MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 10, No. 2, Spring 2008, pp. 236–256
ISSN 1523-4614 | EISSN 1526-5498 | 08 | 1002 | 0236

informs®
DOI 10.1287/msom.1070.0165

Dynamic Inventory Management with Learning About the Demand Distribution and Substitution Probability

Li Chen
TrueDemand, Inc., Los Gatos, California 95032, li@tdemand.com

Erica L. Plambeck
Graduate School of Business, Stanford University, Stanford, California 94305, elp@stanford.edu

A well-known result in the Bayesian inventory management literature is: If lost sales are not observed, the Bayesian optimal inventory level is larger than the myopic inventory level (one should "stock more" to learn about the demand distribution). This result has been proven in other studies under the assumption that inventory is perishable, so the myopic inventory level is equal to the Bayesian optimal inventory level with observed lost sales. We break that equivalence by considering nonperishable inventory. We prove that with nonperishable inventory, the famous "stock more" result is often reversed to "stock less," in that the Bayesian optimal inventory level with unobserved lost sales is lower than the myopic inventory level. We also prove that making lost sales unobservable increases the Bayesian optimal inventory level; in this specific sense, the famous "stock more" result of other studies generalizes to the case of nonperishable inventory.

When the product is out of stock, a customer may accept a substitute or choose not to purchase. We incorporate learning about the probability of substitution. This reduces the Bayesian optimal inventory level in the case that lost sales are observed. Reducing the inventory level has two beneficial effects: to observe and learn more about customer substitution behavior and (for a nonperishable product) to reduce the probability of overstocking in subsequent periods.

Finally, for a capacitated production-inventory system under continuous review, we derive maximum likelihood estimators (MLEs) of the demand rate and probability that customers will wait for the product. (Accepting a raincheck for delivery at some later time is a common type of substitution.) We investigate how the choice of base-stock level and production rate affect the convergence rate of these MLEs. The results reinforce those for the Bayesian, uncapacitated, periodic review system.

Key words: Bayesian inventory management; unknown demand distribution; unobserved lost sales; substitution probability; Bayesian dynamic programming; optimal inventory control; maximum likelihood estimator; make-to-stock queue
History: Received: May 13, 2005; accepted: January 3, 2007. Published online in *Articles in Advance* December 17, 2007.

1. Introduction

Imagine a retailer purchasing inventory for an innovative product. If the retailer runs out of stock, customers may accept a substitute or choose not to purchase. The optimal inventory level for the product depends on both the demand distribution and substitution probability. Furthermore, the choice of inventory level affects the sales data and customer behavior observed by the retailer, and thus shapes the retailer's opportunity for learning about the demand distribution and substitution probability. This paper provides insights and guidelines for dynamically managing the inventory of an innovative product while learning about the demand distribution and substitution probability.

The early literature on Bayesian inventory management for a single item with unknown demand distribution assumes that the item is nonperishable and customer demand is observed perfectly (Scarf 1959, Azoury 1985, Lovejoy 1990). These papers focus on methods to reduce the computational complexity of the statistical inventory control problem by exploring the conjugate prior distribution structure and the state-space reduction technique. A more recent stream of research on this topic has put more emphasis on the derivation of qualitative insights about how the presence of Bayesian learning affects the optimal inventory decision. Assuming the item is perishable, Harpaz et al. (1982) recognize that when lost sales are not observed, one should initially increase the inventory level to learn more about the demand distribution. Further assuming that demand has an exponential distribution with gamma prior on the mean, Lariviere and Porteus (1999) derive a closed-form expression for the Bayesian optimal inventory level and confirm that it exceeds the Bayesian myopic inventory level in every period. Ding et al. (2002) and Lu et al. (2005, 2006) extend this "stock more" result to perishable inventory systems with a general continuous demand distribution. Bensoussan et al. (2005) study a similar problem with Markovian demand.

Following this literature, we say that the "Bayesian myopic inventory policy" maximizes the current-period expected profit, pretending, in the case of nonperishable inventory, that the unsold unit can be returned at the (discounted) original cost. The expectation is taken with respect to updated prior distribution that incorporates all past demand observations. Using this policy, the myopic decision maker fails to account for the potential benefit of gathering information now to improve future-period profits and also fails to account for inventory interaction between periods. However, the myopic decision maker does update her choice of inventory level in every period based on past demand observations. We will subsequently use the terms "myopic inventory level" and "Bayesian myopic inventory level" interchangeably.

The myopic inventory policy has two properties that are important for our analysis. First, in the case of perishable inventory, assuming that the decision maker observes lost sales, the myopic inventory policy is identical to the Bayesian optimal inventory policy (Lariviere and Porteus 1999, Ding et al. 2002). Second, in managing nonperishable inventory with an infinite horizon, the "optimal inventory policy without Bayesian updating is a stationary policy with inventory level equal to the initial myopic inventory level" (Heyman and Sobel 1984, p. 66, Proposition 3-1). Therefore, assuming the planning horizon is sufficiently long for the nonperishable inventory case, all statements below regarding the myopic inventory level in the initial period are also true when one substitutes the optimal inventory level without Bayesian learning. Specifically, when the myopic inventory level is lower (higher) than the Bayesian optimal inventory level, one should stock more (stock less) to account for the opportunity to learn about the demand distribution and substitution probability.

We extend the Bayesian inventory management literature in three directions. The first is to work with general discrete demand distributions. Most of the literature on Bayesian inventory management assumes a continuous demand distribution. We find the assumption of discrete demand simplifies many of our proofs.

Second, we relax the perishability assumption, which destroys the equivalence between the myopic inventory level and the Bayesian optimal inventory level with observed lost sales. We show that the "stock more" result in Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002) is often reversed to "stock less" when the Bayesian optimal inventory level with unobserved lost sales is compared with the myopic inventory level. That is, with unobserved lost sales and nonperishable inventory, one should often stock less than the myopic inventory level to reduce the risk of overstocking in the subsequent periods. However, when the comparison is made between the Bayesian optimal inventory level with unobserved lost sales and the Bayesian optimal inventory level with observed lost sales, we show that the "stock more" result in Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002) generalizes to the case of nonperishable inventory. That is, making lost sales unobservable increases the Bayesian optimal inventory level, regardless of whether inventory is perishable or nonperishable.

The latter insight is useful for computing effective inventory levels in practice. Exact computation of the Bayesian optimal inventory level with unobserved lost sales is tractable only within a limited newsvendor distribution family (Lariviere and Porteus 1999). In contrast, with observed lost sales, the Bayesian optimal inventory level is relatively easy to compute using the results in Azoury (1985) and Lovejoy (1990), which apply for a broad class of demand distributions. Therefore, our result that observing lost sales reduces the Bayesian optimal inventory level provides a useful lower bound for heuristic management of nonperishable inventory with unobserved lost sales.

The third direction in which we extend the Bayesian inventory management literature is to consider the effect of learning about the probability that a customer will accept a substitute product. Intuitively, one learns more about the substitution probability by lowering the inventory level and thus forcing more customers to consider the substitute product. We show that this intuitive insight is valid when lost sales are observed, but is not generally true when lost sales are unobserved.

In §2, we follow Nahmias and Smith (1994) and model customers' decisions to accept a substitute when the product is out of stock by a series of independent and identically distributed (i.i.d.) Bernoulli trials. Increasing the substitution probability reduces the expected underage cost and thus reduces the optimal inventory level. Assuming that the substitution probability and unknown parameter of the demand distribution are subject to a joint prior belief, we make comparisons among the Bayesian optimal inventory level with observed lost sales and with unobserved lost sales and the myopic inventory level. The results are summarized in Table 1.

For analytic tractability, we assume that the substitute product is always available and the substitution decision is observed. Many researchers have addressed joint inventory management for multiple products under customer substitution without the added complexity of Bayesian learning. This work can be categorized into papers that assume centralized control (Parlar and Goyal 1984, Ernst and Kouvelis 1999, Noonan 1995, Agrawal and Smith 1998, Smith and Agrawal 2000, Rajaram and Tang 2001, Mahajan and van Ryzin 2001a) and papers with inventory competition between multiple retailers (Parlar 1988, Wang and Parlar 1994, Lippman and McCardle 1997, Mahajan and van Ryzin 2001b, Netessine and Rudi 2003). Mahajan and van Ryzin (1999) provide a more detailed description of the literature on inventory management under customer substitution.

Table 1 **Comparisons Between Bayesian Optimal Inventory Levels with Unobserved and Observed Lost Sales and the Myopic Inventory Level**

	Perishable inventory	Nonperishable inventory
Estimate demand parameter θ only	$y^O = y^M$, $y^U \geq y^O$, $y^U \geq y^M$.	$y^O \leq y^M$, $y^U \geq y^O$, $y^U \gtrless y^M$.
Estimate demand parameter θ and substitution probability p	$y^O \leq y^M$, $y^U \gtrless y^O$, $y^U \gtrless y^M$.	$y^O \leq y^M$, $y^U \gtrless y^O$, $y^U \gtrless y^M$.

Notes. y^O denotes the Bayesian optimal inventory level with observed lost sales, y^U the Bayesian optimal inventory level with unobserved lost sales, and y^M the myopic inventory level. "$\gtrless$" indicates that the relationship could be either "greater than" or "less than" depending on the cost structure and relative uncertainties.

For complex systems with multiple products or capacity constraints, the Bayesian approach becomes intractable and maximum likelihood estimators (MLEs) are commonly employed. For example, for a multiproduct inventory system with a Poisson demand process and periodic observation of inventory levels, Anupindi et al. (1998) apply the expectation-maximization (EM) algorithm to compute MLEs for the demand rate for each product and substitution probabilities. For a single item with capacitated production and two distribution channels, Armony and Plambeck (2005) investigate the systematic errors in the MLEs for demand rate, and the reneging rate and the consequent errors in capacity investment caused by unsuspected duplicate ordering.

In §3, we consider a capacitated production-inventory system ($M/M/1$ make-to-stock queue with balking) under continuous review. We compute MLEs for the demand rate λ and probability p that when the product is out of stock, a customer will choose to wait for the product rather than balk. Accepting a raincheck for delivery of the product at some later time is a common type of substitution. For any finite time horizon, when lost sales (customers that balk) are not observed, the accuracy of the MLE for λ is monotonically increasing in the base-stock inventory level and production rate, but the accuracy of p is not. With observed lost sales, the accuracy of the MLE for λ does not depend on the base-stock level, and the accuracy of the MLE for p is monotonically decreasing in the base-stock level and production rate. It follows that with observed lost sales, one should initially reduce the inventory level or production rate to improve the accuracy of the MLE for p, which, in turn, will improve future decision making. These results for the capacitated production-inventory system reinforce the results for Bayesian inventory management summarized in Table 1.

2. Bayesian Inventory Management

This section considers periodic review Bayesian inventory management for a product with unknown demand distribution and unknown stockout-based substitution probability. Without imposing any distributional assumptions, we compare the Bayesian optimal inventory level with observed lost sales, the Bayesian optimal inventory level with unobserved lost sales, and the myopic inventory level. We first study the relatively simple case with learning only about the demand distribution, then study the case with learning about the demand distribution and substitution probability.

2.1. Model

The product will be stocked and sold for N periods. At the beginning of each period i ($i = 1, \ldots, N$), the inventory manager selects an inventory level for the product. The inventory level is achieved immediately after the decision (we are assuming a negligible delivery lead time from the supplier). For each unit of the product, the production cost is c and the selling price is r, with $r > c > 0$. At the end of each period, a unit holding cost h is charged to any leftover stock. As a base case, we assume that the inventory is nonperishable and can be used to satisfy demand in subsequent periods. At the end of the selling season, i.e., in period $N+1$, we assume that any unsold inventory will be salvaged at a unit value of c_s, with $c_s \le c$. If the demand during a period exceeds the inventory level, the manager is charged a stockout penalty q per unit of shortage. In addition, each customer who arrives when the product is out of stock is offered a substitute product. If a customer accepts the substitute, then the manufacturer receives a contribution margin of m per unit from selling the substitute. To avoid triviality, we assume that $r - c \ge m$, which guarantees that selling the product from inventory makes economic sense. (Accepting a rain check for delivery at the end of the period is a common type of substitution. In this case, the margin m is the product selling price r less the production cost c and any expediting/handling cost.) The objective of the inventory manager is to maximize total discounted expected profit. We denote the discount factor by δ ($0 \le \delta < 1$).

The demand in each period is independently and identically generated by a general nonnegative discrete demand distribution. The probability mass function of the demand distribution is denoted by $f(\xi \mid \theta)$, $\xi = 0, 1, 2, \ldots,$ where θ is an unknown parameter (or vector of unknown parameters) with $\theta \in \Theta$. In a period when stockout occurs, each unit of excess demand generates an independent Bernoulli trial: Each customer is willing to accept the substitute with probability p and becomes a lost sale with probability $1-p$.

Let z denote the starting inventory, y the chosen inventory level, and X the random total quantity sold, including sales of the substitute product caused by stockout (x will be used to denote the realization of total quantity sold).

For the moment, suppose that the manager observes all customer demand (including any lost sales). In other words, for each period, the manager observes not only the total sales x, but also the demand realization ξ. Given values of θ and p, and an inventory level of y, the likelihood of observing demand realization ξ and sales x in a period is

$$f^y(\xi, x \mid \theta, p) = \begin{cases} f(\xi \mid \theta) & \text{if } \xi = x < y, \\ \dbinom{\xi - y}{x - y} p^{x-y}(1-p)^{\xi - x} f(\xi \mid \theta) & \\ & \text{if } \xi \ge x \ge y, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Analogously, if lost sales are not observed (i.e., the manager only observes sales of the regular product and, in the event of stockout, the number of customers that accept the substitute), the likelihood of observing total sold quantity x can be written as

$$f_X^y(x \mid \theta, p) = \begin{cases} f(x \mid \theta) & \text{if } x < y, \\ \displaystyle\sum_{\xi = x}^{\infty} \binom{\xi - y}{x - y} p^{x-y}(1-p)^{\xi - x} f(\xi \mid \theta) & \text{if } x \ge y. \end{cases} \quad (2)$$

When $p = 0$, (2) becomes

$$f_X^y(x \mid \theta) = \begin{cases} f(x \mid \theta) & \text{if } x < y, \\ \displaystyle\sum_{\xi = y}^{\infty} f(\xi \mid \theta) & \text{if } x = y, \\ 0 & \text{otherwise,} \end{cases}$$

which is the discrete demand version of the lost-sales case (without substitution) studied in Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002). On the other hand, if $p=1$, we have $f_X^y(x \mid \theta, p) = f(x \mid \theta)$ for all $x \geq 0$, which means that the quantity sold in a period is an exact observation of the demand. Likelihood functions (1) and (2) will be used to update the posterior distribution in the observed and unobserved lost-sales cases, respectively.

The expected profit for a single period as a function of the demand parameter θ and substitution probability p is (by straightforward algebra)

$$\begin{aligned} &rE[\min(y,\xi)] + mE[(X-y)^+] - c(y-z) \\ &\quad - hE[(y-\xi)^+] - qE[(\xi-y)^+] \\ &\quad\quad = cz - (h+c)y + (r+h)\bar{\xi}(\theta) \\ &\quad\quad\quad - (r+h+q-mp)\sum_{\xi=y}^{\infty}(\xi-y)f(\xi \mid \theta), \end{aligned}$$

where $\bar{\xi}(\theta)$ is the mean demand given density $f(\xi \mid \theta)$.

Let $v_i(z \mid \theta, p)$ denote the maximum total discounted expected profit over periods $i, \ldots, N$ with z units of on-hand inventory at the beginning of period i as a function of the demand parameter θ and substitution probability p. We assume that the on-hand inventory is zero at the beginning of Period 1. Using a well-known simplifying technique (see Heyman and Sobel 1984, p. 79), we can write the optimality equations as

$$\begin{aligned} &v_1(0 \mid \theta, p) \\ &\quad = \max_{y\geq 0}\left\{R(y \mid \theta, p) + \delta \sum_{\xi=0}^{\infty} v_2((y-\xi)^+ \mid \theta, p) f(\xi \mid \theta)\right\}, \\ &v_i(z \mid \theta, p) \\ &\quad = \max_{y\geq z}\left\{R(y \mid \theta, p) + \delta \sum_{\xi=0}^{\infty} v_{i+1}((y-\xi)^+ \mid \theta, p) f(\xi \mid \theta)\right\} \\ &\qquad\qquad\qquad\qquad \text{for } i = 2, \ldots, N-1, \\ &\qquad v_N(z \mid \theta, p) = \max_{y\geq z}\{R_N(y \mid \theta, p)\}, \end{aligned} \tag{3}$$

where

$$\begin{aligned} R(y \mid \theta, p) &= -(\tilde{h}+c)y + (r+\tilde{h})\bar{\xi}(\theta) \\ &\quad - (r+\tilde{h}+q-mp)\sum_{\xi=y}^{\infty}(\xi-y)f(\xi \mid \theta), \\ R_N(y \mid \theta, p) &= -(\tilde{h}_N+c)y + (r+\tilde{h}_N)\bar{\xi}(\theta) \\ &\quad - (r+\tilde{h}_N+q-mp)\sum_{\xi=y}^{\infty}(\xi-y)f(\xi \mid \theta), \end{aligned}$$

with $\tilde{h} = h - \delta c$ and $\tilde{h}_N = h - \delta c_s$.

Now assume that, starting in Period 1, θ and p are not known but are subject to certain independent prior distributions, i.e., the joint prior distribution $\pi_1(\theta, p) = \pi_1'(\theta) \cdot \pi_1'(p)$, with $\pi_1'(\theta)$ and $\pi_1'(p)$ being the marginal density for θ and p, respectively. Given a joint prior distribution $\pi_i(\theta, p)$ for period i, the posterior distribution calculated from the demand information observed in period i is used as the prior distribution $\pi_{i+1}(\theta, p)$ for the subsequent period $i+1$. Hence, given a specific Bayesian prior updating scheme, the optimality equations for the Bayesian inventory management can be expressed as

$$\begin{aligned} &v_1(0, \pi_1) \\ &= \max_{y\geq 0} E_{\pi_1(\theta,p)}\left\{R(y \mid \theta, p) + \delta \sum_{\xi=0}^{\infty} v_2((y-\xi)^+, \pi_2) f(\xi \mid \theta)\right\}, \\ &v_i(z, \pi_i) = \max_{y\geq z} E_{\pi_i(\theta,p)}\bigg\{R(y \mid \theta, p) \\ &\qquad\qquad + \delta \sum_{\xi=0}^{\infty} v_{i+1}((y-\xi)^+, \pi_{i+1}) f(\xi \mid \theta)\bigg\}, \\ &\qquad\qquad\qquad\qquad i = 2, \ldots, N-1, \\ &\quad v_N(z, \pi_N) = \max_{y\geq z} E_{\pi_N(\theta,p)}\{R_N(y \mid \theta, p)\}. \end{aligned} \tag{4}$$

From the optimality equations (4), it is also useful to denote the objective function for each period as for $i = 1, \ldots, N-1$,

$$\begin{aligned} &G_i(y, \pi_i) \\ &= E_{\pi_i(\theta,p)}\left\{R(y \mid \theta, p) + \delta \sum_{\xi=0}^{\infty} v_{i+1}((y-\xi)^+, \pi_{i+1}) f(\xi \mid \theta)\right\}, \end{aligned}$$

and

$$G_N(y, \pi_N) = E_{\pi_N(\theta,p)}\{R_N(y \mid \theta, p)\}.$$

For each period i, the maximum is attained at some finite integer value(s) of y (because the expected profit goes to $-\infty$ as y goes to ∞). We denote the Bayesian optimal inventory level with unobserved lost sales as y^U, the Bayesian optimal inventory level with observed lost sales as y^O, and the myopic inventory level as y^M. We designate the same superscripts "U," "O," and "M" to the value functions $v_i(\cdot,\cdot)$ and objective functions $G_i(\cdot,\cdot)$ for these three cases, respectively. Below are detailed descriptions of the Bayesian updating schemes with observed and unobserved lost sales.

2.1.1. Observed Lost-Sales Case. Let us denote the marginal distribution of the joint prior distribution $\pi_i(\theta,p)$ for period i by $\pi_i'(\theta)$ and $\pi_i'(p)$ for θ and p, respectively. Given observations of (ξ,x) and inventory level y for period i, the posterior distribution density $\pi_{i+1}(\theta,p\mid\xi,x,y,\pi_i)$ can be updated using the likelihood function (1) according to Bayes' rule:

$$\pi_{i+1}(\theta,p\mid\xi,x,y,\pi_i) = \begin{cases} \dfrac{f(\xi\mid\theta)\pi_i(\theta,p)}{\int_\Theta f(\xi\mid\theta')\pi_i'(\theta')\,d\theta'} & \text{if } \xi=x\le y,\\[2ex] \dfrac{\binom{\xi-y}{x-y}p^{x-y}(1-p)^{\xi-x}f(\xi\mid\theta)\pi_i(\theta,p)}{\int_0^1\int_\Theta\binom{\xi-y}{x-y}(p')^{x-y}(1-p')^{\xi-x}f(\xi\mid\theta')\pi_i(\theta',p')\,d\theta'\,dp'} & \text{if } \xi>x=y,\\[2ex] \dfrac{\binom{\xi-y}{x-y}p^{x-y}(1-p)^{\xi-x}f(\xi\mid\theta)\mathbf{1}\{p>0\}\pi_i(\theta,p)}{\int_{0^+}^1\int_\Theta\binom{\xi-y}{x-y}(p')^{x-y}(1-p')^{\xi-x}f(\xi\mid\theta')\pi_i(\theta',p')\,d\theta'\,dp'} & \text{if } \xi\ge x>y. \end{cases} \tag{5}$$

Because $\pi_1(\theta,p)=\pi_1'(\theta)\cdot\pi_1'(p)$, substituting this relation into (5), we have

$$\pi_2(\theta,p\mid\xi,x,y,\pi_1) = \begin{cases} \pi_2'(\theta\mid\xi,\pi_1'(\theta))\cdot\pi_1'(p) & \text{if } \xi=x\le y,\\[1ex] \pi_2'(\theta\mid\xi,\pi_1'(\theta))\cdot\pi_2'(p\mid\xi-y,x-y,\pi_1'(p)) & \text{if } \xi\ge x\ge y, \end{cases}$$

where

$$\pi_2'(\theta\mid\xi,\pi_1'(\theta)) = \frac{f(\xi\mid\theta)\pi_1'(\theta)}{\int_\Theta f(\xi\mid\theta')\pi_1'(\theta')\,d\theta'} \quad \text{and}$$

$$\pi_2'(p\mid\xi-y,x-y,\pi_1'(p)) = \begin{cases} \dfrac{\binom{\xi-y}{x-y}p^{x-y}(1-p)^{\xi-x}\pi_1'(p)}{\int_0^1\binom{\xi-y}{x-y}(p')^{x-y}(1-p')^{\xi-x}\pi_1'(p')\,dp'}, & \text{if } \xi\ge x=y,\\[2ex] \dfrac{\binom{\xi-y}{x-y}p^{x-y}(1-p)^{\xi-x}\mathbf{1}\{p>0\}\pi_1'(p)}{\int_{0^+}^1\binom{\xi-y}{x-y}(p')^{x-y}(1-p')^{\xi-x}\pi_1'(p')\,dp'}, & \text{if } \xi\ge x>y. \end{cases}$$

Hence, the updated prior distributions of θ and p for Period 2 are also independent. By induction, the updated prior distributions of θ and p of all subsequent periods are independent. The Bayesian updating of θ and p can thus be performed separately based on their respective marginal distributions. This clean updating process is a result of the observability of lost-sales information: Full information is available for each parameter-updating scheme and thus prevents them from tangling up. For ease of notation, in the rest of the paper either $\pi_{i+1}'(p\mid\xi-y,x-y)$ or $\pi_{i+1}'(p)$ will be shorthand for $\pi_{i+1}'(p\mid\xi-y,x-y,\pi_i'(p))$ whenever the meaning is clear from the context.

2.1.2. Unobserved Lost-Sales Case. Now let us consider the unobserved lost-sales case. For period i, let $\pi_i(\theta,p)$ be the joint prior distribution density of θ and p and y be the inventory level selected. If total sales x is observed in period i, the posterior distribution density of θ and p is given by the following equation, according to Bayes' rule and likelihood function (2):

$$\pi_{i+1}(\theta,p\mid x,y,\pi_i(\theta,p)) = \begin{cases} \dfrac{f_X^y(x\mid\theta,p)\pi_i(\theta,p)}{\int_0^1\int_\Theta f_X^y(x\mid\theta',p')\pi_i(\theta',p')\,d\theta'\,dp'} & \text{if } x\le y,\\[2ex] \dfrac{f_X^y(x\mid\theta,p)\mathbf{1}\{p>0\}\pi_i(\theta,p)}{\int_{0^+}^1\int_\Theta f_X^y(x\mid\theta',p')\pi_i(\theta',p')\,d\theta'\,dp'} & \text{if } x>y. \end{cases} \tag{6}$$

For ease of notation, in the rest of the paper we write $\pi_{i+1}(\theta,p\mid x,y,\pi_i(\theta,p))$ with the shorthand expression $\pi_{i+1}(\theta,p\mid x,y)$ or $\pi_{i+1}(\theta,p\mid x)$ in the case of $x<y$ (because the posterior is not affected by y) whenever the meaning is clear from the context. As we can see from likelihood function (2) and updating scheme (6), unlike the observed lost-sales case, the posteriors for θ and p are mixed up once $x\ge y$ is observed.

2.2. Updating Demand Parameter θ Only

In this section, we consider the case that only demand parameter θ is updated and the substitution probability p is assumed known. When lost sales are not observed, the general consensus in the literature is that one should stock more than the myopic inventory level to increase the probability of observing an exact demand realization (Harpaz et al. 1982, Lariviere and Porteus 1999, Ding et al. 2002). However, this insight is derived based on the assumption that inventory is perishable (cannot be carried over to the next period). We relax the perishability assumption in this section.

We first show that this stock more result remains valid when the Bayesian optimal inventory level with unobserved lost sales is compared to that with observed lost sales. However, when the Bayesian optimal inventory level with unobserved lost sales is compared to the myopic inventory level, we show examples in which the opposite result (stock less) holds. In these examples, one should reduce the inventory level to mitigate the risk of overstocking in the subsequent periods.

Define $\Delta g(x,\cdot) = g(x+1,\cdot) - g(x,\cdot)$ for any function $g(x,\cdot)$. It is easy to verify that for the case with observed lost sales, we have, for $i = 1,\ldots,N-1$,

$$\Delta G_i^O(y, \pi_i'(\theta)) = E_{\pi_i'(\theta)}\left\{\Delta R(y \mid \theta, p) + \delta \sum_{x=0}^{y} \Delta v_{i+1}^O(y-x, \pi_{i+1}'(\theta \mid x)) f(x \mid \theta, p)\right\}. \quad (7)$$

It is straightforward to verify through backward induction that $G_i^O(y, \pi_i'(\theta))$ are concave in y (see Scarf 1959 for the continuous demand distribution case).

For the case with unobserved lost sales, we have, for $i = 1,\ldots,N-1$,

$$\Delta G_i^U(y, \pi_i'(\theta)) = E_{\pi_i'(\theta)}\left\{\Delta R(y \mid \theta, p) + \delta \sum_{x=0}^{y} \Delta v_{i+1}^U(y-x, \pi_{i+1}'(\theta \mid x)) f(x \mid \theta, p) + \delta\left\{\sum_{x=y}^{\infty} v_{i+1}^U(0, \pi_{i+1}'(\theta \mid x, y+1)) f_X^{y+1}(x \mid \theta, p) - \sum_{x=y}^{\infty} v_{i+1}^U(0, \pi_{i+1}'(\theta \mid x, y)) f_X^{y}(x \mid \theta, p)\right\}\right\}. \quad (8)$$

The concavity of $G_i^U(y, \pi_i'(\theta))$ is difficult to establish for general distributions. However, we can establish the following result:

Proposition 1. *For $i = 2,\ldots N$, $y \geq 0$, and any $\pi'(\theta)$, the following holds*:

$$E_{\pi'(\theta)}\left\{\sum_{x=y}^{\infty} v_i^U(0, \pi_i'(\theta \mid x, y+1)) f_X^{y+1}(x \mid \theta, p)\right\} \geq E_{\pi'(\theta)}\left\{\sum_{x=y}^{\infty} v_i^U(0, \pi_i'(\theta \mid x, y)) f_X^{y}(x \mid \theta, p)\right\}; \quad (9)$$

in other words, the expected future value function is increasing in the current-period inventory level y (expectation taken when demand exceeds y).

The proof of Proposition 1 is in the appendix. Proposition 1 shows that if p is known, then, conditional on the event that demand exceeds the current-period inventory level y, stocking an additional unit above y will always increase the discounted expected profit in future periods—because the manager will observe one more unit of customer demand.

If the product were perishable, the second terms of (7) and (8) would be zero. Hence, by Proposition 1, the stock more result would follow. The case with nonperishable inventory is complicated by these additional two terms. Nevertheless, we have the following unequivocal result.

Theorem 1. *Suppose that the substitution probability p is known and lost sales are not observed. Regardless of whether the inventory is perishable or nonperishable, given the same prior $\pi_i'(\theta)$, the Bayesian optimal inventory level with unobserved lost sales is greater than the Bayesian optimal inventory level with observed lost sales, i.e., $y_i^U \geq y_i^O$.*

Proof. We first consider the case with nonperishable inventory. To prove $y_i^U \geq y_i^O$, it suffices to show that $\Delta G_i^U(y, \pi_i'(\theta)) \geq \Delta G_i^O(y, \pi_i'(\theta))$, for all $y \geq 1$. With Proposition 1, it thus suffices to show that $\Delta v_{i+1}^U(y-x, \pi_{i+1}'(\theta \mid x)) \geq \Delta v_{i+1}^O(y-x, \pi_{i+1}'(\theta \mid x))$. We show it by backward induction. It is easy to verify that given the same prior distribution $\pi_N'(\theta)$, these above claims hold. Now, assume that these are true for case $i+1$; i.e., given any prior $\pi_{i+1}'(\theta)$, we have $y_{i+1}^U \geq y_{i+1}^O$ and $\Delta G_{i+1}^U(y, \pi_i'(\theta)) \geq \Delta G_{i+1}^O(y, \pi_i'(\theta))$ for all $y \geq 1$. A key observation is that given the same prior $\pi_i'(\theta)$ at the beginning of period i, the updated posterior $\pi_{i+1}'(\theta \mid x)$ after observing sales x (with $x < y$) for both the observed and unobserved lost-sales cases remains identical. Hence, we can successfully apply the induction assumption of period $i+1$. Specifically, we consider three cases:

Case 1. $y^U_{i+1} \geq y^O_{i+1} > z$, we have $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) = \Delta v^O_{i+1}(z, \pi'_{i+1}(\theta)) = 0$.

Case 2. $y^U_{i+1} \geq z > y^O_{i+1}$, we have $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) = 0 \geq \Delta v^O_{i+1}(z, \pi'_{i+1}(\theta))$.

Case 3. $z > y^U_{i+1} \geq y^O_{i+1}$. Notice that since $y^O_{i+1} \leq z$, by concavity of $G^O_{i+1}(\cdot, \cdot)$, we have $\Delta G^O_{i+1}(z', \cdot) \leq 0$ for all $z' \geq z$. Hence, the optimal decision in the observed lost-sales case is to place zero order. As a result,

$$\Delta v^O_{i+1}(z, \pi'_{i+1}(\theta)) = \Delta G^O_{i+1}(z, \pi'_{i+1}(\theta)) \leq 0. \tag{10}$$

For the unobserved lost-sales case, however, $G^U_{i+1}(\cdot, \cdot)$ is *not* necessarily concave. Therefore, there exist two scenarios when starting inventory in period $i+1$ is z: (a) do not order, or (b) place an order to an inventory level greater than z.

In scenario (a), we have

$$\begin{aligned}\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) &= v^U_{i+1}(z+1, \pi'_{i+1}(\theta)) - v^U_{i+1}(z, \pi'_{i+1}(\theta)) \\ &\geq G^U_{i+1}(z+1, \pi'_{i+1}(\theta)) - G^U_{i+1}(z, \pi'_{i+1}(\theta)) \\ &= \Delta G^U_{i+1}(z, \pi'_{i+1}(\theta)),\end{aligned}$$

where the inequality follows from the fact that $v^U_{i+1}(z+1, \pi'_{i+1}(\theta)) \geq G^U_{i+1}(z+1, \pi'_{i+1}(\theta))$ and $v^U_{i+1}(z, \pi'_{i+1}(\theta)) = G^U_{i+1}(z, \pi'_{i+1}(\theta))$. By the induction assumption: $\Delta G^U_{i+1}(y, \cdot) \geq \Delta G^O_{i+1}(y, \cdot)$ for all $y \geq 1$, and the observation that $\Delta v^O_{i+1}(z, \cdot) = \Delta G^O_{i+1}(z, \cdot)$ for $z > y^O_{i+1}$, we conclude that $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) \geq \Delta v^O_{i+1}(z, \pi'_{i+1}(\theta))$.

In scenario (b), we have $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) = 0$, which, again, means $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) \geq \Delta v^O_{i+1}(z, \pi'_{i+1}(\theta))$ because the latter is nonpositive as shown in (10).

Therefore, we have shown that $\Delta v^U_{i+1}(z, \pi'_{i+1}(\theta)) \geq \Delta v^O_{i+1}(z, \pi'_{i+1}(\theta))$ for all $z \geq 0$. With Proposition 1, we conclude $\Delta G^U_i(y, \cdot) \geq \Delta G^O_i(y, \cdot)$, for all $y \geq 1$, and by concavity of $\Delta G^O_i(y, \cdot)$, we obtain $y^U_i \geq y^O_i$. This completes the induction proof.

For the case with perishable inventory, we first replace $\tilde{h}$ and $\tilde{h}_N$ in (3) with h. We know that $G^O_i(y, \pi'_i(\theta))$ is concave in y, and

$$\Delta G^O_i(y, \pi_i) = E_{\pi'_i(\theta)}\{\Delta R(y \mid \theta, p)\}.$$

By Proposition 1 and the perishable inventory assumption, we immediately have $\Delta G^U_i(y, \pi'_i(\theta)) \geq \Delta G^O_i(y, \pi'_i(\theta))$, for all $y \geq 1$, and hence, $y^U_i \geq y^O_i$. □

Theorem 1 generalizes the observation in Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002) that making lost sales unobservable increases the Bayesian optimal inventory level to allow for nonperishable inventory and partial lost sales ($0 \leq p \leq 1$). In general, lack of lost-sales information induces a Bayesian inventory manager to increase the inventory level to observe more exact demand.

Our derivation is based on a general discrete demand distribution. An analogous result holds for general continuous demand distributions, but the analogous proof requires more notation and is more complex.

Theorem 1 provides a relatively easy-to-compute lower bound on the Bayesian optimal inventory level for systems with unobserved lost sales and nonperishable inventory setting (which is only tractable within a limited newsvendor distribution family; see Lariviere and Porteus 1999). Ideally, one can leverage the Bayesian optimal inventory level with observed lost sales (which is tractable for much broader distribution families; see Azoury 1985 and Lovejoy 1990) to obtain approximate solutions to the much harder unobserved lost-sales case. Designing efficient computational methods to determine the Bayesian optimal inventory level is beyond the scope of this paper, but will be a natural extension following the line of this research work.

Now let us compare the Bayesian optimal inventory level with the myopic inventory level. By the definition of myopic inventory level given in §1, we immediately have, for $i = 1, \ldots, N-1$,

$$\Delta G^M_i(y, \pi'_i(\theta)) = E_{\pi'_i(\theta)}\{\Delta R(y \mid \theta, p)\}. \tag{11}$$

When the inventory is perishable, we know that the Bayesian optimal inventory level with observed lost sales is equivalent to the myopic inventory level. Hence, by Theorem 1, the Bayesian optimal inventory level with unobserved lost sales is greater than the myopic inventory level.

When the inventory is nonperishable, comparing (7) with (11) and observing that $\Delta v^O_{i+1}(z, \pi'_{i+1}(\theta \mid x)) \leq 0$ for all $z \geq 0$, we find that the Bayesian optimal inventory level with observed lost sales is less than the myopic inventory level, i.e., $y^O_i \leq y^M_i$ for all i.

If we compare (8) with (11), on the one hand, as shown in Proposition 1, the inventory manager wants to stock more than the myopic inventory level to observe exact demand if she knows the demand is going to be high. On the other hand, in case the realized demand is low, the inventory manager may be left with excess inventory after updating with the low-demand observation. In the latter case, the manager would prefer to stock less than the myopic inventory level (because $\delta\sum_{x=0}^{y}\Delta v_{i+1}^{U}(y-x, \pi_{i+1}'(\theta\mid x))f(x\mid\theta,p)$, the second term of (8) is negative due to the fact that $\Delta v_{i+1}^{U}(z,\pi_{i+1}'(\theta\mid x))\leq 0$ for all $z\geq 0$). There is a trade-off between stocking more (to observe more demand information) and stocking less (to avoid excess inventory) when the product is nonperishable and the Bayesian optimal inventory level with unobserved lost sales is compared with the myopic inventory level.

The next result provides guidance about the conditions under which one should stock more than the myopic inventory level to account for the opportunity to learn about the demand distribution.

Proposition 2. *Suppose that the substitution probability p is known and lost sales are not observed. If there exists a period j ($j<N$) in which for a given prior $\pi'(\theta)$ the Bayesian optimal inventory level with unobserved lost sales is greater than the myopic inventory level, i.e., $y_j^U\geq y_j^M$, then for any period $i<j$, given the same starting prior $\pi'(\theta)$, $y_i^U\geq y_i^M$.*

Proof. Based on the myopic policy definition, it is straightforward that given the same prior $\pi'(\theta)$, $y_i^M = y_j^M$ for any $i<j$. Thus, for any $y\leq y_{j-1}^M = y_j^M\leq y_j^U$, from (8) we have $\Delta G_{j-1}^U(y,\pi'(\theta))\geq 0$, which means the optimal inventory y_{j-1}^U in period $j-1$ must be greater than y_{j-1}^M. By induction, this result holds for any period $i<j$. □

Proposition 2 tells us that increasing the planning horizon tends to make the initial Bayesian optimal inventory level with unobserved lost sales rise above the myopic inventory level. However, with shorter horizons such as $N=2$ periods, we commonly observe the "stock less" result.

2.2.1. Numerical Examples. Suppose that we have $N=2$ periods to sell a nonperishable product. For simplicity, we assume that $p=0$, i.e., all unmet demand is lost, and that the demand in each period follows an independent geometric distribution with an unknown parameter θ, subject to a beta prior with parameters α and β. It is straightforward to show that the predictive demand distribution for the first period is $f_1(x\mid\alpha,\beta)=\alpha\Gamma(\alpha+\beta)\Gamma(\beta+x)/\Gamma(\beta)\Gamma(\alpha+\beta+x+1)$, where $\Gamma(\cdot)$ is the gamma function. If $\alpha>1$, which we assume henceforth, then the mean of the predictive demand distribution equals $\beta/(\alpha-1)$ (see Johnson et al. 1992, p. 243). Furthermore, if the demand realized in the first period is below the inventory level, i.e., $x<y$, the posterior for unknown parameter θ is beta distribution with parameters $\alpha'=\alpha+1$ and $\beta'=\beta+x$; otherwise, if demand is censored at y, the posterior is gamma with parameters $\alpha'=\alpha$ and $\beta'=\beta+y$.

To illustrate the effect of the perishable inventory assumption on Bayesian inventory decisions, we assume that inventory is nonperishable but can be salvaged at the end of first period for c_s^1 ($c_s^1\leq c$) per unit and at the end of second period for $c_s=c$ per unit. For $c_s^1=c$, this is equivalent to the perishable inventory case. When $c_s^1\leq\delta c_s-h$, one would never salvage inventory at the end of Period 1, so this is equivalent to the nonperishable inventory case with no salvaging option.

Given the salvaging option at the end of the first period, the optimal inventory policy for Period 2 (which is also the last period) is an (L,U)-type, i.e., the manager would order up to L if $z_2<L$ (z_2 is the on-hand inventory at the beginning of Period 2); order nothing if $L\leq z_2\leq U$; and salvage inventory down to U if $z_2>U$.

The optimal Period 1 inventory level for either the Bayesian or myopic case cannot be expressed in closed form and is not necessarily equal to the order-up-to level L of Period 2, but it can be easily determined numerically in this two-period example. In particular, we assume that selling price $r=10$, purchase cost $c=c_s=2$, holding cost $h=1$, shortage cost $q=8$, and discount factor $\delta=1$. The salvage value c_s^1 is varied between one and two, with $c_s^1=1$ equivalent to the nonperishable case and $c_s^1=2$ equivalent to the perishable case.

Figure 1 shows the Bayesian optimal inventory level with observed and unobserved lost sales and myopic inventory level for Period 1 with a beta prior $\alpha=2$, $\beta=4$. The coefficient of variation of this beta prior is $\sqrt{\beta/\alpha(\alpha+\beta+1)}=0.5345$, implying a relatively high degree of uncertainty about the parameter θ. (The resulting predictive demand mean for

Figure 1 **Impact of First-Period Salvage Value on the Optimal Inventory-Level Decisions**

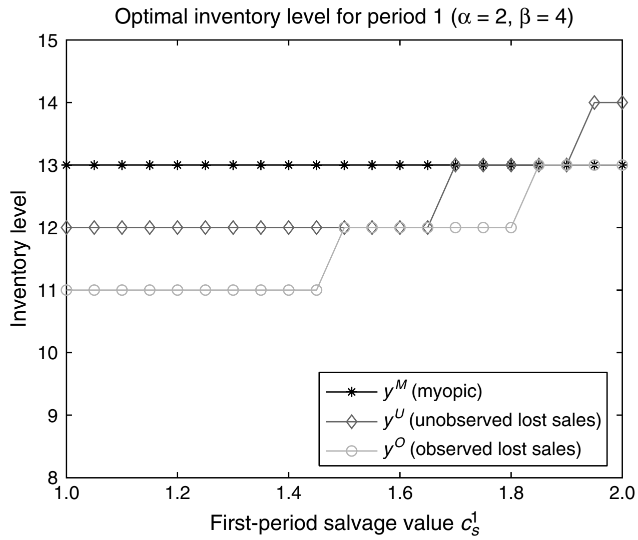

Period 1 is $\beta/(\alpha-1)=4$.) Figure 1 illustrates the analytic result in Theorem 1: The Bayesian optimal inventory level with unobserved lost sales is weakly greater than that with observed lost-sales case. Figure 1 also shows that the Bayesian optimal inventory level with unobserved lost sales y^U is less than the myopic level y^M in most of the parameter region. Specifically, this stock less result holds for all salvage values $c_s^1 \in [1, 1.9]$ and is reversed to stock more only when $c_s^1 > 1.9$, which essentially corresponds to the case of perishable inventory. This illustrates that the penalty of overstocking (even with the first-period salvage option) tends to dominate the informational benefit of observing exact demand by increasing inventory.

This stock less result also holds for the exponential demand distribution. (Following Lariviere and Porteus 1999, we focus on the exponential distribution to obtain a closed-form expression for the Bayesian optimal inventory level with unobserved lost sales.)

Proposition 3 (Chen 2005). *Suppose that the demand distribution is exponential and the product is nonperishable with the close-out salvage value equal to the original cost* ($c_s = c$). *For a two-period problem with censored demand observation, the Bayesian optimal inventory level with unobserved lost sales is less than the myopic inventory level, i.e.,* $y_1^U < y_1^M$.

Derived in a two-period setting, this stock less result is the opposite of the stock more obtained in Lariviere and Porteus (1999) for the same model but with perishable inventory. The proof is in the first author's unpublished PhD thesis (Chen 2005).

In conclusion, the stock less result is remarkably robust in our examples. Often, the Bayesian inventory manager will have less inventory than the *myopic* level because the potential cost of overstocking in subsequent periods (because demand is lower than initially anticipated) outweighs the potential informational gains from increasing inventory to observe a more exact realization of demand.

2.3. Updating Both Demand Parameter θ and Substitution Probability p

In this section, we study the case in which both demand parameter θ and substitution probability p are updated according to the Bayes rule. To differentiate this case from the cases studied in the previous section, we use the superscript "$O2$" and "$U2$" to indicate the cases of updating both θ and p with observed and unobserved lost sales, respectively. Because the myopic inventory level in this case is equivalent to the previous section if the known p value is replaced by the prior mean of p in (11), we keep using the superscript "M" to indicate the myopic case in this section.

We first assume that the manager can observe any lost sales. This assumption is valid, for instance, when customer orders arrive through a call center or over the Internet and detailed order data is recorded. We have, for $i=1,\dots,N-1$,

$$\begin{aligned}\Delta G_i^{O2}(y,\pi_i)\\ =E_{\pi_i(\theta,p)}\Bigg\{&\Delta R(y\mid\theta,p)\\ &+\delta\sum_{\xi=0}^{y}\Delta v_{i+1}^{O2}(y-\xi,\pi'_{i+1}(\theta)\cdot\pi'_i(p))f(\xi\mid\theta)\\ &+\delta\sum_{\xi=y+1}^{\infty}\big\{E_{K\mid\xi-y-1}\big\{v_{i+1}^{O2}(0,\pi'_{i+1}(\theta)\\ &\qquad\cdot\pi'_{i+1}(p\mid\xi-y-1,K))\big\}\\ &-E_{K\mid\xi-y}\big\{v_{i+1}^{O2}(0,\pi'_{i+1}(\theta)\cdot\pi'_{i+1}(p\mid\xi-y,K))\big\}\big\}\\ &\cdot f(\xi\mid\theta)\Bigg\},\end{aligned}\tag{12}$$

where $E_{K\mid n}\{\cdot\}$ denotes expectation taken over random variable K with binomial distribution (n,p).

Proposition 4. *Suppose that lost sales are observed. The ongoing expected value function for the next period increases with the number of substitution trials observed in the current period. For any period $i \in 2, \ldots, N$, number of trials $n \geq 0$ and prior $\pi'_{i-1}(p)$:*

$$E_{\pi'_{i-1}(p)} E_{K\mid n+1}\{v_i^{O2}(z, \pi'_i(\theta) \cdot \pi'_i(p \mid n+1, K, \pi'_{i-1}))\} \geq E_{\pi'_{i-1}(p)} E_{K\mid n}\{v_i^{O2}(z, \pi'_i(\theta) \cdot \pi'_i(p \mid n, K, \pi'_{i-1}))\}.$$

This result is derived based on the fact that when lost sales are observed, updating θ and p can be separated. The intuition behind it is that the greater the number of substitution trials, the greater the expected informational benefits from Bayesian updating. The proof together with two auxiliary lemmas, is given in the appendix. Applying Proposition 4 to (12), we immediately see that the last term is negative, which means one is induced to stock less inventory to increase the number of customer substitution trials.

Because both $G_i^M(y, \pi_i)$ and $G_i^O(y, \pi_i)$ are concave in y, the optimal inventory levels y_i^M and y_i^O can be uniquely determined by the first-order difference conditions. In contrast, $G_i^{O2}(y, \pi_i)$ is not necessarily concave in y. Nevertheless, we have the following unequivocal result.

Theorem 2. *Suppose that lost sales are observed. Regardless of whether the product is perishable or nonperishable, for any $i \in 1, \ldots, N$, given the same prior π_i, the Bayesian optimal inventory level with updating of both demand parameter θ and substitution probability p is less than the Bayesian optimal inventory level with updating only of demand parameter θ, i.e., $y_i^{O2} \leq y_i^O$.*

Proof. We first prove for the case of nonperishable inventory. We need to show $\Delta G_i^{O2}(y, \pi_i) \leq \Delta G_i^O(y, \pi_i)$, for all $y \geq 1$. With Proposition 4, it suffices to show that $\Delta v_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) \leq \Delta v_{i+1}^O(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p))$ for all $z \geq 0$. We show this by backward induction. It is easy to verify that, given the same prior distribution $\pi_N(\theta, p)$, the result holds. Now assume that this is true for case $i+1$, i.e., given any prior $\pi_{i+1}(\theta, p)$, we have $\Delta G_{i+1}^{O2}(y, \pi_{i+1}) \leq \Delta G_{i+1}^O(y, \pi_{i+1})$, for all $y \geq 1$, and $y_{i+1}^{O2} \leq y_{i+1}^O$. Consider three cases:

Case 1. $z < y_{i+1}^{O2} \leq y_{i+1}^O$, we have $\Delta v_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) = \Delta v_{i+1}^O(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) = 0$.

Case 2. $y_{i+1}^{O2} \leq z < y_{i+1}^O$, we have $\Delta v_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) \leq 0 = \Delta v_{i+1}^O(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p))$.

Case 3. $y_{i+1}^O \leq z$, we first show that one would not order in the full updating ($O2$) case. Notice that because $y_{i+1}^O \leq z$, by concavity of $G_{i+1}^O(\cdot, \cdot)$, we have $\Delta G_{i+1}^O(z', \cdot) \leq 0$ for all $z' \geq z$. Hence, by induction assumption, $\Delta G_{i+1}^{O2}(z', \cdot) \leq \Delta G_{i+1}^O(z', \cdot) \leq 0$ for all $z' \geq z$. Therefore, the optimal inventory levels for both ($O2$) and (O) are z in this case. As a result, for the prior $\pi'_{i+1}(\theta) \cdot \pi'_i(p)$, we have $\Delta v_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) = \Delta G_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) \leq \Delta G_{i+1}^O(z', \pi'_{i+1}(\theta) \cdot \pi'_i(p)) = \Delta v_{i+1}^O(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p))$.

Therefore, we have shown that $\Delta v_{i+1}^{O2}(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p)) \leq \Delta v_{i+1}^O(z, \pi'_{i+1}(\theta) \cdot \pi'_i(p))$ for all $z \geq 0$. With Proposition 4, we have $\Delta G_i^{O2}(y, \pi_i) \leq \Delta G_i^O(y, \pi_i)$, for all $y \geq 1$, and by concavity of $\Delta G_i^O(y, \cdot)$, we have $y_i^{O2} \leq y_i^O$.

For the perishable inventory case, we only need to replace $\tilde{h}$ and $\tilde{h}_N$ in (3) with h (the unit disposal/salvage cost of the perishable product). Because no inventory is carried over to subsequent periods, the interaction between periods is based purely on information updating. We have

$$\Delta G_i^O(y, \pi_i) = \Delta G_i^M(y, \pi_i) = E_{\pi_i(\theta, p)}\{\Delta R(y \mid \theta, p)\},$$

and

$$\begin{aligned}
&\Delta G_i^{O2}(y, \pi_i) \\
&= E_{\pi_i(\theta, p)}\Bigg\{\Delta R(y \mid \theta, p) + \delta \sum_{\xi=y+1}^{\infty} \Big\{E_{K\mid \xi-y-1}\{v_{i+1}^{O2}(0, \pi'_{i+1}(\theta) \\
&\qquad \cdot \pi'_{i+1}(p \mid \xi-y-1, K))\} \\
&\qquad - E_{K\mid \xi-y}\{v_{i+1}^{O2}(0, \pi'_{i+1}(\theta) \\
&\qquad \cdot \pi'_{i+1}(p \mid \xi-y, K))\}\Big\} f(\xi \mid \theta)\Bigg\},
\end{aligned}$$

where $\tilde{h}$ and $\tilde{h}_N$ are replaced by h (the unit salvage cost of the perishable product). By arguments analogous to the proof of Proposition 4, we can show that for the perishable product case,

$$E_{\pi'_{i-1}(p)} E_{K\mid n+1}\{v_i^{O2}(0, \pi'_i(\theta) \cdot \pi'_i(p \mid n+1, K, \pi'_{i-1}))\} \geq E_{\pi'_{i-1}(p)} E_{K\mid n}\{v_i^{O2}(0, \pi'_i(\theta) \cdot \pi'_i(p \mid n, K, \pi'_{i-1}))\},$$

for $i = 2, \ldots, N$, $n \geq 0$, and any $\pi'(p)$. Hence, we immediately have

$$\Delta G_i^{O2}(y, \pi_i) \leq \Delta G_i^O(y, \pi_i) = \Delta G_i^M(y, \pi_i).$$

By concavity of $\Delta G_i^O(y, \cdot)$, we have $y_i^{O2} \leq y_i^O$. □

Theorem 2 confirms the intuition that to learn about the substitution probability, a forward-looking manager should reduce the inventory level to observe more trials of customer substitution behavior. From the previous section, we know that $y_i^O \le y_i^M$. Hence, by Theorem 2, we immediately have $y_i^{O2} \le y_i^M$; i.e., the Bayesian optimal inventory level with observed lost sales (updating both parameters θ and p) is less than the myopic inventory level.

Theorem 2 is derived based on a general discrete demand distribution. Assuming discrete rather than continuous demand, as in most of the Bayesian inventory literature, we can explicitly capture the customer substitution behaviors as independent Bernoulli trials and thus incorporate learning about the probability that customers will accept a substitute.

In summary, we have observed two beneficial effects from inventory reduction. First, reducing the inventory level increases the number of observations of customer substitution behavior and thus improves estimation of p (hence, $y_i^{O2} \le y_i^O$). Second, reducing the inventory level reduces the risk of having excess inventory in the next period in the case of observing poor sales in the current period (hence, $y_i^O \le y_i^M$).

Eppen and Iyer (1997) study Bayesian inventory management of a fashion product (nonperishable during the selling season) with observed lost sales and allow for salvaging in any period. For a class of demand distributions that include normal, Poisson, and negative binomial, they prove that the optimal policy is of the form order-up-to L and salvage-down-to U where $L \le U$. When we allow for salvaging in every period in our model with general discrete demand and unknown substitution probability, the optimal policy is not necessarily of the (L, U) type because the expected total reward function for either ordering or salvaging is not necessarily concave in the inventory level y. The optimal inventory level y_i^* for period i after ordering or salvaging is a function, possibly complex, of the inventory remaining at the end of period $i-1$. Nevertheless, Theorem 2 holds when we allow for salvaging in each period (see the online appendix[1] The proof follows the same basic line of argument as the one above but is, of course, more complex.

[1] An online appendix to this paper is available on the *Manufacturing & Service Operations Management* website (http://msom.pubs.informs.org/ecompanion.html).

Now let us consider the case when lost sales are not observed. The first-order difference of the objective function for period i is given by

$$\begin{aligned}\Delta G_i^{U2}(y, \pi_i) &= E_{\pi_i(\theta, p)}\Bigg\{\Delta R(y \mid \theta, p) \\ &\quad + \delta \sum_{x=0}^{y} \Delta v_{i+1}^{U2}(y-x, \pi_{i+1}(\theta, p \mid x)) f(x \mid \theta, p) \\ &\quad + \delta \Bigg\{\sum_{x=y}^{\infty} v_{i+1}^{U2}(0, \pi_{i+1}(\theta, p \mid x, y+1)) f_X^{y+1}(x \mid \theta, p) \\ &\quad - \sum_{x=y}^{\infty} v_{i+1}^{U2}(0, \pi_{i+1}(\theta, p \mid x, y)) f_X^{y}(x \mid \theta, p)\Bigg\}\Bigg\}. \end{aligned} \tag{13}$$

The second term in (13) is negative because $\Delta v_{i+1}^{U2}(z, \pi_{i+1}(\theta, p \mid x)) \le 0$ for all $z \ge 0$, but the sign of the last term is no longer definite because Proposition 1 no longer holds when we update θ and p. As a result, making lost sales unobservable may either increase or decrease the Bayesian optimal inventory level (this is proven by the numerical examples immediately below). Qualitatively, from the observed lost-sales case, we have shown that one should stock less inventory to better estimate p. However, in §2.2, we have shown that one would stock more to better estimate θ when lost sales are not observed. Having a better estimate of θ is also helpful in inferring p from observing the number of customers that chose the substitute (but not lost sales). The following specially constructed numerical examples and analytic result for the case of exponential demand shed some light on the conditions under which one should stock more versus less.

2.3.1. Special Examples. This subsection provides examples in which demand parameter θ is known, but one may either stock more or stock less than the myopic inventory level to estimate the substitution probability p.

Consider a two-period problem in which cost parameters are set as $r = 10$, $c = 0$, $m = 10$, $q = 0$, and $h = 3$. The unknown substitution probability p takes value in $\{0.2, 0.9\}$, with a prior belief of $\Pr(p = 0.2) = w$ and

Figure 2 **Impact of Prior Belief Parameter on the Optimal Inventory-Level Decisions**

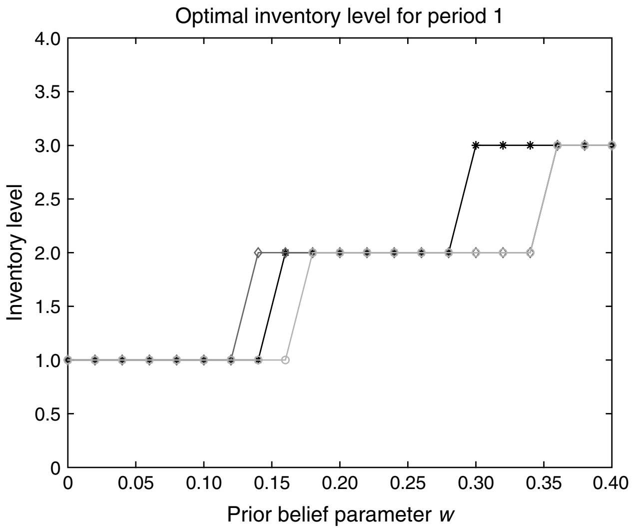

$\Pr(p = 0.9) = 1 - w$. The known discrete demand distribution is given by the following.

Demand ξ	0	1	2	3
Probability $f(\xi)$	0.1	0.3	0.1	0.5

Figure 2 shows the Bayesian optimal inventory level with observed and unobserved lost sales as well as the myopic inventory level in the first period where the prior belief parameter w varies from 0 to 0.4. From this example, we see that the Bayesian optimal inventory level with unobserved lost sales is greater than or equal to the observed lost-sales case in all w values. However, the Bayesian inventory with unobserved lost sales can be either greater than (when $w = 0.14$) or less than (when $w = 0.30$) the myopic inventory level. One would stock more with unobserved lost sales in this example because of the need for reducing ambiguity about the number of customers who arrive after stockout. In particular, in the event that the number sold equals the inventory level, one cannot tell whether the realized demand was perfectly matched to the inventory level or whether it exceeded the inventory level but no customers were willing to substitute. In the former case, one did not learn anything about the customer substitution probability, and thus the posterior on p remains the same as the prior, while in the latter case, the posterior is shifted downward because one learns that customers are not likely to substitute. Without observing lost sales, one cannot discern which of these cases occurred, and thus it may be profitable to increase the inventory level to separate them—especially when the probability is high for the critical point where demand realization equals the myopic inventory level. From the above example, we make the following remark:

Remark 1. Suppose that the demand distribution is known but the substitution probability is not known. If lost sales are not observed, in general, the Bayesian optimal inventory level may be either greater than or less than the myopic inventory level.

For the case when the demand parameter θ and substitution probability p are unknown, analysis of the unobserved lost-sales case is generally intractable. However, the two-period model with exponential demand distribution introduced in Proposition 3 of §2.2.1 is tractable. For this special case, we can identify the conditions under which the Bayesian optimal inventory level is greater than the myopic inventory level.

Proposition 5 (Chen 2005). *Suppose that the demand distribution is exponential and the product is perishable. For a two-period problem with censored demand observation and unknown $p \in \{0, 1\}$, depending on the prior uncertainty and cost parameters, either of the following two claims hold exclusively*:

(a) *For any prior probability $w \in [0, 1)$ for p, the Bayesian optimal inventory level (with updating of both demand parameter θ and substitution probability p) is greater than the myopic inventory level, i.e., $y_1^{U2} > y_1^M$.*

(b) *There exists a prior probability $w^* \in (0, 1)$ for p, subject to $y_1^{U2} > y_1^M$ when $0 \le w < w^*$, and $y_1^{U2} < y_1^M$ when $w^* < w < 1$. Furthermore, w^* is decreasing in m (the contribution margin of the substitute product).*

The proof is in the first author's unpublished PhD thesis (Chen 2005). Proposition 5 says that two scenarios can happen in this special setting: (a) The Bayesian inventory manager will stock more than the myopic level to learn about demand, regardless of her beliefs about the substitution probability (the "stock more" effect dominates the "stock less" effect); or (b) the Bayesian inventory manager may stock more or less, depending on the prior belief of the substitution probability. If $p = 0$ is likely, she stocks more because demand is more likely to be censored, and vice versa. Furthermore, as the contribution margin on the substitute increases, the Bayesian inventory manager will tend to stock less than the myopic inventory level because learning the substitution probability will increase expected profit in the following period.

In summary, when lost sales are not observed, the sign and magnitude of the deviation of the Bayesian optimal inventory level from the myopic inventory level depend on the value of active learning (i.e., the expected gain in future profit obtained by deviating from the myopic inventory level). This, in turn, depends on the prior uncertainty and cost structure of the underlying model. For instance, when the product is perishable, if the stockout penalty is high, the stock more effect tends to dominate the stock less effect because with high service level the chance of stockout is low; hence, the benefit of estimating the substitution probability becomes negligible.

3. Maximum Likelihood Estimation for a Capacitated System

In this section, we consider a system with constrained capacity, operated according to a base-stock policy. Customers arrive according to a Poisson process with rate λ. Let $Z(t)$ denote the inventory level in the system at time t. If $Z(t)$ is positive, the arriving customer purchases one unit of the product. Otherwise, with probability p the customer will join the queue to purchase the product at a later time, and with probability $1-p$ he or she will "balk" or leave the system without making a purchase. The inventory level is reviewed in continuous time, and when it falls below the base-stock level B, production starts. The production lead times are i.i.d. exponential with rate μ. Therefore, the system is essentially an $M/M/1$ make-to-stock queue with state-dependent arrival rates: It is λ if $Z(t) > 0$ and λp otherwise. The process $Z(t)$ is a continuous-time Markov chain with a state space of $\{\ldots, -2, -1, 0, 1, \ldots, B\}$. We denote the process in the steady state by $Z(\infty)$ and define $\rho = \lambda/\mu$.

We begin by showing how operating costs and the optimal base-stock level vary with the probability p that customers will wait. The proofs of these results are elementary; hence, they are omitted.

Proposition 6. *The steady-state distribution of $Z(t)$ is given by*

$$P(Z(\infty)=z) = \begin{cases} \dfrac{(1-\rho)(1-\rho p)}{1-\rho p-(1-p)\rho^{B+1}}\rho^{B-z} & \text{if } 0 \le z \le B, \\[2ex] \dfrac{(1-\rho)(1-\rho p)}{1-\rho p-(1-p)\rho^{B+1}}\rho^{B-z}p^{-z} & \text{if } z \le 0. \end{cases}$$

The steady-state probability that the product is out of stock is

$$P(Z(\infty) \le 0) = \frac{(1-\rho)\rho^B}{1-\rho p-(1-p)\rho^{B+1}}.$$

Furthermore, $P(Z(\infty) \le 0)$ is decreasing in the base-stock level B and increasing in the load factor ρ.

Let h be the holding cost rate and r the foregone revenue on each of the lost sales. The system's long-run average cost rate is given by

$$\begin{aligned} C(B,p) &= hE(Z(\infty)^+) + r\lambda(1-p)P(Z(\infty) \le 0) \\ &= h\frac{(1-\rho p)(\rho^{B+1}+B-(B+1)\rho)}{(1-\rho)(1-\rho p-(1-p)\rho^{B+1})} \\ &\quad + r\frac{\lambda(1-p)(1-\rho)\rho^B}{1-\rho p-(1-p)\rho^{B+1}}. \end{aligned}$$

The following proposition gives the optimal base-stock level $B^*(p)$ as a function of the waiting probability p:

Proposition 7. *The optimal base-stock level $B^*(p) = \arg\min_{B\in\{0,1,2,\ldots\}}[C(B,p)]$ is characterized by*

(1) $g(B^*) \le r(\mu-\lambda)/h$, *and*

(2) $g(B^*+1) \ge r(\mu-\lambda)/h$,

where $g(B) = (1-\rho p)(\rho^{-B}-1)/(1-p)(1-\rho) - B$.

Furthermore, $B^*(p)$ is decreasing in p, and $C(B^*(p),p)$ is decreasing in p.

Bayesian optimization is intractable in this continuous-review queueing model. Instead, we derive MLEs and show how their rate of convergence varies with the base-stock level. Qualitatively, by modifying the base-stock level to increase the rate of convergence of the MLEs, the system manager rapidly "learns" about the demand rate and the probability that a customer is willing to wait.

3.1. Unobserved Lost-Sales Case

Suppose that when the product is out of stock, the manager observes customers that join the queue, but does not observe customers that balk. Let $\tau_1(t)$ be the amount of time the product is out of stock during $[0, t]$, and $N_1(t)$ the number of customers that join the queue and wait while the product is out of stock during $[0, t]$. Finally, let $\Rightarrow$ denote weak convergence. Proofs for Propositions 8 and 9 are given in the appendix.

Proposition 8. *If λ is known, then the MLE for p is given by $\hat{p}(t) = N_1(t)/\lambda\tau_1(t)$. Furthermore, $t^{1/2}(\hat{p}(t) - p) \Rightarrow \sqrt{p/\lambda P(Z(\infty) \le 0)} \cdot N(0, 1)$ as $t \to \infty$, where $N(0, 1)$ is the standard normal distribution.*

Proposition 8 establishes that the accuracy of the MLE $\hat{p}(t)$ is increasing in $P(Z(\infty) \le 0)$. Because $P(Z(\infty) \le 0)$ is decreasing in B and increasing in ρ, as shown in Proposition 6, we conclude that the accuracy of $\hat{p}(t)$ is decreasing in B and increasing in ρ. In other words, lowering the initial base-stock level B or reducing the production rate μ will lead to more accurate estimation of p. Hence, the manager should set the base-stock level below $B^*(p)$ during an experimenting period to increase the MLE accuracy.

When λ is unknown, one can derive similar MLEs and the weak convergence results. Let $N_2(t)$ denote the number of arrivals that occur while the product is in stock during $[0, t]$, and let $\tau_2(t)$ denote the amount of time that the product is in stock during $[0, t]$.

Proposition 9. *The MLEs for λ and p are given by*

$$\hat{\lambda}(t) = N_2(t)/\tau_2(t) \quad \textit{and} \quad \hat{p}(t) = N_1(t)\tau_2(t)/N_2(t)\tau_1(t),$$

respectively. Furthermore, as $t \to \infty$,

$$t^{1/2}(\hat{\lambda}(t) - \lambda) \Rightarrow \sqrt{\frac{\lambda}{P(Z(\infty) > 0)}} N(0, 1) \quad \textit{and}$$

$$t^{1/2}(\hat{p}(t) - p) \Rightarrow \sqrt{\frac{p}{\lambda P(Z(\infty) \le 0)} + \frac{p^2}{\lambda P(Z(\infty) > 0)}} N(0, 1).$$

Proposition 9 establishes that increasing the base-stock level B (or the production rate μ) will increase the accuracy of $\hat{\lambda}(t)$, which is analogous to the stock more result in Theorem 1. However, the accuracy of $\hat{p}(t)$ is not monotonic in the base-stock level B (or the production rate μ). For small B (or μ), increasing B (or μ) increases the accuracy of $\hat{p}(t)$ (by improving the estimate of λ). As B (or μ) becomes very large, $\hat{p}(t)$ becomes less accurate because little waiting behavior is observed. This resembles the insights we obtained in §2.3.1.

3.2. Observed Lost-Sales Case

Suppose that all customer arrivals are observed. Then, the MLEs for λ and p become $\hat{\lambda}(t) = A(t)/t$ and $\hat{p}(t) = N_1(t)/A_1(t)$, where $A(t)$ is the number of customers that arrive during $[0, t]$, and $A_1(t)$ is the number of customers that arrive while the product is out of stock and must decide whether to wait or balk. Clearly, the choice of base-stock level does not affect the maximum likelihood estimation of λ. Furthermore, reducing the base-stock level B improves the accuracy of $\hat{p}(t)$ by increasing the number of observations of the decision to wait or balk. This reinforces the stock less result obtained in Theorem 2 of §2.3.

4. Concluding Remarks

A well-known result in Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002) is that for a perishable product with unobserved lost sales, one should *stock more* when learning about the demand distribution. In contrast, we show that the Bayesian optimal inventory level may be lower than the myopic inventory level (i.e., in some cases one should *stock less*) to learn about the substitution probability, or for a nonperishable product. We also prove that making lost sales unobservable increases the Bayesian optimal inventory level; in this specific sense, the famous stock more result of Harpaz et al. (1982), Lariviere and Porteus (1999), and Ding et al. (2002) generalizes to the case of nonperishable inventory.

Our results are for a single product, but serve as a building block toward analysis of systems with multiple, interacting products. In general, customers may substitute among multiple alternatives when their top choice is not available, and the inventory manager must estimate the probability of substitution among multiple products. Retailers will soon be able to use RFID to continuously observe when a product is out of stock. If the consequent decisions by customers to backorder or substitute an alternative product can be captured, both Theorem 2 and Proposition 8 suggest that with this visibility retailers should lower inventory levels to learn more about whether customers will wait or substitute an alternative product. Insofar as the decision of customers to make substitutions can be observed, our general Bayesian analysis for the single-item problem can be extended to problems involving dynamic estimation of multiple substitution probabilities in consumer choice models such as the multinomial logit model. Effective heuristics are needed in practice. For example, Kok and Fisher (2007) propose a heuristic for estimating multi-item demand and substitution and optimizing the product assortment; in a supermarket application, this heuristic achieves a 50% profit improvement.

Our results are most insightful for inventory and capacity management for an innovative product like the Toyota Prius hybrid. Critically, operations managers must estimate customers' willingness to wait as well as the demand distribution. Toyota has profited from maintaining low capacity and inventory levels and learning that customers will wait in excess of six to eight months for the Prius. Interestingly, Toyota has allocated car inventory within the United States such that customers must wait longer on the coasts than in other parts of the country. This likely reflects regional differences in customers' willingness to wait (Garennes 2004, Nauman 2004).

Acknowledgments

The authors thank an anonymous senior editor and two referees as well as Mor Armony, Peter Glynn, Robert Phillips, and Evan Porteus for helpful insights and suggestions. This research was supported by Nth Orbit and by the National Science Foundation under Grant DMI-0239840.

Appendix

Lemma A1. *For any given $\theta \in \Theta$, $p \in [0,1]$, and $y \ge 0$, $f_X^y(x \mid \theta, p)$ and $f_X^{y+1}(x \mid \theta, p)$ satisfy the following relationship*:

$$f_X^y(x \mid \theta, p) = \begin{cases} f_X^{y+1}(x \mid \theta, p) & \text{for } x < y, \\ f_X^{y+1}(x \mid \theta, p) + (1-p) \cdot f_X^{y+1}(x+1 \mid \theta, p), & \text{for } x = y, \\ p \cdot f_X^{y+1}(x \mid \theta, p) + (1-p) \cdot f_X^{y+1}(x+1 \mid \theta, p), & \text{for } x > y. \end{cases}$$

Proof. By definition (2) and some elementary algebraic manipulations. □

Proof of Proposition 1. We prove the case of $p > 0$; the case of $p = 0$ can be easily verified following the same steps. Because p is known, we write $f_X^y(x \mid \theta)$ as shorthand for $f_X^y(x \mid \theta, p)$. The sequence of predictive sales densities $\{f_{X_i}^{y_i}(x), i = 1, 2, \ldots, N\}$ satisfies

$$f_{X_i}^{y_i}(x) = \begin{cases} \int_\Theta f_X^{y_i}(x \mid \theta) \pi_1'(\theta)\, d\theta & i = 1, \\ \int_\Theta f_X^{y_i}(x \mid \theta) \pi_i'(\theta \mid x_{i-1}, y_{i-1}, \pi_{i-1}'(\theta))\, d\theta & i = 2, \ldots, N. \end{cases}$$

By backward induction, using Lemma A1, we first show that the following holds: for $i = 2, \ldots, N$, $z \ge 0$, and any $\pi'(\theta)$, if $x = y$,

$$\begin{aligned} &E_{\pi'(\theta)}\{v_i^U(z, \pi_i'(\theta \mid x, y, \pi'(\theta))) f_X^y(x \mid \theta)\} \\ &\quad \le E_{\pi'(\theta)}\{v_i^U(z, \pi_i'(\theta \mid x, y+1, \pi'(\theta))) f_X^{y+1}(x \mid \theta) \\ &\qquad + (1-p) \cdot v_i^U(z, \pi_i'(\theta \mid x+1, y+1, \pi'(\theta))) \\ &\qquad \cdot f_X^{y+1}(x+1 \mid \theta)\}; \end{aligned} \tag{14}$$

and if $x > y$,

$$\begin{aligned} &E_{\pi'(\theta)}\{v_i^U(z, \pi_i'(\theta \mid x, y, \pi'(\theta))) f_X^y(x \mid \theta)\} \\ &\quad \le E_{\pi'(\theta)}\{p \cdot v_i^U(z, \pi_i'(\theta \mid x, y+1, \pi'(\theta))) f_X^{y+1}(x \mid \theta) \\ &\qquad + (1-p) \cdot v_i^U(z, \pi_i'(\theta \mid x+1, y+1, \pi'(\theta))) \\ &\qquad \cdot f_X^{y+1}(x+1 \mid \theta)\}. \end{aligned} \tag{15}$$

For case $i = N$, we verify (15) as follows.

$$\begin{aligned} &E_{\pi'(\theta)}\{v_N^U(z, \pi_i'(\theta \mid x, y, \pi'(\theta))) f_X^y(x \mid \theta)\} \\ &= \max_{y' \ge z} \left\{ \int_\Theta R_N(y', \theta) f_X^y(x \mid \theta) \pi'(\theta)\, d\theta \right\} \\ &= \max_{y' \ge z} \left\{ \int_\Theta R_N(y', \theta) \{p \cdot f_X^{y+1}(x \mid \theta) + (1-p) \cdot f_X^{y+1}(x+1 \mid \theta)\} \pi'(\theta)\, d\theta \right\} \\ &\le p \cdot \max_{y' \ge z} \left\{ \int_\Theta R_N(y', \theta) f_X^{y+1}(x \mid \theta) \pi'(\theta)\, d\theta \right\} \\ &\quad + (1-p) \cdot \max_{y' \ge z} \left\{ \int_\Theta R_N(y', \theta) f_X^{y+1}(x+1 \mid \theta) \pi'(\theta)\, d\theta \right\} \\ &= E_{\pi'(\theta)}\{p \cdot v_i^U(z, \pi_i'(\theta \mid x, y+1, \pi'(\theta))) f_X^{y+1}(x \mid \theta) \\ &\qquad + (1-p) \cdot v_i^U(z, \pi_i'(\theta \mid x+1, y+1, \pi'(\theta))) \\ &\qquad \cdot f_X^{y+1}(x+1 \mid \theta)\}. \end{aligned} \tag{16}$$

where (16) follows from identities from Lemma A1. Now, assume the result holds for case $i+1$. We check for case i:

$$\begin{aligned} &E_{\pi'(\theta)}\{v_i^U(z, \pi_i'(\theta \mid x, y, \pi'(\theta))) f_X^y(x \mid \theta)\} \\ &= \max_{y' \ge z} \left\{ \int_\Theta \left\{ R(y', \theta) + \delta \sum_{x'=0}^{\infty} v_{i+1}^U((y'-x')^+, \pi_{i+1}'(\theta \mid x', y' \mid x, y, \pi')) \cdot f_X^{y'}(x' \mid \theta) \right\} f_X^y(x \mid \theta) \pi'(\theta)\, d\theta \right\} \end{aligned}$$

$$
\begin{aligned}
&= \max_{y' \ge z} \bigg\{ \int_\Theta R(y', \theta) f_X^y(x \mid \theta) \pi'(\theta)\, d\theta \\
&\qquad + \delta \sum_{x'=0}^{\infty} \int_\Theta v_{i+1}^U\big((y'-x')^+, \pi'_{i+1}(\theta \mid x', y' \mid x, y, \pi')\big) \\
&\qquad\qquad \cdot f_X^{y'}(x' \mid \theta) f_X^y(x \mid \theta) \pi'(\theta)\, d\theta \bigg\} \\
&= \max_{y' \ge z} \bigg\{ \int_\Theta R(y', \theta) f_X^y(x \mid \theta) \pi'(\theta)\, d\theta + \delta \sum_{x'=0}^{\infty} E_{\pi'_i(\theta \mid y', x', \pi')} \\
&\qquad \cdot \big\{ v_{i+1}^U((y'-x')^+, \pi'_{i+1}(\theta \mid x, y, \pi'_i)) f_X^y(x \mid \theta) \big\} \bigg\} \\
&\le \max_{y' \ge z} \bigg\{ \int_\Theta R(y', \theta) f_X^y(x \mid \theta) \pi'(\theta)\, d\theta + \delta \sum_{x'=0}^{\infty} E_{\pi'_i(\theta \mid y', x', \pi')} \\
&\qquad \cdot \big\{ p \cdot v_{i+1}^U((y'-x')^+, \pi'_{i+1}(\theta \mid x, y+1, \pi'_i)) f_X^{y+1}(x \mid \theta) \\
&\qquad + (1-p) \cdot v_{i+1}^U((y'-x')^+, \pi'_{i+1}(\theta \mid x+1, y+1, \pi'_i)) \\
&\qquad\qquad \cdot f_X^{y+1}(x+1 \mid \theta) \big\} \bigg\} \\
&= \max_{y' \ge z} \bigg\{ \int_\Theta R(y', \theta) \big\{ p \cdot f_X^{y+1}(x \mid \theta) + (1-p) \cdot f_X^{y+1}(x+1 \mid \theta) \big\} \pi'(\theta)\, d\theta \\
&\qquad + \delta \sum_{x'=0}^{\infty} p \int_\Theta v_{i+1}^U\big((y'-x')^+, \pi'_{i+1}(\theta \mid x', y' \mid x, y+1, \pi')\big) \\
&\qquad\qquad \cdot f_X^{y'}(x' \mid \theta) f_X^{y+1}(x \mid \theta) \pi'(\theta)\, d\theta \\
&\qquad + \delta \sum_{x'=0}^{\infty} (1-p) \int_\Theta v_{i+1}^U\big((y'-x')^+, \\
&\qquad\qquad \pi'_{i+1}(\theta \mid x', y' \mid x+1, y+1, \pi')\big) \\
&\qquad\qquad \times f_X^{y'}(x' \mid \theta) f_X^{y+1}(x+1 \mid \theta) \pi'(\theta)\, d\theta \bigg\} \\
&\le p \cdot \max_{y' \ge z} \bigg\{ \int_\Theta \bigg\{ R(y', \theta) \\
&\qquad + \delta \sum_{x'=0}^{\infty} v_{i+1}^U\big((y'-x')^+, \pi'_{i+1}(\theta \mid x', y' \mid x, y+1, \pi')\big) \\
&\qquad\qquad \cdot f_X^{y'}(x' \mid \theta) \bigg\} \cdot f_X^{y+1}(x \mid \theta) \pi'(\theta)\, d\theta \bigg\} \\
&\quad + (1-p) \cdot \max_{y' \ge z} \bigg\{ \int_\Theta \bigg\{ R(y', \theta) \\
&\qquad + \delta \sum_{x'=0}^{\infty} v_{i+1}^U\big((y'-x')^+, \pi'_{i+1}(\theta \mid x', y' \mid x+1, y+1, \pi')\big) \\
&\qquad\qquad \cdot f_X^{y'}(x' \mid \theta) \bigg\} \cdot f_X^{y+1}(x+1 \mid \theta) \pi'(\theta)\, d\theta \bigg\} \\
&= E_{\pi'(\theta)} \big\{ p \cdot v_i^U(z, \pi'_i(\theta \mid x, y+1, \pi'(\theta))) f_X^{y+1}(x \mid \theta) \\
&\qquad + (1-p) \cdot v_i^U\big(z, \pi'_i(\theta \mid x+1, y+1, \pi'(\theta))\big) \cdot f_X^{y+1}(x+1 \mid \theta) \big\}.
\end{aligned} \tag{17}
$$

Note that inequality (17) follows from induction assumption. Hence, (15) holds. Similarly, we can show that (14) also holds.

Apply (14) and (15) to the right-hand side of (9), we have

$$
\begin{aligned}
&E_{\pi'(\theta)} \bigg\{ \sum_{x=y}^{\infty} v_i^U(0, \pi'_i(\theta \mid x, y)) f_X^y(x \mid \theta) \bigg\} \\
&\le E_{\pi'(\theta)} \bigg\{ v_i^U(0, \pi'_i(\theta \mid y, y+1)) f_X^{y+1}(y \mid \theta) + (1-p) \\
&\qquad \cdot v_i^U(0, \pi'_i(\theta \mid y+1, y+1)) f_X^{y+1}(y+1 \mid \theta) \\
&\qquad + \sum_{x=y+1}^{\infty} \big\{ p \cdot v_i^U(0, \pi'_i(\theta \mid x, y+1)) f_X^{y+1}(x \mid \theta) + (1-p) \\
&\qquad\qquad \cdot v_i^U(0, \pi'_i(\theta \mid x+1, y+1)) \cdot f_X^{y+1}(x+1 \mid \theta) \big\} \bigg\} \\
&= E_{\pi'(\theta)} \bigg\{ \sum_{x=y}^{\infty} v_i^U(0, \pi'_i(\theta \mid x, y+1)) f_X^{y+1}(x \mid \theta) \bigg\}.
\end{aligned}
$$

Hence, inequality (9) holds. □

Lemma A2. *For $i = 2, \ldots, N$, $n \ge 0$, and any $\pi'(p)$ the following holds*:

$$
\begin{aligned}
&E_{\pi'(p)} E_{K \mid n+1} \big\{ v_i^{O2}(z, \pi'_i(\theta) \cdot \pi'_i(p \mid n+1, K, \pi')) \big\} \\
&= E_{\pi'(p)}\{p\} \cdot E_{\pi''(p)} E_{K \mid n} \big\{ v_i^{O2}(z, \pi'_i(\theta) \cdot \pi'_i(p \mid n, K, \pi'')) \big\} \\
&\quad + E_{\pi'(p)}\{1-p\} \cdot E_{\pi'''(p)} E_{K \mid n} \big\{ v_i^{O2}\big(z, \pi'_i(\theta) \cdot \pi'_i(p \mid n, K, \pi''')\big) \big\},
\end{aligned}
$$

where $\pi''(p) = p\pi'(p) / \int p\pi'(p)\, dp$, *and* $\pi'''(p) = (1-p)\pi'(p) / \int (1-p)\pi'(p)\, dp$.

Proof. The right-hand side of the equation can be expressed as

$$
\begin{aligned}
&\int p\pi'(p)\, dp \cdot \sum_{k=0}^{n} v_i^{O2}\bigg(z, \pi'_i(\theta) \frac{p^{k+1}(1-p)^{n-k}\pi'(p)}{p^{k+1}(1-p)^{n-k}\pi'(p) dp}\bigg) \\
&\qquad \cdot \frac{\int \binom{n}{k} p^{k+1}(1-p)^{n-k}\pi'(p)\, dp}{\int p\pi'(p)\, dp} + \int (1-p)\pi'(p)\, dp \\
&\qquad \cdot \sum_{k=0}^{n} v_i^{O2}\bigg(z, \pi'_i(\theta) \frac{p^k(1-p)^{n-k+1}\pi'(p)}{\int p^k(1-p)^{n-k+1}\pi'(p)\, dp}\bigg) \\
&\qquad \cdot \frac{\int \binom{n}{k} p^k(1-p)^{n-k+1}\pi'(p)\, dp}{\int (1-p)\pi'(p)\, dp} \\
&= \sum_{k=0}^{n} v_i^{O2}\bigg(z, \pi'_i(\theta) \frac{p^{k+1}(1-p)^{n-k}\pi'(p)}{\int p^{k+1}(1-p)^{n-k}\pi'(p)\, dp}\bigg) \\
&\qquad \cdot \int \binom{n}{k} p^{k+1}(1-p)^{n-k}\pi'(p)\, dp \\
&\quad + \sum_{k=0}^{n} v_i^{O2}\bigg(z, \pi'_i(\theta) \frac{p^k(1-p)^{n-k+1}\pi'(p)}{\int p^k(1-p)^{n-k+1}\pi'(p)\, dp}\bigg) \\
&\qquad \cdot \int \binom{n}{k} p^k(1-p)^{n-k+1}\pi'(p)\, dp \\
&= \sum_{k=1}^{n+1} v_i^{O2}\bigg(z, \pi'_i(\theta) \frac{p^k(1-p)^{n-k+1}\pi'(p)}{\int p^k(1-p)^{n-k+1}\pi'(p)\, dp}\bigg)
\end{aligned}
$$

$$
\begin{aligned}
&\cdot \int \binom{n}{k-1} p^k (1-p)^{n-k+1} \pi'(p)\, dp \\
&+ \sum_{k=0}^{n} v_i^{O2}\left(z,\, \pi_i'(\theta) \frac{p^k (1-p)^{n-k+1}\pi'(p)}{\int p^k (1-p)^{n-k+1}\pi'(p)\, dp}\right) \\
&\cdot \int \binom{n}{k} p^k (1-p)^{n-k+1} \pi'(p)\, dp \\
&= \sum_{k=0}^{n+1} v_i^{O2}\left(z,\, \pi_i'(\theta) \frac{p^k (1-p)^{n-k+1}\pi'(p)}{\int p^k (1-p)^{n-k+1}\pi'(p)\, dp}\right) \\
&\cdot \int \binom{n+1}{k} p^k (1-p)^{n-k+1} \pi'(p)\, dp \\
&= E_{\pi'(p)} E_{K\,|\,n+1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n+1, K, \pi'))\},
\end{aligned}
$$

where the identity $\binom{n+1}{k} = \binom{n}{k-1} + \binom{n}{k}$ is used in the second-to-last equality. □

Lemma A3. *For $i = 2, \ldots, N$, $n \geq 0$, and any $\pi'(p)$, if*

$$
E_{\pi'(p)} E_{K|1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,1, K, \pi'))\} \geq v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi'),
$$

then

$$
\begin{aligned}
&E_{\pi'(p)} E_{K\,|\,n+1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n+1, K, \pi'))\} \\
&\quad \geq E_{\pi'(p)} E_{K\,|\,n}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n, K, \pi'))\}.
\end{aligned}
$$

Proof. By induction, the case of $n = 0$ holds by the lemma condition. Now, assume it is true for case $n-1$. By Lemma A2, we have

$$
\begin{aligned}
&E_{\pi'(p)} E_{K\,|\,n+1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n+1, K, \pi'))\} \\
&= E_{\pi'(p)}\{p\} \cdot E_{\pi''(p)} E_{K\,|\,n}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n, K, \pi''))\} \\
&\quad + E_{\pi'(p)}\{1-p\} \cdot E_{\pi'''(p)} E_{K\,|\,n}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n, K, \pi'''))\} \\
&\geq E_{\pi'(p)}\{p\} \cdot E_{\pi''(p)} E_{K\,|\,n-1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n-1, K, \pi''))\} \\
&\quad + E_{\pi'(p)}\{1-p\} \cdot E_{\pi'''(p)} E_{K|n-1}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p|n-1, K, \pi'''))\} \\
&= E_{\pi'(p)} E_{K\,|\,n}\{v_i^{O2}(z,\, \pi_i'(\theta)\cdot \pi_i'(p\,|\,n, K, \pi'))\}.
\end{aligned}
$$

The inequality follows from the induction assumption and the last equality follows from Lemma A2. □

Proof of Proposition 4. The sequence of predictive demand densities $\{f_i(\xi),\ i = 1, 2, \ldots, N\}$ satisfies

$$
f_i(\xi) = \begin{cases} \int_\Theta f(\xi\,|\,\theta)\pi_1'(\theta)\, d\theta & i = 1, \\ \int_\Theta f(\xi\,|\,\theta)\pi_i'(\theta\,|\,\xi_{i-1}, \pi_{i-1}'(\theta))\, d\theta & i = 2, \ldots, N. \end{cases}
$$

Now prove by backward induction. It is straightforward to verify that

$$
\begin{aligned}
&E_{\pi_{N-1}'(p)} E_{K\,|\,1}\{v_N^{O2}(z,\, \pi_N'(\theta)\cdot \pi_N'(p\,|\,1, K, \pi_{N-1}'))\} \\
&\quad \geq v_N^{O2}(z,\, \pi_N'(\theta)\cdot \pi_{N-1}').
\end{aligned}
$$

Hence, by Lemma A3, we have

$$
\begin{aligned}
&E_{\pi_{N-1}'(p)} E_{K\,|\,n+1}\{v_N^{O2}(z,\, \pi_N'(\theta)\cdot \pi_N'(p\,|\,n+1, K, \pi_{N-1}'))\} \\
&\quad \geq E_{\pi_{N-1}'(p)} E_{K\,|\,n}\{v_N^{O2}(z,\, \pi_N'(\theta)\cdot \pi_N'(p\,|\,n, K, \pi_{N-1}'))\}.
\end{aligned}
$$

Now, assume it is true for the case $i+1$. By Lemma A2, we have

$$
\begin{aligned}
&E_{\pi_{i-1}'(p)} E_{K|1}\{v_i^{O2}(z, \pi_i'(\theta)\cdot \pi_i'(p|1, K, \pi_{i-1}'))\} \\
&= \int p\pi_{i-1}'(p)\, dp \cdot \max_{y\geq z} E_{\pi_i'(\theta)\pi_i'(p|1,1,\pi_{i-1}')} \\
&\quad \cdot \Big\{R(y\,|\,\theta, p) + \delta \sum_{\xi=0}^{y} v_{i+1}^{O2}(y-\xi,\, \pi_{i+1}'(\theta)\cdot \pi_i'(p)) f(\xi\,|\,\theta) \\
&\qquad + \delta \sum_{\xi=y+1}^{\infty} E_{K|\xi-y-1}\{v_{i+1}^{O2}(0,\, \pi_{i+1}'(\theta) \\
&\qquad\qquad \cdot \pi_{i+1}'(p\,|\,\xi-y-1, K, \pi_i'))\} f(\xi\,|\,\theta)\Big\} \\
&\quad + \int (1-p)\pi_{i-1}'(p)\, dp \cdot \max_{y\geq z} E_{\pi_i'(\theta)\pi_i'}(p|1, 0, \pi_{i-1}') \\
&\quad \cdot \Big\{R(y\,|\,\theta, p) + \delta \sum_{\xi=0}^{y} v_{i+1}^{O2}(y-\xi,\, \pi_{i+1}'(\theta)\cdot \pi_i'(p)) f(\xi\,|\,\theta) \\
&\qquad + \delta \sum_{\xi=y+1}^{\infty} E_{K|\xi-y-1}\{v_{i+1}^{O2}(0,\, \pi_{i+1}'(\theta) \\
&\qquad\qquad \cdot \pi_{i+1}'(p\,|\,\xi-y-1, K, \pi_i'))\} f(\xi\,|\,\theta)\Big\} \\
&\geq \max_{y\geq z}\Big\{E_{\pi_i'(\theta)\pi_{i-1}'(p)}\{R(y\,|\,\theta, p)\} \\
&\qquad + \delta \sum_{\xi=0}^{y} E_{\pi_{i-1}'(p)} E_{K|1}]\{v_{i+1}^{O2}(y-\xi,\, \pi_{i+1}'(\theta)\cdot \pi_i'(p))\} f_i(\xi) \\
&\qquad + \delta \sum_{\xi=y+1}^{\infty} \Big\{\int p\pi_{i-1}'(p)\, dp \times E_{\pi_i'\left(p|1,1,\pi_{i-1}'\right)} E_{K|\xi-y-1} \\
&\qquad\qquad \cdot \{v_{i+1}^{O2}(z,\, \pi_{i+1}'(\theta)\cdot \pi_{i+1}'(p|\xi-y-1, K, \pi_i'))\} \\
&\qquad\qquad + \int (1-p)\pi_{i-1}'(p)\, dp \times E_{\pi_i'(p|1,0,\pi_{i-1}')} E_{K|\xi-y-1} \\
&\qquad\quad \cdot \{v_{i+1}^{O2}(z,\, \pi_{i+1}'(\theta)\cdot \pi_{i+1}'(p\,|\,\xi-y-1, K, \pi_i'))\}\Big\} f_i(\xi)\Big\} \\
&= \max_{y\geq z}\Big\{E_{\pi_i'(\theta)\pi_{i-1}'(p)}\{R(y\,|\,\theta, p)\} + \delta \sum_{\xi=0}^{y} E_{\pi_{i-1}'(p)} E_{K|1} \\
&\qquad \cdot \{v_{i+1}^{O2}(y-\xi,\, \pi_{i+1}'(\theta)\cdot \pi_i'(p))\} f_i(\xi) + \delta \sum_{\xi=y+1}^{\infty} E_{\pi_{i-1}'(p)} \\
&\qquad \cdot E_{K|\xi-y}\{v_{i+1}^{O2}(z,\, \pi_{i+1}'(\theta)\cdot \pi_{i+1}'(p\,|\,\xi-y, K, \pi_{i-1}'))\} f_i(\xi)\Big\}
\end{aligned}
$$

$$\geq \max_{y\geq z}\bigg\{E_{\pi_i'(\theta)\pi_{i-1}'(p)}\{R(y\mid\theta,p)\} + \delta\sum_{\xi=0}^{y} E_{\pi_{i-1}'(p)}\cdot\{v_{i+1}^{O2}(y-\xi,\pi_{i+1}'(\theta)\cdot\pi_{i-1}'(p))\}f_i(\xi) + \delta\sum_{\xi=y+1}^{\infty} E_{\pi_{i-1}'(p)}E_{K\mid\xi-y-1}\{v_{i+1}^{O2}(z,\pi_{i+1}'(\theta) \cdot\pi_{i+1}'(p\mid\xi-y-1,K,\pi_{i-1}'))\}f_i(\xi)\bigg\}$$

$$=\max_{y\geq z}E_{\pi_i'(\theta)\pi_{i-1}'(p)}\bigg\{R(y\mid\theta,p)+\delta\sum_{\xi=0}^{y}v_{i+1}^{O2}(y-\xi,\pi_{i+1}'(\theta)\cdot\pi_{i-1}'(p)) \cdot f(\xi\mid\theta)+\delta\sum_{\xi=y+1}^{\infty}E_{K\mid\xi-y-1}\{v_{i+1}^{O2}(z,\pi_{i+1}'(\theta) \cdot\pi_{i+1}'(p\mid\xi-y-1,K,\pi_{i-1}'))\}f(\xi\mid\theta)\bigg\}$$

$$=v_i^{O2}(z,\pi_i'(\theta)\cdot\pi_{i-1}'(p)).$$

The last inequality follows from the induction assumption and the equality right before that is by Lemma A2. Applying Lemma A3 again, we conclude that for $n\geq 0$,

$$E_{\pi_{i-1}'(p)}E_{K\mid n+1}\{v_i^{O2}(z,\pi_i'(\theta)\cdot\pi_i'(p\mid n+1,K,\pi_{i-1}'))\} \geq E_{\pi_{i-1}'(p)}E_{K\mid n}\{v_i^{O2}(z,\pi_i'(\theta)\cdot\pi_i'(p\mid n,K,\pi_{i-1}'))\}. \quad \square$$

Proof of Proposition 8. $\hat{p}(t)=N_1(t)/\lambda\tau_1(t)$ can be obtained by solving the first-order condition of the likelihood function (given that λ is known) over period $[0,t]$. We focus on showing the weak convergence result of $\hat{p}(t)$. First, note that $t^{1/2}(N_1(t)/\lambda\tau_1(t)-p)=t^{-1/2}(N_1(t)-\lambda p\tau_1(t))/\lambda(\tau_1(t)/t)$. Because $\lim_{t\to\infty}\tau_1(t)/t=P(Z(\infty)\leq 0)$ almost surely (a.s.), it suffices to show that $t^{-1/2}(N_1(t)-\lambda p\tau_1(t))\Rightarrow\sqrt{\lambda pP(Z(\infty)\leq 0)}\cdot N(0,1)$, as $t\to\infty$.

Define $M(t)=N_1(t)-\lambda p\tau_1(t)=N_1(t)-\int_0^t\lambda p\mathbf{1}\{Z(s)\leq s\}\,ds$, which is a martingale. Scale $M(t)$ as $M_\epsilon(t)=\epsilon^{1/2}(N_1(t/\epsilon)-\int_0^{t/\epsilon}\lambda p\mathbf{1}\{Z(s)\leq s\}\,ds)$. Now we show that the quadratic variation of $M_\epsilon(t)$, denoted by $[M_\epsilon,M_\epsilon](t)$, converges to $\lambda pP(Z(\infty)\leq 0)$ as $\epsilon\to 0$. Note that

$$[M_\epsilon,M_\epsilon](t)=\lim_{m\to\infty}\sum_{j=1}^{m}\big(M_\epsilon(jt/m)-M_\epsilon((j-1)t/m)\big)^2 =\epsilon\cdot\lim_{m\to\infty}\sum_{j=1}^{m}\Big(N(jt/m\epsilon)-N((j-1)t/m\epsilon) -\lambda p\int_{(j-1)t/m\epsilon}^{jt/m\epsilon}\mathbf{1}\{Z(s)\leq 0\}\,ds\Big)^2.$$

Utilize the fact that

$$\lim_{m\to\infty}\sum_{j=1}^{m}\big(N(jt/m\epsilon)-N((j-1)t/m\epsilon)\big)^2 =\lim_{m\to\infty}\sum_{j=1}^{m}\big(N(jt/m\epsilon)-N((j-1)t/m\epsilon)\big)=N(t/\epsilon).$$

Also,

$$\lim_{m\to\infty}\sum_{j=1}^{m}2\lambda p\int_{(j-1)t/m\epsilon}^{jt/m\epsilon}\mathbf{1}\{Z(s)\leq 0\}\,ds\cdot(N(jt/m\epsilon)-N((j-1)t/m\epsilon)) \leq\lim_{m\to\infty}\sum_{j=1}^{m}2\lambda p\cdot(t/m\epsilon)\cdot(N(jt/m\epsilon)-N((j-1)t/m\epsilon)) =\lim_{m\to\infty}2\lambda ptN(t/\epsilon)/m\epsilon=0.$$

Similarly, $\lim_{m\to\infty}\sum_{j=1}^{m}(\lambda p\int_{(j-1)t/m\epsilon}^{jt/m\epsilon}\mathbf{1}\{Z(s)\leq 0\}\,ds)^2=0$. Hence, we have $[M_\epsilon,M_\epsilon](t)=\epsilon N(t/\epsilon)$.

Now, let $\widetilde{M}_{\epsilon_n}(t)=\epsilon_n(N(t/\epsilon_n)-\int_0^{t/\epsilon_n}\lambda p\mathbf{1}\{Z(s)\leq 0\}\,ds)$, with $\epsilon_n\to 0$. By Theorem 2.18 of Hall and Heyde (1980), it is easy to verify that $\lim_{n\to\infty}\epsilon_n\cdot\widetilde{M}_{\epsilon_n}(t)=0$. Because ϵ_n can be any real positive decreasing sequence to zero, we have $\lim_{\epsilon\to 0}\epsilon\cdot\widetilde{M}_\epsilon(t)=0$ a.s. Hence,

$$\lim_{\epsilon\to 0}\epsilon\cdot N(t/\epsilon)=\lim_{\epsilon\to 0}\epsilon\lambda p\int_0^{t/\epsilon}\mathbf{1}\{Z(s)\leq 0\}\,ds =\lambda pt\lim_{\epsilon\to 0}\int_0^{t/\epsilon}\mathbf{1}\{Z(s)\leq 0\}\,ds/(t/\epsilon) =\lambda ptP(Z(\infty)\leq 0).$$

The last equality follows from the property of continuous-time Markov chain. Also, check condition that for $T>0$,

$$\lim_{\epsilon\to 0}E\Big\{\sup_{t\leq T}|M_\epsilon(t)-M_\epsilon(t^-)|\Big\} =\lim_{\epsilon\to 0}\epsilon^{1/2}E\Big\{\sup_{t\leq T}|N(t/\epsilon)-N(t^-/\epsilon)|\Big\}\leq\lim_{\epsilon\to 0}\epsilon^{1/2}=0.$$

Thus, by the martingale central limit theorem (Theorem 7.1.4 of Ethier and Kurtz 1986), we conclude that $M_\epsilon(t)\Rightarrow\sqrt{\lambda pP(Z(\infty)\leq 0)}\cdot B(t)$, as $\epsilon\to 0$ (where $B(t)$ is the standard Brownian motion process). Now, if we let $t=1$, we have $\epsilon^{1/2}(N_1(1/\epsilon)-\int_0^{1/\epsilon}\lambda p\mathbf{1}\{Z(s)\leq s\}\,ds)\Rightarrow\sqrt{\lambda pP(Z(\infty)\leq 0)}\cdot N(0,1)$, as $\epsilon\to 0$, which completes our proof for the convergence result. $\square$

Proof of Proposition 9. $\hat{\lambda}(t)=N_2(t)/\tau_2(t)$, and $\hat{p}(t)=N_1(t)\tau_2(t)/N_2(t)\tau_1(t)$ can be obtained by solving the first-order conditions of the likelihood function. By similar proof steps of Proposition 8, it is straightforward to show $t^{1/2}(\hat{\lambda}(t)-\lambda)\Rightarrow\sqrt{\lambda/P(Z(\infty)>0)}N(0,1)$, as $t\to\infty$. Now we focus on showing the convergence result for $\hat{p}(t)$. First, note that $t^{1/2}(N_1(t)\tau_2(t)/N_2(t)\tau_1(t)-p)=t^{-1/2}[(\tau_2(t)/\tau_1(t))\cdot(N_1(t)-\lambda p\tau_1(t))+p(N_2(t)-\lambda\tau_2(t))]/(N_2(t)/t)$. Note that $\lim_{t\to\infty}N_2(t)/t=\lambda P(Z(\infty)>0)$ a.s. Let $\gamma=P(Z(\infty)>0)/P(Z(\infty)\leq 0)$. Define $M(t)=\gamma(N_1(t)-\lambda p\tau_1(t))+p(N_2(t)-\lambda\tau_2(t))$. Hence, it suffices to show that

$$t^{-1/2}M(t)\Rightarrow\sqrt{\lambda p\gamma P(Z(\infty)>0)+\lambda p^2P(Z(\infty)>0)}N(0,1),$$

as $t\to\infty$.

Notice that $M(t)$ is a martingale. Scale $M(t)$ as

$$M_\epsilon(t) = \epsilon^{1/2}\left\{\gamma\left(N_1(t/\epsilon) - \lambda p\int_0^{t/\epsilon}\mathbf{1}\{Z(s)\le 0\}\,ds\right) + p\left(N_2(t/\epsilon) - \lambda\int_0^{t/\epsilon}\mathbf{1}\{Z(s)>0\}\,ds\right)\right\}.$$

Checking the quadratic variation of $M_\epsilon(t)$, we have

$$\begin{aligned}[M_\epsilon(t), M_\epsilon(t)] &= \lim_{m\to\infty}\epsilon\sum_{j=1} m\left\{\gamma\left(N_1(jt/m\epsilon) - N_1((j-1)t/m\epsilon) - \lambda p\int_{(j-1)t/m\epsilon}^{jt/m\epsilon}\mathbf{1}\{Z(s)\le 0\}\,ds\right)\right.\\ &\quad\left. + p\left(N_2(jt/m\epsilon) - N_2((j-1)t/m\epsilon) - \lambda\int_{(j-1)t/m\epsilon}^{jt/m\epsilon}\mathbf{1}\{Z(s)\le 0\}\,ds\right)\right\}^2\\ &= \gamma^2\epsilon N_1(t/\epsilon) + p^2\epsilon N_2(t/\epsilon).\end{aligned}$$

Now, by the Law of Large Numbers of the martingale, we have

$$\begin{aligned}&\lim_{\epsilon\to 0}\gamma^2\epsilon N_1(t/\epsilon) + p^2\epsilon N_2(t/\epsilon)\\ &\quad= \lim_{\epsilon\to 0}\gamma^2\epsilon\lambda p\int_0^{t/\epsilon}\mathbf{1}\{Z(s)\le 0\}\,ds + p^2\epsilon\lambda\int_0^{t/\epsilon}\mathbf{1}\{Z(s)>0\}\,ds\\ &\quad= [\lambda p\gamma P(Z(\infty)>0) + \lambda p^2 P(Z(\infty)>0)]t.\end{aligned}$$

Also, verify that for $T>0$, $\lim_{\epsilon\to 0}E\{\sup_{t\le T}|M_\epsilon(t) - M_\epsilon(t^-)|\} = 0$. Thus, by the martingale central limit theorem (Theorem 7.1.4 of Ethier and Kurtz 1986), we have

$$M_\epsilon(t) \Rightarrow \sqrt{\lambda p\gamma P(Z(\infty)>0) + \lambda p^2 P(Z(\infty)>0)}\cdot B(t),$$

as $\epsilon\to 0$. Letting $t=1$, we immediately obtain the convergence result for $\hat{p}(t)$. □

References

Agrawal, N., S. Smith. 1998. Optimal retail assortments for substitutable items purchased in set. *Naval Res. Logist. Quart.* **50**(7) 793–822.

Anupindi, R., M. Dada, S. Gupta. 1998. Estimation of consumer demand with stockout based substitution: An application to vending machine products. *Marketing Sci.* **17**(4) 406–423.

Armony, M., E. Plambeck. 2005. The impact of duplicate orders on demand estimation and capacity investment. *Management Sci.* **51**(10) 1505–1518.

Azoury, K. S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* **31** 1150–1160.

Bensoussan, A., M. Cakanyildirim, S. P. Sethi. 2005. A multi-period newsvendor problem with partially observed demand. *Math. Oper. Res.* Forthcoming.

Chen, L. 2005. Optimal information acquisition, inventory control, and forecast sharing in operations management. Unpublished PhD thesis, Stanford University, Stanford, CA.

Ding, X., M. L. Puterman, A. Bisi. 2002. The censored newsvendor and the optimal acquisition of information. *Oper. Res.* **50**(3) 517–527.

Eppen, G. D., A. V. Iyer. 1997. Improved fashion buying with Bayesian updates. *Oper. Res.* **45**(6) 805–819.

Ernst, R., P. Kouvelis. 1999. The effects of selling package goods on inventory decisions. *Management Sci.* **45**(8) 1142–1155.

Ethier, S. N., T. G. Kurtz. 1986. *Markov Processes: Characterization and Convergence*. John Wiley & Sons, New York.

Garennes, C. 2004. Area car buyers line up to purchase hybrids. *News-Gazette* (June 2), Champaign-Urbana, IL.

Hall, P., C. C. Heyde. 1980. *Martingale Limit Theory and Its Application*. Academic Press, New York.

Harpaz, G., W. Y. Lee, R. L. Winkler. 1982. Learning, experimentation, and the optimal output decisions of a competitive firm. *Management Sci.* **28**(6) 589–603.

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

Johnson, N. L., S. Kotz, A. W. Kemp. 1992. *Univariate Discrete Distributions*, 2nd ed. John Wiley & Sons, New York.

Kok, G. A., M. L. Fisher. 2007. Demand estimation and assortment optimization under substitution: Methodology and application. *Oper. Res.* **55** 1001–1021.

Lariviere, M. A., E. L. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* **45**(3) 346–363.

Lippman, S. A., K. F. McCardle. 1997. The competitive newsboy. *Oper. Res.* **45**(1) 54–65.

Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* **36**(6) 724–738.

Lu, X., J. S. Song, K. Zhu. 2005. On "The censored newsvendor and the optimal acquisition of information." *Oper. Res.* **53**(6) 1024–1026.

Lu, X., J. S. Song, K. Zhu. 2006. Analysis of perishable inventory systems with censored data. *Oper. Res.* Forthcoming.

Mahajan, S., G. van Ryzin. 1999. Retail inventories and consumer choice. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, Norwell, MA, 491–551.

Mahajan, S., G. van Ryzin. 2001a. Stocking retail assortments under dynamic consumer substitution. *Oper. Res.* **49**(3) 334–351.

Mahajan, S., G. van Ryzin. 2001b. Inventory competition under dynamic consumer choice. *Oper. Res.* **49**(5) 646–657.

Nahmias, S., S. A. Smith. 1994. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Sci.* **40**(5) 582–596.

Nauman, M. 2004. Buyers line up and wait for the Toyota gas-electric hybrid. *Mercury News* (August 26). San Jose, CA, 1c–2c.

Netessine, S., N. Rudi. 2003. Centralized and competitive inventory models with demand substitution. *Oper. Res.* **51**(2) 329–335.

Noonan, P. S. 1995. When customers choose: A multi-product, multi-location newsboy model with substitution. Working paper, Emory University, Atlanta, GA.

Parlar, M. 1988. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Res. Logist. Quart.* **35**(3) 397–409.

Parlar, M., S. K. Goyal. 1984. Optimal ordering decisions for two substitutable products with stochastic demand. *OPSEARCH* **21**(1) 1–15.

Rajaram, K., C. S. Tang. 2001. The impact of product substitution on retail merchandising. *Eur. J. Oper. Res.* **135**(3) 582–601.

Scarf, H. E. 1959. Bayes solution of the statistical inventory problem. *Ann. Math. Statist.* **30** 490–508.

Smith, S. A., N. Agrawal. 2000. Management of multi-item retail inventory system with demand substitution. *Oper. Res.* **48**(1) 50–64.

Wang, Q., M. Parlar. 1994. A three-person game theory model arising in stochastic inventory control theory. *Eur. J. Oper. Res.* **76**(1) 83–97.