

Managing Fashion Goods Inventories:

Nicholas Petruzzi

University of Illinois at Urbana–Champaign

George E. Monahan

University of Illinois at Urbana–Champaign

Abstract

The proliferation of both online and bricks and mortar outlet stores underscores the observation that secondary markets are readily accessible to retailers of short-life-cycle products. These secondary markets provide recourse channels for retailers to sell excess inventory of out of favor items at reduced prices when overstocking occurs in a primary market. We study the problem of determining when a retailer should terminate its primary selling season by selling remaining inventory on a secondary market. The retailer has a single opportunity to procure prior to a primary selling season consisting of multiple periods. Demand in each period is random, but correlated. At the end of each period, any remaining inventory incurs a holding cost. Then, based upon the current level of inventory and the cumulative demand-to-date, the retailer decides either to terminate the primary selling season by selling all or part of the remaining inventory on a secondary market, or to extend the current primary selling season by another period. We develop structural properties of the optimal policy for determining when to terminate the primary selling season, and we develop corresponding implications for procurement.

MANAGING FASHION GOODS INVENTORIES: DYNAMIC RECOURSE FOR RETAILERS WITH OUTLET STORES

NICHOLAS C. PETRUZZI ¹

GEORGE E. MONAHAN ²

DECEMBER 2002

Abstract

The proliferation of both online and bricks and mortar outlet stores underscores the observation that secondary markets are readily accessible to retailers of short-life-cycle products. These secondary markets provide recourse channels for retailers to sell excess inventory of out-of-favor items at reduced prices when overstocking occurs in a primary market. We study the problem of determining when a retailer should terminate its primary selling season by selling remaining inventory on a secondary market. The retailer has a single opportunity to procure prior to a primary selling season consisting of multiple periods. Demand in each period is random, but correlated. At the end of each period, any remaining inventory incurs a holding cost. Then, based upon the current level of inventory and the cumulative demand-to-date, the retailer decides either to terminate the primary selling season by selling all or part of the remaining inventory on a secondary market, or to extend the current primary selling season by another period. We develop structural properties of the optimal policy for determining when to terminate the primary selling season, and we develop corresponding implications for procurement.

¹Corresponding author. Assistant Professor, Department of Business Administration, University of Illinois at Urbana-Champaign, 1206 South Sixth Street Champaign, IL 61820. Phone: 217-333-9578. e-mail: petruzzi@uiuc.edu.

²Professor, Department of Business Administration, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, IL 61820. Phone: 217-333-8270. e-mail: gmonahan@uiuc.edu.

1 Introduction

The proliferation of both online and bricks and mortar outlet stores underscores the observation that secondary markets are readily accessible to retailers of short-life-cycle products. These secondary markets provide recourse channels for retailers to sell excess inventory of out-of-favor items at reduced prices when overstocking occurs in a primary market. One key reason for such overstocking is that original procurement decisions for fashion goods typically occur well before the actual selling season begins; thus, they often are subject to unreliable demand forecasts. Yet, once such items are introduced to the market and actual demand patterns are observed, demand projections can be revised and recourse actions can be exercised.

The identification and implementation of recourse options when there is only one procurement opportunity have received much attention in the literature, thanks to the heightened awareness of the benefits of supply-chain coordination. Considerable academic attention, for example, has been given to the design of contract parameters between a buyer and a seller to govern postponement alternatives. Postponement is a strategy in which a buyer is afforded the procurement flexibility either to place a tentative order with a supplier only to adjust that order at a later time once more accurate market information becomes available or to delay placing an order to achieve the same effect. Mark-down pricing is another strategy that has the effect of adjusting an original procurement decision. Effectively, mark-down pricing helps coordinate market demand to inflexible supply. We investigate a related strategy in this paper: We analyze a problem from the perspective of a retailer who has the option of terminating its primary selling season by salvaging its inventory in a secondary market if, at any time during the season, the retailer determines that the demand in the primary market, which is characterized by fashion-conscious consumers, does not warrant the continued stocking of the item.

The focus of this paper is two-fold: First, what are the structural properties that define an optimal termination strategy when salvage recourse for remaining inventory exists? And second, what are the corresponding implications on the retailer's original procurement decision? We answer these questions within the following modeling framework: A retailer has a single opportunity to procure prior to the start of a primary selling season that consists of multiple periods. Demand in each period is random but is influenced by fashion, meaning that demand in any period is positively correlated with cumulative past demand. At the end of each period, any remaining inventory incurs a holding cost. Then, based upon the current level of inventory

and the cumulative demand-to-date, the retailer decides either to terminate the primary selling season by selling its remaining inventory on a secondary market, or to extend the current primary selling season by another period. If the retailer decides to continue for another period, demand in the primary market is realized, a new level of inventory is obtained, and the decision regarding termination of the primary selling season is revisited.

Our model makes a contribution to the literature primarily by recognizing the existence of salvage channels as a recourse alternative to retailers of fashion goods, and by incorporating the economic potential of such an alternative into the retailer's procurement decision. In addition, by explicitly accounting for a salvage opportunity, which represents a form of *downstream* recourse available to the retailer, our model complements both the literature addressing mark-down pricing, which provides an alternative form of downstream recourse, and the literature addressing postponement, which provides various forms of *upstream* recourse.

The remainder of this paper is organized as follows. In Section 2, we position the paper in relation to the existing literature. In Section 3, we specify the formal mathematical model and formulate the retailer's dynamic decision problem. In Section 4, we establish the structure of the optimal termination policy, and in Section 5 we discuss implications for procurement. Section 6 contains a numerical example that illustrates the manner in which an optimal termination policy can be computed. It also illustrates how certain parameter values may influence the optimal termination and the optimal procurement decisions. We conclude with a discussion of the applicability and scope of our model in Section 7.

2 Relationship to the Literature

There are two key aspects of our model, the specification of two non-overlapping market segments and the incorporation of dynamic recourse, that, collectively, either generalize, extend, or complement models from four streams of literature. To provide context for the positioning of our paper in relation to these streams, we first provide a more detailed description of the two basic elements of our model.

Fundamental to our model is the specification of two non-overlapping markets, each of which is represented by a distinct demand characterization. In particular, we specify a primary market that consists of consumers who value fashion over price, and a secondary market consisting of consumers who value price over fashion. Correspondingly, demand in the primary mar-

ket is positively influenced by cumulative past demand, whereas demand in the secondary market is immune to such contagious effects. Essentially, the primary market is constructed to represent a suitable target for a fashion item that turns out to be a “winner”; whereas the secondary market represents a fall-back target for a fashion item that turns out to be a “loser”.

The incorporation of dynamic recourse in our model also is important because a retailer’s decision to target a fashion-conscious market need not be an irrevocable one. If, at any time, the retailer determines that its fashion product is a “loser”, which is a determination that the retailer can reconsider periodically as it observes demand over time, it can react by effectively switching target markets from the primary market to the secondary market.

One stream of literature directly related to our work is that on combined forecasting and inventory management of style goods when demand growth is explicitly considered. Relevant models include those developed by Ravindran (1972), Hartung (1973), Kurawarwala and Matsuo (1996), and Hertz and Schaffir (1960). Ravindran’s model, in particular, can be viewed as a special case of the one we formulate. Ravindran analyzes a single-market inventory model in which both stocking quantity and season length are chosen to maximize total expected profit over the season, given that demand is specified as a non-homogenous Poisson process with rate proportional to cumulative demand-to-date. However, both decisions of his model are fixed at the start of the selling season; consequently, his model does not incorporate dynamic recourse. In contrast, our model generalizes the demand process while retaining the contagious effect of past demand; and it treats the decision to terminate the primary selling season as a dynamic one.

Hartung (1973) incorporates a demand process similar to the process specified by Ravindran (1972). However, unlike Ravindran, Hartung develops a dynamic model that uses on-going demand observations to refine the forecast of demand and to enhance the management of inventory. Hartung’s model nevertheless differs from our model in the form of the recourse available once an updated forecast of demand is obtained: Hartung assumes a fixed season length and allows for periodic procurement opportunities; we assume a single procurement opportunity and allow for the termination of the primary selling season. Kurawarwala and Matsuo (1996) and Hertz and Schaffir (1960) also consider time-dependent demand models within the context of managing a fashion good. However, they emphasize the development and impact of practical forecasting methodologies while we emphasize the benefit and impact of dynamic recourse.

A second, closely-related literature stream concerns two variations of multiple-period inventory models. Porteus (1990) provides an excellent over-

view of this general body of literature. However, one variation that is related to this paper concerns time-dependent demand. Karlin (1960) presents an infinite horizon, multi-period inventory model with time-dependent demands. Production decisions are made each period and, since the time horizon is infinite, disposal is not an issue. He establishes the optimality of a base stock policy. In addition, he shows that if demand increases stochastically over time, the optimal base stock values also increase. In our model, current period demand is an explicit function of past demand and is not specified as a general stochastic ordering. Furthermore, we have only a single procurement decision that is made at the beginning of the planning horizon. Each period, the decision is not how much to make but whether or not to continue selling in the primary market.

The second variation of multi-period inventory models that is related to the model of this paper comprises those models that allow for salvaging (or disposal). Representative papers include Eppen and Iyer (1997b) and Lovejoy (1990), both of which incorporate Bayesian learning, and Fukuda (1961), which does not. Our model differs from this stream in the following respects. Each of these papers considers a multi-procurement business setting in which demand does not grow over time and the available salvage option may be exercised more than once. Effectively, in these papers, salvaging complements procurement as a mechanism to maintain a desired inventory level: At the beginning of a new period, if the inventory level is below the desired stocking level for the period, additional units are procured. However, if the inventory level is above the desired stocking level, excess units are salvaged. In contrast, we consider a single-procurement business setting in which demand does grow over time and the available salvage option may be exercised only once, to exit one market segment and to enter an alternative one. We make this assumption to avoid the possibility of having the same product available to both markets, for two different prices, simultaneously. Such a possibility may be undesirable. As more retailers turn to outlet stores for salvaging opportunities, the easier it becomes for primary-market customers to notice if an item is being offered simultaneously for multiple prices, thereby resulting either in an erosion of the primary market or in an irreversible damage to good will.

A third germane literature stream investigates dynamic pricing as a form of recourse for selling a fixed stock of items over a finite horizon. Such dynamic models represent a special set of problems called *perishable asset revenue management (PARM)* models. As a class, PARM models effectively prescribe methods for maximizing expected revenues by managing fixed inventories at different prices. McGill and van Ryzin (1999) provide an insight-

ful survey of this literature. Papers of this type that are most closely related to our model include Feng and Gallego (1995), Bitran and Mondschein (1997), Gallego and van Ryzin (1994), and Zhao and Zheng (2000). Feng and Gallego's (1995) model, in particular, is similar to ours in the sense that it explicitly includes the formulation of a case in which a retailer may make only one price change during the course of a selling season, which is analogous to our specification of a single opportunity to switch between market segments. However, like the other papers identified in this literature stream, it differs from our model in two respects: one, it specifies demand as a Poisson process; and two, it does not assume any element of demand growth. More recently, Feng and Gallego (2000) extended their earlier model to include a case in which demand is characterized by Markovian, time-dependent intensities that depend on cumulative demand. Nevertheless, they approach the problem within the context of optimal control theory and they treat the starting inventory level as an exogenous parameter. In contrast, we generalize the demand process, employ a dynamic-programming approach, and treat the starting inventory level as a decision variable.

A fourth literature stream analyzes postponement strategies, which comprise techniques for delaying final procurement decisions in hopes of obtaining more accurate demand information. Tsay et al. (1999) provide a timely and comprehensive review of this area. Representative papers include Eppen and Iyer (1997a), Bassok and Anupindi (1997), Tsay and Lovejoy (1999), Iyer and Bergen (1997), and Fisher and Raman (1996). Although these papers focus on the supplier-retailer relationship, thereby adopting a perspective different from the one we adopt in our model, they are important to note because our model, which focuses on a form of downstream recourse available to the retailer, complements these papers, which instead focus on various forms of upstream recourse available to the retailer.

3 Model and Notation

A primary selling season consists of at most T time periods. Prior to the beginning of the season, an initial stock of Q units is procured at a per-unit cost of c and is made available for sale in the primary market at a per-unit selling price of r . At the beginning of each subsequent period in the primary selling season, any remaining inventory is assessed a holding cost of h per unit, and then a decision is made either to face the current period demand in the primary market or to terminate the selling season. Any residual demand not satisfied in the primary market incurs a per-unit lost sales penalty of

π . There are two reasons why residual demand can go unsatisfied: either because the primary season is *not* terminated but demand exceeds available supply (in which case demand in the primary market is involuntarily lost), or because the primary selling season *is* terminated (in which case demand in the primary market is voluntarily forfeited).

If, at the beginning of period t , the decision to terminate is made, any remaining stock is sold in a secondary market at a price of $v_t < r$ per unit. In other words, we assume that the salvage function associated with terminating the primary selling season is proportional to the inventory remaining when the decision to terminate is made. This assumption directly applies to a situation in which the salvage price is exogenous (because, for example, it depends on the *function* of the good, not on the *style* of the good), and the salvage demand is ample relative to the primary demand. This assumption would also apply if the retailer's supplier offered a liberal returns policy, namely, one in which the retailer can return any unused stock, at any time, for a pre-specified, per-unit credit. In such a case, the supplier serves as the retailer's salvage market. We discuss implications and challenges of a non-linear salvage function in Section 7.

We assume that demand in the primary market, although random, is *contagious*; that is, it is an increasing function of cumulative past demand. (In our model, sales and realized demand as of period t are equivalent if the inventory available for sale at the beginning of period t is positive. Thus, whenever the retailer is faced with a termination decision, total cumulative sales is equivalent to total realized demand.) Our assumption that single-period demand is increasing in cumulative past demand is consistent with related models involving the management of fashion goods (e.g., Ravindran (1972) and Hartung (1973)), and applies to situations in which potential demand is large relative to the number of customers that already have made purchases. It also is prevalent in marketing models in which word of mouth effects are explicitly incorporated into the specification of periodic demand (e.g., Monahan and Nti (1988), Mahjan and Muller (1979), and Horsky and Simon (1983)), and in diffusion models of pricing that ignore saturation (e.g., Bass (1980) and Bass and Bultez (1982)).

To formalize the notion of contagious demand, we let X_t denote the period- t random demand in the primary market. We consider two general specifications of X_t : the *multiplicative* case, defined by $X_t = \mu(s_t)Z_t$, and the *additive* case, defined by $X_t = \mu(s_t) + Z_t$. In both cases, the Z_t are independent, identically distributed random variables with a known cumulative distribution function (cdf) Φ and a corresponding probability density function (pdf) ϕ ; and $\mu(s_t)$ is known and is strictly increasing in s_t , where s_t is the

cumulative demand-to-date. Note that the multiplicative case is a generalization of the classic *percent-done* formulation (Hartung (1973)), which refers to the special case in which $\mu(s_t) = s_t$. The additive case is a generalization of ARMA (auto-regressive/moving average) models.

Given these specifications, we adopt the following definitions:

- $F(x|s_t)$ cdf of X_t .
- $f(x|s_t)$ pdf of X_t .
- Y_t random variable representing cumulative demand in the primary market as of the *end* of period t ; i.e., $Y_t \equiv s_t + X_t$.
- $G(y|s_t)$ cdf of Y_t ; i.e., $G(y|s_t) = F(y - s_t|s_t)$.
- $g(y|s_t)$ pdf of Y_t ; i.e., $g(y|s_t) = f(y - s_t|s_t)$.

There are two parts to the retailer's decision problem. First, the retailer must choose Q , the amount of stock to procure prior to the beginning of the primary selling season. And second, the retailer must decide after each period whether to continue operating in the primary market or to terminate the primary selling season by salvaging its remaining inventory on the secondary market. The decision to continue or to terminate the primary selling season is a dynamic one, made each time the retailer observes cumulative past demand in the primary market and compares the result to its remaining inventory.

Accordingly, we next construct a dynamic program that captures the total expected profit associated with following an optimal termination policy for as long as the retailer continues to operate in the primary market. To that end, let $W_t(s_t, Q)$ be the maximum expected profit that the retailer can generate from period t onward, given that Q units are procured prior to the start of the primary selling season, and given that, as of the beginning of period t , the primary selling season has not been terminated and a total of s_t units have been demanded (and sold) to date. Then the procurement decision prior to the start of the primary selling season is the solution to the following optimization problem:

$$\max_Q \{W_1(0, Q) - cQ\}. \quad (1)$$

To develop the recursive structure of $W_t(s_t, Q)$, suppose that it currently is the beginning of period t of the primary selling season and s_t units have been demanded to date. It follows that $Q - s_t$ units remain in inventory. However, the retailer's expected profit for the period depends on the decision to terminate or not to terminate the primary selling season. If the decision is made to terminate the primary selling season, the retailer salvages its remaining inventory on the secondary market, forfeits all residual demand

from period t to T in the primary market, and the dynamic-decision problem ends. The corresponding period- t contribution to expected profit is

$$v_t \cdot (Q - s_t) - \pi \mathbb{E} \left[\sum_{i=t}^T X_i | s_t \right].$$

The last term in this expression represents what total expected demand would have been *if* the retailer continued to operate in the primary market (without stocking out) through the end of the selling season. Consequently, this is the demand that the retailer forfeits if the primary season is terminated.

If the decision is made not to terminate the primary selling season, the retailer faces demand in period t in the primary market, which results in a period- t contribution to total expected profit of

$$\begin{aligned} M(s_t, Q) &= r \mathbb{E} [\text{Period-}t \text{ Sales}(s_t, Q)] - h \mathbb{E} [\text{Period-}t \text{ Leftovers}(s_t, Q)] \\ &\quad - \pi \mathbb{E} [\text{Period-}t \text{ Shortages}(s_t, Q)] \\ &= (r + \pi) \cdot (Q - s_t) - (r + \pi + h) \int_0^{Q-s_t} F(x|s_t) dx - \pi \mathbb{E} [X_t | s_t], \end{aligned}$$

but either of two outcomes is possible. On the one hand, if it turns out that $X_t \geq Q - s_t$, then the dynamic decision problem still ends and the retailer loses all residual demand from period $t + 1$ to T because all of the available stock is depleted. On the other hand, if $X_t < Q - s_t$, then

- i. $s_{t+1} = s_t + x_t$, where x_t is the realized demand in period t ;
- ii. $Q - s_{t+1}$ is the stock available for sale in period $t + 1$;
- iii. the decision to terminate or not to terminate the primary selling season repeats.

Thus, the termination decision in period t depends on the future stream of single-period contributions to expected profit, which, in turn, depends on the length of the sequence of non-termination decisions. Formally, given the definitions of $W_t(s_t, Q)$ and $M(s_t, Q)$, $W_t(s_t, Q)$ satisfies the following

recursion for $t = 1, \dots, T$:

$$\begin{aligned}
W_t(s_t, Q) = \max & \left\{ v_t \cdot (Q - s_t) - \pi \mathbb{E} \left[\sum_{i=t}^T X_i | s_t \right], M(s_t, Q) \right. \\
& - \pi \int_{Q-s_t}^{\infty} \mathbb{E} \left[\sum_{i=t+1}^T X_i | (s_t + x) \right] f(x | s_t) dx \\
& \left. + \int_0^{Q-s_t} W_{t+1}(s_t + x, Q) f(x | s_t) dx \right\}, \tag{2}
\end{aligned}$$

where $W_{T+1}(s_{T+1}, Q) = \max\{v_{T+1} \cdot (Q - s_{T+1}), 0\}$. In other words, the optimal decision at the beginning of period t , given that the primary selling season has not been terminated as of that time, is to face the demand in the primary market in period t if and only if

$$\begin{aligned}
M(s_t, Q) - \pi \int_{Q-s_t}^{\infty} \mathbb{E} \left[\sum_{i=t+1}^T X_i | (s_t + x) \right] f(x | s_t) dx \\
+ \int_0^{Q-s_t} W_{t+1}(s_t + x, Q) f(x | s_t) dx > v_t \cdot (Q - s_t) - \pi \mathbb{E} \left[\sum_{i=t}^T X_i | s_t \right].
\end{aligned}$$

For convenience in making that assessment, we define the following three functions:

$$L_t(s_t, Q) = (r + \pi - v_t)(Q - s_t) - (r + \pi + h - v_{t+1}) \int_0^{Q-s_t} F(x | s_t) dx, \tag{3}$$

$$J_t(s_t, Q) = L_t(s_t, Q) + \int_0^{Q-s_t} V_{t+1}(s_t + x, Q) f(x | s_t) dx, \tag{4}$$

and

$$V_t(s_t, Q) = \max\{0, J_t(s_t, Q)\}. \tag{5}$$

Using these definitions, the recursion in (2) can be rewritten as:

$$W_t(s_t, Q) = v_t \cdot (Q - s_t) - \pi \mathbb{E} \left[\sum_{i=t}^T X_i | s_t \right] + V_t(s_t, Q), \tag{6}$$

where $V_{T+1}(s_{T+1}, Q) = 0$; and, for $t = 1, \dots, T$, $V_t(s_t, Q)$ is determined by (5), (4), and (3).

As a result of this transformation, the analysis of the retailer's procurement decision in (1) now corresponds to the problem of finding the Q that maximizes

$$V_1(0, Q) - (c - v_1) \cdot Q - \pi E \left[\sum_{i=1}^T X_i | s_1 = 0 \right]. \quad (7)$$

Furthermore, the analysis of the retailer's decision problem at the beginning of period t (for $t = 2, \dots, T$) corresponds to the determination of whether or not $J_t(s_t, Q) > 0$. If this test is satisfied, then the retailer's optimal decision in period t is to face period- t demand in the primary market; if this test is not satisfied, then the retailer's optimal decision in period t is to terminate the primary selling season and salvage its remaining stock in the secondary market.

The transformed version of the retailer's dynamic-decision problem is equivalent to the following scenario. If the retailer faces demand in the primary market in period $t - 1$, then at the end of that period, the retailer temporarily agrees to sell all of the resulting leftovers to a salvage dealer for the price of v_t per-unit. Then the retailer assesses whether or not it would be profitable, in expectation, to sell in the primary market for another period. If the assessment results in a "yes", then the retailer backs out of its agreement to sell its leftovers to the salvage dealer, effectively buying back all of its stock at a cost of v_t per-unit. Thus, the retailer's dynamic-decision scenario can be thought of as a multi-period, "take-it-or-leave-it" newsvendor problem in which the period- t purchase cost is v_t per unit, and the corresponding end-of-period net holding cost is $h - v_{t+1}$ per unit, but in which the purchase quantity is fixed: the "newsvendor" may not deviate from the pre-specified purchase quantity, it may only decide against making the purchase at all. This interpretation of the problem helps explain the results of an illustration provided in the next section.

4 Optimal Termination Policy

In this section, we establish that the retailer's optimal termination policy is unique; thus, it can be determined efficiently using standard search techniques. Moreover, we demonstrate that the optimal policy is, in general, a *sales-target termination policy*, which is a policy that is characterized by a single critical level of cumulative demand-to-date for each period, which we call a *sales target*. In a special case, the optimal policy is a *monotone sales-target termination policy*, which is a policy in which the sales targets increase over

time. Thus, if cumulative demand as of the beginning of any period fails to meet the sales target associated with that period, then it behooves the retailer to terminate the primary selling season and to cut its losses by salvaging its remaining inventory on the secondary market. In other words, the sales targets provide benchmarks for the retailer to ascertain whether its product is a “winner” or a “loser”.

Lemma 4.1. *For a given Q , the retailer’s optimal termination policy for period T , the last possible period of the primary selling season, can be characterized uniquely by the sales target s_T^* .*

Proof. First, we define the function $U(s_t, Q)$ as follows:

$$U(s_t, Q) \equiv \int_{s_t}^Q (Q - y)g(y|s_t) dy = \int_{s_t}^Q G(y|s_t) dy. \quad (8)$$

Note that $U(s_t, Q)$ indicates the expected number of leftovers at the end of any period t in which the retailer operates in the primary market, given that s_t is the cumulative demand-to-date as of the beginning of this period (i.e., given that $Q - s_t$ is the inventory available for sale in period t).

Applying (8) and the definition of $G(y|s_t)$ to (3) yields

$$L_t(s_t, Q) = (r + \pi - v_t)(Q - s_t) - (r + \pi + h - v_{t+1})U(s_t, Q), \quad (9)$$

which implies

$$\frac{\partial L_t(s_t, Q)}{\partial s_t} = -(r + \pi - v_t) - (r + \pi + h - v_{t+1})\frac{\partial U(s_t, Q)}{\partial s_t}. \quad (10)$$

Next, for mathematical convenience, we define the function $P(s_t, Q, y)$ as follows:

$$P(s_t, Q, y) \equiv (Q - y)g(y|s_t) + (Q - s_t)\frac{\partial G(y|s_t)}{\partial s_t}. \quad (11)$$

Note that $P(s_t, Q, y) < -(y - s_t)g(y|s_t) < 0$ for all $y > s_t$ since, from Lemma A.1 in the Appendix, $\partial G(y|s_t)/\partial s_t \leq -g(y|s_t)$. Thus, (9)-(11) establishes the following inequality for all t :

$$L_t(s_t, Q) + (Q - s_t)\frac{\partial L_t(s_t, Q)}{\partial s_t} = -(r + \pi + h - v_{t+1})\int_{s_t}^Q P(s_t, Q, y) dy > 0. \quad (12)$$

Notice from (4) that $J_T(s_T, Q) = L_T(s_T, Q)$; thus, the retailer’s optimal decision at the beginning of period T , given that the primary selling season has not been terminated as of that time, is to face period- T demand in the

primary market if and only if $L_T(s_T, Q) > 0$. Since (12) indicates that, for all $s_T < Q$, $\partial L_T(s_T, Q)/\partial s_T > 0$ for any s_T that satisfies $L_T(s_T, Q) = 0$, it follows that no more than one value of s_T can satisfy $L_T(s_T, Q) = 0$. Therefore, there exists a unique value of s_T , namely s_T^* , that is such that $L_T(s_T, Q) > 0$ if and only if $s_T > s_T^*$. In other words, s_T^* denotes the retailer's sales target for period T . \square

The existence and uniqueness of an optimal sales target termination policy now can be extended to all periods of the primary selling season by induction.

Theorem 4.1. *For a given Q , the retailer's optimal termination policy is characterized by the sequence of sales targets, s_1^*, \dots, s_T^* , where, for all t , s_t^* is the unique value of s_t that satisfies $J_t(s_t, Q) = 0$.*

Proof. The proof follows by induction on t , given the following two induction hypotheses:

- i. $\frac{\partial J_t(s_t, Q)}{\partial s_t} > -(r + \pi - v_t)$; and
- ii. $J_t(s_t, Q) + (Q - s_t) \frac{\partial J_t(s_t, Q)}{\partial s_t} > 0$.

If $t = T$, then $J_T(s_T, Q) = L_T(s_T, Q)$; thus, induction hypothesis (i) follows from (10) since $\partial U(s_T, Q)/\partial s_T < 0$ by Lemma A.2 in the Appendix; and induction hypothesis (ii) follows since it is a restatement of (12). Assume, then, that induction hypotheses (i) and (ii) are true for $t = i + 1$ and consider the case $t = i$.

Induction hypothesis (i). Let $\bar{s}_i \equiv \max\{s_i, s_{i+1}^*\}$, where s_{i+1}^* satisfies $J_{i+1}(s_{i+1}^*, Q) = 0$. For $s_{i+1} < Q$, induction hypothesis (ii) implies that $\partial J_{i+1}(s_{i+1}, Q)/\partial s_{i+1} > 0$ for all s_{i+1} that satisfy $J_{i+1}(s_{i+1}, Q) < 0$. Thus, s_{i+1}^* is unique. Moreover, $J_{i+1}(s_{i+1}, Q) > 0$ if and only if $s_{i+1} > s_{i+1}^*$. Therefore, from (4) and (5) and integration by parts, we can write:

$$J_i(s_i, Q) = L_i(s_i, Q) - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} G(y|s_i) dy, \quad (13)$$

which implies

$$\frac{\partial J_i(s_i, Q)}{\partial s_i} = \frac{\partial L_i(s_i, Q)}{\partial s_i} - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} \frac{\partial G(y|s_i)}{\partial s_i} dy \quad (14)$$

because $G(\bar{s}_i|s_i) = 0$ if $\bar{s}_i = s_i$ and $d\bar{s}_i/ds_i = 0$ if $\bar{s}_i = s_{i+1}^*$. From Lemma A.1, $\partial G(y|s_i)/\partial s_i < 0$, and, from induction hypothesis (i), $\partial J_{i+1}(y, Q)/\partial y > -(r + \pi - v_{i+1})$. Thus, (14) implies

$$\begin{aligned} \frac{\partial J_i(s_i, Q)}{\partial s_i} &> \frac{\partial L_i(s_i, Q)}{\partial s_i} + (r + \pi - v_{i+1}) \int_{\bar{s}_i}^Q \frac{\partial G(y|s_i)}{\partial s_i} dy \\ &= \left[-(r + \pi - v_i) - (r + \pi + h - v_{i+1}) \frac{\partial U(s_i, Q)}{\partial s_i} \right] \\ &\quad + (r + \pi - v_{i+1}) \left[\frac{\partial U(s_i, Q)}{\partial s_i} - \frac{\partial U(s_i, \bar{s}_i)}{\partial s_i} \right] \\ &= -(r + \pi - v_i) - h \frac{\partial U(s_i, Q)}{\partial s_i} - (r + \pi - v_{i+1}) \frac{\partial U(s_i, \bar{s}_i)}{\partial s_i} \\ &> -(r + \pi - v_i), \end{aligned}$$

where the second line follows from (10) and Lemma A.2, and the last line follows from Lemma A.2.

Induction hypothesis (ii). From (4), (5), and induction hypothesis (ii),

$$\begin{aligned} J_i(s_i, Q) &= L_i(s_i, Q) + \int_{\bar{s}_i}^Q J_{i+1}(y, Q) g(y|s_i) dy \\ &> L_i(s_i, Q) - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} (Q - y) g(y|s_i) dy. \end{aligned}$$

Moreover, from (14),

$$(Q - s_i) \frac{\partial J_i(s_i, Q)}{\partial s_i} = (Q - s_i) \frac{\partial L_i(s_i, Q)}{\partial s_i} - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} (Q - s_i) \frac{\partial G(y|s_i)}{\partial s_i} dy.$$

Therefore, applying (11) and (12),

$$\begin{aligned} J_i(s_i, Q) + (Q - s_i) \frac{\partial J_i(s_i, Q)}{\partial s_i} &> L_i(s_i, Q) + (Q - s_i) \frac{\partial L_i(s_i, Q)}{\partial s_i} - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} P(s_i, Q, y) dy \\ &= (r + \pi + h - v_{i+1}) \int_{s_i}^Q [-P(s_i, Q, y)] dy - \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial y} P(s_i, Q, y) dy. \end{aligned}$$

However, since $P(s_i, Q, y) < 0$ and $s_i \leq \bar{s}_i$, and since $\partial J_{i+1}(y, Q)/\partial y >$

$-(r + \pi - v_{i+1})$ (by induction hypothesis (i)), this reduces to

$$J_i(s_i, Q) + (Q - s_i) \frac{\partial J_i(s_i, Q)}{\partial s_i} > \int_{\bar{s}_i}^Q [-P(s_i, Q, y)] \left[(r + \pi + h - v_{i+1}) + \frac{\partial J_{i+1}(y, Q)}{\partial y} \right] dy > 0.$$

Induction hypothesis (ii) indicates that, for all $s_t < Q$, $\partial J_t(s_t, Q) / \partial s_t > 0$ for any s_t that satisfies $J_t(s_t, Q) = 0$. Therefore, for any given Q , there exists a unique value of s_t , namely s_t^* , that satisfies $J_t(s_t, Q) = 0$, and that is such that $J_t(s_t, Q) > 0$ if and only if $s_t > s_t^*$. \square

Theorem 4.1 means that, if the retailer has not terminated the primary selling season by the beginning of period t , then, at the beginning of period t , the retailer need only compare the cumulative demand-to-date with s_t^* to determine its optimal course of action. If cumulative demand-to-date meets the target (i.e., $s_t \geq s_t^*$), then it is optimal for the retailer to continue in the primary market for another period; if not, then it is optimal for the retailer to terminate the primary selling season and to salvage its remaining inventory ($Q - s_t$) on the secondary market. Moreover, computation of the sales targets is a simple matter because Theorem 4.1 provides verification that the recursive functions of the dynamic program specified by (3)–(6), though not necessarily concave (the embedded recursion includes the term $\max\{0, J_t(s_t, Q)\}$), are sufficiently-well behaved that efficient search algorithms will yield the appropriate sales targets, for a given Q .

The optimality of a sales-target termination policy is consistent with a similar result demonstrated by Feng and Gallego (1995) for a situation in which a retailer can make a single price change sometime during the course of a selling season. Our model is analogous to Feng and Gallego's model in the sense that a decision to terminate the primary selling season at the beginning of period t is similar in construct to a one-time decision to change the selling price from r to v_t . However, our model differs from Feng and Gallego's model because it applies to a periodic-review setting, the magnitude of the one-time price reduction is time dependent, and demand growth is incorporated into a general demand distribution.

It is interesting to note that Theorem 4.1 holds regardless of the relationship between the v_t 's, as long as the natural restriction $v_t < r + \pi$ is satisfied for all t . However, if v_t is constant over time, then a *monotone* sales-target termination policy is optimal.

Theorem 4.2. *If $v_t = v$ for $t = 1, \dots, T + 1$, then, for a given Q , $s_1^* < s_2^* < \dots < s_T^*$.*

Proof. The proof is by induction and is similar to the proof of Theorem 4.1. See the Appendix for details. \square

The monotonicity property dictated by Theorem 4.2, although true for the special case in which v_t is constant over time, is not true in general. The following example illustrates this observation.

Example of non-monotone sales targets. For the purpose of this illustration, let

$$\gamma_t = v_t - v_{t+1}, \quad (15)$$

$$\alpha_t = r + \pi - v_t, \quad \text{and} \quad (16)$$

$$\beta_t = h + \gamma_t. \quad (17)$$

These definitions, applied to (9), imply

$$L_t(s_t, Q) = \alpha_t \cdot (Q - s_t) - (\alpha_t + \beta_t)U(s_t, Q), \quad \text{and} \quad (18)$$

$$L_{t-1}(s_t, Q) = L_t(s_t, Q) + \gamma_t U(s_t, Q) - \gamma_{t-1}(Q - s_t). \quad (19)$$

Theorem 4.1 implies that $J_{T-1}(s, Q) < 0$ if and only if $s < s_{T-1}^*$. Thus, $J_{T-1}(s_T^*, Q) < 0$ is sufficient to ensure that $s_T^* < s_{T-1}^*$. Suppose, then, that $X_t = sZ_t$ and $\Phi(z) = 1 - \exp(-z/\lambda)$ for $\lambda > 0$. This implies $Y_t = s + sZ_t$, which implies $G(y|s) = 1 - \exp(-(y - s)/(\lambda s))$. Accordingly, from (8),

$$U(s_t, Q) = \int_{s_t}^Q G(y|s_t) dy = (Q - s_t) - \lambda s_t G(Q|s_t).$$

This, together with (18), imply

$$\lambda s_T^* G(Q|s_T^*) = (Q - s_T^*) - U(s_T^*, Q) = \frac{\beta_T}{\alpha_T} U(s_T^*, Q)$$

since, from Theorem 4.1, s_T^* satisfies $L_T(s_T^*, Q) = 0$. Therefore, from (4), (5), (18), and (19),

$$\begin{aligned} J_{T-1}(s_T^*, Q) &= \left[\gamma_T - \gamma_{T-1} \frac{\alpha_T + \beta_T}{\alpha_T} \right] U(s_T^*, Q) \\ &\quad + \int_{s_T^*}^Q [\alpha_T \cdot (Q - y) - (\alpha_T + \beta_T)U(y, Q)] g(y|s_T^*) dy \end{aligned}$$

$$\begin{aligned}
&= \left[\gamma_T - \gamma_{T-1} \frac{\alpha_T + \beta_T}{\alpha_T} - \beta_T \right] U(s_T^*, Q) \\
&\quad + (\alpha_T + \beta_T) \int_{s_T^*}^Q \lambda s_T^* G(Q|\gamma) g(\gamma|s_T^*) d\gamma \\
&< \left[\gamma_T - \gamma_{T-1} \frac{\alpha_T + \beta_T}{\alpha_T} - \beta_T \right] U(s_T^*, Q) + (\alpha_T + \beta_T) \lambda s_T^* G(Q|s_T^*) \\
&= \left[\gamma_T - \gamma_{T-1} \frac{\alpha_T + \beta_T}{\alpha_T} + \frac{\beta_T^2}{\alpha_T} \right] U(s_T^*, Q) \\
&= [h\beta_T - (\gamma_{T-1} - \gamma_T)(\alpha_T + \beta_T)] \frac{U(s_T^*, Q)}{\alpha_T}.
\end{aligned}$$

Thus, if v_{T-1} , v_T , and v_{T+1} are such that

$$\gamma_{T-1} - \gamma_T \geq h\beta_T / (\alpha_T + \beta_T), \quad (20)$$

then $J_{T-1}(s_T^*, Q) < 0$, which implies that $s_T^* < s_{T-1}^*$. So, for example, if $v_{T-1} = 2$, $v_T = 1$, $v_{T+1} = 0.5$, and $h = 1.5$, then any $r + \pi \geq 5$ results in $s_T^* < s_{T-1}^*$.

This illustration, and (20) in particular, demonstrate that a future sales target (e.g., s_T^*) can actually be lower than a nearer-term sales target (e.g., s_{T-1}^*) if the salvage values are “convex enough” over time. Intuitively, it is natural to expect that v_t is non-increasing in t . Correspondingly, when compared to $L_{t-1}(s, Q)$, $L_t(s, Q)$ can have associated with it a lower per-unit salvage value (i.e., $v_{t+1} < v_t$), but it also can have associated with it a lower per-unit “purchase cost” (i.e., $v_t < v_{t-1}$). Consequently, if the purchase-cost discount (i.e., $\gamma_{t-1} = v_{t-1} - v_t$) is deep enough to offset the salvage-value discount (i.e., $\gamma_t = v_t - v_{t+1}$), then $L_t(s, Q) > L_{t-1}(s, Q)$ so that, everything else being equal, the single-period contribution associated with period t will be greater than the single-period contribution associated with period $t - 1$. If that benefit is large enough to further offset any expected gain associated with period $t - 1$ that results from period $t - 1$ having one extra potential future selling period as compared to period t , then period t becomes a more attractive starting point than period $t - 1$, everything else being equal. In such a scenario, the prudent choice is to establish a test for continuing operations in the primary market that is more stringent for period $t - 1$ than it is for period t ; in other words, $s_{t-1}^* > s_t^*$.

5 Implications for Procurement

In this section we develop and discuss two key results involving the retailer's original procurement quantity. Specifically, we establish that each of the sales targets that constitute the optimal termination policy is increasing in Q , and that, in turn, the optimal Q is strictly greater than a benchmark Q , which we define as the optimal Q for an analogous situation in which termination of the primary selling season prior to the end of period T is not an option.

Theorem 5.1. $\{s_1^*, \dots, s_T^*\}$, the sales targets that constitute the retailer's optimal termination policy, are each strictly increasing in Q , the original procurement quantity.

Proof. The proof of this result is in the Appendix. \square

Theorem 5.1 indicates that a higher procurement quantity results in more aggressive sales targets. This suggests that, everything else being equal, an increase in the initial procurement quantity decreases the likelihood that the sales targets will be reached, thereby increasing the likelihood that the retailer will terminate the primary selling season earlier in the season. Thus, the ultimate determination of whether the retailer's fashion product is a "winner" or a "loser" in the primary market is based on a relative comparison between demand and supply, not on an absolute level of demand.

Although Theorem 5.1 establishes that each s_t^* is well-behaved as a function of Q , finding Q^* , the retailer's optimal procurement quantity, still can be computationally intensive because of the sequence of embedded functions of Q that is implicit in (6). (We compute Q^* for a variety of parameter values in the numerical example in the section that follows.) Nevertheless, additional insight comes from comparing Q^* against a benchmark procurement quantity.

Theorem 5.2. *The retailer's optimal procurement quantity, Q^* , is at least as large as the benchmark quantity Q_b , which is defined as the optimal procurement quantity for an otherwise equivalent decision scenario in which the retailer does not have the flexibility to terminate the primary selling season.*

Proof. Suppose that s_t is the cumulative demand-to-date as of the beginning of period t . Define $N_t(s_t, Q)$ as follows:

$$N_t(s_t, Q) = L_t(s_t, Q) + \int_{s_t}^Q J_{t+1}(y, Q)g(y|s_t) dy.$$

Note, from (4) and (5), that $N_t(s_t, Q)$ represents the conditional value of $J_t(s_t, Q)$, given that the retailer commits to operating in the primary market in period $t + 1$, regardless of the observed level of cumulative demand as of the beginning of period $t + 1$.

By applying the definition of $N_t(s_t, Q)$ and Theorem 4.1 to (4) and (5), we can write:

$$J_t(s_t, Q) = L_t(s_t, Q) + \int_{\bar{s}_t}^Q J_{t+1}(y, Q)g(y|s_t) dy = N_t(s_t, Q) - \int_{s_t}^{\bar{s}_t} J_{t+1}(y, Q)g(y|s_t) dy,$$

where $\bar{s}_t \equiv \max\{s_t, s_{t+1}^*\}$. For a given s_t , $J_{t+1}(\bar{s}_t, Q)[d\bar{s}_t/dQ] = 0$ because $J_{t+1}(\bar{s}_t, Q) = 0$ if $\bar{s}_t = s_{t+1}^*$, and $d\bar{s}_t/dQ = 0$ if $\bar{s}_t = s_t$. Therefore, for a given s_t ,

$$\frac{\partial J_t(s_t, Q)}{\partial Q} = \frac{\partial N_t(s_t, Q)}{\partial Q} - \int_{s_t}^{\bar{s}_t} \frac{\partial J_{t+1}(y, Q)}{\partial Q} g(y|s_t) dy \geq \frac{\partial N_t(s_t, Q)}{\partial Q}. \quad (21)$$

The inequality in (21) follows because $\partial J_{t+1}(y, Q)/\partial Q < J_{t+1}(y, Q)/(Q - y)$ (from induction hypothesis (ii) provided in the proof of Theorem 5.1), and $J_{t+1}(y, Q) < 0$ for all $y < s_{t+1}^*$ (from Theorem 4.1).

The inequality in (21) implies that, for a given s_t , the value of Q that maximizes $J_t(s_t, Q)$ is not less than the value of Q that maximizes $N_t(s_t, Q)$. That is,

$$\arg \max_Q J_t(s_t, Q) \geq \arg \max_Q N_t(s_t, Q).$$

Thus,

$$\begin{aligned} \arg \max_Q \left\{ J_1(0, Q) - (c - v_1)Q - \pi E \left[\sum_{i=1}^T X_i \right] \right\} \\ \geq \arg \max_Q \left\{ N_1(0, Q) - (c - v_1)Q - \pi E \left[\sum_{i=1}^T X_i \right] \right\}. \end{aligned}$$

To interpret this inequality, consider, from (7), that

$$Q^* = \arg \max_Q \left\{ V_1(0, Q) - (c - v_1)Q - \pi E \left[\sum_{i=1}^T X_i \right] \right\},$$

where, presumably, Q^* is such that $V_1(0, Q) = J_1(0, Q)$ because otherwise, the retailer's optimal solution would be *not* to introduce its product to the

primary market, and hence, not to procure any units at all. But, recall that $N_1(0, Q) - (c - v_1)Q - \pi E\left[\sum_{i=1}^T X_i\right]$ represents the total expected profit over the decision-problem horizon, given that the retailer commits to operating in the primary market in period 2, regardless of the observed level of cumulative demand as of the beginning of period 2. Thus, the result follows by induction. \square

Intuitively, the flexibility to terminate the primary selling season upon concluding that its product is a “loser” decreases the retailer’s cost of having leftovers in the sense that the termination flexibility allows the retailer to cut its losses by salvaging eventual leftovers early, rather than first letting holding costs accrue and then salvaging the leftovers later.

Theorems 5.1 and 5.2 provide additional insight by underscoring the benefits of coordinating operations and marketing decisions. To illustrate, consider a scenario in which marketing is charged with setting sales quotas, and operations is charged with procuring stock. Further suppose that operations and marketing decisions are made sequentially rather than in a coordinated fashion so that operations leads by choosing a stocking quantity without considering the effect of sales quotas, and then marketing follows by setting sales quotas in response to the stocking quantity chosen. In such a scenario, operations will under-procure (Theorem 5.2); hence marketing will respond by setting its sales targets too low (Theorem 5.1). Consequently, the sales targets will be easier to achieve as compared to the sales targets associated with a coordinated strategy, which could result in the retailer over-staying its welcome in the primary market.

6 Numerical Example

In this section, we further explore the economic rationale behind the sales-target termination policy through a numerical example. Toward that end, we also use the example to illustrate how managers could determine the sales targets and the optimal procurement policy.

For the purpose of this example, let $X_t = \mu(s_t)Z_t$, where $\mu(s_t) = B + \rho s_t$ for any $B > 0$ and $s_1 \equiv 0$. This implies that $Y_t = s_t + (B + \rho s_t)Z_t$; thus, given that z is the realized value of Z_t , $\mu(s_{t+1}|z) = \mu(s_t)(1 + \rho z)$. In addition, let $v_t = v$ for all t , and define $\alpha = r + \pi - v$ and $\theta = \alpha/(\alpha + h)$. Given these specifications, note that $B\lambda$ can be interpreted as an expected demand base, and $\rho\lambda$ can be interpreted as a (constant) per-period growth rate of expected demand ($\rho\lambda = (E[X_{t+1}] - E[X_t])/E[X_t]$). Note also that α represents the

per-unit net revenue associated with a sale in the primary market since, for every unit sold, r is collected, π is avoided, and v is forfeited. Thus, θ represents the trade-off between the opportunity cost of forfeited sales in the primary market (α) and the explicit cost of holding inventory for sale (h). As h increases relative to α , θ decreases.

To develop a procedure for computing the sales targets and the optimal procurement quantity for this example, we define the transformation $K_t = (Q - s_t)/\mu(s_t)$. Note that K_t represents the period- t stocking factor (Petruzzi and Dada (1999)), and can be updated from period t to period $t + 1$ as follows: $K_{t+1}|z = (Q - s_{t+1})/\mu(s_{t+1}) = (K_t - z)/(1 + \rho z)$, where z denotes the realized value of Z_t , and $K_1 = Q/B$. This transformation is convenient because $L_t(s_t, Q)$ becomes multiplicatively separable when the substitution $Q = s_t + \mu(s_t)K_t$ is applied. To demonstrate this, notice from (3),

$$L_t(s_t, Q) = L_t(s_t, s_t + \mu(s_t)K_t) = (\alpha + h)\mu(s_t)\tilde{L}_t(K_t)$$

where

$$\tilde{L}_t(K_t) = \left[\theta K_t - \int_0^{K_t} \Phi(z) dz \right].$$

By induction, $J_t(s_t, Q)$ also is separable and simplifies to

$$J_t(s_t, Q) = J_t(s_t, s_t + \mu(s_t)K_t) = (\alpha + h)\mu(s_t)\tilde{J}_t(K_t)$$

where

$$\tilde{J}_t(K_t) = \tilde{L}_t(K_t) + E \left[\max \left\{ 0, (1 + \rho z)\tilde{J}_{t+1} \left(\frac{K_t - z}{1 + \rho z} \right) \right\} \right].$$

As a result, $J_t(s_t, Q) = 0$ if and only if $\tilde{J}_t(K_t) = 0$. Thus, if we define K_t^* as the value of K_t that satisfies $\tilde{J}_t(K_t) = 0$, then the sales targets can be recovered as a function of Q from the definition of K_t as follows:

$$s_t^*(Q) = (Q - BK_t^*)/(1 + \rho K_t^*). \quad (22)$$

Similarly, from (6) and the definition of K_1 ,

$$\begin{aligned} \max_Q \{W_1(0, Q) - cQ\} &= \max_{K_1} \{W_1(0, BK_1) - cBK_1\} \\ &= B \cdot \max_K \left\{ \tilde{W}_1(K_1) - \pi E \left[\sum_{i=1}^T X_i | s_1 = 0 \right] \right\}, \end{aligned}$$

where

$$\widetilde{W}_1(K_1) = (\alpha + h)\widetilde{J}_1(K_1) - (c - v)K_1.$$

Therefore, if we define K_1^{opt} as the value of K_1 that maximizes $\widetilde{W}_1(K_1)$, then the retailer's optimal procurement quantity can be written as $Q^* = BK_1^{\text{opt}}$. Finally, substituting this into (22) provides expressions for the optimal sales targets:

$$s_1^*(Q^*) = \left(\frac{K_1^{\text{opt}} - K_t^*}{1 + \rho K_t^*} \right) B. \quad (23)$$

To implement this procedure, we let $\Phi(z) = 1 - \exp(-z/\lambda)$; and we arbitrarily set $B = 1$, $\lambda = 1$, $T = 3$, $r + \pi = 100$, and $c = 60$. Then, we varied v and h relative to c and solved a variety of problem instances. In particular, we let $v = 10, 30, 50$ (i.e., we let $\alpha = r + \pi - v = 80, 60, 40$), and we let $h = 10, 20, 30, 40, 50, 60$. These parameter values are summarized in Table 1. Consequently, we experimented with 18 problem instances with corresponding values of θ ranging from 0.40 to 0.89. Finally, we varied ρ , using values 10, 30, 50, 70, and 90. In total, we computed the optimal initial procurement and the optimal sales-target termination policy for $18 \times 5 = 90$ sets of parameter values. It took just under 4 hours to do this on a 550 MHz desk top computer running Windows 2000.

Parameter	Value(s)
B	1
λ	1
T	3
$r + \pi$	100
c	60
v	10, 30, 50
$\alpha = r + \pi - v$	90, 80, 50
h	10, 20, 30, 40, 50, 60
ρ	10, 30, 50, 70, 90

Table 1: Parameter values used in the numerical example.

The computations were done in Microsoft Excel, using a combination of built-in tools and modules coded in Visual Basic. For each problem instance, we first observed that $\widetilde{W}_1(K_1)$ was unimodal in K_1 , and then determined the values of K_1^{opt} and each of the K_t^* using optimization by simulation. Specifically, we first generated an empirical distribution comprising a common string of 150 independent observations of Z (from $\Phi(z) = 1 - \exp(-z)$, the stipulated cdf). This string of observations was then used to simulate the

computation of the relevant expected values embedded in the $\tilde{J}_t(K_t)$ functions. Finally, the value of $\tilde{J}_1(K)$ was computed recursively in terms of $\tilde{J}_2(K)$ and $\tilde{J}_3(K)$. The number of observations of Z used in the simulation is a parameter that can be altered, but the time required to compute K_t^* for one set of parameter values is exponential in the number of observations. We tested the robustness of the results that follow with respect to the number of observations by increasing the number of observations from 150 to 500 and running a subset of the numerical experiments. In all cases, while the magnitude of the values changed slightly, the relationships among the values that were computed persisted; e.g., if we report that a certain value is not monotonic in a particular parameter, this lack of monotonicity was also observed when 500 observations were used to simulate the expectations.

Then, we used K_1^{opt} and each of the K_t^* to recover the retailer's optimal procurement quantity and optimal sales targets as described above. Table 2 shows the values of Q^* when $v = 50$ for all of the values of h and ρ that were used.

		ρ				
		10	30	50	70	90
h	10	36.34	201.91	495.59	926.54	1486.00
	20	9.28	33.15	64.08	100.32	139.99
	30	3.50	6.32	9.66	14.04	17.96
	40	1.91	2.49	2.77	3.07	3.23
	50	1.31	1.45	1.60	1.68	1.70
	60	0.97	1.09	1.15	1.17	1.18

Table 2: Q^* when $v = 50$.

Table 3 shows the values for s_2^* and s_3^* , the sales targets used to determine if the retailer should operate in the primary market in periods 2 and 3, respectively.

From these results, we made several qualitative observations, which are summarized in Table 4. For any row in the body of the table, the variable in the first column is increased. The resulting change in Q^* , s_2^* , and s_3^* are noted in the next three columns. Up and down arrows indicate that the value in that column always increases or decreases in the variable associated with that row for every other set of parameter values. A ? indicates that the variable in that column was not monotonic in the variable in that row.

Several of the qualitative results are not surprising:

- As h , the per unit cost of holding an item in stock in the primary

		ρ				
		10	30	50	70	90
h	10	0.000	0.031	0.039	0.042	0.044
	20	0.000	0.000	0.001	0.003	0.004
	30	0.000	0.000	0.000	0.000	0.000
	40	0.000	0.000	0.000	0.000	0.000
	50	0.000	0.000	0.000	0.000	0.000
	60	0.000	0.000	0.000	0.000	0.000

(a) s_2^*

		ρ				
		10	30	50	70	90
h	10	0.422	0.940	1.427	1.906	2.385
	20	0.135	0.249	0.308	0.353	0.388
	30	0.024	0.041	0.049	0.057	0.060
	40	0.000	0.004	0.005	0.006	0.005
	50	0.000	0.000	0.000	0.000	0.000
	60	0.000	0.000	0.000	0.000	0.000

(b) s_3^* Table 3: s_2^* and s_3^* when $\nu = 50$.

market, increases, it is optimal to procure fewer items.

- As ρ , the rate of growth of demand, increases, it is optimal to procure more.
- As ν , the salvage value, increases, there is a higher incentive to salvage earlier (i.e., the sales targets in each period increase).

In addition, however, we also observed some qualitative properties that are less obvious:

- Higher inventory carrying costs may *decrease* sales targets.

From Tables 3(a) and (b), both s_2^* and s_3^* are nonincreasing in h for each value of ρ . This means that, all else being equal, there is *less* incentive to salvage as the cost of holding inventory in the primary market

Increase Only	Q^*	s_2^*	s_3^*
v	?	↑	↑
h	↓	↓	↓
ρ	↑	↑	?

Table 4: Summary of qualitative results observed in the numerical experiment.

increases. What we observe here is the confounding interactions between Q^* and s_t^* . If the procurement quantity Q is held constant, we would expect increases in the holding cost to *raise* the sales targets, thereby increasing the incentive to salvage. However, as noted earlier, a higher holding cost also results in a lower procurement quantity, which *decreases* the incentive to salvage (Theorem 4.2). What is interesting in this example is that although a higher value of h creates an explicit incentive to salvage, that incentive is more than offset by the implicit dampening of the salvage incentive that is due to a correspondingly lower Q^* . The net result of the higher value of h is an incentive *not* to salvage. In effect, a higher value of h discourages the retailer from planning to have leftovers early in the season to hold for potential sale later in the season. Consequently, if it turns out that the retailer *does* have leftovers, those leftovers almost surely will be very low relative to expected demand. In that case, the retailer is as likely to sell the remainder of the stock in the primary market as she is to sell it in the salvage market; thus, her best recourse at that point is *not* to salvage.

- Sales targets do not necessarily increase as ρ , the demand growth rate, increases.

From Table 3(b), if $h = 40$, s_3^* first increases in ρ but then begins to decrease. Seemingly, this is another example of the confounding interactions between Q^* and s_t^* . Consider: on the one hand, higher values of ρ suggest that, everything else being equal, the retailer can be more confident that leftovers will sell in the primary market. Thus, she can be more willing to accept the risks associated with the primary market by setting a less aggressive (i.e., lower) sales target. On the other hand, a higher value of ρ also results in a higher value of Q^* , since both the mean and standard deviation of each period's demand

increases with ρ . But, a higher value of Q^* requires a more aggressive sales target. Thus, the net effect that changes in ρ have on the sales targets is, in general, ambiguous.

7 Scope and Applicability

We have considered the problem of determining *how much* a retailer of a fashion product should procure, and *when* the retailer should terminate its primary selling season when demand in the primary market is contagious and salvage recourse exists. We show that the optimal termination policy is characterized by a general sequence of sales targets, which extends similar results in the mark-down pricing literature. We also show that the sequence of sales targets is monotone (increasing) for the special case in which the per-unit salvage value is independent of the termination time. In addition, we demonstrate that the sales targets are strictly increasing as a function of the original procurement quantity, and that the optimal procurement quantity is higher when salvage recourse is dynamic. These results also provide insight on the benefits of coordinating procurement and marketing decisions.

We conclude with a discussion of possible extensions and applications of the model.

1. Stochastic v_t .

We have assumed that v_1, \dots, v_{T+1} are deterministic, but this can easily be relaxed. If instead, the v_t 's are random variables, then the analysis and insights presented in this paper still hold after replacing the v_t 's with their expected values, as long as the probability functions characterizing the v_t 's are independent of other model parameters and decisions. Thus, the assumption of deterministic v_t 's is an innocuous one.

2. Non-linear salvage function.

The assumption of a linear salvage function is clearly a strong one. In effect, it means that if the retailer salvages its remaining inventory at the beginning of period t , then v_t , the per-unit salvage value received, does not depend on $Q - s_t$, the amount of inventory salvaged. This might well be the case if the retailer's product, though unique in style (which is valued by fashion-conscious consumers), is not unique in function (which is valued by price-conscious consumers). In such a scenario, the inventory provided by the retailer to the salvage market would represent only a small fraction

of the overall supply to the price-conscious consumers so the effect of the retailer's salvage quantity on the salvage price would be negligible.

Alternatively, a non-linear salvage function could be constructed from different assumptions, but then technical challenges arise. For example, suppose that the v_t 's were endogenous so that period- t demand in the salvage market depended on v_t as follows: $D_t(v_t) = a_t - b_tv_t$. Then, the retailer would set v_t to maximize $\Psi_t(v_t|Q - s_t) = v_t \cdot \min\{Q - s_t, a_t - b_tv_t\}$, its period- t salvage function, for each t . Accordingly, for each t ,

$$v_t^* = \begin{cases} \frac{a_t - (Q - s_t)}{b_t} & \text{if } Q - s_t < a_t/2 \\ \frac{a_t}{2b_t} & \text{if } Q - s_t \geq a_t/2, \end{cases}$$

which implies that

$$\Psi_t(v_t^*|Q - s_t) = \begin{cases} \frac{(a_t - (Q - s_t))(Q - s_t)}{b_t} & \text{if } Q - s_t < a_t/2 \\ \frac{a_t^2}{4b_t} & \text{if } Q - s_t \geq a_t/2. \end{cases}$$

Hence, the retailer's optimal salvage function would be non-decreasing and concave as a function of the inventory that remained when the primary selling season is terminated. Similarly, if the v_t 's were exogenous, but period- t demand in the salvage market were stochastic (and denoted by D_t), then the retailer's expected salvage function would be

$$v_t E[\min\{Q - s_t, D_t\}],$$

which also is non-decreasing and concave in $Q - s_t$. In such scenarios, a single sales target associated with each period is not necessarily optimal. This suggests the possibility that the retailer's optimal termination policy would be more generally represented by a cumulative "demand range" termination policy, thereby indicating that the retailer should operate in the primary market for another period only if cumulative demand-to-date were not less than a lower-bound target and not more than an upper-bound target. Nonetheless, similar insights to those developed here should apply.

3. Supply chain interpretation.

Finally, our model could be used as a basis for analyzing supply chain considerations. For example, one question that can be investigated is whether

or not the retailer's supplier should consider offering an "any time" returns policy in which the retailer would be allowed to return any unused stock, at any time, to the supplier for a pre-specified, per-unit credit. In such a scenario, the supplier would represent the retailer's salvage market, and the supplier's problem would be to determine the v_t 's to quote to the retailer prior to procurement. The supplier's trade-off would be as follows: on the one hand, the expected number of returns from an any-time policy would be no less than the expected number of returns from an "at time t " policy, but on the other hand, the retailer's procurement quantity if offered an any-time policy would be no less than its procurement quantity if offered an "at time t " policy (from Theorem 5.2).

Acknowledgements

We are grateful to the Associated Editor and the two referees for their insightful comments and suggestions.

A Appendix

We begin by establishing two lemmas that are used in several proofs.

Lemma A.1. $\partial G(y|s_t)/\partial s_t \leq -g(y|s_t) < 0$.

Proof. Multiplicative Case. In the multiplicative case, $X_t = \mu(s_t)Z_t$. Thus,

$$G(y|s_t) = P\{\mu(s_t)Z_t \leq y - s_t\} = \Phi\left(\frac{y - s_t}{\mu(s_t)}\right) \quad (24)$$

and

$$g(y|s_t) = \frac{1}{\mu(s_t)} \phi\left(\frac{y - s_t}{\mu(s_t)}\right). \quad (25)$$

Therefore,

$$\frac{\partial G(y|s_t)}{\partial s_t} = -\left[1 + \frac{(\partial\mu(s_t)/\partial s_t)}{\mu(s_t)}(y - s_t)\right] g(y|s_t) < -g(y|s_t)$$

where the inequality follows from the assumption that $\partial\mu(s_t)/\partial s_t > 0$.

Additive Case. In the additive case, $X_t = \mu(s_t) + Z_t$. Thus,

$$G(y|s_t) = P\{\mu(s_t) + Z_t \leq y - s_t\} = \Phi(y - s_t - \mu(s_t)) \quad (26)$$

and

$$g(y|s_t) = \phi(y - s_t - \mu(s_t)). \quad (27)$$

Therefore,

$$\frac{\partial G(y|s_t)}{\partial s_t} = - \left[1 + \frac{\partial \mu(s_t)}{\partial s_t} \right] g(y|s_t) < -g(y|s_t)$$

where, again, the inequality follows from the assumption that $\partial \mu(s_t)/\partial s_t > 0$. \square

Lemma A.2. *The following relationships are true:*

- i. $\frac{\partial U(s_t, Q)}{\partial Q} = G(Q|s_t) > 0$; and
- ii. $\frac{\partial U(s_t, Q)}{\partial s_t} = \int_{s_t}^Q \frac{\partial G(y|s_t)}{\partial s_t} dy < 0$.

Proof. The relationships follow from (8) and Lemma A.1. \square

Proof of Theorem 4.2. It suffices to show that, for all t , $J_{t-1}(s_t, Q) > J_t(s_t, Q)$ since Theorem 4.1 implies that $J_t(s_t^*, Q) = J_{t-1}(s_{t-1}^*, Q) = 0$ and that $J_{t-1}(s_t, Q) > 0$ if and only if $s_t > s_{t-1}^*$. Consider, then, the following two induction hypotheses:

- i. $J_{t-1}(s_t, Q) > J_t(s_t, Q)$; and
- ii. $s_{t-1}^* < s_t^*$.

Note that if $v_t = v$ for all t , then $L_t(s_t, Q) = L(s_t, Q)$ for all t (from (9)). Thus, if $t = T$, then (4), (5), and Theorem 4.1 imply that $J_{T-1}(s_T, Q) > L(s_T, Q) = J_T(s_T, Q)$; thus, $J_{T-1}(s_T^*, Q) > J_T(s_T^*, Q) = 0$, which implies that $s_{T-1}^* < s_T^*$. This establishes that the induction hypotheses are true for $t = T$. Now, assume that induction hypotheses (i) and (ii) are true for $t = i + 1$ and consider the case $t = i$.

Induction hypothesis (i). From (4):

$$\begin{aligned} J_{i-1}(s_i, Q) &= L(s_i, Q) + \int_{\max\{s_i, s_i^*\}}^Q J_i(y, Q) g(y|s_i) dy \\ &\geq L(s_i, Q) + \int_{\max\{s_i, s_{i+1}^*\}}^Q J_i(y, Q) g(y|s_i) dy, \end{aligned}$$

where the inequality follows because $J_i(\mathcal{y}, Q) > 0$ for $\mathcal{y} > s_i^*$ by Theorem 4.1, and $s_{i+1}^* > s_i^*$ by induction hypothesis (ii). Therefore,

$$J_{i-1}(s_i, Q) > L(s_i, Q) + \int_{\max\{s_i, s_{i+1}^*\}}^Q J_{i+1}(\mathcal{y}, Q) g(\mathcal{y}|s_i) d\mathcal{y} = J_i(s_i, Q), \quad (28)$$

which follows because $J_i(\mathcal{y}, Q) > J_{i+1}(\mathcal{y}, Q)$ by induction hypothesis (i).

Induction hypothesis (ii). From (28), $J_{i-1}(s_i^*, Q) > J_i(s_i^*, Q) = 0$, which implies that $s_{i-1}^* < s_i^*$. \square

Proof of Theorem 5.1. The proof follows by induction on t , given the following induction hypotheses:

- i. $\frac{\partial J_t(s_t, Q)}{\partial Q} < (r + \pi - v_t)$;
- ii. $J_t(s_t, Q) - (Q - s_t) \frac{\partial J_t(s_t, Q)}{\partial Q} > 0$; and
- iii. $\frac{ds_t^*}{dQ} > 0$.

If $t = T$, then $J_T(s_T, Q) = L_T(s_T, Q)$. Thus, from (9) and (8), $\partial J_T(s_T, Q)/\partial Q < (r + \pi - v_T)$. And, from (9) and Lemma A.2,

$$L_T(s_T, Q) - (Q - s_T) \frac{\partial L_T(s_T, Q)}{\partial Q} = (r + \pi + h - v_{T+1}) [(Q - s_T)G(Q|s_T) - U(s_T, Q)] > 0, \quad (29)$$

where the inequality follows because $U(s_T, Q) = \int_{s_T}^Q G(\mathcal{y}|s_T) d\mathcal{y} < (Q - s_T)G(Q|s_T)$. Since $L_T(s_T^*, Q) = J_T(s_T^*, Q) = 0$ by Theorem 4.1, (29) implies that

$$\left. \frac{\partial J_T(s_T, Q)}{\partial Q} \right|_{s_T=s_T^*} < \frac{J_T(s_T^*, Q)}{Q - s_T^*} = 0.$$

Moreover, induction hypothesis (ii) from the proof of Theorem 4.1 implies that

$$\left. \frac{\partial J_T(s_T, Q)}{\partial s_T} \right|_{s_T=s_T^*} > -\frac{J_T(s_T^*, Q)}{Q - s_T^*} = 0.$$

Therefore,

$$\frac{ds_T^*}{dQ} = -\frac{\partial J_T(s_T, Q)/\partial Q|_{s_T=s_T^*}}{\partial J_T(s_T, Q)/\partial s_T|_{s_T=s_T^*}} > 0. \quad (30)$$

This establishes that the induction hypotheses are true for $t = T$. Now assume that they are true for $t = i + 1$, and consider the case $t = i$.

Induction hypothesis (i). Let $\bar{s}_i = \max \{s_i, s_{i+1}^*\}$. Since s_{i+1}^* is increasing in Q , by induction hypothesis (iii), let Q_i be defined as the unique value of Q , for a given s_i , that satisfies $s_{i+1}^* = s_i$. Then, $s_i < s_{i+1}^*$ if and only if $Q > Q_i$. Thus, for $Q \leq Q_i$, $\bar{s}_i = s_i$, which implies that $d\bar{s}_i/dQ = 0$. And, for $Q > Q_i$, $\bar{s}_i = s_{i+1}^*$, which implies that $J_{i+1}(\bar{s}_i, Q) = 0$. Therefore, for all Q ,

$$J_{i+1}(\bar{s}_i, Q) \frac{d\bar{s}_i}{dQ} = 0$$

Moreover, the definition of \bar{s}_i applied to (4) yields

$$J_i(s_i, Q) = L_i(s_i, Q) + \int_{\bar{s}_i}^Q J_{i+1}(y, Q) g(y|s_i) dy, \quad (31)$$

which implies

$$\frac{\partial J_i(s_i, Q)}{\partial Q} = \frac{\partial L_i(s_i, Q)}{\partial Q} + \int_{\bar{s}_i}^Q \frac{\partial J_{i+1}(y, Q)}{\partial Q} g(y|s_i) dy. \quad (32)$$

Since $\partial J_{i+1}(y, Q)/\partial Q < r + \pi - v_i$ by induction hypothesis (i), (32) implies

$$\begin{aligned} \frac{\partial J_i(s_i, Q)}{\partial Q} &< \frac{\partial L_i(s_i, Q)}{\partial Q} + (r + \pi - v_{i+1}) \int_{\bar{s}_i}^Q g(y|s_i) dy \\ &= (r + \pi - v_i) - (r + \pi + h - v_{i+1})G(Q|s_i) \\ &\quad + (r + \pi - v_{i+1}) [G(Q|s_i) - G(\bar{s}_i|s_i)] \\ &= (r + \pi - v_i) - hG(Q|s_i) - (r + \pi - v_{i+1})G(\bar{s}_i|s_i) \\ &< (r + \pi - v_i). \end{aligned}$$

Induction hypothesis (ii). From (31) and (32),

$$\begin{aligned}
& J_i(s_i, Q) - (Q - s_i) \frac{\partial J_i(s_i, Q)}{\partial Q} \\
&= L_i(s_i, Q) - (Q - s_i) \frac{\partial L_i(s_i, Q)}{\partial Q} \\
&\quad + \int_{\bar{s}_i}^Q \left[J_{i+1}(y, Q) - (Q - s_i) \frac{\partial J_{i+1}(y, Q)}{\partial Q} \right] g(y|s_i) dy \\
&= (r + \pi + h - v_{i+1}) \int_{s_i}^Q (y - s_i) g(y|s_i) dy \\
&\quad + \int_{\bar{s}_i}^Q \left[J_{i+1}(y, Q) - (Q - s_i) \frac{\partial J_{i+1}(y, Q)}{\partial Q} \right] g(y|s_i) dy.
\end{aligned}$$

But, $J_{i+1}(y, Q) > (Q - y) \partial J_{i+1}(y, Q) / \partial Q$ by induction hypothesis (ii), $\partial J_{i+1}(y, Q) < r + \pi - v_{i+1}$ by induction hypothesis (i), and $s_i \leq \bar{s}_i$ by the definition of \bar{s}_i . Therefore,

$$\begin{aligned}
& J_i(s_i, Q) - (Q - s_i) \frac{\partial J_i(s_i, Q)}{\partial Q} > \\
& \int_{\bar{s}_i}^Q \left[r + \pi + h - v_{i+1} - \frac{\partial J_{i+1}(y, Q)}{\partial Q} \right] (y - s_i) g(y|s_i) dy > 0.
\end{aligned}$$

Induction hypothesis (iii). Since $J_i(s_i^*, Q) = 0$ by Theorem 4.1, induction hypothesis (ii) implies

$$\left. \frac{\partial J_i(s_i, Q)}{\partial Q} \right|_{s_i=s_i^*} < \frac{J_i(s_i^*, Q)}{Q - s_i^*} = 0.$$

Moreover, induction hypothesis (ii) from the proof of Theorem 4.1 implies

$$\left. \frac{\partial J_i(s_i, Q)}{\partial s_i} \right|_{s_i=s_i^*} > -\frac{J_i(s_i^*, Q)}{Q - s_i^*} = 0.$$

Therefore,

$$\frac{ds_i^*}{dQ} = -\frac{\partial J_i(s_i, Q) / \partial Q|_{s_i=s_i^*}}{\partial J_i(s_i, Q) / \partial s_i|_{s_i=s_i^*}} > 0.$$

□

References

- Bassok, Y. and R. Anupindi. 1997. Analysis of supply contracts with total minimum commitment. *IIE Transactions* **29** 373-381.
- Bass, F. and A. Bultez. 1982. A note on optimal strategic pricing of technological innovations. *Marketing Science* **1** 371-378.
- Bass, F. 1980. The relationship between diffusion rates, experience curves, and demand elasticities for consumer durables. *Journal of Business* **53** 551-567.
- Bitran, G. R. and S. V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Science* **43** 64-79.
- Eppen, G. D. and A. V. Iyer. 1997a. Backup agreements in fashion buying — the value of upstream flexibility. *Management Science* **43** 1469-1484.
- Eppen, G. D. and A. V. Iyer. 1997b. Improved fashion buying with Bayesian updates. *Operations Research* **45** 805-819.
- Feng, Y. and G. Gallego. 1995. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science* **41** 1371-1391.
- Feng, Y. and G. Gallego. 2000. Perishable asset revenue management with markovian time dependent demand intensities. *Management Science* **46** 941-956.
- Fisher, M. and A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* **44** 87-99.
- Fukuda, Y. 1961. Optimal disposal policies. *Naval Research Logistics Quarterly* **8** 221-227.
- Gallego, G. and G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40** 999-1020.
- Hartung, P. H. 1973. A simple style goods inventory model. *Management Science* **19** 1452-1458.
- Hertz, D. and K. Schaffir. 1960. A forecasting method for management of seasonal-style goods inventories. *Operations Research* **8** 45-52.

- Horsky, D. and L. Simon. 1983. Advertising and the diffusion of new products. *Marketing Science* **2** 1-17.
- Iyer, A. V. and M. V. Bergen. 1997. Quick response in manufacturer-retailer channels. *Management Science* **43** 559-570.
- Karlin, S. 1960. Dynamic inventory policy with varying stochastic demands. *Management Science* **6** 231-258.
- Kurawarwala, A. A. and H. Matsuo. 1996. Forecasting and inventory management of short life-cycle products. *Operations Research* **44** 131-150.
- Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Science* **36** 724-738.
- Mahjan, V. and E. Muller. 1979. Innovation diffusion and new product growth models in marketing. *Journal of Marketing* **43** 55-68.
- McGill, J. I. and G. J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* **33** 233-256.
- Monahan, G. E. and K. O. Nti. 1988. Optimal advertising and pricing for new products with repeat purchases. In T. M. Devinney, ed., *Issues in Pricing: Theory and Research*, 145-170. Lexington Books, Lexington, MA.
- Petruzzi, N. C. and M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research* **74** 183-194.
- Porteus, E. L. 1990. Stochastic inventory theory. In D. Heyman and M. Sobel, eds., *Handbooks in Operations Research and Management Science, Vol 2: Stochastic Models*, 605-652. North Holland, Amsterdam.
- Ravindran, A. 1972. Management of seasonal style-goods inventories. *Operations Research* **20** 265-275.
- Tsay, A. A. and W. S. Lovejoy. 1999. Quantity flexibility contracts and supply chain performance. *Manufacturing & Service Operations Management* **1** 89-111.
- Tsay, A. A., S. Nahmias, and N. Agrawal. 1999. Modeling supply chain contracts: A review. In S. Tayur, R. Ganeshan, and M. Magazine, eds., *Quantitative Models for Supply Chain Management*, 299-336. Kluwer Academic Publishers.
- Zhao, W. and Y. Zheng. 2000. Optimal dynamic pricing for perishable assets. *Management Science* **46** 375-388.