = 1.0 (e/s)	8	23	$ 915.30
cv = 0.75 (e/s)	6	17	$ 931.03
cv = 1.4 (m/d)	15	45	$ 901.85
cv = 1.2 (m/d)	10	30	$ 914.21
cv = 1.0 (m/d)	9	27	$ 925.54
cv = 0.75 (m/d)	6	18	$ 937.84
cv = 0.50 (m/d)	4	13	$ 947.76

of capital, maintenance, insurance, loss, and damage, annualized holding cost rates amount to approximately 20% of the initial retail price, while high service levels (fill rates) are ensured by setting the backlogging cost rates at 100 times the holding cost rates. Finally, all scenarios assume an emergency procurement cost rate $\bar{c} = 10c$ and in accordance with the firm's experience, a salvage value $s = 0.8c$ for any inventory remaining at the end of the horizon.

Table I summarizes all parameters for the base scenarios pertaining to the dress and the skirt. Note that the selected fixed price is the price that maximizes expected single period profits. The remaining figures gauge the impact of each of the model's parameters on the optimal base stock and list price. More specifically, the graphs display, as a function of the number of periods remaining until the end of the horizon, the optimal base stock and list price. (We report these values for each period until they converge to their limit values.) Further, we report for two sets of scenarios the percentage increase in profits achieved by the dynamic price policy as opposed to a fixed price policy, when five periods remain in the horizon, and assuming zero starting inventory. Further, we report (in Table II) the long-run average profit per period as well as the number of periods and the total computational time (in seconds) required until the optimal base stock/list price combination and the average profit per period converge (to less than 10^{-2} units from their limit values). For the scenarios where only markdowns are permitted, long-run average profit depends on the initial price. For these scenarios we choose the optimal fixed price as the initial price. Finally, we illustrate the price monotonicity in the initial inventory level, as proven in Theorem 2 (see Figure 3). We present only the results for the dresses; those for the skirts follow similar patterns.

Figures 4a–4b display the results for our base scenario as well as a range of scenarios with an alternative value for the coefficient of variation of weekly demands (specifically, $c.v.$ varies between 0.12 and 1.4). Figures 5a–5b exhibit the above results for our base scenario and a range of scenarios with an alternative value for π (π = 14.55, 8.55 and 4.05, corresponding with a $\pi/(\pi + h)$ ratio of 0.97, 0.95, and 0.90, respectively).

In Figures 6a–6b we investigate the impact of different price elasticities by modifying the slope b of the demand functions $\delta(p)$ to $b = -1$ and $b = -5$. The demand

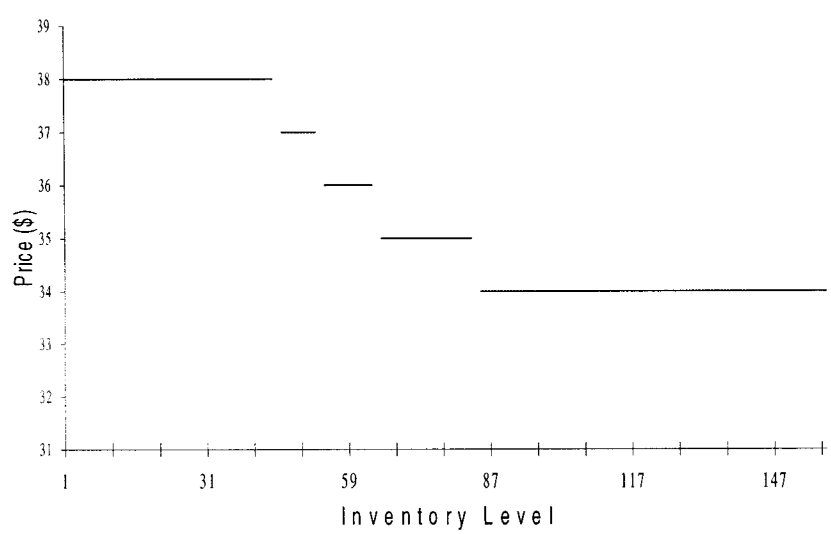

Figure 3. Price trajectory for dress base scenario, with one period remaining.

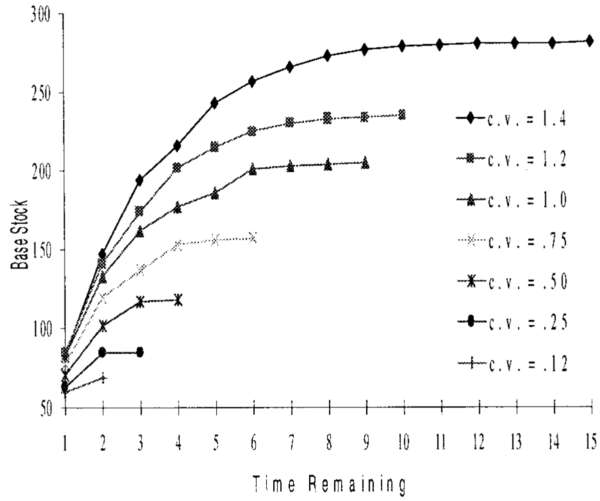

Figure 4a. Optimal base stock levels for varying levels of demand uncertainty, as function of remaining time.

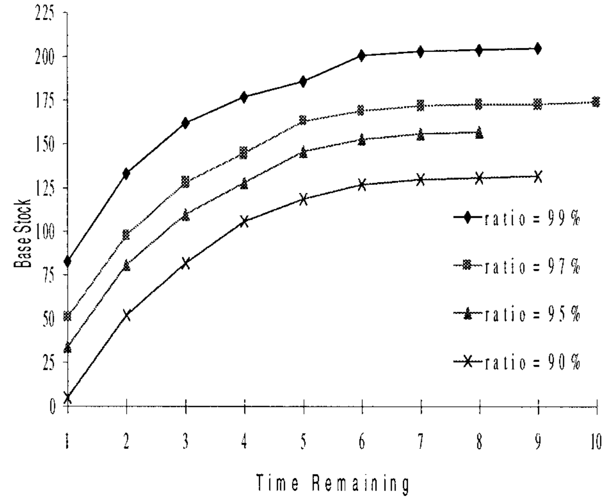

Figure 5a. Optimal base stock levels for varying values of $\pi/(\pi + h)$ ratio as function of remaining time.

functions continue to go through the point (p^0, Q) as discussed above. In Figures 7a–7b we illustrate the impact of the *shape* of the distribution of the random terms ϵ_t, by considering two alternative scenarios in which ϵ_t has a Geometric and Poisson distribution, respectively. (Since the Geometric and Poisson distributions depend on a single parameter, we choose this parameter to match the expectation only.)

Next, Figures 8a–8b consider the impact of requiring excess demand to be met by an emergency procurement, under the parameters of the base scenarios as well as (up to) three alternative scenarios with $c.v. = 1.4$, 1.2, and 0.75. Figures 9a–9b consider the case where only markdowns are permitted, and display the impact of the coefficient of variation which is varied over the same range as in the case of bi-directional price changes. In Figure 9b we compare the optimal list prices for systems in which only markdowns are permitted to those which allow for bi-directional price changes.

Figures 10a–10b gauge the impact of seasonalities. Toward this end we used the actual 21 *seasonality* factors $\{\gamma_t\}$ used by the firm. We then consider the impact of increasing the amplitude of the seasonality cycle by a factor of 3.5, i.e., we use the factors $\hat{\gamma}_t = 1 + 3.5(\gamma_t - 1)$, see Figure 2. Figures 11a–11b display the percentage increase in profits achieved by a dynamic pricing policy as compared with a fixed price policy for the set of scenarios represented by Figures 4 and 8, respectively.

Finally, in Table III we assess the benefits of continuous (weekly) replenishments, by comparing, for the 21-week sales season experienced by the firm, the profits generated by the base scenario with those generated by scenarios in which (i) only one order can be placed at the beginning of

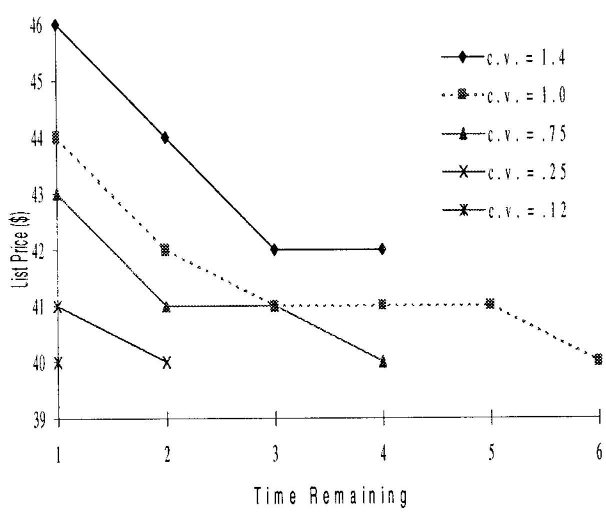

Figure 4b. Optimal list prices for varying levels of demand uncertainty, as function of remaining time.

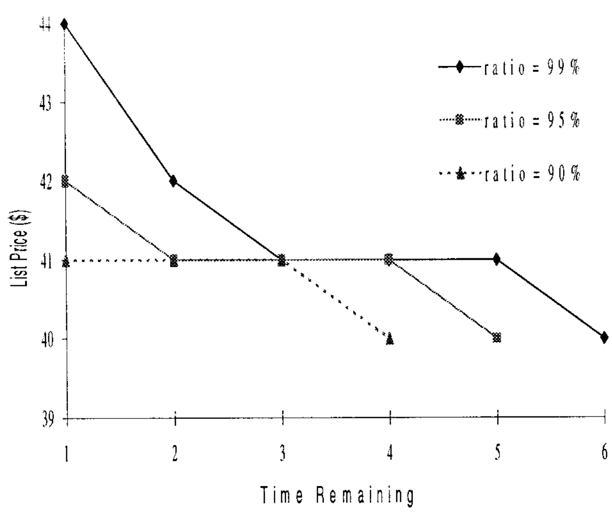

Figure 5b. Optimal list prices for varying values of $\pi/(\pi + h)$ ratio as function of remaining time.

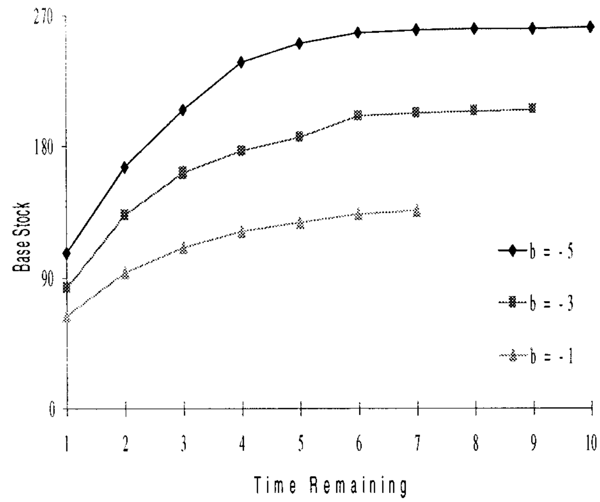

Figure 6a. Optimal base stock levels for varying demand elasticities as a function of remaining time.

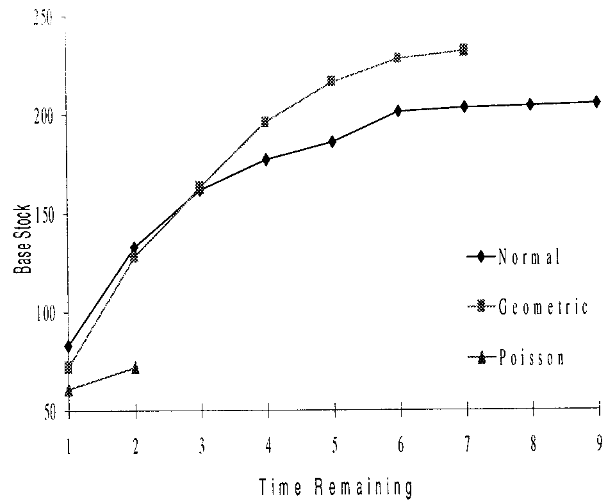

Figure 7a. Optimal base stock levels for varying demand distributions as a function of remaining time.

the season, and (ii) at most two orders can be received, in periods 1 and 11. A full backlogging model is somewhat unrealistic for models that assume that orders may only be placed at limited times, since it assumes that all unfilled demand will be filled by subsequent orders placed with the usual associated per unit order cost, c. Instead, for these limited ordering models we assume that all unfilled demand is met by an emergency procurement in the period in which it occurs.

The results in Figures 4a and 4b demonstrate that the optimal base stock level decreases as one approaches the end of the planning horizon, while the optimal list price increases. The latter is in sharp contrast with markdown strategies employed when a single initial procurement of inventory is used to cover the complete season. As reported in Table II, in these stationary scenarios the optimal base stock/list price combinations as well as the estimate of the long-run average profit converge rather rapidly; convergence occurs always at a horizon length n less than or equal to fifteen, and often already at $n \leq 5$. Not surprisingly, convergence occurs faster when system randomness is reduced.

We also note that the optimal list price increases with the degree of system uncertainty. This may be explained as follows. For any given list price, the required safety stocks as well as expected shortage costs increase as the coefficient of variation of weekly demands increases. Thus it is beneficial to respond to an increase in the $c.v.$ value by increasing the price, thus reducing both the mean and standard deviation of weekly demands in the same proportion. For similar reasons, as exhibited in Figure 5b, the optimal list price is increasing in π. On the other hand, we observe that the increases in the optimal base stocks and

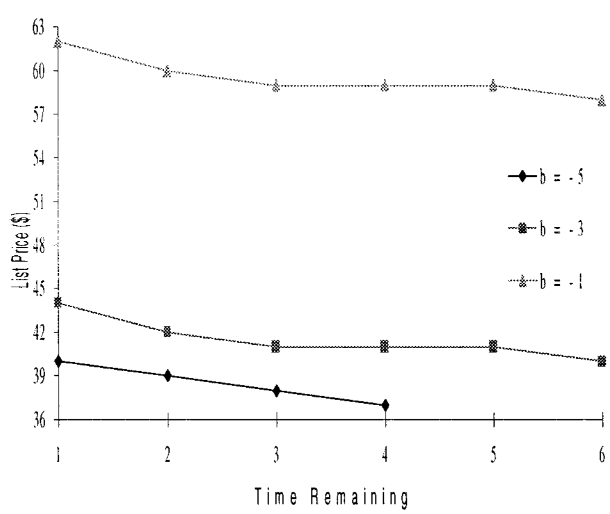

Figure 6b. Optimal list prices for varying demand elasticities as a function of remaining time.

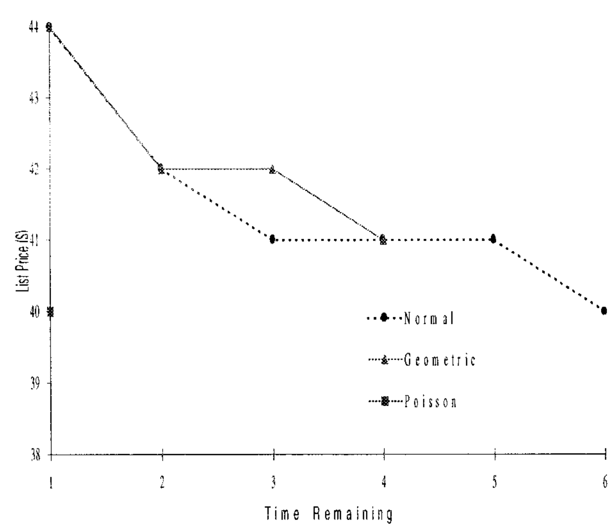

Figure 7b. Optimal list prices for varying demand distributions as a function of remaining time.

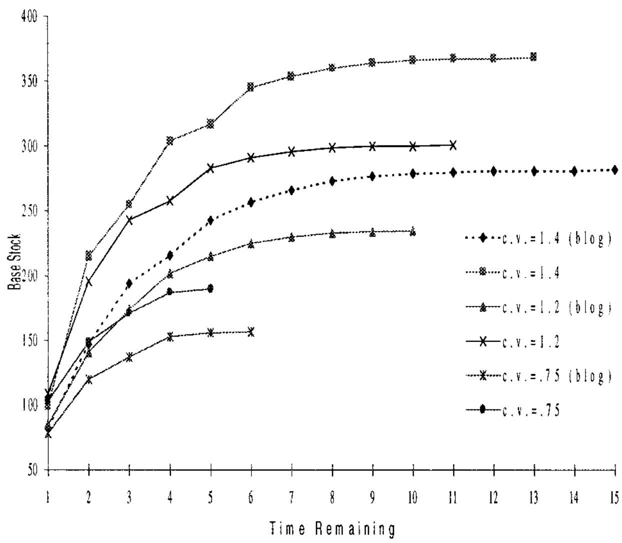

Figure 8a. Comparison of optimal base stock levels for systems with no backlogging and those with full backlogging, as a function of remaining time.

list prices, as a function of $c.v.$ or π, is modest and vanishes as the length of the planning horizon increases.

Figure 11a shows that even in a stationary environment significant benefits accrue from a dynamic pricing strategy compared to the fixed price strategy. The variable profit enhancements may amount to up to 2.25% when five or fewer periods remain until the end of the planning horizon (e.g., the end of the sales season). As observed in the general literature on yield management, in the retail sector these differences may have very large impacts on bottom line profit figures. Fisher and Raman (1996) report that in the ski wear industry, for example, a cost reduction in the amount of 1% of sales, results in an increase of profits by

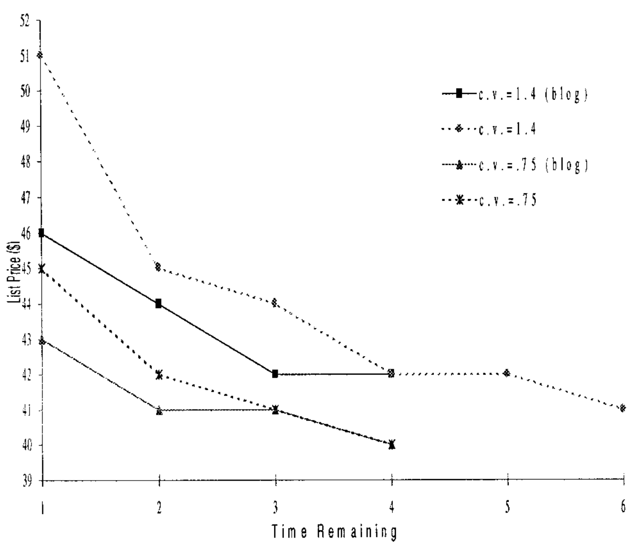

Figure 8b. Comparison of optimal list prices for systems with no backlogging and those with full backlogging, as a function of remaining time.

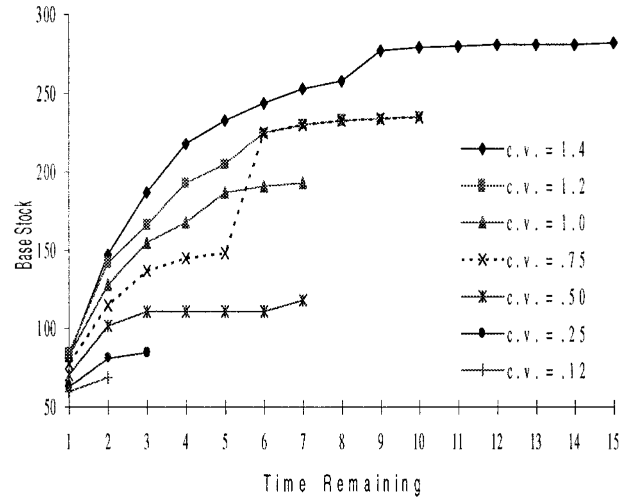

Figure 9a. Optimal base stock levels for systems with markdown-only price policy, as a function of remaining time.

60%. Consistent with Theorem 8, showing that in the long run a fixed price strategy is optimal among all markdown-only strategies, the relative benefits decrease as the length of the planning horizon increases. As can be expected, the benefits of a dynamic pricing strategy increase as the degree of uncertainty (i.e., the value of $c.v.$) increases. Finally, consistent with standard inventory models with a fixed price level, we observe that the optimal base stock level increases and expected profit values decrease as either the value of $c.v.$ or the value of π increase. Figure 3 exhibits, for the base scenario, the optimal price level as a function of the starting inventory in the final period of the horizon.

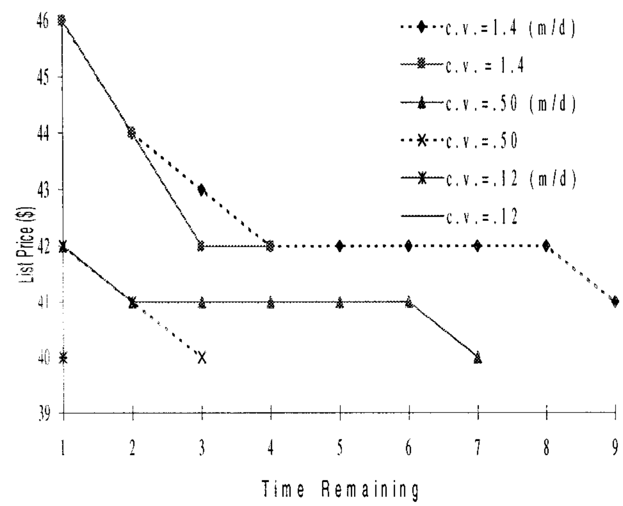

Figure 9b. Comparison of optimal list prices for systems with markdown-only and bi-directional price change policies, as a function of remaining time.

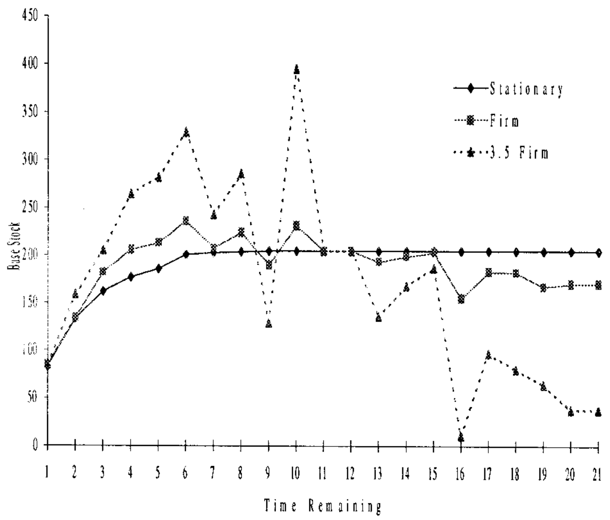

Figure 10a. Optimal base stock levels for varying levels of demand nonstationarity as function of remaining time.

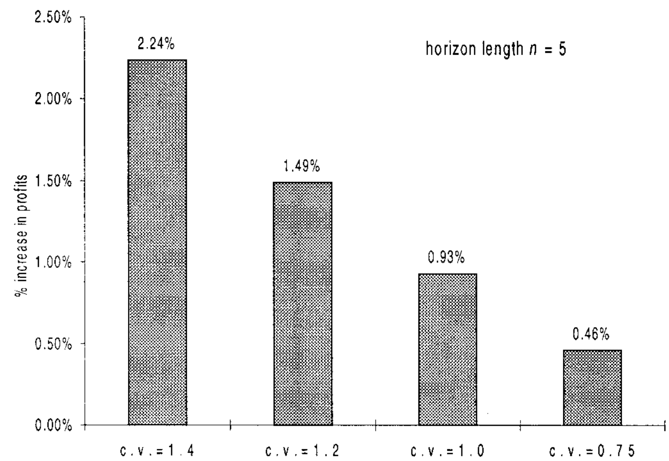

Figure 11a. Percentage difference between profit under fixed price and dynamic price policies.

Figures 6a–6b exhibit the impact of different price elasticities in the demand functions. If the slope of the demand function decreases (in absolute value) from 3 to 1, the optimal list price converges to 58 instead of 40 (as the length of the planning horizon increases); the lower price sensitivity is exploited to increase the price sharply, at the expense of incurring a relatively modest decrease in demand volume. The latter does however result in a systematic decrease of the optimal base stock levels (from 205 to 136, when $n \geq 9$). Conversely, if the slope of the demand functions increases from 3 to 5 (in absolute value), the optimal response consists of a markdown from 40 to 37 (when $n \geq 6$) so as to capture a relatively large increase in demand volume. Notice that for $n \geq 10$, the 7.5% price markdown causes an approximately 28% increase in the mean and standard deviation of weekly demand, and a similar increase in the optimal base stock level. Since the demand functions are rotated around the point whose price component is the optimal fixed price for the original demand curve, one observes that expected profits increase both when the slope of the demand functions increases and when it decreases.

Figures 7a–7b exhibit the impact of the shape of the distribution. The scenario in which weekly demands follow a Poisson distribution, is best compared with the corresponding scenario in Figures 4a–4b where $c.v. = 0.12$; under a price of 40, the mean demand per week is for 54 units in *all* scenarios of Figures 4a–7b. Under a Poisson distribution this implies that the standard deviation of weekly demand is in the amount of 7.3, hence a value of $c.v. = 0.136$ (approximately). Looking at Table II, the scenario with $c.v. = 0.12$ and that where demand follows the above Poisson distribution, both exhibit exceedingly fast convergence to the optimal steady-state parameters.

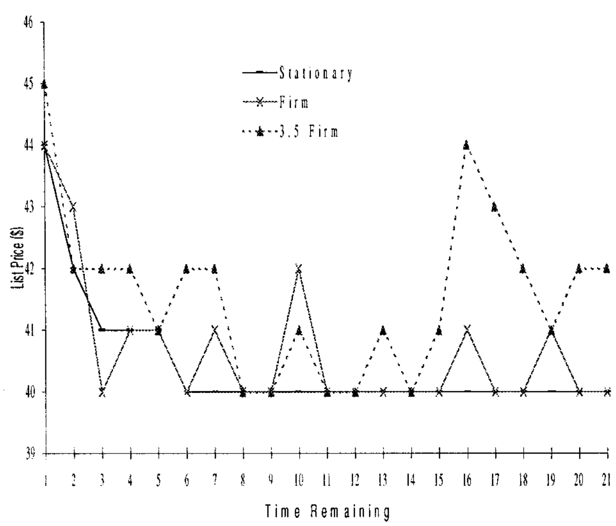

Figure 10b. Optimal list prices for varying levels of demand nonstationarity as function of remaining time.

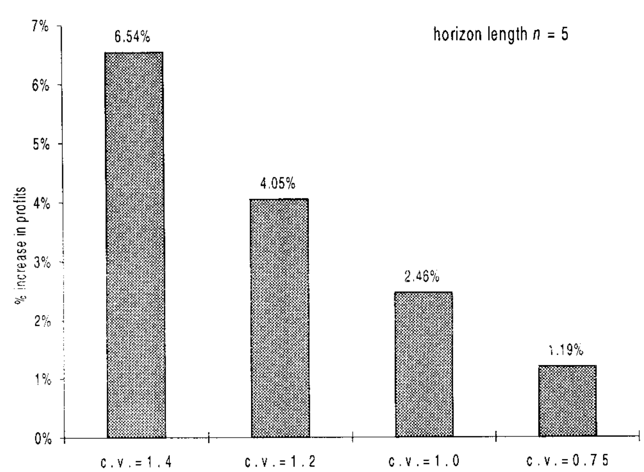

Figure 11b. Percentage difference in profits under fixed price and dynamic price policies for systems with no backlogging.

Table III
Limited Ordering; Excess Demand Met by Emergency Procurement

Period		Base Case	Bi-directional Pricing		Markdowns only	
			One Order	Two Orders	One Order	Two Orders
$n = 10$	(y^*_{10}, p^*_{10})	(242, 41)		(605, 41)		(679, 45)
	$v^*_{10}(0)$	\$ 8531.69		\$ 7761.69		\$ 6411.08
	y^*_{10} (fixed)	256		874		
	$v^*_{10}(0)$ (fixed)	\$ 8435.83		\$ 6370.60		
$n = 21$	(y^*_{21}, p^*_{21})	(242, 41)	(1105, 41)	(754, 41)	(1578, 40)	(8245, 44)
	$v^*_{21}(0)$	\$18599.98	\$15944.80	\$17022.15	\$12866.81	\$15119.51
	y^*_{21} (fixed)	256	1571	1028		
	$y^*_{21}(0)$ (fixed)	\$18502.42	\$12766.07	\$15091.93		

Comparing Figures 4a and 4b to Figures 7a and 7b, respectively, one notes that the optimal base stock levels and list prices are almost identical in the considered pair of scenarios. The same can be said for long-run average profits in this pair of scenarios; see Table II. The scenario with geometrically distributed weekly demands is comparable to the base case since, under the price of 40, the $c.v.$ value equals 0.9906. Nevertheless we observe that the steady-state optimal base stock level is 13% larger and the long-run average profit value about 2% lower. The geometric distribution has a considerably fatter tail than the corresponding normal distribution, thus requiring additional safety stocks to ensure comparable service, resulting in lower expected profits.

The "emergency procurements" scenarios in Figures 8a–8b show similar convergence characteristics and similar patterns of the optimal base stock levels and list prices $\{(y^*_t, p^*_t): t = 1, \ldots, T\}$ as a function of the value of $c.v.$, as compared to the corresponding scenarios with full backlogging. In comparing the parameter values and long run profits themselves one notices, in Figure 8a and Table II, that larger optimal base stock levels are needed and lower expected profits are obtained. As can be expected, the differences increase with the degree of uncertainty in the system, i.e., with the value of $c.v.$ Finally, Figure 11b demonstrates that a dynamic price policy enhances profits by as much as 6.5% for a horizon with $n = 5$.

Figures 9a–9b gauge the impact of a "markdowns only" restriction on the price strategy. As observed above, in a stationary environment a fixed price policy behaves as well as any dynamic price policy for sufficiently large planning horizons. This explains why for large planning horizons in a stationary environment, the "markdowns only" restriction is without significant loss of optimality. On the other hand, when $n \leq 5$ and $c.v. = 1.0$ ($c.v. = 1.4$), a loss in expected profits of at least 1% (2%) is observed. We also note that the optimal list prices are somewhat higher under the "markdowns only" scenarios. This occurs because under "markdowns only" the space of feasible future prices is enlarged by adopting a higher current price while in the case of bi-directional price changes the range of future prices is *unaffected* by current choice. Finally, as with the scenarios permitting bi-directional price changes, one observes (in Figure 9b) that the optimal list price increases as one approaches the end of the planning horizon. However, in the case of "markdowns only," future (higher) list prices are unachievable once a (lower) current optimal list price is implemented.

Figures 10a and 10b exhibit the impact of seasonalities in the demand pattern on the optimal base stock levels and list prices. As expected, one observes that the optimal base stock levels fluctuate significantly from period to period. Moreover, the optimal list prices are also adjusted upwards and downwards as one progresses over the seasonal cycle. As with the stationary scenarios, one observes that the loss in expected profits over the 21-week cycle under the "markdowns only" restriction is of the order of 0.5%. Recall that under stationary scenarios and continuously available replenishment opportunities, the benefits of a dynamic pricing strategy, while significant at first, vanish as the length of the planning horizon is increases. Under seasonal fluctuations, the benefits of a dynamic price strategy are considerably larger and persist even as the length of the planning horizon is increased. Expected profits over the complete season of $n = 21$ periods are increased by 0.5% (0.8%) for the "Firm" ("3.5 Firm") seasonality pattern.

We have simulated the 21 week selling season scenario with non-stationary demand using the seasonality factors experienced by the firm and allowing bi-directional price changes. In a sample of 500 replicas, we observed that price changes were implemented on average in 12.18 of the 21 periods, while the order-up-to level is modified in 14.84 periods, on average. The average number of periods in which the implemented and order-up-to level or price is different from the optimal base stock or list price combination is 1.55 and 4.74, respectively. Finally, in those weeks in which a price change is implemented, the average absolute price change amounts to \$1.33, and in those weeks in which the order-up-to level deviates from the optimal base stock level the average deviation is in the amount of 25.27 units. These simulations demonstrate the importance of a dynamic pricing strategy and the degree of fluctuation in order-up-to levels that are experienced under dynamic pricing schemes.

Table III exhibits the magnitude of the benefits of weekly replenishment opportunities as opposed to settings where demand over a complete season must be covered by one or two procurements only. For example, one observes that total expected profits decrease by 14.8% when the entire season must be covered by a single order. Observe that the loss in expected profits under limited replenishment opportunities (i.e., one or two replenishment orders) is significantly smaller when a dynamic pricing strategy is adopted as compared to the case where a fixed price is employed. In other words, the benefits of a dynamic pricing strategy are remarkably large in those setting where only a single or a limited number of replenishment opportunities can be arranged. For example, total expected profits over 21 weeks can be improved by 19.9% because of the adoption of a dynamic pricing strategy when a single replenishment opportunity exists, and by 11.8% when two replenishment orders can be placed. We also observe that under limited replenishment opportunities, the benefits of bi-directional price changes become extensive: compared with the "markdowns only" case, expected profits over the season increase by 19.3% when a single replenishment opportunity exists, and by 11.18% when two replenishment orders can be placed. Note that if prices can be increased over the course of the season, a 30% lower initial procurement quantity suffices to cover the entire season as compared to the "markdowns only" option.

APPENDIX

Proof of Theorem 3. To prove Theorem 3 we first need the following lemma.

Lemma 2. *Let* $f:\Re \times [\underline{v}, \bar{v}] \to \Re : (y, v) \to f(y, v)$ denote a jointly concave function with isotone differences. Define $g(x, u) = \max_{\{y \geqslant x, v \geqslant u, \underline{v} \leqslant v \leqslant \bar{v}\}} f(y, v)$ and assume that for all pairs (x, u) the maximum is achieved in some point (y^0, v^0). Then *g is jointly concave with isotone differences*.

Proof. Define $F(y, u) = \max_{\{v \geqslant u, \underline{v} \leqslant v \leqslant \bar{v}\}} f(y, v)$. We show that $F(y, u)$ is jointly concave with isotone differences. Since $g(x, u) = \max_{y \geqslant x} F(y, u)$, g shares these properties as well. (Note that in view of the existence of finite maximizers, the mapping transforming the function f into F is identical to that transforming F into g.) For any given value of y, the concave function $f(y, \cdot)$ clearly has a maximizer on the bounded interval $[\underline{v}, \bar{v}]$. Let $v^*(y)$ denote the smallest such maximizer. Since $f(\cdot, \cdot)$ has isotone differences, it follows that $v^*(y)$ is nondecreasing in y. (See, e.g., Heyman and Sobel 1984, Theorem 8-4.) Thus, for any $y_1 > y_2$ we have $v^*(y_1) \geqslant v^*(y_2)$. Consider the difference function $F(y_1, u) - F(y_2, u)$:

$$F(y_1, u) - F(y_2, u) = \begin{cases} f(y_1, v^*(y_1)) - f(y_2, v^*(y_2)) \\ \quad \text{if } u < v^*(y_2) \leqslant v^*(y_1), \\ f(y_1, v^*(y_1)) - f(y_2, u) \\ \quad \text{if } v^*(y_2) \leqslant u \leqslant v^*(y_1), \\ f(y_1, u) - f(y_2, u) \\ \quad \text{if } v^*(y_1) \leqslant u. \end{cases} \tag{26}$$

(26) follows from the concavity of $f(y, \cdot)$. We conclude that $F(y_1, u) - F(y_2, u)$ is constant for $u \leqslant v^*(y_2)$, and increases from $f(y_1, v^*(y_1)) - f(y_2, v^*(y_2))$ to $f(y_1, v^*(y_1)) - f(y_2, v^*(y_1))$ as u increases from $v^*(y_2)$ to $v^*(y_1)$. (Note that $f(y_2, \cdot)$ is nonincreasing for $u \geqslant v^*(y_2)$.) Thereafter, i.e., for $u > v^*(y_1)$, the difference function $F(y_1, u) - F(y_2, u) = f(y_1, u) - f(y_2, u)$ increases in view of f having isotone differences. We conclude that F has isotone differences, while its concavity is immediate from that of the function f. As mentioned, the same properties thus carry over to the function g. □

Proof of Theorem 3. (a) By induction: J_1 is jointly concave since it coincides with the J_1 function in the model with bi-directional price changes. The concavity and monotonicity properties of $V_1(\cdot, \cdot)$ are thus straightforwardly satisfied. Now fix $t = 2, \dots, T$ and assume $V^*_{t-1}(\cdot, \cdot)$ has the desired concavity and monotonicity properties. We then show that J_t and V^*_t have the desired properties as well. We first show that for any given value of ϵ_t, the function $V^*_{t-1}(y - d_t(p', \epsilon_t), p')$ is jointly concave in (y, p'). For a fixed value of ϵ_t, let (y_1, p'_1) and (y_2, p'_2) denote two points in $\Re^2$:

$$\begin{aligned} &V^*_{t-1}\left(\frac{y_1 + y_2}{2} - d_t\left(\frac{p'_1 + p'_2}{2}, \epsilon_t\right), \frac{p'_1 + p'_2}{2}\right) \\ &\quad \geqslant V^*_{t-1}\left(\frac{y_1 + y_2}{2} - \frac{1}{2} d_t(p'_1, \epsilon_t) \right. \\ &\qquad\qquad \left. - \frac{1}{2} d_t(p'_2, \epsilon_t), \frac{p'_1 + p'_2}{2}\right) \\ &\quad \geqslant \tfrac{1}{2} V^*_{t-1}(y_1 - d_t(p'_1, \epsilon_t), p'_1) \\ &\qquad + \tfrac{1}{2} V^*_{t-1}(y_2 - d_t(p'_2, \epsilon_t), p'_2), \end{aligned}$$

where the first inequality follows from the concavity of d_t in p and the nonincreasingness of V^*_{t-1} in its first argument, and the second inequality follows from the concavity of V^*_{t-1}. This implies that $\mathsf{E}_{\epsilon_t} V^*_{t-1}(y - d_t(p', \epsilon_t), p')$ is jointly concave as well, and since the first three terms to the right of (13) were shown to be jointly concave (see the proof of Theorem 1), the same applies to the function J_t. The concavity and monotonicity properties of $V^*_t(\cdot, \cdot)$ again follow immediately.

(b) Let $\hat{J}_t(y, p')$ denote the expected total net profits in periods t to 1 when adopting, in period t, in the model with bi-directional price changes, an inventory level y and price p' and making optimal decisions thereafter. Clearly, $J_t(y, p') \leqslant \hat{J}_t(y, p')$ and by Theorem 1(b), $J_t(y, p') = O(|y|^\rho)$, $V^*_t(x, p) = O(|x|^\rho)$, and $\lim_{|y| \to \infty} J_t(y, p') = -\infty$. The existence of a finite maximizer thus follows from the concavity of J_t.

(c) We show, by induction, that the function $J_t(y, p')$ has antitone differences for all $t = 1, \dots, T$. The remainder of the proof is identical to that of Theorem 2. The submodularity of J_1 was demonstrated in the proof of Theorem 2. Assume now that J_{t-1} is submodular for some $t = 2, \dots, T$. Since by parts (a) and (b) of this theorem J_{t-1} is concave and the maximization problems in (12) have finite maximizers, it follows from Lemma 2 that $V^*_{t-1}(\cdot, \cdot)$ has antitone differences as well. (Apply Lemma 2 with $u = -p$, $v = -p'$ and $f(y, v) = V^*_{t-1}(y, -v)$.) Fix a value for ϵ_t and choose an arbitrary quadruple (y_1, y_2, p_1, p_2) with $y_1 > y_2$ and $p_1 > p_2$. Note that

$$\begin{aligned} &V^*_{t-1}(y_1 - d_t(p_1, \epsilon_t), p_1) - V^*_{t-1}(y_2 - d_t(p_1, \epsilon_t), p_1) \\ &\leqslant V^*_{t-1}(y_1 - d_t(p_2, \epsilon_t), p_1) \\ &\quad - V^*_{t-1}(y_2 - d_t(p_2, \epsilon_t), p_1) \\ &\leqslant V^*_{t-1}(y_1 - d_t(p_2, \epsilon_t), p_2) \\ &\quad - V^*_{t-1}(y_2 - d_t(p_2, \epsilon_t), p_2). \end{aligned}$$

The first inequality follows from the concavity of V^*_{t-1} in its first argument and the fact that $d_t(p_1, \epsilon_t) \leqslant d_t(p_2, \epsilon_t)$ by Assumption 1. The second inequality follows from the fact that V^*_{t-1} has antitone differences. We conclude that $V^*_{t-1}(y - d_t(p', \epsilon_t), p')$ has antitone differences in y and p' and the same property therefore applies to $\mathsf{E}_{\epsilon_t} V^*_{t-1}(y - d_t(p', \epsilon_t), p')$. Since the first three terms in (13) are submodular as well, it follows that $J_t(y, p')$ is submodular, thus completing the induction step. □

Proof of Theorem 5. (a) The transformed model has nonpositive one-step expected profits. In particular, $\hat{v} \leqslant 0$ for all $t = 1, 2, \ldots$. In view of Proposition 9.17 in Bertsekas and Shreve (1978), it suffices to verify that for all $t = 1, 2, \ldots$ and all λ the sets

$$\begin{aligned} U_t(x, \lambda) = \{(y, p) : y \geqslant x, \max(p_{\min}, c) \\ \leqslant p \leqslant p_{\max} \text{ and } \hat{J}_t(y, p) \geqslant -\lambda\} \end{aligned}$$

are compact subsets of $\Re^2$. Note by the definition of M and $\hat{v}_t \leqslant 0$ that for $t = 1, 2, \ldots$

$$-cy - G(y, p) \geqslant \hat{J}_t(y, p). \tag{27}$$

Since $G(y, p)$ is convex (by Assumption 5) it is continuous, so that by Assumption 2, $\lim_{|y|\to\infty}[cy + G(y, p)] = \infty$ *uniformly* on the bounded interval $[p_{\min}, p_{\max}]$. This implies the existence of constants $\bar{y}(\lambda)$ and $\underline{y}(\lambda)$ such that for $y < \underline{y}(\lambda)$ and $y > \bar{y}(\lambda)$, $-cy - G(y, p) < -\lambda$. Thus by (27), $U_t(x, \lambda) \subseteq \{(y, p) : \max(p_{\min}, c) \leqslant p \leqslant p_{\max}$ and $\underline{y}(\lambda) \leqslant y \leqslant \bar{y}(\lambda)\}$ is a bounded set. By Theorem 1(a), J_t is concave and hence continuous and so is $\hat{J}_t = J_t - M(1 - \alpha^{t+1})/(1 - \alpha)$, which guarantees that the sets $U_t(x, \lambda)$ are closed and hence compact.

(b) It is immediate from Proposition 9.8 in Bertsekas and Shreve (1978) that $\hat{v}$ and $\hat{J}$ satisfy the optimality Equation (14) in the transformed model. Since $v^* = \hat{v} + M/(1 - \alpha)$ and $J^* = \hat{J} + M/(1 - \alpha)$, v^* and J^* satisfy the optimality equation in the original model.

(c) The concavity properties of J^* and v^* and the fact that J^* has antitone differences, follow from the corresponding properties of J_t and v^*_t (see Theorem 1(a) and the proof of Theorem 2) as well as the fact that $J^* = \lim_{t\to\infty} J_t$ and $v^* = \lim_{t\to\infty} v^*_t$; see part (a). Also, applying the proof of Theorem 1(a) to the transformed model we get $\lim_{|y|\to\infty} \hat{J}_t(y, p) = -\infty$ for all $p \in [p_{\min}, p_{\max}]$ and $\hat{J} \leqslant \hat{J}_t \leqslant \hat{J}_{t-1} \leqslant \cdots \leqslant \hat{J}_1$ as can be verified by induction. Thus, $\lim_{|y|\to\infty} J^*(y, p) = \lim_{|y|\to\infty} \hat{J}(y, p) + M/(1 - \alpha) = -\infty$, which, by the concavity of J^* implies that J^* has a finite maximizer.

It remains to be shown that $v^*(x) = O(|x|^{\rho+1})$. Clearly,

$$v^*(x) \leqslant \frac{M}{1 - \alpha}, \tag{28}$$

the maximum present value of profits when all negative (cost) components are ignored. Furthermore, let $v^*(x|p_{\min})$ denote the expected maximum infinite horizon profit when starting in state x in the basic inventory model with *static* prices which arises when adopting a constant price $p = p_{\min}$ over the entire planning horizon. Thus, $v^*(x) \geqslant v^*(x|p_{\min})$.

$$\begin{aligned} &v^*(x|p_{\min}) - \frac{\alpha p_{\min} \mathsf{E} d(p_{\min}, \epsilon)}{1 - \alpha} \\ &\quad = -c(y^* - x) - (1 - \alpha)^{-1} \\ &\qquad \cdot [G(y^*, p_{\min}) + \alpha c \mathsf{E} d(p_{\min}, \epsilon)], \quad \text{if } x \leqslant y^*. \end{aligned} \tag{29}$$

Let $\pi_0 = \Pr[d(p_{\min}, \epsilon) = 0] < 1$ by Assumption 1. Note that under the optimal order-up-to policy in the model with a fixed price $p_{\min}$, the present value of the expected order costs is bounded by the value that arises when in each period last period's demand is ordered. Thus, for $x > y^*$,

$$\begin{aligned} &v^*(x|p_{\min}) - \frac{\alpha p_{\min} \mathsf{E}(p_{\min}, \epsilon)}{1 - \alpha} \\ &\quad \geqslant -c(1 - \alpha)^{-1} \mathsf{E} d(p_{\min}, \epsilon) \\ &\qquad - (1 - \pi_0)^{-1} \sum_{l=y}^{x} G(l, p) \\ &\qquad - \alpha(1 - \alpha)^{-1} G(y^*, p_{\min}). \end{aligned} \tag{30}$$

We conclude from (29) and (30) that $v^*(x) = O(|x|^{\rho+1})$.

(d) Immediate from the proof of Theorem 1(b).

(e) Fix $t = 1, 2, \ldots$. Since $\hat{J}(y, p) \leqslant \hat{J}_t(y, p) \leqslant -G(y, p)$ for all (y, p), $-\lambda \doteq \hat{J}(y^*, p^*) \leqslant \hat{J}_t(y^*_t, p^*_t) \leqslant -G(y^*_t, p^*_t)$. By the proof of part (a), there exist two constants, $\underline{y}(\lambda)$ and $\bar{y}(\lambda)$, whose values are independent of α, such that $\underline{y}(\lambda) \leqslant y^*_t \leqslant \bar{y}(\lambda)$. In other words, the sequence $\{(y^*_t, p^*_t)\}$ is bounded and thus has at least one limit point (y^*, p^*). Clearly, $\underline{y} \leqslant y^* \leqslant \bar{y}$. Recall that $\lim_{t\to\infty} J_t = J$ (part (a)) and that (y^*_t, p^*_t) is a maximizer of the concave function J_t, so that the vector 0 is a subgradient of J_t in (y^*, p^*). It follows from Rockefellar (1970, Theorem 24.5) that 0 is a subgradient of J in the point (y^*, p^*) so that (y^*, p^*) is a maximizer of J. □

REFERENCES

Amihud, Y., H. Mendelson. 1983. Price smoothing and inventory. *Rev. Econom. Studies* **50** 87–98.

Aviv, Y., A. Federgruen. 1995. The value iteration method for countable state Markov decision processes. To appear in *OR Letters*.

——, ——. 1997. Stochastic inventory models with limited production capacity and periodically varying demands. *Probability Engrg. Informational Sci.* **11** 107–135.

Balvers, R. J., T. F. Cosimano. 1990. Actively learning about demand and the dynamics of price adjustment. *Econom. J.* **100** 882–898.

Bassok, Y., R. Anupindi. 1997. Analysis of supply contracts with total minimum commitment. *IIE Trans.* **29** 373–381.

——, ——. 1995. Analysis of supply contracts with forecasts and flexibility. Working Paper. Northwestern University, Dept. of Industrial Engineering and Operations Research, Evanston, IL.

Bertsekas, D. P., S. E. Shreve. 1978. *Stochastic Optimal Control*. Academic Press, New York.

Bitran, G. R., S. V. Mondschein. 1993. Perishable product pricing: an application to the retail industry. Working Paper. Massachusetts Institute of Technology, Cambridge, MA.

——, ——. 1995. Periodic pricing of seasonal products in retailing. *Management Sci.* **43** 64–79.

Braden, D. J., S. S. Oren. 1994. Nonlinear pricing to produce information. *Marketing Sci.* **13** 310–326.

Eliashberg, J., R. Steinberg. 1991. Marketing-production joint decision making. J. Eliashberg, J. D. Lilien, eds. *Management Science in Marketing*, Volume 5 of *Handbooks in Operations Research and Management Science*, North Holland, Amsterdam.

Eppen, G. D., A. V. Iyer. 1995. Improved fashion buying with Bayesian updates. Working Paper. Graduate School of Business, University of Chicago, Chicago, IL.

——, ——. 1997. Backup agreements in fashion buying—the value of upstream flexibility. Management Sci. **43** 1469–1484.

Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87–99.

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

——, ——. 1994b. A multi-product dynamic pricing problem and its applications to network yield management. *Oper. Res.* **45** 24–41.

Grossman, S., R. E. Kihlstrom, L. J. Mirman. 1977. A Bayesian approach to the production of information and learning by doing. *Rev. Econom. Studies* **44** 533–547.

Heching, A., G. van Ryzin, G. Gallego. 1999. A theoretical and empirical investigation of markdown pricing in fashion retailing. Working Paper. Graduate School of Business, Columbia University, New York.

Hempenius, A. L. 1970. *Monopoly with Random Demand*. Universitaire Pers Rotterdam, Rotterdam, The Netherlands.

Heyman, D. P., M. J. Sobel. 1984. *Stochastic Models in Operations Research Volume II*. McGraw-Hill, New York.

Karlin, S., C. R. Carr. 1962. Prices and optimal inventory policy. *Studies in Applied Probability and Management Science*. K. Arrow, S. Karlin, H. Scarf, eds. Stanford University Press, Stanford, CA.

Kunreuther, H., L. Schrage. 1973. Joint pricing and inventory decisions for non-seasonal items. *Econometrica* **39** 173–175.

Lau, A. H., H. Lau. 1988. The newsboy problem with price-dependent demand distribution. *IIE Trans.* **20** 168–175.

Lee, H. L., S. Nahmias. 1993. Single product single location models. S. C. Graves, A. H. G. Rinnooy Kan, P. H. Zipkin, eds. *Logistics of Production and Inventory*. *Handbooks in Operations Research and Management Science*, Elsevier Science Publishers B.V., Amsterdam, The Netherlands.

Li, L. 1988. A stochastic theory of the firm. *Math. Oper. Res.* **13** 447–466.

McLennan, A. 1984. Price dispersion and incomplete learning in the long run. *J. Econom. Dynamics Control* **7** 331–347.

Mills, E. S. 1959. Uncertainty and price theory. *Quart. J. Econom.* **73** 116–130.

——. 1962. *Price, Output and Inventory Policy*. John Wiley, New York.

Polatoglu, L. H. 1991. Optimal order quantity and pricing decisions in single period inventory systems. *Internat. J. Production Econom.* **23** 175–185.

Porteus, E. L. 1990. Stochastic inventory theory. D. P. Heyman, M. J. Sobel, eds. *Handbooks in OR and MS*, *Vol. 2*. Elsevier Science Publishers B.V., Amsterdam, The Netherlands.

Rajan, A. Rakesh, R. Steinberg. 1992. Dynamic pricing and ordering decisions by a monopolist. *Management Sci.* **38** 240–262.

Rockafellar, R. T. 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ.

Rothschild, M. 1974. A two-armed bandit theory of market pricing. *J. Econom. Theory* **9** 185–202.

Sennott, L. I. 1989. Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Oper. Res.* **37** 626–633.

Thomas, L. J. 1974. Price and production decisions with random demand. *Oper. Res.* **22** 513–518.

Thowsen, G. T. 1975. A dynamic, nonstationary inventory problem for a price/quantity setting firm. *Naval Res. Logist. Quart.* **22** 461–476.

Topkis, D. 1978. Minimizing a submodular function on a lattice. *Oper. Res.* **26** 305–321.

Veinott, A. Jr. 1966. On the optimality of (s, S) inventory policies: new conditions and a new proof. *SIAM J. Appl. Math.* **14** 1067–1083.

Wagner, H. M., T. M. Whitin. 1958. Dynamic problems in the theory of the firm. *Naval Res. Logist. Quart.* **7** 7–12.

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

Young, L. 1978. Price, inventory and the structure of uncertain demand. *New Zealand Oper. Res.* **6** 157–177.

Zabel, E. 1970. Monopoly and uncertainty. *Rev. Econom. Studies* **37** 205–219.

——. 1972. Multiperiod monopoly under uncertainty. *J. Econom. Theory* **5** 524–536.