The Censored Newsvendor and the Optimal Acquisition of Information

Xiaomei Ding Martin L. Puterman

Faculty of Commerce and Business Administration

University of British Columbia

2053 Main Mall

Vancouver, B.C., Canada V6T 1Z2

December 15, 1998

Abstract

This paper investigates the effect of demand censoring on the optimal policy in newsvendor inventory models with general parametric demand distributions and unknown parameter values. We show that the newsvendor problem with *observable* lost sales reduces to a sequence of single-period problems while the newsvendor problem with *unobservable* lost sales requires a dynamic analysis. Using a Bayesian Markov decision process approach we show that the optimal inventory level in the presence of partially observable demand is *higher* than when demand is completely observed. We explore the economic rationality for this observation and illustrate it with numerical examples.

Key words: Inventory, Bayesian Markov decision processes, lost sales, demand estimation, censoring

Inspite of the extensive research on inventory models, there still remain some practical issues that have not received due consideration. One of them is demand estimation and its effect on optimal policies. Most results in stochastic inventory assume a known demand distribution. However in practice, this is rarely the case and demand distribution estimation is required. The most common approach is to observe some data, estimate unknown parameters by maximum likelihood and choose inventory policies assuming that the demand distribution parmeters equal the estimated values. This practical recipe often leads to sub-optimal policies and ignores the opportunity to tradeoff *short term* optimality for more efficient information gathering and *long term* optimality. This problem is further exacerbated by *censoring* of the observed demand that occurs when there are lost sales and no backordering. This paper is motivated by a practical and ubiquitous problem facing retail inventory managers; how to optimally set inventory levels in the presence of an unknown demand distribution when demand observations are censored by the inventory on hand. We will refer to this situation as *demand censored by the inventory level* (DCI).

There is extensive research on finding and characterizing optimal policies in lost sales inventory systems assuming the demand distribution is *known*. When the demand distribution is *unknown*, most research on lost sales inventory systems has concerned demand estimation but ignored its effect on optimal policies. We believe that Conrad (1976) was the first to explicitly distinguish sales and demand in such systems. He investigated the effect of DCI on Poisson demand estimation, and proposed an unbiased maximum likelihood estimate (MLE) of the Poisson parameter. Nahmias (1994) considered a DCI model with normal demands and proposed a procedure for sequentially updating estimates of the normal parameters. Agrawal and Smith (1996) suggested that the negative binomial distribution provides a better fit than either the normal or Poisson distributions to discrete data and developed a parameter estimation method for base stock inventory systems in which sales are truncated at a constant level. In her PhD dissertation, Ding (1999) develops Bayesian and maximum likelihood estimates of the parameters of Poisson and zero-inflated Poisson demand distributions in DCI and fully observable newsvendor and (s, S) inventory systems. She notes an interesting feature of the (s, S) system is that the censoring level varies as the stock decreases during lead time.

When the demand distribution is not known with certainty, it is desirable to update the demand distribution as new data becomes available while avoiding storage of all historical data. This is naturally implemented in a Bayesian fashion. Karlin (1960), Scarf (1960) and Iglehart (1964) studied dynamic inventory policy updating when the demand density has unknown parameters and is a member of the exponential and range families. They showed that an adaptive critical value (or order-up-to) policy is optimal where the critical value depends on the history path through a sufficient statistic. Scarf (1960) and Iglehart (1964) also discussed the asymptotic behavior of the optimal adaptive policy. Azoury (1985) extended the work to more general demand distributions. Azoury (1984) also investigated the effect of dynamic Bayesian demand updating on optimal order quantities. She concluded that Bayesian demand updating, compared to non-Bayesian method, yields a more flexible optimal policy by allowing updates of the order quantities in future periods. Although the Bayesian approach provides a rigorous framework for dynamic demand updating, it is generally difficult to implement because of extensive computational demands. Lovejoy (1990) showed that a simple inventory policy based on a critical fractile can be optimal or near-optimal in some inventory models. He also gave two numerical examples to illustrate the performance of the simple myopic policies.

This paper investigates the DCI effect of unobservable lost sales on the optimal policy by using a Bayesian Markov decision process (BMDP) (van Hee, 1978) formulation which allows demand distribution updating and policy updating as information is gathered. However, a serious disadvantage of Bayesian MDP model is its probability distribution valued state variable and computational intractability. To reduce this complexity and capture the essence of the problem we focus on the newsvendor inventory system. While the model is conceptually simple, it has practical significance in the fashion goods and high technology industries where products are characterized by rapid obsolescence, large lot sizes and long lead times. Eppen and Iyer (1997) described a related inventory problem facing merchandising managers in fashion industry in which managers have the option of dumping a portion of stocks before the end of the season. They used a newsvendor heuristic to simplify computation.

From the DCI perspective, *the relevant feature of the newsvendor model is that demand in excess of the inventory level is not observed.* The main result of this paper is that the combined effect of an unknown demand distribution and unobservable lost sales results in higher optimal order quantities than in the fully observable demand case. That is, it is optimal to order set the censoring (inventory) level higher in early periods to obtain additional information about the demand distribution that can lead to better decisions in latter periods. We believe that this paper might be the only example in control theory which explicitly illustrates the trade off between *information* and *optimality*. Further, we obtain this result with no assumptions about the structure of the demand distribution other than it is continuous, non-negative and from a parametric family. The key result described above is established in Section 3. Sections 1 and 2, introduce and discuss three versions of newsvendor problem in progressively more realistic, yet more complicated settings. In Section 4, we provide economic justification for our result and give numerical examples in Section 5. Concluding comments appear in Section 6.

1 Three Newsvendor Problems

In a newsvendor (NV) model, at each decision epoch the NV selects an inventory level for a perishable good and then observes demand. If demand is less than the order the quantity, some goods are left over while if the demand exceeds the order quantity, sales are lost. The NV disposes of ending inventory, receives a salvage value and faces the identical problem with perhaps different costs and demand distributions at the next decision epoch. The objective is to a choose an inventory level that minimizes expected total cost over the planning horizon. Refer to Gallego and Moon (1993) for a thorough discussion of the newsvendor problem and its variants.

We characterize three versions of the newsvendor problems on the basis of the available demand information as follows:

i) The traditional newsvendor problem in which the demand distribution is assumed known;

ii) A newsvendor problem with an unknown demand distribution in which demand is completely observed even when demand exceeds sales;

iii) A newsvendor problem with an unknown demand distribution in demand is censored by the inventory level.

We will assume throughout this paper that the demand distribution belongs to a parametric family with unknown parameter values. We assume positive continuous demand so that our attention won't be diverted by technicalities arising from discrete demand.

1.1 The Traditional Newsvendor Problem

The costs in an NV problem are a variable ordering cost per unit, c, a salvage value of h per unit and a penalty p per unit short. It is reasonably and commonly assumed that $h < c < p$.

The demand in period n is represented an iid random variable X_n, that has probability density function (pdf) $m(x)$ with support $R_+ = [0, \infty)$. When the NV sets the inventory level at y units and customers demand x units, the cost to the NV in is:

$$r(x, y) = \begin{cases} cy - h(y - x) & \text{if } x \leq y \\ cy + p(x - y) & \text{if } x \geq y \end{cases} .$$

The corresponding expected cost is

$$R(y) = \int_{R_+} r(x, y)m(x)dx = cy - h\int_0^y (y - x)m(x)dx + p\int_y^\infty (x - y)m(x)dx. \qquad (1)$$

Obviously, $R(y)$ is continuous and convex in y.

1.2 The Bayesian Newsvendor with Observable Lost Sales

We now formulate an NV model in which the demand distribution parameters are unknown but in which the NV can observe *the total demand* at each decision epoch. In most applications this is unrealistic since only sales are observed but this model might apply when orders arrive through a call center or over the internet and accurate records of orders are retained. In a lighter vein, the newsvendor may stay around after stockout and observe how many papers she would have sold.

Assume that the demand at decison epoch n, X_n is observable and that the demand in different periods are independent and identically distributed (*iid*). In each period, X_n is generated by a probability distribution with known density $f(\cdot|\theta)$ and unknown parameter (or vector of parameters) $\theta \in \Theta$. Given a prior distribution $\pi_n(\theta)$ and a demand observation x_n, the posterior distribution $\pi_{n+1}(\theta|x_n)$ is given by

$$\pi_{n+1}(\theta|x_n) = Prob(\theta|\pi_n, X_n = x_n) = \frac{f(x_n|\theta)\pi_n(\theta)}{\int_\Theta f(x_n|\theta')\pi_n(\theta')d\theta'} \quad (2)$$

We may view the sequence of posterior probabilities as a distribution valued stochastic process which is the basis for dynamic demand distribution updating. When the demand x_n in decision epoch n is observed, the prior π_n is updated to the posterior π_{n+1}, and the posterior at n becomes the prior for $n+1$.

Correspondingly, the marginal demand distribution process $\{m_n(x), n = 1, \ldots, N\}$ satisfies

$$m_n(x) = \int_\Theta f(x|\theta)\pi_n(\theta|x_{n-1})d\theta. \quad (3)$$

The marginal is the Bayesian estimate of the demand distribution in period n. The cost structure remains the same as in the traditional newsvendor problem.

We formulate this newsvendor problem as a Bayesian MDP (BMDP) model (van Hee (1978)) as follows. Let N denote the finite number of decision epochs.

Decision Epochs : The points of time at which the NV chooses the inventory level.

States: The state summarizes the relevant information available to the NV at decision epoch $n, n = 1, 2, \ldots, N$. Let S_n denote the state space at decision epoch n; $S_n \subseteq \{f : \ f$ is a probability distribution on Θ $\}$ and $S_1 \subseteq$ {all appropriate priors}. The state space represents the set of all prior distributes when $n = 1$ and posterior distributions when $n = 2, \ldots, N$. Since we do not investigate the effect of prior specification here, we assume a fixed prior so that $s_1 = \pi_1(\theta)$. Then S_n is the set of all possible posteriors corresponding to the given prior $\pi_1(\theta)$.

Actions: Since any inventory level can be chosen in at any decision epoch, the action set $A_s = [0, \infty)$ for each $s \in S_n$.

Expected Costs: The Bayesian expected cost with prior distribution respect π_n and inventory level y_n is given by

$$\begin{aligned} R_B(\pi_n, y_n) &= E_{\pi_n}[r(X, y_n)] \\ &= cy_n - h\int_0^{y_n}(y_n - x)m_n(x)dx + p\int_{y_n}^{\infty}(x - y_n)m_n(x)dx \end{aligned}$$

where $m_n(x)$ is defined in (3), and y_n is the inventory level (action) in state π_n. We assume that $R_B(\pi_{N+1}, y_{N+1}) = 0$. As noted above, $R_B(\pi_n, y_n)$ is continuous and convex in y_n.

Transition probabilities: At decision epoch n, the transition probability (density) for the BMDP may be specified for any $y_n \in [0, \infty)$ as

$$p(\pi_{n+1}|\pi_n, y_n) = m_n(x_n) \text{ where } \pi_{n+1}(\theta|x_n) = \frac{f(x_n|\theta)\pi_n(\theta)}{\int_\Theta f(x_n|\theta')\pi_n(\theta')d\theta'}$$

since x_n occurs with probability $m_n(x_n)$. The transition probability equals 0 for posteriors that cannot be attained from $\pi_n(\theta)$. Note that since the demand is completely observable, the transition probability is *independent* of the action y_n.

Optimality Equations: Under the total expected reward criterion, the optimality equations are given by

$$u_n(\pi_n) = \min_{y_n \in R_+}\{R_B(\pi_n, y_n) + \int_{R_+} u_{n+1}(\pi_{n+1}(\cdot|x))m_n(x)dx\} \tag{4}$$

for $n = 1, \ldots, N$ with the boundary condition

$$u_{N+1}(\pi_{N+1}) = 0$$

for all values of π_{n+1}.

1.3 The Bayesian Newsvendor with Unobservable Lost Sales

In this model, the demand is censored by the inventory level and the NV observes the sales x_n *but not the demand* in period n. That is $x_n = \min(X_n, y_n)$, where y_n denotes the inventory level chosen at decision epoch n and X_n the demand between decision epochs n and $n+1$. Demand is exactly observed when sales are less than the order quantity, that is, $X_n = x_n$ when $x_n < y_n$; and the demand is *censored* at the order quantity when sales equal y_n, that is , $X_n \geq x_n$ when $x_n = y_n$. Therefore, by comparing the sales and the order quantities, we can conclude if the observed demand equals sales or is at least as great as sales. A BMDP formulation follows:

States: The state space S_n has the same structure as in Section 1.2, but the posterior probability $\pi'_{n+1}(\theta|x_n)$ depends on both the sales observation and the order quantity. It is given by

$$\pi'_{n+1}(\theta|x_n) \equiv \begin{cases} \pi_{n+1}(\theta|x_n) & \text{if } x_n < y_n \\ \\ \pi^c_{n+1}(\theta|y_n) & \text{if } x_n = y_n \end{cases} . \tag{5}$$

Figure 1 illustrates the calculation of the posterior given different sales observations in the case of a discrete distribution. As we can see, when the demand in n is fully observed, that is when $x_n < y_n$, the state at $n+1$ becomes $\pi'_{n+1}(\theta|x_n) = \pi_{n+1}(\theta|x_n)$ as defined in (2). On the other hand, if the demand is censored at the order quantity, that is when $x_n = y_n$, the state becomes $\pi'_{n+1}(\theta|x_n) = \pi^c_{n+1}(\theta|y_n)$ which is given by

$$\pi^c_{n+1}(\theta|y_n) = Prob(\theta|\pi'_n, X_n \geq y_n) \equiv \frac{\int_{y_n}^{\infty} f(x|\theta)dx\, \pi'_n(\theta)}{\int_\Theta \int_{y_n}^{\infty} f(x|\theta')\pi'_n(\theta')dx\, d\theta'} . \tag{6}$$

We use the superscript c to denote censoring. Note in case of stockouts, lost sales is taken into account by integrating $f(x|\theta)$ over the range (y_n, ∞) corresponding to the probability of a censored demand.

Actions: Same as in Section 1.2.

Costs: The Bayesian expected cost with respect to π'_n is

$$\begin{aligned} R_B(\pi'_n, y_n) &= E_{\pi'_n}[r(X, y_n)] \\ &= cy_n - h\int_0^{y_n}(y_n - x)m'_n(x)dx + p\int_{y_n}^{\infty}(x - y_n)m'_n(x)dx \end{aligned}$$

where $m'_n(x) = \int_\Theta f(x|\theta)\pi'_n(\theta|x_{n-1})d\theta$. Since π'_n is calculated differently with an exact demand observation than with a censored observation as indicated in (5), the marginal demand distribution m'_n can be distinguished into the corresponding two cases as well, i.e.

$$m'_n(x) \equiv \begin{cases} m_n(x) & \text{if } \pi'_n(\theta|x_{n-1}) = \pi_n(\theta|x_{n-1}) \text{ i.e., when } x_{n-1} < y_{n-1} \\ \\ m^c_n(x) & \text{if } \pi'_n(\theta|x_{n-1}) = \pi^c_n(\theta|y_{n-1}) \text{ i.e., when } x_{n-1} = y_{n-1} \end{cases} . \tag{7}$$

Transition probabilities:

$$p(\pi'_{n+1}|\pi'_n, y_n) = \begin{cases} m'_n(x_n) & \text{if } \pi'_{n+1}(\theta|x_n) = \pi_{n+1}(\theta|x_n) \text{ i.e., when } x_n < y_n \\ \\ 1 - M'_n(y_n) & \text{if } \pi'_{n+1}(\theta|x_n) = \pi^c_{n+1}(\theta|y_n) \text{ i.e., when } x_n = y_n \end{cases}$$

where $M'_n(y_n) = Prob(X_n \leq y_n)$. The transition probability equals 0 otherwise. Note that in contrast to the observed demand case, transition probabilities *depend on actions.*

Optimality Equations: The optimality equations in the unobservable case become

$$\begin{aligned} u_n(\pi'_n) &= \min_{y_n \in R_+}\{R_B(\pi'_n, y_n) + \int_0^{y_n} u_{n+1}(\pi_{n+1}(\cdot|x))m'_n(x)dx \\ &\quad + u_{n+1}(\pi^c_{n+1}(\cdot|y_n))[1 - M'_n(y_n)]\} \end{aligned} \tag{8}$$

for $n = 1, \ldots, N$, with the same boundary condition as in Section 1.2.

2 Optimal Policies for the Three Newsvendor Problems

2.1 Traditional Newsvendor Problem

This problem decomposes into a series of identical one period problems and y^N, the optimal inventory level at any decision epoch satisfies

$$(c-h)M(y^N) = (p-c)[1-M(y^N)]. \tag{9}$$

Equation (9) can be derived by calculus or marginal analysis (Silver, Pyke and Peterson (1998)). Solving it yields the well known newsvendor policy:

$$y^N = M^{-1}(k) \tag{10}$$

where $k = \frac{p-c}{p-h}$, $0 < k < 1$, is referred to as the *critical fractile*; and $M(x) = Prob(X \leq x)$.

2.2 The Bayesian Newsvendor with Observable Lost Sales

We obtain the optimal inventory level y_n^{BN} by solving the optimality equations (4). Since u_n depends on history only through the current state π_n, and π_n is updated by the *observed* demands independent of past actions, the BMDP reduces to a sequence of single period NV problems, in which the demand distribution is the updated marginal demand distribution. The optimality equations become

$$u_n(\pi_n) = \min_{y_n \in R_+} R_B(\pi_n, y_n) \quad \text{for } n = 1, \ldots, N \tag{11}$$

which are linked only through the Bayesian information structure (state dynamics). For this version of the newsvendor problem, the optimal order quantity varies between decision epochs because the demand distribution is updated after observing demand however the ordering policy has no effect on the state transitions. We refer to the optimal policy y_n^{BN} for $n = 1, \ldots, N$ as the Bayesian Newsvendor (BN) policy. Since the optimal policy for both this and the traditional *Single Period Equivalent Newsvendor* (SPEN) problems.

We obtain the optimal decision rule for period n by solving (9) with M replaced by M_n. Therefore, the optimal policy is obtained by the following procedure:

i) After observing x_{n-1}, use (2) to update $\pi_{n-1}(\theta|x_{n-2})$ to $\pi_n(\theta|x_{n-1})$.

ii) Calculate $M_n(y)$ by integrating (3) over x.

iii) Compute $y_n^{BN} = M_n^{-1}(k)$.

iv) Increment n by 1 and return to i.

When the prior is degenerate at θ^0, corresponding to knowing the value of θ with certainty, this model reduces to the traditional newsvendor model. Lovejoy (1993) gives an example of a prior distribution for which it is optimal to set the inventory level to zero.

2.3 The Bayesian Newsvendor with Unobservable Lost Sales

Unobservable lost sales complicates the problem greatly. Because the transition probabilities depend on the actions, the demand distribution updating now depends on the entire history. As a result, the newsvendor problem with unobservable demand can no longer be reduced to a sequence of single-period problems and the SPEN policy need not be optimal. To find optimal policies we solve (8) by backward induction. Unfortunately, analytic solutions do not appear to be readily available. Instead we attempt to establish a structural result and carry out some numerical calculations. An alternative approach we do not consider here is to assume a parametric form the demand distribution and seek accompanying closed form characterizations of the optimal policy. In the remainder of this paper, we explore the relation between optimal in the fully observable newsvendor and DCI version. Althought the Bayesian newsvendor policy need not be optimal, we will show that it provides a lower bound on the optimal policy.

3 The Key Result

In this section we establish for a two period NV problem with unobservable lost sales and continuous demand that the optimal policy (y_1^*, y_2^*) satisfies $y_1^* \geq y_1^{BN}$ and $y_2^* = y_2^{BN}$. The reason for this is as follows. At decision epoch 1, the newsvendor might want to sacrifice cost minimization so that she can obtain extra information about the demand distribution in order to make a better decision at epoch 2. The proof of this result when $n = 2$ is very intricate and it is not apparent how to extend this result to $n > 2$ at this time. We conjecture that for finite-horizon problem $N > 2$ that $y_i^* \geq y_i^{BN}$ for $i = 1, 2, \ldots, N-1$; and $y_N^* = y_N^{BN}$.

Now onto the technical analysis. Assume $N = 2$. Since the estimate of the demand distribution is updated after the demand in period 1 is observed, the demand process is $\{X_1, X_2\}$ with X_1 having probability density function (pdf) $m_1(x)$, X_2 having pdf $m_2'(x)$ where $m_2'(x)$ is the marginal density corresponding to the posterior $\pi_2'(\theta|x_1)$ given by (5). To avoid technicalities, we consider non-negative continuous demand only.

Some notation follows:

$m_2(x|x_1)$: the marginal of X_2 corresponding to state $\pi_2(\theta|x_1)$ which is the posterior when the demand in period 1 is observed to be x_1 (i.e. $X_1 = x_1$).

$m_2^c(x|y_1)$: the marginal of X_2 corresponding to state $\pi_2^c(\theta|y_1)$ which is the posterior when the demand in period 1 is censored at the order quantity y_1 (i.e. $X_1 \geq y_1$).

$M_2(x|x_1) = Prob(X_2 \leq x|X_1 = x_1)$, $\quad M_2^c(x|y_1) = Prob(X_2 \leq x|X_1 \geq y_1)$.

y_2^e : optimal policy for period 2 with state $\pi_2(\theta|x_1)$.

y_2^c : optimal policy for period 2 with state $\pi_2^c(\theta|y_1)$.

Note, as mentioned earlier, $m_2(x|x_1)$ and $m_2^c(x|y_1)$ depend on x_1, y_1 through $\pi_2(\theta|x_1)$ and $\pi_2^c(\theta|y_1)$ respectively (refer to (7)). We will use the following technical lemma based on the implicit function theorem which combines results appearing in Bartle (1964, p.260-262).

Lemma 1. *Let $g(x, y) = 0$ be a function from $R \times R$ to R which satisfies:*

i) $g(a, b) = 0$;

ii) the partial derivatives g_x and g_y exist and are continuous in a neighborhood of (a, b);

iii) the partial derivative g_y is not zero at (a, b).

Then there exists function $\phi(x)$ satisfying $g(x, \phi(x)) = 0$ which is continuously differentiable in a neighborhood of (a, b) with

$$\phi'(x) = -\frac{g_x}{g_y}$$

where the partial derivatives are evaluated at $(x, \phi(x))$.

We provide definitions of stochastic ordering and likelihood ratio ordering from Ross (1983).

[**Stochastic Ordering**] *We say that the random variables X is stochastic larger than the random variable Y, written $X \geq_{st} Y$ if*

$$P(X > a) \geq P(Y > a) \qquad \text{for all } a.$$

[**Likelihood Ratio Ordering**] *Let X and Y denote continuous non-negative random variables having respectively densities f and g. We say that X is larger than Y in the likelihood ratio* *sense, and write $X \geq_{LR} Y$ if*

$$\frac{f(x)}{g(x)} \leq \frac{f(y)}{g(y)} \qquad \textit{for all} \quad x \leq y.$$

That is $X \geq_{LR} Y$ if the ratio of their respective densities, $f(x)/g(x)$, is nondecreasing in x.

The next proposition gives the stochastic ordering property of demand distributions that are updated using censored observations and exact observations respectively. It states that a censored demand observation results in a marginal demand distribution which is stochastically greater than that from an exact observation at the same level.

Proposition 1. *Suppose X_1 and X_2 are non-negative iid random variables with density $f(x|\theta)$, and suppose $f(x|\theta)$ is likelihood ratio increasing in θ. Let $\pi_1(\theta)$ be given and suppose $\pi_2'(\theta|x_1)$ is given by (5). Then for any x_1, $X_2|[X_1 \geq x_1]$ is stochastically greater than*

$X_2|[X_1 = x_1]$ or equivalently, $M_2^c(x|x_1) \leq M_2(x|x_1)$.

Proof. Our proof consists of three steps. First, we show that the posterior $\pi_2(\theta|x_1)$ is stochastically increasing in x_1. Second, we show $\pi_2^c(\theta|x_1) >_{st} \pi_2(\theta|x_1)$. Then we show that this implies the stochastic ordering of the marginals.

[*Step 1*] At any x_1 and x_1' such that $x_1 < x_1'$, we have

$$\pi_2(\theta|x_1) = \frac{f(x_1|\theta)\pi_1(\theta)}{m_1(x_1)}, \quad \pi_2(\theta|x_1') = \frac{f(x_1'|\theta)\pi_1(\theta)}{m_1(x_1')}.$$

We will show in the following that $\int_t^\infty \pi_2(\theta|x_1')d\theta \geq \int_t^\infty \pi_2(\theta|x_1)d\theta$ for any t.

Since

$$\int_t^\infty \pi_2(\theta|x_1')d\theta - \int_t^\infty \pi_2(\theta|x_1)d\theta = \int_t^\infty \frac{f(x_1'|\theta)\pi_1(\theta)}{m_1(x_1')}d\theta - \int_t^\infty \frac{f(x_1|\theta)\pi_1(\theta)}{m_1(x_1)}d\theta \quad (12)$$

and the common denominator of the right hand side of (12) is obviously positive, we only need to show that the numerator is greater than or equal to 0. The numerator is,

$$\int_t^\infty f(x_1'|\theta)\pi_1(\theta)d\theta m_1(x_1) - \int_t^\infty f(x_1|\theta)\pi_1(\theta)d\theta m_1(x_1') =$$
$$\int_t^\infty f(x_1'|\theta)\pi_1(\theta)d\theta \int_0^\infty f(x_1|\theta')\pi_1(\theta')d\theta' - \int_t^\infty f(x_1|\theta)\pi_1(\theta)d\theta \int_0^\infty f(x_1'|\theta')\pi_1(\theta')d\theta'.$$

The right hand side of the above equation can be reduced to,

$$\int_t^\infty f(x_1'|\theta)\pi_1(\theta)d\theta \int_0^t f(x_1|\theta')\pi_1(\theta')d\theta' - \int_t^\infty f(x_1|\theta)\pi_1(\theta)d\theta \int_0^t f(x_1'|\theta')\pi_1(\theta')$$
$$= \int_{\theta=t}^\infty \int_{\theta'=0}^t [f(x_1'|\theta)f(x_1|\theta') - f(x_1|\theta)f(x_1'|\theta')]\pi_1(\theta)\pi_1(\theta')d\theta' d\theta.$$

Since $f(x|\theta)$ is likelihood ratio increasing in θ ,

$$f(x_1'|\theta)f(x_1|\theta') \geq f(x_1|\theta)f(x_1'|\theta')$$

for any $x_1 < x_1'$. Therefore, $\pi_2(\theta|x_1)$ increases stochastically in x_1.

[*Step 2*]

$$\pi_2^c(\theta|x_1) = \frac{\int_{x_1}^\infty f(x|\theta)dx\pi_1(\theta)}{\int_\Theta \int_{x_1}^\infty f(x|\theta)dx\pi_1(\theta)d\theta} = \frac{\int_{x_1}^\infty m_1(x)\frac{f(x|\theta)\pi_1(\theta)}{m_1(x)}dx}{\int_{x_1}^\infty m_1(x)dx}.$$

That is,

$$\pi_2^c(\theta|x_1) = \frac{\int_{x_1}^\infty m_1(x)\pi_2(\theta|x)dx}{\int_{x_1}^\infty m_1(x)dx}. \tag{13}$$

Note the right hand side of (13) can be considered as a weighted average of $\pi_2(\theta|x)$ for $x > x_1$. We showed in step 1, $\pi_2(\theta|x)$ is stochastically increasing in x. Therefore, $\pi_2^c(\theta|x_1) >_{st} \pi_2(\theta|x_1)$ for arbitrary x_1.

[*Step 3*] We give the marginals in the following

$$M_2(x|x_1) = \int_0^x \int_\Theta f(t|\theta)\pi_2(\theta|x_1)d\theta dt \ . \tag{14}$$

By changing the order of integral, (14) becomes,

$$M_2(x|x_1) = \int_\Theta F(x|\theta)\pi_2(\theta|x_1)d\theta = E_\theta[F(x|\theta)] \tag{15}$$

where θ has density $\pi_2(\theta|x_1)$.

Similarly,

$$M_2^c(x|x_1) = \int_\Theta F(x|\theta)\pi_2^c(\theta|x_1)d\theta = E_{\theta'}[F(x|\theta')] \tag{16}$$

where θ' has density $\pi_2^c(\theta|x_1)$.

The condition that $f(x|\theta)$ is increasing in likelihood ratio as θ increases implies that $f(x|\theta)$ is also stochastically increasing in θ since the likelihood ratio ordering is stronger than stochastic ordering. Therefore, $1 - F(x|\theta)$ is non-decreasing in θ. By Proposition 8.1.2 in Ross, (1983, p. 252), we have

$$E_\theta[1 - F(x|\theta)] \leq E_{\theta'}[1 - F(x|\theta')] \tag{17}$$

if θ' is stochastically greater than θ. It follows from (17) that

$$M_2(x|x_1) \geq M_2^c(x|x_1).$$

Proposition 2 establishes a further result about the optimal policy based on the stochastic ordering property in Proposition 1. It establishes that censoring in the first period results in a greater optimal order quantity for the second period.

Proposition 2. *Let y_2^e be the optimal order quantity for period 2 if in period 1, demand is observed to be x_1, i.e. $X_1 = x_1$. Similarly, let y_2^c be the optimal order quantity if demand in period 1 is censored at x_1, i.e. $X_1 \geq x_1$. Under the hypothesis of Proposition 1, $y_2^e \leq y_2^c$.*

Proof. Since $u_3(\pi_3') = R_B(\pi_3', y_3) = 0$ by the assumption of boundary condition,

$$u_2(\pi_2') = \min_{y_2}\{R_B(\pi_2', y_2) + 0\}. \tag{18}$$

As we have pointed out earlier, $R_B(\pi_2', y_2)$ is convex in y_2. So $u_2(\pi_2')$ is minimized at

$$y_2^* = y_2^{BN} = [M_2'(k|x_1)]^{-1}$$

where $M_2'(x|x_1) = \begin{cases} M_2(x|x_1) & \text{corresponding to } \pi_2'(\cdot|x_1) = \pi_2(\cdot|x_1),\ i.e. \text{ when } X_1 = x_1 \\ M_2^c(x|x_1) & \text{corresponding to } \pi_2'(\cdot|x_1) = \pi_2^c(\cdot|x_1),\ i.e. \text{ when } X_1 \geq x_1 \end{cases}.$

That is,

$$y_2^* = \begin{cases} y_2^e = [M_2(k|x_1)]^{-1} & \text{if } X_1 = x_1 \\ \\ y_2^c = [M_2^c(k|x_1))]^{-1} & \text{if } X_1 \geq x_1. \end{cases} \tag{19}$$

Since both $M_2(x|x_1)$ and $M_2^c(x|x_1)$ are monotonically increasing in x, using the result of the proposition, it is obvious that

$$y_2^c = [M_2^c(k|x_1)]^{-1} \geq y_2^e = [M_2(k|x_1)]^{-1}. \tag{20}$$

The key result is given in the following.

Theorem A. *Suppose the hypotheses of Proposition 1 hold and in addition $u_2(\pi_2'(\cdot|x_1))$ is differentiable in x_1. Then for the Bayesian newsvendor problem with unobservable lost sales*

$$y_1^* \geq y_1^{BN} \qquad \textit{and} \quad y_2^* = y_2^{BN}.$$

Proof. We will start from period 2 and move backward to period 1 in computing the total expected cost.

We have (18) from the boundary condition. Recall $\pi_2'(\cdot|x_1)$ is calculated differently depending on whether X_1 is fully observed or censored (refer to (5)). Hence,

$$u_2(\pi_2'(\cdot|x_1)) \equiv \begin{cases} u_2(\pi_2(\cdot|x_1)) & \text{if } \pi_2'(\cdot|x_1) = \pi_2(\cdot|x_1),\ i.e. \text{ when } X_1 = x_1 < y_1 \\ u_2(\pi_2^c(\cdot|y_1)) & \text{if } \pi_2'(\cdot|x_1) = \pi_2^c(\cdot|y_1),\ i.e. \text{ when } X_1 \geq x_1 = y_1 \end{cases}. \quad (21)$$

Since $\pi_2(\cdot|x_1)$ can be considered as a function of π_1 and x_1 (refer to (2)), and $\pi_2^c(\cdot|y_1)$ as a function of π_1 and y_1 (refer to (6)), we redefine $u_2(\pi_2'(\cdot|x_1))$ as

$$u_2(\pi_2'(\cdot|x_1)) \equiv \begin{cases} u_2(\pi_1, x_1) & \text{if } x_1 < y_1 \\ u_2^c(\pi_1, y_1) & \text{if } x_1 = y_1 \end{cases}.$$

Omitting π_1 as it is fixed, we further simplify the above to

$$u_2(\pi_2'(\cdot|x_1)) \equiv \begin{cases} u_2(x_1) & \text{if } x_1 < y_1 \\ u_2^c(y_1) & \text{if } x_1 = y_1 \end{cases}.$$

We omit π_1 through out the proof.

Obviously, the optimality equation at $n = 1$ is

$$u_1(\pi_1) = \min_{y_1}\{R_B(y_1) + \int_0^{y_1} u_2(x_1)m_1(x_1)dx_1 + u_2^c(y_1)[1 - M_1(y_1)]\} \quad (22)$$

in which the second term at the right hand side is the expected cost at $n = 2$ if X_1 is fully observed (i.e. $X_1 < y_1$) and the third term is the expected cost if X_1 is censored at y_1 (i.e. $X_1 \geq y_1$).

We omit the subscript of x_1 for convenience. Let

$$\begin{aligned} J(y_1) &= R_B(y_1) + \int_0^{y_1} u_2(x)m_1(x)dx + u_2^c(y_1)[1 - M_1(y_1)] \\ &\equiv R_B(y_1) + I(y_1) \quad (23) \end{aligned}$$

then, the optimal policy is obtained from

$$y_1^* \in arg \min_{y_1} J(y_1).$$

To show $y_1^* \geq y_1^{BN} = M_1^{-1}(k)$, it suffices to show that $\frac{dJ(y_1)}{dy_1}|_{y_1=y_1^{BN}} \leq 0$ where

$$\begin{aligned} \frac{dJ(y_1)}{dy_1} &= \frac{dR_B(y_1)}{dy_1} + \frac{dI(y_1)}{dy_1} \qquad (24) \\ &= (c-p) + (p-h)M_1(y_1) + \frac{du_2^c(y_1)}{dy_1}[1 - M_1(y_1)] - (u_2^c(y_1) - u_2(y_1))m_1(y_1). \end{aligned}$$

Since $\frac{dR_B(y_1)}{dy_1}|_{y_1=y_1^{BN}} = 0$, we only need to show that at y_1^{BN},

$$\frac{dI(y_1)}{dy_1} = \frac{du_2^c(y_1)}{dy_1}[1 - M_1(y_1)] - (u_2^c(y_1) - u_2(y_1))m_1(y_1) \leq 0. \qquad (25)$$

As a matter of fact, we will show that this is true at any value of y_1.

We first write out $u_2(y_1)$ and $u_2^c(y_1)$ below.

$$\begin{aligned} u_2(y_1) &= R_B(\pi_2, y_2^e) = E_{\pi_2}[r(X, y_2^e)] \\ &= p\int_{y_2^e}^{\infty} xm_2(x|y_1)dx + h\int_0^{y_2^e} xm_2(x|y_1)dx \\ &= p[\int_{y_2^c}^{\infty} xm_2(x|y_1)dx + \int_{y_2^e}^{y_2^c} xm_2(x|y_1)dx] \\ &\quad + h[\int_0^{y_2^c} xm_2(x|y_1)dx - \int_{y_2^e}^{y_2^c} xm_2(x|y_1)dx] \qquad (26) \end{aligned}$$

Similarly,

$$\begin{aligned} u_2^c(y_1) &= R_B(\pi_2^c, y_2^c) = E_{\pi_2^c}[r(X, y_2^c)] \\ &= p\int_{y_2^c}^{\infty} xm_2^c(x|y_1)dx + h\int_0^{y_2^c} xm_2^c(x|y_1)dx. \qquad (27) \end{aligned}$$

Subtracting (26) from (27) gives

$$u_2^c(y_1) - u_2(y_1) = p\int_{y_2^c}^{\infty} x[m_2^c(x|y_1) - m_2(x|y_1)]dx + h\int_0^{y_2^c} x[m_2^c(x|y_1) - m_2(x|y_1)]dx$$

$$
\begin{aligned}
& - (p-h) \int_{y_2^e}^{y_2^c} x m_2(x|y_1) dx \\
= \quad & D_1 - (p-h) \int_{y_2^e}^{y_2^c} x m_2(x|y_1) dx, \qquad (28)
\end{aligned}
$$

where

$$
D_1 = p \int_{y_2^c}^{\infty} x[m_2^c(x|y_1) - m_2(x|y_1)]dx + h \int_0^{y_2^c} x[m_2^c(x|y_1) - m_2(x|y_1)]dx
$$

with $m_2(x|y_1)$, $m_2^c(x|y_1)$ given below

$$
m_2(x|y_1) = \int_\Theta f(x|\theta)\pi_2(\theta|x_1)d\theta = \frac{\int_\Theta f(x|\theta)f(y_1|\theta)\pi_1(\theta)d\theta}{m_1(y_1)}; \qquad (29)
$$

$$
m_2^c(x|y_1) = \frac{m_1(x) - \int_\Theta f(x|\theta)F(y_1|\theta)\pi_1(\theta)d\theta}{1 - M_1(y_1)}. \qquad (30)
$$

We now calculate $\frac{du_2^c(y_1)}{dy_1}$. From (27), we have the following by Leibniz's rule.

$$
\begin{aligned}
\frac{du_2^c(y_1)}{dy_1} &= p \int_{y_2^c}^{\infty} x \frac{dm_2^c(x|y_1)}{dy_1} dx - p y_2^c m_2^c(y_2^c|y_1) \frac{dy_2^c}{dy_1} \\
&\quad + h \int_0^{y_2^c} x \frac{dm_2^c(x|y_1)}{dy_1} dx + h y_2^c m_2^c(y_2^c|y_1) \frac{dy_2^c}{dy_1} \\
&= D_2 - (p-h) y_2^c m_2^c(y_2^c|y_1) \frac{dy_2^c}{dy_1} \qquad (31)
\end{aligned}
$$

where

$$
D_2 = p \int_{y_2^c}^{\infty} x \frac{dm_2^c(x|y_1)}{dy_1} dx + h \int_0^{y_2^c} x \frac{dm_2^c(x|y_1)}{dy_1} dx
$$

with

$$
\frac{dm_2^c(x|y_1)}{dy_1} = \frac{m_1(y_1)[m_1(x) - \int_\Theta f(x|\theta)F(y_1|\theta)\pi_1(\theta)d\theta]}{[1 - M_1(y_1)]^2} - \frac{\int_\Theta f(x|\theta)f(y_1|\theta)\pi_1(\theta)d\theta}{1 - M_1(y_1)}. \qquad (32)
$$

Now, replace (28) and (31) in (25), we can rewrite $\frac{dI(y_1)}{dy_1}$ by,

$$
\begin{aligned}
\frac{dI(y_1)}{dy_1} &= \{[1 - M_1(y_1)]D_2 - m_1(y_1)D_1\} + \{m_1(y_1)(p-h) \int_{y_2^e}^{y_2^c} x m_2(x|y_1) dx \\
&\quad - [1 - M_1(y_1)](p-h) y_2^c m_2^c(y_2^c|y_1) \frac{dy_2^c}{dy_1}\}. \qquad (33)
\end{aligned}
$$

From (32),

$$[1 - M_1(y_1)]\frac{dm_2^c(x|y_1)}{dy_1} = \frac{m_1(y_1)m_1(x)}{1 - M_1(y_1)} - \frac{\int_\Theta f(x|\theta)F(y_1|\theta)\pi_1(\theta)d\theta}{1 - M_1(y_1)}m_1(y_1)$$
$$- \int_\Theta f(x|\theta)f(y_1|\theta)\pi_1(\theta)d\theta. \quad (34)$$

And from (30),

$$m_1(y_1)m_2^c(x|y_1) = \frac{m_1(y_1)m_1(x)}{1 - M_1(y_1)} - \frac{\int_\Theta f(x|\theta)F(y_1|\theta)\pi_1(\theta)d\theta}{1 - M_1(y_1)}m_1(y_1). \quad (35)$$

Using (34) and (35), we have

$$[1 - M_1(y_1)]\frac{dm_2^c(x|y_1)}{dy_1} - m_1(y_1)[m_2^c(x|y_1) - m_2(x|y_1)] =$$
$$- \int_\Theta f(x|\theta)f(y_1|\theta)\pi_1(\theta)d\theta + m_1(y_1)m_2(x|y_1) = 0.$$

Since the integrand is zero,

$$[1 - M_1(y_1)]D_2 - m_1(y_1)D_1 = 0.$$

Therefore (33) reduces to

$$\frac{dI(y_1)}{dy_1} = m_1(y_1)(p-h)\int_{y_2^e}^{y_2^c} x m_2(x|y_1)dx$$
$$- [1 - M_1(y_1)](p-h)y_2^c m_2^c(y_2^c|y_1)\frac{dy_2^c}{dy_1}. \quad (36)$$

Note $M_2^c(x_2|x_1)$ is a function of both x_2 and x_1. It is a function of x_1 through π_2^c. Since $M_2^c(y_2^c|y_1) = k$ implicitly defines the function $y_2^c(y_1)$, it follows from Lemma 1 that

$$\frac{dy_2^c}{dy_1} = - \frac{\frac{\partial M_2^c(y_2^c|y_1)}{\partial y_1}}{m_2^c(y_2^c|y_1)}.$$

This gives

$$[1 - M_1(y_1)]y_2^c m_2^c(y_2^c|y_1)\frac{dy_2^c}{dy_1} = -[1 - M_1(y_1)]y_2^c\frac{\partial M_2^c(y_2^c|y_1)}{\partial y_1}. \quad (37)$$

Since

$$M_2^c(y_2^c|y_1) = \int_0^{y_2^c} m_2^c(x|y_1)dx = \frac{M_1(y_2^c) - \int_\Theta F(y_1|\theta)F(y_2^c|\theta)\pi_1(\theta)d\theta}{1 - M_1(y_1)},$$

(37) becomes

$$-[1 - M_1(y_1)]y_2^c \frac{\partial M_2^c(y_2^c|y_1)}{\partial y_1} = y_2^c m_1(y_1)[M_2(y_2^c|y_1) - M_2^c(y_2^c|y_1)]. \tag{38}$$

By Lemma 2, we have $y_2^e \leq y_2^c$ so that

$$\begin{aligned} m_1(y_1)\int_{y_2^e}^{y_2^c} x m_2(x|y_1)dx &\leq m_1(y_1)y_2^c \int_{y_2^e}^{y_2^c} m_2(x|y_1)dx \\ &= m_1(y_1)y_2^c[M_2(y_2^c|y_1) - M_2(y_2^e|y_1)]. \end{aligned} \tag{39}$$

From (19), we know $M_2^c(y_2^c|y_1) = M_2(y_2^e|y_1) = k$. Replacing (38) and (39) in (36), we have the following inequality

$$\frac{dI(y_1)}{dy_1} \leq (p-h)y_2^c[m_1(y_1)(M_2(y_2^c|y_1) - k) - m_1(y_1)(M_2(y_2^c|y_1) - k)] = 0$$

which gives $\frac{dJ(y_1)}{dy_1}|_{y_1^{BN}=M_1^{-1}(k)} \leq 0$. That is, $u_1(\pi_1)$ is non-increasing at y_1^{BN} completing the proof.

4 The Effect of Lost Sales on the Optimal Policy

We've shown that when lost sales are unobservable, the optimal policy at the first decision epoch is as least as great as the BN policy, while in the second period the BN policy is optimal. In this section, we provide some economic justification for this result.

4.1 The Optimal Policy for a Two Period Bayesian Newsvendor with Unobservable Lost Sales

Because of the terminal condition, the BN policy is optimal at the last decision epoch. We now develop a rationale for the relationship $y^* \geq y^{BN}$ at the first decision epoch.

Theorem B. *Under the hypotheses in Theorem A, the optimal order quantity y_1^* in the first period of a two period newsvendor problem with unobservable lost sales satisfies*

$$(c-h)M_1(y_1) =$$
$$[1-M_1(y_1)][p - \frac{du_2^c(y_1)}{dy_1} + (u_2^c(y_1) - u_2(y_1))\frac{m_1(y_1)}{1-M_1(y_1)} - c]. \quad (40)$$

Proof. Because of the convexity of the cost function $J(y_1)$ (defined in (23)), we can obtain the optimal order quantity by solving for y_1 from setting $\frac{dJ(y_1)}{dy_1} = 0$ (refer to (24)). Setting $\frac{dJ(y_1)}{dy_1} = 0$ and rearrangeing terms gives (40).

We now use (40) to interpret the result in Theorem A. For the newsvendor problem with observable lost sales, the optimal policy at decision epoch 2 satisfies

$$(c-h)M_2(y_2) = [1-M_2(y_2)](p-c) \quad (41)$$

where M_2 is the marginal demand CDF. For the NV model with unobservable lost sales, let

$$p'(y_1) = p - \frac{du_2^c(y_1)}{dy_1} + (u_2^c(y_1) - u_2(y_1))\frac{m_1(y_1)}{1-M_1(y_1)} . \quad (42)$$

Then (40) may be expressed as

$$(c-h)M_1(y_1) = [1-M_1(y_1)][p'(y_1) - c]. \quad (43)$$

This shows that we can view $p'(y_1)$ as a *policy dependent variable penalty cost.* In Section 3 we established that

$$\frac{du_2^c(y_1)}{dy_1}[1 - M_1(y_1)] - (u_2^c(y_1) - u_2(y_1))m_1(y_1) \leq 0.$$

Dividing this expression by $1 - M_1(y_1)$ establishes that $p'(y_1) \geq p$ for any y_1. Since the solution of (43) is increasing in $p'(\cdot)$ it follows from the observation that $p'(\cdot) \geq p$ that the optimal inventory level at decision epoch 1 is as least as great as the Bayesian newsvendor policy y_1^{BN}.

4.2 Interpretation of the Policy Dependent Penalty

Let

$$i(y_1) = (u_2^c(y_1) - u_2(y_1))\frac{m_1(y_1)}{1 - M_1(y_1)} - \frac{du_2^c(y_1)}{dy_1} \tag{44}$$

so that, $p'(y_1) = p + i(y_1)$. The penalty $p'(y_1)$ consists of the shortage penalty p and a positive term which varies with policy y_1. To understand how $i(y_1)$ relates to lost sales and information loss, we consider the following scenario.

If $x_1 = y_1$, that is, the demand is censored, the optimal policy y_2^c for period 2 gives the value function u_2^c. Now, suppose we have the option of purchasing additional information regarding whether the demand x_1 lies above or below $y_1 + \Delta y$ where Δy is small and positive. Buying this information when $X_1 = x_1$ gives an exact observation of the total demand in period 1 if $x_1 < y_1 + \Delta y$ and a censored observation at a slightly higher level if $x_1 = y_1 + \Delta y$. Correspondingly, the value functions in these two cases are $u_2 + \Delta u(x)$ and $u_2^c + \Delta u^c$ respectively. Thus the value of this information is the difference between the expected cost without the information and the expected cost with it. Table 1 summarizes these costs and probabilities.

From Table 1, we calculate the expected cost if we acquire the additional information to be

$$\int_{y_1}^{y_1+\Delta y}[u_2+\Delta u(x)]\frac{m_1(x)}{1-M_1(y_1)}dx+(u_2^c+\Delta u^c)\frac{1-M_1(y_1)-\Delta M}{1-M_1(y_1)} \tag{45}$$

where $\Delta M=\int_{y_1}^{y_1+\Delta y}m_1(x)dx$. The expected value of the information is, therefore,

$$c(y_1,\Delta y)=u_2^c-\{\int_{y_1}^{y_1+\Delta y}[u_2+\Delta u(x)]\frac{m_1(x)}{1-M_1(y_1)}dx+(u_2^c+\Delta u^c)\frac{1-M_1(y_1)-\Delta M}{1-M_1(y_1)}\}.$$

Letting $\Delta y\downarrow 0$, we obtain the marginal expected value of information at y_1,

$$\begin{aligned}\lim_{\Delta y\to 0}\frac{c(y_1,\Delta y)}{\Delta y} &= \lim_{\Delta y\to 0}\frac{u_2^c-[(u_2^c+\Delta u^c)\frac{1-M_1(y_1)-\Delta M}{1-M_1(y_1)}+(u_2+\Delta u)\frac{\Delta M}{1-M_1(y_1)}]}{\Delta y}\\ &= \lim_{\Delta y\to 0}\{u_2^c\frac{\Delta M}{[1-M_1(y_1)]\Delta y}-\frac{\Delta u^c}{\Delta y}(1-\frac{\Delta M}{1-M_1(y_1)})\\ &\quad -(u_2+\Delta u)\frac{\Delta M}{[1-M_1(y_1)]\Delta y}\}.\end{aligned} \tag{46}$$

Since,

$$\lim_{\Delta y\to 0}\frac{\Delta u^c}{\Delta y}\;\frac{\Delta M}{1-M_1(y_1)}=0$$

$$\lim_{\Delta y\to 0}\frac{\Delta u}{\Delta y}\;\frac{\Delta M}{1-M_1(y_1)}=0$$

(46) becomes

$$\lim_{\Delta y\to 0}\frac{c(y_1,\Delta y)}{\Delta y}=(u_2^c(y_1)-u_2(y_1))\frac{m_1(y_1)}{1-M_1(y_1)}-\frac{du_2^c(y_1)}{dy_1}=i(y_1). \tag{47}$$

From this argument we conclude that the new penalty $p'(y_1)$ consists of the shortage penalty p and the *marginal expected value of information* $i(y_1)$. Hence this problem too can be viewed as an SPEN, but computing the penalty requires backwards induction. The augmented penalty cost accounts for the gain in future expected total return through the acquisition of additional information in the current period by setting the inventory level higher.

5 Examples

To gain insight into the results of the previous section as well as to learn more about the effects of censoring, we present and discuss some numerical examples. The analysis of these examples shows that the optimal policy may require *strictly* higher inventory levels than the optimal BN policy at the first decision epoch.

We have assumed continuous demand so far. For discrete demand, the optimal policy satisfies

$$y_1^* = \min_{y_1 \in Z}\{y_1 : (c-h)M_1(y_1) \geq [1 - M_1(y_1)][p'(y_1) - c]\} \tag{48}$$

where Z is the set of non-negative integers and

$$y_2^* = y_2^{BN} = \min_{y_2 \in Z}\{y_2 : (c-h)M_2(y_2) \geq [1 - M_2(y_2)](p - c)\}.$$

In this section, we give three examples of a two period newsvendor problem with unobservable lost sales. We assume that the demand has a Poisson distribution with unknown parameter λ and assume a gamma distribution prior for λ denoted Gamma(α, β). If historical data is available, the gamma parameters can be specified so that the expectation of demand is close to the mean of the demand data. Note that the gamma distribution is the conjugate prior for the Poisson. We vary the costs and/or the prior parameter values over the three examples.

Example 1. The costs and demand distribution assmptions are as follows:

Costs: $c = 1,\ h = 0.25,\ p = 1.5$; this gives $k = \frac{p-c}{p-h} = 0.4$.

Demand: $X|\lambda \sim Poisson(\lambda); \quad \pi_1(\lambda) = Gamma(0.4, 10)$.

This gives $E(X) = 4$ with respect to π_1.

We enumerate policies for $n = 1$, and calculate the optimal BN policy for $n = 2$. Results in Table 2 show that the BN policy $y_1^{BN} = 1$ is optimal at $n = 1$ in the two period problem and gives the minimum total expected cost 11.6763. Note that the optimal inventory level in period two depends on the observed sales in period 1. If the sales equal 0, the optimal inventory level in period 2 is 0 and if the sales in period 1 equals 1 unit, the optimal inventory level in period 2 is 3 units.

Example 2. In this example, we change the costs while leaving the prior distribution assumptions unchanged.

Costs: $c = 1,\ h = 0.5,\ p = 2$; and $k = \frac{p-c}{p-h} = 0.67$.

Demand: $X|\lambda \sim Poisson(\lambda); \quad \pi_1(\lambda) = Gamma(0.4, 10)$.

The policies and their costs are given in Table 3. In contrast to the first example, the optimal order quantity in period 1, $y_1^* = 5$, exceeds the BN policy $y_1^{BN} = 3$. In this case it is optimal to set the censoring level two units higher to acquire additional information regarding the demand distribution. The total expect cost for the optimal policy is 13.2126 while that from using the BN policy in both periods is 13.3709.

It is also informative to observe some features of the marginal demand distributions and optimal policies. If we were to use the BN policy at decision epoch 1, we would set the inventory level at 3 units in which case there is a .3887 probability of a censored observation. If the sales is observed at this level than the optimal policy would set the inventory level at 10 units in the next period. This contrasts to the situation when sales equal 2 units and the next inventory level is set to 3 units. Plots in Ding (1999) confirm this dramatic shift to the right in the marginal demand distribution resulting from sales of 3 units as opposed to 2 units. Note also that when the optimal policy of 5 units is used in period 1, the censoring probability drops to .2744. In this case, sales of 3 units in period 1 leads to an optimal order quantity of 4 units in the second period.

Example 3. In this example, we use the same costs as in Example 2 but change the prior distribution parameters.

Costs: $c = 1,\ h = 0.5,\ p = 2$; and $k = \frac{p-c}{p-h} = 0.67$.

Demand: $X|\lambda \sim Poisson(\lambda); \quad \pi_1(\lambda) = Gamma(1.2, 8)$.

This gives $E(X) = 9.6$ with respect to π_1. In Table 4, we only give the enumerated policies and corresponding expected total costs for the two periods. We omit y_2^* to avoid the long list of all possible historical paths at each y_1. As in Example 2, the optimal inventory level in period 1 exceeds that specified by the BN policy, namely $y_1^* = 12$ and $y_1^{BN} = 11$.

These examples show that under the parameter choices in Examples 2 and 3 the optimal policy will call for a *strictly greater* inventory level in the first period than would be selected by the BN policy.

6 Conclusion

This paper has established that when the demand is unknown, that there is a tradeoff between short term optimality and information acquisition. While we have concentrated on a two period newsvendor problem we are confident that similar observations may be made in more complex systems.

We view this paper as a first and general step to the analysis of DCI inventory models. Several important questions remain which we describe below:

i) Extension of the monotonicity result to $N > 2$ in general and for specific distributions. (Some numerical examples for a 3-period newsvendor problem support with our conjecture about the N-period problem);

ii) Efficient numerical calculations of the optimal order quantity for continuous demand;

iii) Application of the BMDP to base stock and (s,S) inventory systems with lost sales and unknown demand;

iv) Establishment of the finiteness of the optimal order quantity y_1^* .

ACKNOWLEDGEMENTS

We thank Professor Harry Joe, Department of Statistics, University of British Columbia for helpful comments on this paper. This research is supported by NSERC grant OGP0005527.

References

[1] Agrawal, N. and Smith, S. A., "Estimating Negative Binomial Demand for Retail Inventory Management with Unobservable Lost Sales," *Naval Research Logistics,* Vol. 43, 1996.

[2] Azoury, K. S., "Comparison of Optimal Ordering Levels of Bayesian and Non-Bayesian Inventory Models," *Management Science,* Vol. 30, No. 8, 1984.

[3] Azoury, K. S., "Bayes Solution to Dynamic Inventory Models Under Unknown Demand Distribution," *Management Science,* Vol. 31, No. 9, 1985.

[4] Bartle, R. G., *The Elements of Real Analysis,* John Wiley & Sons, 1964.

[5] Conrad, S. A., "Sales Data and the Estimation of Demand," *Operational Research Quarterly,* Vol. 27, No. 1, 1976.

[6] Ding, X. , "Estimation and Optimization in Discrete Inventory Models," Unpublished PhD Dissertation, The Faculty of Commerce, University of British Columbia, Vancouver, Canada, 1999.

[7] Eppen, G. D. and Iyer, A. V., "Improved Fashion Buying with Bayesian Updates," *Operations Research,* Vol. 45, No. 6, 1997.

[8] Gallego, G. and Moon, I. "The Distribution Free Newsvendor Problem: Review and Extensions," *Journal of Operational Research Society,* Vol. 44, No. 8, 1993.

[9] Iglehart, D. L., "The Dynamic Inventory problem with Unknown Demand Distribution," *Management Science,* Vol. 10, No. 3, 1964.

[10] Karlin, S., "Dynamic Inventory Policy with Varying Stochastic Demands," *Management Science,* Vol. 6, No. 3, 1960.

[11] Lovejoy, W. S., "Myopic Policies for Some Inventory Models with Uncertain Demand Distribution," *Management Science,* Vol. 36, No. 6, 1990.

[12] Lovejoy, W. S., "Suboptimal Policies, with Bounds, for Parameter Adaptive Decision Processes," *Operations Research,* Vol. 41, No. 3, 1993.

[13] Nahmias, S., "Demand Estimation in Lost Sales Inventory Systems," *Naval Research Logistics,* Vol. 41, 1994.

[14] Ross, S., *Stochastic Processes,* John Wiley & Sons, 1983.

[15] Scarf, H., "Some Remarks on Bayes Solutions to the Inventory Problem," *Naval Research Logistics Quarterly,* Vol. 7, 1960.

[16] Silver E. A., Pyke, D. F. and Peterson, R. *Inventory Management and Production Planning and scheduling, Third Edition, ch. 10,* John Wiley & Sons, 1998.

[17] van Hee, K. M., *Bayesian Control of Markov Chains,* Mathematisch Centrum, 1978.