Style Goods Pricing with Demand Learning

Alper Sen
Department of Industrial Engineering
Bilkent University
Bilkent, Ankara 06800, Turkey
alpersen@bilkent.edu.tr

Alex X. Zhang
Hewlett Packard Laboratories
1501 Page Mill Road, MS 1U 2
Palo Alto, CA 94304, USA
alex_zhang@hp.com

Abstract

For many industries (e.g., apparel retailing) managing demand through price adjustments is often the only tool left to companies once the replenishment decisions are made. A significant amount of demand uncertainty can be resolved using the early sales information. In this study, a Bayesian model is developed to summarize sales information and pricing history in an efficient way. This model is incorporated into a periodic pricing model to optimize revenues for a given stock of items over a finite horizon. A computational study is carried out in order to find out the circumstances under which learning is most beneficial. The model is extended to allow for replenishments within the season, in order to understand global sourcing decisions made by apparel retailers. Some of the findings are empirically validated using data from U.S. apparel industry.

1. Introduction

Fashion goods such as ski–apparel or sunglasses are characterized by high degrees of demand uncertainty. Most of the merchandise in this category are new designs. Although some of the demand uncertainty may be resolved using sales history of similar merchandise offered in previous years, most of the uncertainty still remains due to the changing consumer tastes and economic conditions every year. Retailers of these items face long lead times and relatively short selling seasons that force them to order well in advance of the sales season with limited replenishment opportunities during the season. Demand and supply mismatches due to this inflexible and highly uncertain environment result in forced mark–downs or shortages. Frazier [21] estimates that the forced mark–downs average 8 percent of net retail sales in apparel industry, which he states is also an indication of as much as 20 percent in lost sales from stock–outs. He estimates that the overall resulting revenue losses of the industry may be as much as $25 billion.

In 1985, U.S. textile and apparel industry initiated a series of business practices and technological innovations, called Quick Response, to cut down these costs and to be able to compete with foreign industry enjoying lower wages. Quick Response aims to shorten lead times through improvements in production and information technology. As a result, production and ordering decision can be shifted closer to the selling season which will help to resolve some uncertainty. Moreover, additional replenishment opportunities during the season may be created. See Hammond and Kelly [21] for a review of Quick Response and Sen [37] and Sen [38] for reviews of operations and current business practices and trends in the U.S. apparel industry.

Despite the efforts of domestic manufacturers to remain competitive in this industry, retailers are using more and more imports to source their apparel, preferring cost advantage over responsiveness. For most imported apparel and some domestic apparel, managing demand through price adjustments is often the only tool left to retailers once the buying decisions took place. These adjustments are usually in the form of mark–downs in the apparel industry. Fisher et al. [19] note that 25 % of all merchandise sold in department stores in 1990 was sold with mark–downs. Systems that can intelligently decide the timing and magnitude of such mark–downs may help balance the supply and demand and improve the profits of these companies operating with thin margins. Despite enormous amount of data made available to decision makers, such intelligent systems have found limited use in the apparel industry. Recent academic research such as Gallego and van Ryzin [22] and Bitran and Mondschein [6] successfully model dynamic pricing of a given stock of items when the demand is probabilistic and price sensitive. These studies assume that the retailer's estimate of the demand does not change over the course of the season. However, substantial amount of uncertainty about the demand process can be resolved using the early sales information.

The purpose of this paper is to develop a dynamic pricing model that incorporates demand learning. By demand learning, we mean learning by using the early sales information during the selling season as opposed to improving forecasts over time before the start of the season. A considerable portion of demand uncertainty can be eliminated by observing early sales in the apparel industry. A consultant at Dayton Hudson Corp. states "a week after an item hits the floor, a merchant knows whether it's going to be a dog or a best–seller" (Chain Store Age [12]). For our pricing only model, we assume that ordering decision has already been made with the best use of pre–season information and no further replenishment opportunities are available to the retailer. Basically, the model uses a Bayesian approach to update retailer's estimate of a demand parameter. Our model enables us to summarize sales and price history in a direct way to set the problem as a computationally feasible dynamic program. We also conduct a numerical study to analyze the impact of different factors on pricing decisions. First, we study how the accuracy and degree of uncertainty of the initial demand estimates, starting stock levels and price sensitivity of customers impact optimal price paths and expected revenues. We are also interested in finding the conditions under which earlier sales information has the most impact on revenues and whether it is always optimal to use this information. Further, we explore the trade–off between more information and early control in pricing decisions. Finally, we extend the model to account for the possibility of re–ordering during the selling season. This helps us to understand the possible trade–offs for using quicker but more costly domestic manufacturing to achieve such flexibility.

Next, we review literature on Bayesian learning in inventory control and dynamic pricing of fashion goods. We present our basic model in Section 3. Our computational analysis is in Section 4. Section 5 studies the effects of inventory flexibility during the horizon. Section 6 states our conclusions and avenues for future research.

2. Literature Survey

Inventory models that incorporate the updating of demand forecasts have been studied extensively. Most of these models utilize a Bayesian approach to update demand parameters of a periodic inventory model. Demand in one period is assumed to be random with a known distribution but with an unknown parameter (or unknown parameters). This unknown parameter has a prior probability distribution which reflects the initial estimates of the decision maker. Observed sales are then used to find a posterior distribution of the unknown parameter using Bayes' rule. As more observations become available, uncertainty is resolved and the distribution of the demand approaches its true distribution. The prior distribution of the unknown parameter should be such that the posterior distribution is similar to the prior which could be calculated easily. In addition, the demand distribution and the distribution of the unknown parameter should enable the decision maker to summarize information such that a dynamic program to solve the problem is computationally feasible. See DeGroot [15, chapter 9] for such distributions.

Demand learning in inventory theory using a Bayesian approach is first studied by Scarf [32]. He studies a simple periodic inventory problem in which at the beginning of each period the problem is how much to order with the assumption of linear inventory holding, shortage and ordering costs and an exponential family of demand distributions with an unknown parameter. The distribution of the unknown parameter is updated after each period using Bayes' rule. He formulates the problem as a stochastic dynamic program and among other results, shows that the optimal policy is to order up to a critical level and the critical level for each period is an increasing function of the past cumulative demand. Iglehart [25] extends the results of Scarf [32] to account for a range family of distributions and convex inventory holding and shortage costs. Azoury and Miller [3] show that in most cases non–Bayesian order quantities are greater than Bayesian order quantities, but also state that this may not always be true. The dynamic programs used in these studies have two–dimensional state spaces, one for the starting inventory level and one for the cumulative sales. Scarf [33] and Azoury [4] show that the two–dimensional dynamic program can be reduced to one–dimensional for some specific demand distributions.

A particular form of Bayesian approach to demand learning is assuming Poisson demand with an unknown rate in each period. The unknown demand rate's prior distribution is assumed to be Gamma, resulting in an unconditional prior distribution of demand which can be shown to be Negative Binomial. Posterior distributions are also Gamma and Negative Binomial whose parameters can be calculated by using only cumulative demand. These specific distributions are used to model inventory decisions of aircraft spare parts by Brown and Rogers [10]. Popovic [31] extends the model to account for non–constant demand rates.

Demand learning models are most valuable to inventory problems of style goods that are characterized with moderate to extreme degrees of demand uncertainty that is resolvable significantly by observing early sales. Murray and Silver [28] use a Bayesian model in which the purchase probability of homogeneous customers is unknown but distributed priorly with a Beta distribution. This distribution is updated after each period to optimize inventory levels in succeeding periods. Chang and Fyfee [13] present an alternative approach to demand learning. Their model defines the demand in each period as a noise term plus a fraction of total demand which is a random variable whose distribution is revised once the sales information becomes available each period. Bradford and Sugrue [9] use Negative Binomial demand model described earlier to derive optimal inventory stocking policies in a two–period style–goods context.

Fisher and Raman [20] propose a production planning model for fashion goods which uses early sales information to improve forecasts. Their model, which is called Accurate Response, also considers the constraints in the production systems such as production capacity and minimum production quantities. Iyer and Bergen [26] study the Quick Response systems, where the retailers have more information about upcoming demand due to the decreased lead times. They use Bayesian learning to address whether the retailer or the manufacturer wins under such systems. Eppen and Iyer [16] develop a different methodology for Bayesian learning of demand. The demand process is assumed to be one of a set of pure demand processes with discrete prior distribution. This distribution is updated periodically based on Bayes' rule. This demand model is used in a dynamic programming formulation to derive the initial inventory levels and how much to divert periodically to a secondary outlet for a catalog merchandiser. Eppen and Iyer [17] use the same demand model to study the impact of backup agreements on expected profits and inventory levels for fashion goods. Gurnani and Tang [23] study the effect of forecast updating on ordering of seasonal products. Their model allows the retailer to order at two instants before the selling season. The forecast quality may be improved in the second instance, but the cost may either decrease or increase probabilistically.

All of the studies above ignore one crucial aspect of the problem: pricing. In economics literature, Lazear [27] studies clearance sales where he uses Bayesian learning to update the reservation price distribution after observing early sales in the season. However, his model considers the initial and the mark–down prices of a single item and thus lacks the dynamics of price adjustments for a stock of items. Balvers and Casimano [5] incorporate Bayesian learning in pricing models, but they assume a completely flexible supply and ignore inventories that link the pricing decisions. Style goods, on the other hand, face supply inflexibility as a result of short seasons, long lead times and limited production capacities. This characteristic of the problem gave rise to models such as those in Gallego and van Ryzin [22] and Bitran and Mondschein [6] that dynamically price the perishable good over the selling season. Both of these models assume that there is no replenishment opportunity and the only decisions to be made are the timing and magnitude of price changes over the course of the season. Gallego and van Ryzin [22] use a Poisson process for demand where the demand rate depends on the price of the product. Monotonicity results as a function of the remaining stock level and remaining time in the selling season are derived via a dynamic continuous–time model. Among other results, they show that the optimal profit of the deterministic problem, in which demand rates are assumed to be constant, gives an upper bound for the optimal expected profit. For the continuous price case, fixed–price heuristics are shown to be asymptotically optimal. For the discrete price case, a deterministic solution can be used to develop again asymptotically optimal heuristics. Feng and Gallego [18] derive the optimal policy for the two price case. In Bitran and Mondschein's [6] model, the purchase process for a given price is determined by a Poisson process for the store arrival and a reservation price distribution. They show that the model is equivalent to the model in Gallego and van Ryzin [22]. They also show that the loss associated with preferring a discrete–time rather than a continuous–time model is small. Smith and Achabal [34] study clearance pricing in retailing. Their model is deterministic, but incorporates impact of reduced assortment and seasonal changes on demand rates. Petruzzi and Dada [30] consider a periodic review model where the retailer is allowed to order new inventory as well as change the price at each period. However, the stochastic component of their demand model is very specific. If the retailer can fully satisfy the demand in any period, the uncertainty is completely resolved and the remaining problem is a deterministic one. Otherwise, the retailer updates the lower bound for the uncertain component, the remaining problem remains to be a stochastic one, with a new estimate for the uncertain component.

Recently, three closely related papers discuss Bayesian learning in pricing of style goods.

Subrahmanyan and Shoemaker [36] develop a general periodic demand learning model to optimize pricing and stocking decisions. As in Eppen and Iyer [16,17], they use a set of possible demand distribution functions for each period and a discrete prior distribution that tabulates the probability of these possible demand distributions being the true demand distribution. This discrete distribution is updated after each period using the Bayes' rule. The information requirements are extremely large in a general model as updating requires the history of sales, inventory levels and prices in each period. They present computational results on specific demand and price parameters. Bitran and Wadhwa [7] consider only the pricing decisions utilizing the two–phased demand model and discrete–time dynamic programming formulation in Bitran and Mondschein [6]. A Poisson process for store arrival and a reservation price distribution are used to define the purchase process. They assume that uncertainty is involved in a parameter of this reservation price distribution. An updating procedure on this parameter is proposed such that the rate of the purchase process has Gamma priors and posteriors. The methodology allows them to summarize all sales and price information in two variables. They present computational results to show the impact of demand learning on prices and expected profits. Aviv and Pazgal [2] studied a problem where the arrival process is Poisson, the arrival rate has a Gamma distribution and the retailer controls the price continuously. The resulting model is a continuous-time optimal control problem. Among other results, it is shown that initial high variance leads to higher prices and the expected revenues of the optimal pricing policy are compared with expected revenues from several other policies including a fixed price scheme.

3. Model

3.1 Demand Model

Assume that there are N points in time that the pricing decisions can be made. Without loss of generality, assume that each period in consideration is of unit length. The demand in each period has a Poisson distribution. The demand rate is separable and consists of two components: a base demand rate λ, and a multiplier $m(p)$ for the charged price p. The Poisson rate is equal to

$$\lambda(p) = m(p)\lambda.$$

Without loss of generality, we assume $m(1) = 1$. Although our model does not depend on a particular demand function, we use exponential price sensitivity $\lambda(p) = ae^{-\gamma p}$ to find

$$m(p) = e^{-\gamma(p-1)}, \tag{1}$$

in our computational analysis. Exponential price sensitivity and multiplicative demand functions are widely used in practice and research (see Smith and Achabal [34] and Smith et. al [35] for examples). The distribution of demand given the price and base rate is given by,

$$f(x|\lambda, p) = \frac{e^{-m(p)\lambda}[m(p)\lambda]^x}{x!}, \quad \text{for } x = 0, 1, 2, \ldots$$

We assume that there is uncertainty only in the magnitude of demand, but not in the functional form of the price demand relationship (e.g., price elasticity). That is, we assume that we have perfect information about the function $m(p)$ and that the uncertainty of the demand rate for a given price can be totally characterized by the uncertainty in the base demand rate λ. Consequently, observing sales will facilitate learning on the magnitude of demand only.

We assume that λ is distributed as Gamma with parameters α and β. The probability density function for Gamma is given by,

$$f(\lambda) = \frac{\beta^\alpha \lambda^{\alpha-1} e^{-\beta\lambda}}{\Gamma(\alpha)}, \quad \lambda > 0.$$

Then, the *prior* distribution (unconditional of λ) of demand in the first period will have Negative Binomial distribution:

$$D_1 \sim NB\left(\alpha, \frac{\beta}{\beta + m(p_1)}\right),$$

if the price for period 1 is p_1. The probability function is given by,

$$f(x|p_1) = \begin{pmatrix} \alpha + x - 1 \\ x \end{pmatrix} \left(\frac{\beta}{\beta + m(p_1)}\right)^\alpha \left(\frac{m(p_1)}{\beta + m(p_1)}\right)^x, \quad \text{for } x = 0, 1, 2, \ldots$$

This uses the fact that if λ is distributed with Gamma with parameters α and β, $k\lambda$ is distributed with Gamma with parameters α and β/k for any multiplier k. An unconditional distribution that is Negative Binomial is consistent with the high uncertainty involved in fashion goods as the variance to mean ratio of a Negative Binomial random variable is always greater than 1. Additional support is provided in Nahmias and Smith [29] where they discuss the suitability of Negative Binomial demand for a retailer system.

If the realized demand in period 1 is x_1, Bayes' rule implies that the *posterior* distribution of λ will be again Gamma with parameters $\alpha+x_1$ and $\beta+m(p_1)$. Then, the *prior* distribution of demand in the second period will have a Negative Binomial form:

$$D_2 \sim NB\left(\alpha + x_1, \frac{\beta + m(p_1)}{\beta + m(p_1) + m(p_2)}\right),$$

given that the price in second period is p_2.

Likewise, when the realized demand is $x_1, x_2, \ldots, x_{n-1}$ and prices charged are $p_1, p_2, \ldots, p_{n-1}$ in periods $1, 2, \ldots, n-1$, the *prior* distribution of demand in the nth period will also have a Negative Binomial distribution:

$$D_n \sim NB\left(\alpha + \sum_{i=1}^{n-1} x_i, \frac{\beta + \sum_{i=1}^{n-1} m(p_i)}{\beta + \sum_{i=1}^{n} m(p_i)}\right),$$

given that the price in nth period is p_n. Denote cumulative sales prior to period n as $X_{n-1} = \sum_{i=1}^{n-1} x_i$ and cumulative price multipliers prior to period n as $M_{n-1} = \sum_{i=1}^{n-1} m(p_i)$. X_{n-1} and M_{n-1} summarize all the information in periods $1, \ldots, n-1$ and are called the *sufficient statistics* for estimating demand in period n, for a given price p_n.

The unconditional demand distribution for the nth period will have a mean of

$$E[D_n] = \frac{(\alpha + \sum_{i=1}^{n-1} x_i) m(p_n)}{\beta + \sum_{i=1}^{n-1} m(p_i)},$$

which basically means that the sales rate in the nth period is a linear function of sales rate in the earlier $n-1$ periods. This is in fact not surprising. Carlson [11] studies sales data of apparel merchandise from a major department store to see whether the sales rate after a mark–down is predictable. Given an initial price and a mark–down percentage, he shows that past mark–down sales rate is in fact a linear function of pre mark–down sales rate. Our model completely agrees with this empirical result.

It is also worthwhile to see how the mean and variance of the unconditional distribution of demand behaves as n increases. For simplicity of the exposition, assume that the price is equal to 1 throughout the season so that $m(p_i) = 1$ for all i. The expected value and variance of the unconditional demand are given by,

$$\begin{aligned} E[D_n] &= \frac{\alpha + \sum_{i=1}^{n-1} x_i}{\beta + n - 1}, \\ Var[D_n] &= \frac{(\alpha + \sum_{i=1}^{n-1} x_i)(\beta + n)}{(\beta + n - 1)^2}. \end{aligned}$$

It is easy to see that as n approaches infinity, both the mean and variance approach $\overline{x}$, average of x_i, which is the true rate of the Poisson process. We note that the convergence is faster if β is smaller. This corresponds to higher degrees of uncertainty in the decision maker's initial estimate of demand rate, and thus more reliance on actual sales information in estimating future demand.

While our analysis so far assumes that the periods are identical except for the prices charged, our model allows us to permit seasonality and any other extensions as long as the multiplicative nature of the demand function is preserved. That is, as long as we can state the demand rate in period i as

$$\lambda_i(p_i, \tau_i) = m(p_i, \tau_i)\lambda,$$

(where m now is a more general function of price p_i and seasonality factor τ_i) our model is applicable. Also, uneven period lengths are easily accountable. In a situation where the periods have varying lengths and seasonality, the state variable M_{n-1} can be represented as

$$M_{n-1} = \sum_{i=1}^{n-1} m(p_i, \tau_i)\ell_i,$$

where ℓ_i is the length of the period i. The details of the derivation for two periods are given in Section 4.

3.2 Pricing Model

The problem is determining prices in periods $1, \ldots, N$ so that a fixed stock of I_0 items is sold with maximum expected revenue. For simplicity of the presentation, We assume that the inventory holding costs within the selling season are negligible. We note that the it is very easy to relax this assumption in the context of our model.

We use a discrete–time dynamic programming model. Let $V_n(I_{n-1}, X_{n-1}, M_{n-1})$ be the maximum expected revenue from period n through N when the initial inventory is I_{n-1} and the cumulative sales and cumulative price multipliers are X_{n-1} and M_{n-1}, respectively. Note that

$$I_{n-1} = \max\{0, I_0 - X_{n-1}\},$$

and can be dropped from the formulation. But we keep I_{n-1} in our formulation for ease of exposition. Also let p_s be the salvage value for any inventory left unsold beyond period N.

Backward recursion can be written as

$$V_n(I_{n-1}, X_{n-1}, M_{n-1}) =$$

$$\max_{p_n \geq p_s} E\Big[p_n \min\{D_n, I_{n-1}\} + V_{n+1}\Big((I_{n-1} - D_n)^+, X_{n-1} + D_n, M_{n-1} + m(p_n)\Big) \qquad (2)$$
$$\Big| X_{n-1}, M_{n-1} + m(p_n)\Big].$$

Boundary conditions are

$$V_{N+1}(I_N, X_N, M_N) = p_s I_N, \quad \text{for all } I_N, X_N, M_N, \qquad (3)$$
$$V_n(0, X_{n-1}, M_{n-1}) = 0, \quad \text{for all } n, X_{n-1}, M_{n-1}. \qquad (4)$$

First condition states that any left over merchandise has only salvage value when the season ends at the end of period N. The second conditions states that the future expected profits are zero, when there is no merchandise left in stock since re–ordering is not allowed. This property also allows us to avoid the problem of censored of demand information due to unsatisfied demand. In case of excess demand (when the inventory is exhausted), there are no further decisions to be made and no further information about demand is required. The dynamic program can be solved by starting with the Nth period and proceeding backwards.

We solved many problems with different sets of parameters to investigate the structural properties of the optimal policy. In all these problems, we observed that higher sales in earlier periods always translate into higher prices in future periods. The intuition behind this behavior is the following. First, higher sales in earlier periods mean (stochastically) higher demand in future periods because of the Bayesian nature of the demand distributions. Second, higher sales in earlier periods also mean lower left–over inventory for future periods since there are no further replenishment opportunities. Thus, higher sales in earlier periods inflate the expected demand while decreasing the available supply in future periods. This allows the seller to charge higher prices to balance the demand and supply. The second part of the argument (lower inventory calls for higher prices), is formally proved by Chun [14] for Negative Binomial demand. The first part of the argument (stochastically larger demand calls for higher prices), however, is not true in general. See Bitran and Wadhwa [8] for counter examples and certain conditions that are required.

In order to show how the model works, we provide the following example.

Example: The retailer has 30 units to sell in a season of length 1. When the price is set to 1.00, the demand is Poisson with a rate distributed with Gamma with parameters $\alpha = 10$ and $\beta = 0.5$. Thus, mean demand is $\alpha/\beta = 20$ and variance of the demand is $\alpha(\beta+1)/\beta^2 = 60$. Within the season, there are two periods of equal length. The retailer can charge different prices in these periods from a discrete set $\mathcal{P}$=$\{0.50, 0.55, .., 0.95, 1.00\}$.

The price affects the demand in an exponential manner with $\gamma = 3$ described in equation 1. The mean total demand is given as follows with the prices in $\mathcal{P}$.

price	0.50	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00
mean demand	89.6	77.2	66.4	57.2	49.2	42.3	36.4	31.4	27.0	23.2	20.0

The problem is to find the price in the first period and the form of the pricing policy in the second period so as to maximize the total revenues. We solve the problem with the dynamic program given in equations (2-4). The optimal policy is to charge 1.00 in the first period and then charge the following prices in the second period based on the demand realization in the first period.

1st period sales x_1	0-5	6	7-8	9-10	11-12	13-14	15	16-17	18	19-30
2nd period opt. price p_2^*	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00

The resulting optimal expected revenue is 23.248, about 0.775 per unit.

4. Computational Study

We first note that although pricing through a demand learning model is the best the retailer can do, it is not necessarily optimal. The optimal policy depends on the true value of underlying base demand rate. The optimal prices can be computed by using a dynamic programming formulation which uses the true Poisson demand distribution. The performance of the demand learning model depends on how accurate the retailer's initial demand estimate is and how fast the retailer can learn about the true demand process. Note that prior to the start of the season, the retailer assumes that the base demand rate is distributed with Gamma with parameters α and β. The expected value and variance of this random variable are given by,

$$E[\lambda] = \frac{\alpha}{\beta} \text{ and } Var[\lambda] = \frac{\alpha}{\beta^2}.$$

Hence, α/β defines the accuracy of the point estimate. Given a fixed ratio α/β, the magnitude of β (or α) defines the variance of the initial estimate, and hence the decision maker's reliance on her prior beliefs about demand. When β (or α) is large, the retailer is confident about her initial estimate, and she hardly updates her demand estimate based on observed sales. As β (or α) gets smaller, more weight is given to the observed sales in estimating future demand.

We analyze three different model in our computational study. Under the *Perfect Information* model, the true value of the underlying base rate is known, and an optimal policy is derived using Poisson distributed demand with rate $m(p)\lambda$. Under *No Learning* model, the decision maker only knows the values of α and β and an optimal policy is derived using Negative Binomial distribution with parameters α and $\beta/(\beta + m(p))$. The parameters of Negative Binomial distribution is *not* updated as the sales are observed. Under *Learning* model, the decision maker also only knows the values of α and β, however when the optimal policy is derived, parameters of Negative Binomial distribution is updated using observed sales following the learning model as described in Section 2 (more details to be given below).

In order to understand the impact and value of learning, the performance of the policies that are derived under the Learning and No Learning models are also evaluated using Poisson distributed demand with the true value of the base rate. We should note again however that this rate is not revealed to the decision maker before the season (for otherwise, the decision maker would simply use Perfect Information model to maximize its revenues) and thus evaluation of Learning and No Learning models based on the true Poisson rate cannot appropriately guide the decision maker *before* the season. In addition, the performance of the policy under the Learning model is evaluated using Negative Binomial distribution whose parameters are updated based on the observed sales.

Our primary objective in the computational study is to discover the conditions under which the early sales information has the most impact on revenues by comparing the revenues of Learning model with that of No Learning model. While doing this we also generate the optimal revenues for Perfect Information model. We specifically study the impacts of accuracy of the initial estimate, the variance of the initial estimate, price elasticity of demand on all three models.

For the purposes of computational study, we assume that there is only one chance to change the price. The resulting model is a special two–period case of the model described earlier. In addition, we assume that the salvage value is zero. We also note that the assumption that the periods are of equal length can be relaxed as briefly described below. This model is briefly described below.

Demand in one period is again Poisson with rate

$$\lambda(p) = m(p)\lambda,$$

where $m(p) = e^{-\gamma(p-1)}$. Assume again that λ is unknown, but distributed with Gamma with parameters α and β. We assume that the length of the first period is t and the length of the second period is $1-t$. Then *prior* distribution of demand in the first period is Negative Binomial:

$$D_1 \sim NB\Big(\alpha, \frac{\beta}{\beta + m(p_1)t}\Big), \tag{5}$$

given that the price is p_1 in the first period. If the realized demand in period 1 is x_1 and the charged price is p_1, *posterior* distribution of λ would be again Gamma with parameters $\alpha + x_1$ and $\beta + m(p_1)t$. Then the *priori* distribution of demand in the second period is Negative Binomial:

$$D_2 \sim NB\Big(\alpha + x_1, \frac{\beta + m(p_1)t}{\beta + m(p_1)t + m(p_2)(1-t)}\Big), \tag{6}$$

given that the price is p_2 in the second period.

The problem is to set prices p_1 and p_2 such that an initial stock of I_0 items are sold with maximum expected revenue. The problem again can be solved by a dynamic program for each t. Given the price and observed demand in the first period, the second period problem is

$$V_2(I_1, x_1, p_1) = \max_{p_2 \geq 0} E\Big[p_2 \min\{I_1, D_2\}\Big| x_1, p_1, p_2\Big],$$

where the distribution of D_2 given x_1, p_1, p_2 is given by equation (6). Note also that $I_1 = \max\{I_0 - x_1, 0\}$ and $V_2(0, x_1, c_1) = 0$. The problem of the first period is

$$V_1(I_0) = \max_{p_1 \geq 0} E\Big[p_1 \min\{I_0, D_1\} + V_2\Big((I_0 - D_1)^+, D_1, p_1\Big)\Big| p_1\Big],$$

where the distribution of D_1 given p_1 is shown in equation (5).

For No Learning model, the demand distribution for the second periods does not depend on the price or the sales in the first period. While the first period's demand distribution is still given by equation (5), the second period's demand distribution has different parameters although it is still Negative Binomial:

$$D_2 \sim NB\Big(\alpha, \frac{\beta}{\beta + m(p_2)(1-t)}\Big), \tag{7}$$

Then, the dynamic program for No Learning model can be constructed as:

$$V_2(I_1) = \max_{p_2 \geq 0} E\Big[p_2 \min\{I_1, D_2\}\Big| p_2\Big],$$

where the distribution of D_2 given p_2 is given by equation (7). Note again that $I_1 = \max\{I_0 - x_1, 0\}$ and $V_2(0) = 0$. The problem of the first period is

$$V_1(I_0) = \max_{p_1 \geq 0} E\Big[p_1 \min\{I_0, D_1\} + V_2\Big((I_0 - D_1)^+\Big)\Big|p_1\Big],$$

where the distribution of D_1 given p_1 is given by equation (5).

Throughout the computational study, we assume equal periods, i.e., $t = 0.5$ and we allow the first and second prices to be in the set $\{0.50, 0.55, 0.60, .., 0.95, 1.00\}$. We do not put any restrictions on the direction of the price change in the second period, i.e., the second period price can be higher or lower than the first period price.

4.1 The impact of the accuracy of the initial point estimate

In this part of the study, we assess the impact of the initial estimate on profits of Learning and No Learning models in a variety of settings. For the price sensitivity of demand, we use a moderate value, e.g., $\gamma = 3$.

The analysis is done in two steps; first we keep the initial point estimate constant and vary the true rate of the Poisson distribution and later we keep the true rate of the Poisson distribution and vary the initial point estimate. Note that the value of the initial estimate is α/β. In the first part of the analysis, we set $\alpha/\beta = 20$. However in order to study also the impact of decision maker's reliance on the initial estimate, we use two scenarios. In High Variance case, $\alpha = 10$ and $\beta = 0.5$, resulting in a variance of 40 for the gamma distribution (or a coefficient of variation of $1/\sqrt{10}$). In Low Variance case, $\alpha = 40$ and $\beta = 2$ resulting in a variance of 10 for the gamma distribution (or a coefficient of variation of $1/\sqrt{40}$). We also use different values for the starting inventory level, in order to incorporate the impact of imbalance between supply and demand in pricing decisions. This first step of the analysis is summarized in Table 1.

Note that λ is the true Poisson rate when the price is set at the maximum price 1.00. The true Poisson rate takes on values 10, 15, 20, 25 and 30, while the decision maker's initial point estimate is fixed at 20. Note also that optimal policies for No Learning and Learning models are evaluated using Poisson distribution with the true rate. When we compare the profits obtained from No Learning and Learning models, we conclude that learning from observed sales is most beneficial when the initial point estimate is inaccurate and when the variance is high (the decision maker relies less on the initial estimate and is more willing to update its estimate based on observed sales). This gives an opportunity to Learning model

Table 1: The impact of initial estimate, revenues as a function of λ

Initial Inv			λ 10	15	20	25	30
10	Perfect Information		9.0361	9.8697	9.9918	9.9997	10.0000
	High Variance	No Learning	8.7672	9.8647	9.9918	9.9997	10.0000
		Learning	8.9886	9.8564	9.9816	9.9974	9.9996
	Low Variance	No Learning	8.7521	9.8634	9.9918	9.9997	10.0000
		Learning	8.8166	9.8681	9.9914	9.9995	10.0000
20	Perfect Information		14.2552	16.8405	18.6529	19.6623	19.9511
	High Variance	No Learning	11.8433	16.3099	18.6484	19.5302	19.8411
		Learning	12.7448	16.5088	18.5672	19.5032	19.8488
	Low Variance	No Learning	11.8396	16.2913	18.6510	19.5608	19.8675
		Learning	12.1111	16.3755	18.6448	19.5545	19.8660
30	Perfec Information		17.8773	21.7092	24.4823	26.6369	28.3606
	High Variance	No Learning	14.8758	20.9285	24.4795	25.7868	26.0604
		Learning	15.6680	20.9794	24.3064	26.2548	27.1922
	Low Variance	No Learning	14.8758	20.9271	24.4804	25.8100	26.0904
		Learning	15.4997	21.2023	24.4517	25.8500	26.1351

to quickly identify the inaccuracy of the initial estimate and correct the estimate for the second period. The benefits are more pronounced when the true Poisson rate is lower (e.g., $\lambda = 10$) than the initial estimate and the initial inventory levels are high (e.g., $I_0 = 20$ and $I_0 = 30$). Since the maximum allowed price is 1.00, pricing is more instrumental when the demand rate is significantly lower than the initial inventory.

Notice that in 13 cases, No Learning model is performing better than Learning model. These are the cases where the initial estimate is fairly accurate and updating the demand distribution using a random sample can therefore reduce the profits. The reductions are minimal when the variance is low (the decision maker relies more on the initial estimate and is less willing to update its estimate based on observed sales). It should be noted however that the savings due to Learning model when the initial estimate is inaccurate is much higher than the losses due to Learning model when the initial estimate is accurate.

Finally we should note that when the initial inventory is low (i.e., $I_0 = 10$), pricing is not very useful as the maximum price is set at 1.00. Therefore, the difference between the Learning and No Learning models are minimal, and both models can perform very close to the Perfect Information model.

The second step of the analysis is summarized in Table 2. In the second step of the analysis we fixed the true Poisson rate (λ) at 20 and let the initial point estimate (α/β)

Table 2: The impact of initial estimate, revenues as a function of α/β

Initial Inv			α/β 10	15	20	25	30
10	Perfect Information		9.9918	9.9918	9.9918	9.9918	9.9918
	High Variance	No Learning	9.9554	9.9892	9.9918	9.9918	9.9918
		Learning	9.9461	9.9780	9.9816	9.9890	9.9890
	Low Variance	No Learning	9.9647	9.9892	9.9918	9.9918	9.9918
		Learning	9.9645	9.9890	9.9914	9.9918	9.9918
20	Perfect Information		18.6529	18.6529	18.6529	18.6529	18.6529
	High Variance	No Learning	16.2351	17.9570	18.6484	18.5209	18.3379
		Learning	16.3240	17.9525	18.5672	18.5995	18.6227
	Low Variance	No Learning	16.2632	17.9570	18.6510	18.5209	18.2916
		Learning	16.2620	17.9903	18.6448	18.6101	18.4804
30	Perfect Information		24.4823	24.4823	24.4823	24.4823	24.4823
	High Variance	No Learning	20.0839	23.8004	24.4795	23.8294	22.6325
		Learning	20.2463	23.9832	24.3064	24.0655	23.5536
	Low Variance	No Learning	20.1108	23.7998	24.4804	23.8294	22.6138
		Learning	20.1874	23.9254	24.4517	24.0109	23.0534

take on values 10, 15, 20, 25 and 30. In order to eliminate the impact of the variance in the analysis, we fixed the coefficient of variation of the gamma distribution (which is equal to $\sqrt{(\alpha/\beta^2)}/(\alpha/\beta) = 1/\sqrt{\alpha})$ to $1/\sqrt{10}$ for the High Variance case, and to $1/\sqrt{40}$ for Low Variance case.

In addition to results that are similar to those that are obtained in the first step, the second step provides an additional interesting observation. While the maximum revenue is achieved when the estimate is accurate in No Learning model, the same is not necessarily true for Learning model. When the initial inventory is 10 for both High and Low Variance, and when the initial inventory is 20 for High Variance, the maximum revenue is achieved when the decision maker is in fact overestimating the demand. By overestimating the demand, the decision maker is less likely to charge lower than the maximum price in the second based on a random sample.

4.2 The impact of the variance of the initial estimate

In this part of the study, we investigate the impact of the variance of the initial estimate on the performance of Learning and No Learning models. Note again that the variance of the initial estimate reflects the decision maker's reliance on its initial estimate and how much she is willing to update her estimate based on observed sales for Learning model.

The analysis is summarized in Table 3 for an initial inventory level of 20, and Table 4 for an initial inventory level of 30. For both tables, parameter α of the Gamma distribution is taking on values 5, 10, 15, 25, 40 and 80 while the parameter β of the Gamma distribution is taking on values 0.25, 0.5, 0.75, 1.25, 2 and 4, respectively. This keeps the mean of the Gamma distribution constant at 20, while the variance of the Gamma distribution is taking on values 80, 40, 26.67, 16, 10, and 5. The tables show the optimal first period price, expected optimal second period price and optimal expected revenue for Perfect Information model to form a benchmark. As mentioned earlier, No Learning model uses the same Negative Binomial distribution when deciding the first period price and deriving a policy for the second period price, while Learning model uses an updated Negative Binomial distribution for the second period. However, the expected revenues and expected second period prices reported in Table 3 and Table 4 use the true Poisson distribution when taking the expectations. In the less likely case that the initial inventory is totally depleted in the first period, we take the second period price to be 1.00 when calculating expected second period price.

When the initial inventory (I_0) is 20, we note that No Learning and Learning models set the initial price to 1.00 for all variance levels (Table 3). When the initial inventory (I_0) is 20 and the true Poisson rate (λ) is 10, we observe that Perfect Information model sets the initial price to 0.80, significantly lower than No Learning and Learning models. However, as the variance gets higher, Learning model is better able to correct its estimate and thus charges lower prices in the second period. This is in contrast to No Learning model where the second period price and the revenue is insensitive to the variance.

When the initial inventory (I_0) is 20 and the true Poisson rate (λ) is 20, we observe that Perfect Information model sets the initial price to 1.00. The revenues of No Learning and Learning models are also quite close to the optimal revenue obtained in Perfect Information model. However, we note that when the variance is high for Learning model, the decision maker runs the risk of charging a less than optimal price as she may interpret a randomly low demand in the first period as a sign for low demand overall.

When the initial inventory (I_0) is 20 and the true Poisson rate (λ) is 30, we again observe that Perfect Information model sets the initial price to 1.00. Since the demand rate is quite high as compared to the supply, expected optimal second price also needs to be close to 1.00. Similar to the case when the true Poisson rate (λ) is 20, the decision maker still has the risk of charging a less than optimal second period price, based on a randomly low demand in the first period when he uses Learning model. This is especially true for High Variance case.

Table 3: The impact of variance (initial inventory=20)

			α					
			5	10	15	25	40	80
	Learning	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00
	No earning	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00
$\lambda = 10$	Perfect Information	p_1^*	0.80					
		$E[p_2^*]$	0.7383					
		$E[V]$	14.2552					
	Learning	$E[p_2^*]$	0.7128	0.7562	0.7670	0.7922	0.8129	0.8216
		$E[V]$	13.1040	12.7448	12.6378	12.3563	12.1111	12.0021
	No Learning	$E[p_2^*]$	0.8345	0.8345	0.8345	0.8345	0.8349	0.8401
		$E[V]$	11.8433	11.8433	11.8433	11.8433	11.8396	11.7814
$\lambda = 20$	Perfect Information	p_1^*	1.00					
		$E[p_2^*]$	0.9400					
		$E[V]$	18.6529					
	Learning	$E[p_2^*]$	0.9088	0.9198	0.9208	0.9288	0.9330	0.9334
		$E[V]$	18.4570	18.5672	18.5799	18.6217	18.6448	18.6469
	No Learning	$E[p_2^*]$	0.9298	0.9298	0.9298	0.9298	0.9355	0.9400
		$E[V]$	18.6484	18.6484	18.6484	18.6484	18.6510	18.6529
$\lambda = 30$	Perfect Information	p_1^*	1.00					
		$E[p_2^*]$	0.9999					
		$E[V]$	19.9511					
	Learning	$E[p_2^*]$	0.9872	0.9881	0.9882	0.9893	0.9895	0.9895
		$E[V]$	19.8367	19.8488	19.8494	19.8621	19.8660	19.8661
	No Learning	$E[p_2^*]$	0.9863	0.9863	0.9863	0.9863	0.9896	0.9902
		$E[V]$	19.8411	19.8411	19.8411	19.8411	19.8675	19.8737

When the initial inventory (I_0) is 30, we note that No Learning model sets the initial price to 0.85 for all variance levels, while Learning model sets the initial price to 0.90 when the variance is high and to 0.85 when the variance is low (Table 4). The difference between the initial prices of Learning and No Learning models shows that the fact that the decision maker will learn from observed sales may lead the decision maker to different decisions, even before she observes sales.

Table 4 shows results similar to those in Table 3, except that now, Learning model provides significant benefits also when the true Poisson rate (λ) is 30. The value of learning is more pronounced, when the starting inventory level is high, i.e., correcting an underestimation pays more.

Table 4: The impact of variance (initial inventory=30)

			α					
			5	10	15	25	40	80
	Learning	p_1^*	0.90	0.90	0.90	0.90	0.85	0.85
	No Learning	p_1^*	0.85	0.85	0.85	0.85	0.85	0.85
$\lambda = 10$	Perfect Information	p_1^*	0.65					
		$E[p_2^*]$	0.6327					
		$E[V]$	17.8773					
	Learning	$E[p_2^*]$	0.5967	0.6215	0.6314	0.6582	0.6828	0.6995
		$E[V]$	15.9317	15.6680	15.5679	15.2358	15.4997	15.2789
	No Learning	$E[p_2^*]$	0.8345	0.8345	0.8345	0.8345	0.8349	0.8401
		$E[V]$	14.8544	14.8758	14.8758	14.8758	14.8758	14.8748
$\lambda = 20$	Perfect Information	p_1^*	0.85					
		$E[p_2^*]$	0.8565					
		$E[V]$	24.4823					
	Learning	$E[p_2^*]$	0.8035	0.8081	0.8067	0.8091	0.8518	0.8541
		$E[V]$	24.2300	24.3064	24.3227	24.3962	24.4517	24.4685
	No Learning	$E[p_2^*]$	0.8500	0.8500	0.8500	0.8500	0.8520	0.8565
		$E[V]$	24.4795	24.4795	24.4795	24.4795	24.4804	24.4823
$\lambda = 30$	Perfect Information	p_1^*	1.00					
		$E[p_2^*]$	0.9496					
		$E[V]$	28.3606					
	Learning	$E[p_2^*]$	0.9536	0.9536	0.9495	0.9466	0.9738	0.9739
		$E[V]$	27.1888	27.1922	27.1600	27.1293	26.1351	26.1360
	No Learning	$E[p_2^*]$	0.9657	0.9657	0.9657	0.9657	0.9695	0.9712
		$E[V]$	26.0604	26.0604	26.0604	26.0604	26.0904	26.1116

4.3 The impact of price elasticity

In Table 5, the impact of price elasticity of demand is analyzed for Perfect Information, Learning and No Learning models. In this specific analysis, the parameter α of the Gamma distribution is taken as 10 and the parameter β of the Gamma distribution is taken as 0.5 leading to an initial estimate with mean 20. The γ, value which controls the price elasticity of demand, takes on 7 values between 1.0 and 4.0, where $\gamma = 1.0$ models inelastic demand (for this case, the optimal price is 1.00 since reducing the price will not modify demand). We first note that the optimal revenue is an increasing function of the price elasticity of demand for Perfect Information model as the decision maker is better able to manipulate the demand.

For $\lambda = 10$, the decision maker is initially overestimating the demand for both Learning and No Learning models and charges an initial price higher than the optimal price in Perfect

Table 5: The impact of price elasticity (initial inventory=20)

			γ						
			1.0	1.5	2.0	2.5	3.0	3.5	4.0
	Learning	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	No learning	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\lambda = 10$	Perfect Information	p_1^*	1.00	0.75	0.70	0.75	0.80	0.80	0.80
		$E[p_2^*]$	1.0000	0.7338	0.7141	0.7204	0.7383	0.7764	0.8211
		$E[V]$	9.9972	10.8237	12.2038	13.3753	14.2552	14.9601	15.4778
	Learning	$E[p_2^*]$	1.0000	0.8129	0.7509	0.7397	0.7562	0.7665	0.7799
		$E[V]$	9.9972	10.3495	11.1176	12.0069	12.7448	13.4707	14.0700
	No Learning	$E[p_2^*]$	1.0000	0.9106	0.8622	0.8419	0.8345	0.8362	0.8451
		$E[V]$	9.9972	10.1957	10.6685	11.2426	11.8433	12.3988	12.8192
$\lambda = 20$	Perfect Information	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		$E[p_2^*]$	1.0000	0.9730	0.9524	0.9461	0.9400	0.9352	0.9356
		$E[V]$	18.2233	18.2669	18.3928	18.5295	18.6529	18.7620	18.8537
	Learning	$E[p_2^*]$	1.0000	0.9609	0.9375	0.9148	0.9198	0.9151	0.9171
		$E[V]$	18.2233	18.2575	18.3586	18.4402	18.5672	18.6515	18.7302
	No Learning	$E[p_2^*]$	1.0000	0.9851	0.9622	0.9462	0.9298	0.9298	0.9260
		$E[V]$	18.2233	18.2599	18.3843	18.5293	18.6484	18.7539	18.8318
$\lambda = 30$	Perfect Information	p_1^*	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		$E[p_2^*]$	1.0000	1.0000	1.0000	0.9999	0.9999	0.9996	0.9996
		$E[V]$	19.9505	19.9505	19.9506	19.9508	19.9511	19.9516	19.9521
	Learning	$E[p_2^*]$	1.0000	0.9967	0.9932	0.9876	0.9881	0.9849	0.9850
		$E[V]$	19.9505	19.9230	19.8920	19.8410	19.8488	19.8239	19.8261
	No Learning	$E[p_2^*]$	1.0000	0.9989	0.9955	0.9918	0.9863	0.9863	0.9822
		$E[V]$	19.9505	19.9437	19.9171	19.8875	19.8411	19.8424	19.8128

Information model. However, Learning model can partially correct its estimate based on observed sales and improve its revenue by reducing the price in the second period. The revenue of Learning model increases as the price sensitivity increases since the price reductions are more effective with high price sensitivity. As the price sensitivity increases, we observe that the difference between Perfect Information and Learning and the difference between Learning and No Learning also increase as the information is more useful with a more elastic demand.

For $\lambda = 20$, Learning model performs worse than No Learning model for all demand elasticities. This is because the initial estimate is accurate, and the decision maker is better off if she does not change her estimate based on observed sales. We also see that performance of Learning and No Learning models improves as γ increases, which shows that when the demand is accurately estimated, an elastic demand will always help.

For $\lambda = 30$, the decision maker is initially underestimating the demand for both Learning and No Learning models. Note that with Perfect Information model, the initial price needs to be 1.00, and we hardly need a reduction in price in the second period. However, with inaccurate information, both Learning and No Learning models can ask for price reductions in the second period especially when the demand is highly elastic. However, we should see that the relationship between γ and the performances of Learning and No Learning models is not clear to have any further conclusions.

5. Inventory flexibility

The analysis so far assumes that there are no further replenishment opportunities available once the selling season starts. In the apparel industry, this corresponds to the case when the retailer orders from overseas and is not able to order during the season because of the long lead times relative to the selling seasons. Obviously, this limits the retailer's control during the selling season to pricing only, which sharply diminishes its responsiveness. As a result, some retailers are willing to use domestic suppliers and be able to order frequently, even though domestic suppliers are more costly. With domestic suppliers, the retailer is also able to make its initial order much closer to the season, when there is more information, hence less variance, about the demand process.

Some companies are using two (or sometimes even three) different suppliers for the very same product: an off–shore low–cost supplier for the initial large orders, and a domestic high–cost supplier for replenishments during the selling season (Apparel Industry Magazine [1]).

We study the value of this additional flexibility in the context of our pricing model. In a related study, Gurnani and Tang [23] study the impact of forecast improvements by having the flexibility to order at two instances, one of them being closer to the season. Their model differs from ours as they do not consider the possibility of ordering during the season by utilizing a structured learning from observed sales. Also they do not consider any pricing during the season. While their model allows the cost to go up or down as the merchandise is ordered closer to the season, we always assume that the ordering later is more costly reflecting the reality in the apparel industry.

In our model, the off–shore strategy will allow the company to order only once, but possibly with a low unit cost c^o. The domestic strategy will allow the company to order before and during the selling season, but possibly with a high unit cost c^d. The blended strategy, on the other hand, will allow the company to make its initial order at a unit cost c^o, but later replenishments at the unit cost c^d. We assume that there are no other costs involved, the pricing and inventory decisions are made simultaneously at the start of the each period; period lengths are equal for each strategy and the lead time is zero for all strategies.

To be able to compare these three strategies, we need to extend our pricing model to allow for inventory decisions. We suggest the following model.

The problem is determining prices and stock levels in periods $1, \ldots, N$ so that total expected profit is maximized. We use a discrete–time dynamic programming model.

Let $V_n(I_{n-1}, X_{n-1}, M_{n-1})$ be the maximum expected profit from period n through N where the starting inventory is I_{n-1} and the cumulative sales and cumulative price multipliers are X_{n-1} and M_{n-1}, respectively. Also let B_n be the starting inventory level for period n, after the retailer receives its orders. Thus, the retailer acquires $B_n - I_{n-1}$ new units in the beginning of period n. Let p_n be the price set in period n and let a_n be the acquisition cost per unit in period n.

Backward recursion can be written as

$$V_n(I_{n-1}, X_{n-1}, M_{n-1}) =$$

$$\max_{p_n \geq p_s, B_n \geq I_{n-1}} \quad E\Big[c_n(B_n - I_{n-1}) + p_n \min\{D_n, B_n\} \\ + V_{n+1}\Big((B_n - D_n)^+, X_{n-1} + D_n, M_{n-1} + m(p_n)\Big)\Big| X_{n-1}, M_{n-1}, p_n\Big].$$

Boundary conditions are

$$\begin{aligned} V_{N+1}(I_N, X_N, M_N) &= p_s I_N, \quad \text{for all } I_N, X_N, M_N. \\ X_0 = M_0 = I_0 &= 0. \end{aligned}$$

The first condition states that any left over merchandise has only salvage value (p_s) when the season ends at the end of period N. The dynamic program can be solved by starting with the Nth period and proceeding backwards.

For the off–shore strategy, the model can be used with the following acquisition costs.

$$\begin{aligned} c_1 &= c^o, \\ c_n &= \infty, \quad n = 2, \ldots, N. \end{aligned}$$

For the domestic strategy, we simply have

$$c_n = c^d, \quad n = 1, \ldots, N.$$

For the blended strategy, we have,

$$\begin{aligned} c_1 &= c^o, \\ c_n &= c^d, \quad n = 2, \ldots, N. \end{aligned}$$

Let $V^o(c^o, c^d), V^d(c^o, c^d)$ and $V^b(c^o, c^d)$ be the optimal profits for the off–shore, domestic and blended strategies respectively. Without any analytical derivations, it is easy to see the following.

Observation 1 *When the off–shore cost is higher than or equal to the domestic cost (which is not likely), domestic strategy outperforms the off–shore strategy. That is, for* $c^o \geq c^d$, $V^o(c^o, c^d) \leq V^d(c^o, c^d)$.

Intuition: A domestic policy can simply imitate the the optimal off–shore policy by ordering as much as the optimal off–shore policy does in the first period and ordering zero units in later periods. Since the acquisition costs are lower for the domestic orders, this policy generates more profit than the optimal off–shore policy.

Observation 2 *When the off–shore cost is lower than the domestic cost (which is typical), blended strategy outperforms both strategies. That is, for* $c^o < c^d$, $V^b(c^o, c^d) \geq V^o(c^o, c^d)$ *and* $V^b(c^o, c^d) > V^d(c^o, c^d)$.

Intuition: A blended policy can imitate the optimal off–shore policy by simply ordering as much as the off–shore policy does in the first period and ordering zero units in later periods. Since the acquisition costs are the same for blended and off–shore strategies in the first period, this policy generates the same profit with the optimal off–shore profit. Likewise, another blended policy can imitate the optimal domestic policy by simply ordering as much as the optimal domestic policy does in each period. Since first period's acquisition costs are lower for the blended strategy, this policy generates more profit than the optimal domestic policy.

While these comparisons are trivial, a question of interest is under which other circumstances the retailers should favor domestic policies over off–shore policies and under which circumstances the gap between the blended and domestic and off–shore policies are minimal. This is important as acquisition costs may not be the only concern for a retailer. For example, using an additional supplier may involve additional fixed setup costs and complicate the coordination of the sourcing process which disadvantage the blended strategies. Also, in our study we do not consider the inventory holding and other logistics costs that may be incurred within the selling season. Inclusion of inventory holding costs to the model may favor domestic and blended strategies against the off–shore strategy as domestic purchases may be used for frequent replenishments and may reduce inventory levels. However, if unit inventory holding costs are proportional to the unit cost and domestic cost is excessively higher than the import cost, inventory reduction effect will be less apparent.

We use the computational design in Section 4 to answer above questions. Again, the mean demand is $\alpha/\beta = 20$ and we have two periods of equal length. Different from the analysis in Section 4, the starting inventory level is optimized for all strategies. We assume that the maximum price to charge is 1.00. We set the off–shore acquisition cost to 0.5 and vary the domestic acquisition cost to study the effect of acquisition costs on different strategies. We note that in this analysis, we use Learning model as described in Section 3, and the expected profits are evaluated using the Negative Binomial distribution (with initial parameters in the first period and with updated parameters in the second period). We do not use the evaluations based on the true Poisson rate, as this is not available to the decision maker until after the season, and the decision maker makes her sourcing decisions based on her prior beliefs and how she updates her beliefs based on sales during the season. Figure 1 shows the (expected) optimal profits of off–shore, domestic and blended strategies when $\gamma = 2$ and when variance equals 1.5μ or 3μ. The optimal profits are normalized with the profit of the optimal off–shore policy when variance equals to 3 μ. We first note that the optimal off–shore profits do not vary with the domestic acquisition costs. Blended strategies, as shown above, outperform the domestic and off–shore strategies. Clearly, optimal blended and domestic profits decrease with acquisition costs. However, optimal blended profit curves are rather flat, as blended strategies prefers to order more from the off–shore supplier as the domestic supplier becomes more expensive. In fact, optimal blended profits approach optimal off–shore profits as domestic acquisition costs increase. The reduction in profits is more dramatic for domestic policies as they have to live with the expensive domestic suppliers. While domestic policies outperform off–shore policies for low domestic acquisition costs, off–shore policies are favorable as the domestic suppliers become more costly.

For this particular example, domestic and off–shore profit curves intersect when the normalized domestic cost is 1.1 for $\sigma^2 = 3\mu$. This means that the "break–even" point where off–shore profit equals domestic profit is when the unit domestic acquisition cost is 10% more than the unit off–shore acquisition cost. Any unit domestic acquisition cost 10% more

Figure 1: Comparison of off-shore, domestic and blended strategies ($\gamma = 2$)

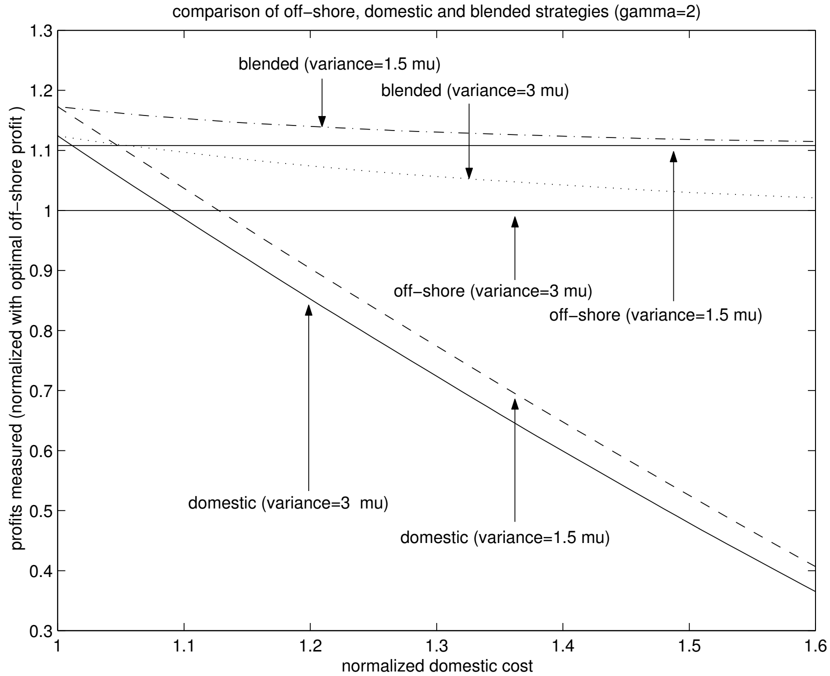

than the unit off–shore acquisition cost will lead the retailers to source their merchandise off–shore.

Another important factor for the efficiency of the off–shore, domestic and blended policies is the variance of the demand process. Typically, apparel retailers choose domestic suppliers for their high fashion content–high variance merchandise, while standard low fashion content–low variance merchandise can be sourced overseas. Also, the policy itself may help to reduce the variance as the domestic strategies can order closer to the season. Figure 2 shows the optimal profits for domestic, off–shore and blended strategies for three different domestic acquisition costs (1.00, 1.06 or 1.12 times the off–shore acquisition cost). Again, the profits are normalized with the optimal off–shore profit for $\sigma^2 = 3\mu$. Note that blended and domestic policies are equivalent when the cost equals 1.00. These policies outperform any other policy. As the variance increases, all profits decrease. The off–shore optimal profit curve is steeper as off–shore policies are subject to more variance since they order only once. Optimal domestic policy when the cost equals 1.06 is inferior to the off–shore policy for low variance levels, but becomes favorable as the variance increases. Note again that the base line is the optimal off–shore profit when the variance is 3 μ. If we can reduce the variance to 1.75 μ by using a domestic policy, even the domestic acquisition costs of 1.06 can be desirable.

Finally, the price sensitivity of the customers also affects the relative efficiency of these

Figure 2: Comparison of off-shore, domestic and blended strategies ($\gamma = 2$)

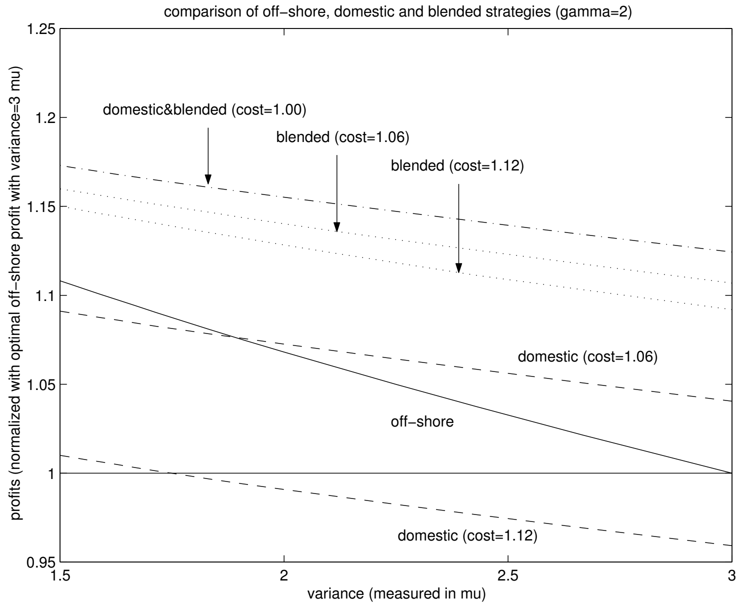

policies. We expect the off–shore strategy to be more sensitive to price sensitivity (γ) as pricing is the only control for such strategy once the season starts. Figure 3 shows the optimal profits for varying levels of γ. Note that when $\gamma = 1$, the demand is inelastic and it is optimal to keep the price at its maximum. Thus, optimal profits at $\gamma = 1$ represent the optimal profits when the only control over the process is through inventory. All profits increase, as the price sensitivity increases. As expected, optimal off–shore profit increases faster with the price sensitivity. The off–shore strategy outperforms the domestic strategy with cost 1.06 when γ is close to 3.

Combining these ideas, we generate the regions in which one strategy is favorable to the other. Figure 4 shows the trade–off curves for off–shore and domestic strategies for $\gamma = 2$ and $\gamma = 3$. For variances and domestic costs on these lines, off–shore and domestic strategies generate the same profit. As domestic cost increases and/or variance decreases, off–shore strategy becomes more desirable and vice versa. Note that, the region for which off–shore strategy is more profitable is larger when $\gamma = 3$, reflecting the increased strength of off–shore strategies with price sensitive demand. In general, as γ increases, the trade–off curve moves to southeast. Although it is difficult to detect visually, we observe that the trade–off curves are concave in variance, possibly becoming flatter as variance gets larger. This means that if the variance gets excessively high, variance differences would have less impact and supplier selection decisions would depend more on cost differences and price sensitivity of demand.

Figure 3: Comparison of off-shore, domestic and blended strategies as a function of price sensitivity

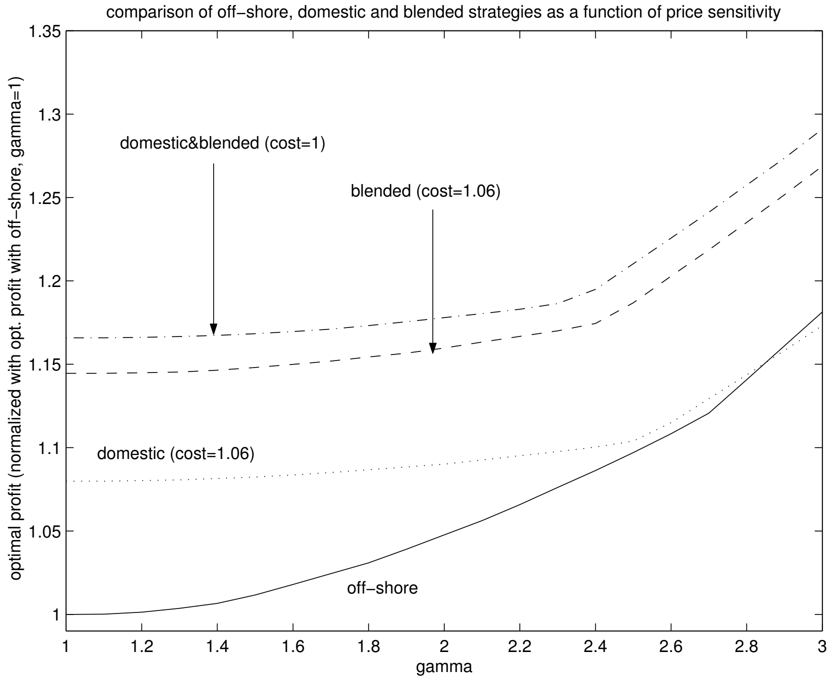

Figure 4: Trade-off curves for off-shore and domestic production

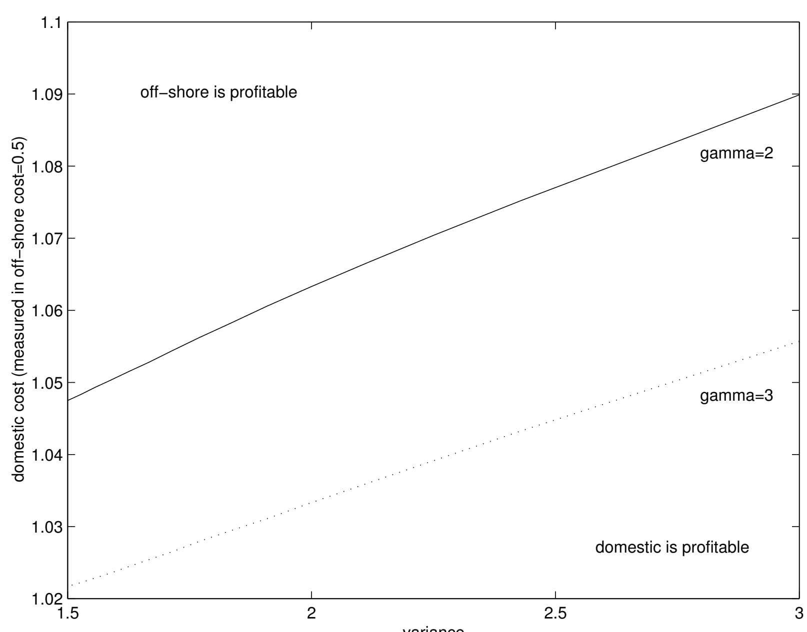

Figure 5: Trade-off between domestic and off-shore production

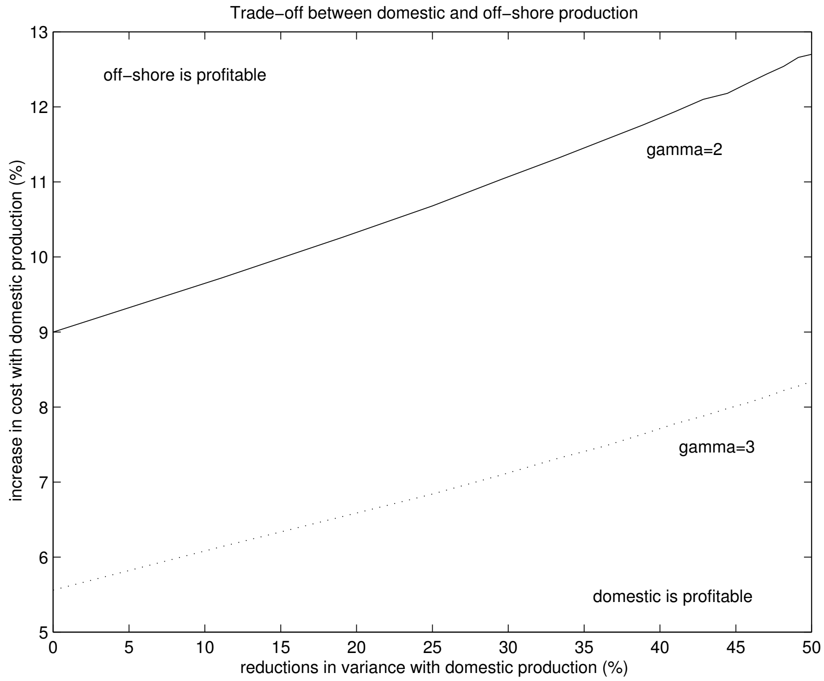

The above analysis is based on the fact that the products are subject to same level of uncertainty under both strategies. However, in most cases, the choice of strategy itself may affect the level of uncertainty. As the retailers are able to order closer to the season with domestic strategies, they are able to know more about the consumer tastes that will shape the demand in the coming season and hence they face a more stable demand when they make their ordering decisions. To incorporate the possible reductions in variance, we choose a base case which is an off–shore strategy with acquisition cost equals 0.5 and variance (σ^2) equals 3 μ. The trade–off curves in Figure 5 shows the increases in cost and reductions in variance with domestic strategy for which domestic and off–shore strategies generate equal profits. For example, for $\gamma = 3$, if the domestic cost is about 7% higher than off–shore cost, the domestic strategy will still result in higher profits, if the variance is reduced by more than 30% as a result. Alternatively, if the variance is reduced about 30%, the domestic strategy will generate higher profits only if the cost does not increase by more than 7%. Again, optimality region is larger for off–shore strategy for more price sensitive demand.

Each retailer faces its own trade–off curve for each apparel item it offers and makes its decision to source it overseas or domestically. An aggregation of these individual decisions determines the market share of imports and domestic production in the domestic market. In Table 6, we provide the average unit import and domestic costs and the market share of imports for selected apparel in the U.S. market in 2002.

Table 6: Market Share and Cost of Imports in Apparel in 2002

	Average Domestic Price[1] (c^d)	Average Import Price[2] (c^o)	Ratio (c^d/c^o)	Market Share of Imports[3]
Men's				
Sweaters	15.27	10.42	1.47	98.3
Swimwear	12.93	4.52	2.86	100.0[4]
Suits	100.64	62.71	1.60	74.0
Women's				
Sweaters	12.24	9.69	1.26	86.1
Swimwear	13.61	6.10	2.23	70.8
Dresses	20.10	9.08	2.21	64.0

Data compiled from U.S. Census Bureau [39]. Prices are per dozen in U.S. dollars.

1 Average cost ($) per unit for manufacturers' shipments.
2 Average cost ($) (cost+insurance+freight) per unit from imports for consumption.
3 Derived by dividing imports for consumption to apparent consumption in the U.S. market.
4 Missing quantity data for 2002 imports is estimated using data from 2001.

We observe that for all product categories, imports have a substantial cost advantage. This is particularly true for men's and women's swim–wear. While this translates into a perfect market domination of imports in men's swim–wear, domestic manufacturers still control 70.8% of the market for women's swim–wear. Similar arguments are valid for sweaters for men and women and suits for men and dresses for women. This shows that the cost advantage is not the only factor in supplier selection. Although it is very difficult to find an aggregate measure for the variance of demand in apparel, we are certainly aware of the importance of fashion in women's apparel. Popular styles and colors change every year, making it very difficult to forecast demand for a particular SKU. From both our computational analysis and industry data, we see that predictability of demand plays a considerable role in sourcing decisions. When the variance effect is less apparent (as in sweaters product category, or as in men's apparel in general as compared to women's apparel), we observe that cost difference is the main driver for such decisions.

6. Conclusion

In this paper, we study the pricing decisions of a perishable products retailer in the existence of demand learning. This is one of the first studies that incorporate "structured" Bayesian updating in the context of pricing for perishable products. The resulting model is computationally feasible and easy to understand and implement. We think that our model is most useful for apparel retailers, as this industry is identified with high levels of uncertainty, most of which can be resolved after observing sales during the earlier weeks of the selling season. Moreover, information required for the application of our model is readily available through point–of–sales scanners.

Through our computational study, we are able to understand the economics of pricing in this context. First, we observe that the optimal price in a given period is a non–decreasing function of sales in the earlier periods when demand learning takes place. Second, we pinpoint the circumstances under which this learning based on observed sales has the most value. We study the impact of the accuracy and the variance of the initial estimate and the price elasticity of demand. Our major finding here is that demand learning is most beneficial when the initial estimate is inaccurate, the demand/supply mismatches necessitates price changes and demand is sensitive to price changes. Finally, we study the impact of an opportunity to procure merchandise during the season, in addition to the up-front procurement before the selling season. This helps us to see how supplier selection decisions are affected by the volatility and price–sensitivity of demand and the procurement costs. We support our conclusions with aggregate data from the apparel industry.

We note that our model with inventory flexibility can be extended to incorporate lead times, inventory holding costs and set–ups (cost and/or time) that may be attached to each purchase. We refrain from doing so, as their effects on costs are fairly trivial and their inclusion may complicate the presentation of the model.

Several avenues for future research are in order. First, the model could be extended to allow for the case where the time to switch price is also a decision variable. This way, one can study the impact of learning on optimal markdown times. An interesting question is whether the Learning model will always delay the pricing decisions in an effort to learn more about the underlying demand. Second, the impact of more than two periods and mark–down only restrictions can be studied. Finally, the model can be modified to incorporate structured learning about the price elasticity of demand.

Bibliography

[1] Apparel Industry Magazine. Quick response mandates today: an assessment of the industry's status and needs after 10 years of implementing QR. March (1997), 45–48.

[2] Aviv, Y. and A. Pazgal. Pricing of short life-cycle products through active learning. Working paper. Washington University, Olin School of Business. 2002.

[3] Azoury, K. S. and B. .L. Miller. A comparison of the optimal ordering levels of Bayesian and non-Bayesian inventory models *Management Science*, **30** (1984), 993–1003.

[4] Azoury, K. S. Bayes solution to dynamic inventory models under unknown demand distributions. *Management Science*, **31** (1985) , 1150–1161.

[5] Balvers, R. J. and T. F. Cosimano. Actively learning about demand and dynamics of price adjustment. *The Economic Journal*, **100** (1990), 882–898.

[6] Bitran, G. R. and S. .V. Mondschein. Periodic pricing of seasonal products in retailing. *Management Science*, **43** (1997), 64–79.

[7] Bitran, G. R. and H. K. Wadhwa. A methodology for demand learning with an application to the optimal pricing of seasonal products. M.I.T. Sloan School of Management, Working Paper #3896–96. 1996.

[8] Bitran, G. R. and H. K. Wadhwa. Some structural properties of the seasonal product pricing problem. M.I.T. Sloan School of Management, Working Paper #3897–96. 1996.

[9] Bradford, J. W. and P. K. Sugrue. A Bayesian approach to the two–period style–goods inventory problem with single replenishment and heterogeneous Poisson demands. *Journal of the Operational Research Society*, **41** (1990), 211–218.

[10] Brown, G. F. and W. F. Rogers. A Bayesian approach to demand estimation and inventory provisioning. *Naval Research Logistics Quarterly*, **19** (1973), 607–624.

[11] Carlson, P. G. Fashion retailing: the sensitivity of rate of sale to mark–down. *Journal of Retailing*, **59** (1983), 67–76.

[12] Chain Store Age. Managing mark–down madness. March 1 (1999), page 108.

[13] Chang, S. H. and D. E. Fyfee. Estimation of forecast errors for seasonal–style–goods sales. *Management Science*, **18** (1971), B-89–B-96.

[14] Chun, Y. H. Optimal pricing and ordering policies for perishable commodities. *European Journal of Operational Research*, **144** (2002), 68–82.

[15] DeGroot, M. H. Optimal statistical decisions. McGraw–Hill Book Company. 1970.

[16] Eppen, G. D. and A. .V. Iyer. Improved fashion buying with Bayesian updates. *Operations Research*, **45** (1997), 805–819.

[17] Eppen, G. D. and A. .V. Iyer. Backup agreements in fashion buying – The value of upstream flexibility. *Management Science*, **43** (1997), 1469–1484.

[18] Feng, Y. and G. Gallego. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, **41** (1995), 1371–1391.

[19] Fisher, M. L., J. H. Hammond, W. R. Obermeyer and A. Raman. Making supply meet demand in an uncertain world. *Harvard Business Review*, **72** (1994), May–June, 83–93.

[20] Fisher, M. and A. Raman. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research*, **44** (1996), 87–99.

[21] Frazier, R. M. Quick Response in Soft Lines. *Discount Merchandiser*, January (1986), 40–46.

[22] Gallego, G. and G. van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, **40** (1994), 999–1020.

[23] Gurnani, H. and C. S. Tang. Optimal ordering decisions with uncertain cost and demand forecast updating. *Management Science*, **45** (1999), 1456-1462.

[24] Hammond, J. H. and M. Kelly. Quick response in the apparel industry. Harvard Business School Teaching Note #9-690-038. 1991.

[25] Iglehart, D. L. The dynamic inventory problem with unknown demand distribution. *Management Science*, **10** (1964), 429–440.

[26] Iyer, A. V. and M. E. Bergen. Quick Response in Manufacturer-Retailer Channels *Management Science*, **43** (1997), 559–570.

[27] Lazear, E. P. Retail pricing and clearance sales. *American Economic Review*, **76** (1986), 14–32.

[28] Murray, G. R. and E. A. Silver. A Bayesian analysis of the style goods inventory problem. *Management Science*, **12** (1966) 785–797.

[29] Nahmias, S. and S. A. Smith. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Science*, **40** (1994), 582–596.

[30] Petruzzi, N. C. and M. Dada. Dynamic pricing and inventory control with learning. *Naval Research Logistics*, **49** (2002), 303–325.

[31] Popovic, J. B. Decision making on stock levels in cases of uncertain demand rate. *European Journal of Operational Research*, **32** (1987), 276–290.

[32] Scarf, H. Bayes solutions to the statistical inventory problem. *Annals of Mathematical Statistics*, **30** (1959), 490–508.

[33] Scarf, H. Some remarks on Bayes solutions to the inventory problem. *Naval Research Logistics Quarterly*, **7** (1960), 591–596.

[34] Smith, S. A. and D. D. Achabal. Clearance pricing and inventory policies for retail chains. *Management Science*, **44** (1998), 285–300.

[35] Smith, S. A., N. Agrawal, and S. H. McIntyre. A discrete optimization model for seasonal merchandise planning. *Journal of Retailing*, **74** (1998), 193–221.

[36] Subrahmanyan, S. and R. Shoemaker. Developing optimal pricing and inventory policies of retailers who face uncertain demand. *Journal of Retailing*, **72** (1996), 7–30.

[37] Sen, A. Inventory and pricing models for perishable products. Ph.D. Dissertation. University of Southern California, Marshall School of Business. 2000.

[38] Sen, A. The U.S. apparel industry: a supply chain review. Working paper. Bilkent University, Department of Industrial Engineering, Ankara, Turkey. 2004.

[39] U.S. Census Bureau. Current Industrial Reports. Apparel. MQ315a. 2002. (http://www.census.gov/cir/www/315/mq315a.html).