

Demand Estimation and Ordering Under Censoring: Stock-Out Timing Is (Almost) All You Need

Aditya Jain

Indian School of Business, Hyderabad 500 032, India, aditya_jain@isb.edu

Nils Rudi

INSEAD, 138676 Singapore, nils.rudi@insead.edu

Tong Wang

NUS Business School, National University of Singapore, 119245 Singapore, tong.wang@nus.edu.sg

Retailers facing uncertain demand can use observed sales to update demand estimates. However, such learning is limited by the amount of inventory carried; when demand exceeds inventory (i.e., when a stock-out event occurs), a retailer in general cannot observe actual demand. We propose using observations on the *timing* of sales occurrences in a Bayesian fashion to learn about demand, and we analyze this learning method for a multiperiod newsvendor setting. We find that, as previously shown with the use of only stock-out *event* observations, the optimal order quantity with timing observations is greater than the optimal order quantity with full demand observations. We prove this result using a novel methodology from the statistics literature on comparison of experiments. Although the optimal over-ordering with timing observations tends to be less than that with only stock-out event observations in most cases, we do observe cases where the opposite is true. Such cases correspond to high demand uncertainty and low margins, where marginal learning from timing observations is significantly higher than using only a stock-out event. In an extensive numerical study we find that, on average and with respect to uncensored demand observations, the use of timing observations eliminates 76.1% of the loss in expected profit from using only stock-out event observations. We show that, for Poisson and normal demand with unknown mean, the proposed learning method is tractable as well as intuitively appealing: the information contained in the timing of sales occurrences is fully captured by a single number—the timing of stock-out. We also investigate checkpoint models in which the newsvendor can make observations only at predetermined times in a period, and illustrate its convergence to the models with timing and stock-out event observations.

Subject classifications: Bayesian inventory; lost sales; censored observations.

Area of review: Operations and Supply Chains.

History: Received September 2012; revisions received November 2013, May 2014; accepted September 2014. Published online in *Articles in Advance* December 4, 2014.

1. Introduction and Literature Review


With increased product variety, shorter product lifetimes, and longer lead times due to increased levels of global sourcing, managing stock-outs has become a major concern for manufacturers and retailers. Not only do stock-outs signify a loss of immediate sales and revenue, they also obscure observations of true demand. Such censoring of demand observations undermines a firm's ability to improve its demand estimates. Yet, even when a stock-out occurs, additional valuable information is readily available that can be utilized to facilitate better demand estimation. The following anecdote, experienced by one of the coauthors, provides an example of such information.

The coauthor, accompanied by a senior supply chain manager of Yijiaxian (a convenience store chain specializing in fresh food), visited one of the stores in Chengdu, China. During the visit, the manager casually asked a salesperson to name the store's best-selling product. The salesperson pointed toward an empty shelf meant for *long beans*.

Looking surprised, the manager questioned this—according to the daily sales report generated by the firm's point-of-sale (POS) system, *cabbages* had consistently outsold *long beans*. The salesperson clarified that although *cabbages* also sold well, its sales numbers were larger because of its higher daily stocking levels. In contrast, *long beans* are stocked in smaller quantities but are sold at a much faster rate; they are usually out of stock within the first few hours of the store opening.

This anecdote points out the critical difference between inferences made from observing only aggregate sales (hereafter, *sales*) and observing the occurrence of sales over time. Relying only on sales data, the manager underestimated demand for *long beans* because the demand following a stock-out is not met and hence not recorded. Yet the salesperson—who observes, in addition, sales over time and the timing of stock-outs—could intuitively use this information to assess the higher demand levels for *long beans*. To substantiate this intuition and highlight the value of

Figure 1. Daily inventory of products *Holly* and *Amber Glow* during the test period.



observing stock-out timing, we offer a more detailed example from our interaction with Mikyajy, a Middle East-based cosmetics brand with its own retail outlets.

In May 2009, Mikyajy launched the *Minx* product line. To improve its knowledge about demand, the company carried out a limited test launch at one store a month prior to the full launch. Two products—a lip-plumping gloss (*Holly*) and a frosted eyeshadow (*Amber Glow*)—shared similar pretest launch forecasts, and 10 units of each product were stocked for the test launch. Figure 1 plots the actual daily inventory levels for the two products over the first 30 days of sales. If one were to consider only sales and the information on whether the product stocked out, then there would be little to differentiate these two products: both had sales of 10 units and were stocked out within a month. However, a careful observer would note that *Holly* stocked out much more rapidly than *Amber Glow* (after only 9 versus 25 days), which indicates higher demand for the former. Realizing this, Mikyajy acted swiftly to secure aggressive subsequent replenishment of *Holly* and was rewarded handsomely; thanks largely to that additional inventory, *Holly* sold about 5 times as many units as *Amber Glow* did in the first six months after the products were launched across nearly a hundred stores.

These examples suggest that managers can significantly improve their learning about demand, and hence make better future ordering decisions, if they keep track of the timing of sales. It is seldom possible to capture precise information on unmet demand, but the timing of sales occurrences prior to stock-out is typically available from the firm's POS system (because it amounts to the timing of sales transactions). The aim of this paper is to substantiate these findings analytically. We consider a parsimonious multiperiod newsvendor model in which lost sales are unobserved and knowledge about demand is updated, in a Bayesian fashion, after each period. For the general case of our model, which models demand as a renewal process, we consider three types of observations: (i) sales along with a binary variable indicating whether a stock-out occurred (hereafter, *stock-out event*); (ii) sales together with

the timing of sales occurrences until stock-out, as motivated by the preceding anecdotes; and (iii) complete observation of all demand occurrences together with its timing. Of these, the uncensored observation type (iii) serves as our benchmark scenario. Of the two censored observation types, type (i) has garnered more of the literature's attention (e.g., Lariviere and Porteus 1999, Ding et al. 2002, Chen 2010). For each of these three observation types, we establish Bayesian updating mechanisms and analyze the optimal inventory decisions.

We prove that the optimal order quantity under observation type (ii), when the timing of sales occurrences is observed, is larger than the optimal order quantity under observation type (iii), when complete demand is observed. This finding is consistent with the optimal order quantity under observation type (i) when only sales and stock-out events are observed (Ding et al. 2002, Lu et al. 2008, Bensoussan et al. 2009). To prove this result, we employ a methodology from the statistics literature on comparison of experiments (see Ginebra 2007). The resulting proof is simpler and more intuitive than the previous proofs of corresponding results. Our analysis uncovers drivers of the value of learning based on censored observations of types (i) and (ii). The “value of learning” is defined as improvements in expected profit, after the current period, that are due to learning about the demand from current period's inventory decision. With both stock-out event and timing observations, the value of learning from a marginal unit of current inventory is positive. However, with timing observations the marginal value of learning is monotonically decreasing in inventory levels, whereas with event observations it is unimodal. So even though the firm optimally over-orders with both types of censored observations, it is not clear a priori which of these two order quantities is greater. We find that the over-ordering is usually much less with timing observations than with stock-out event observations, but there are exceptions. In particular, if demand uncertainty is high and unit margins are small, so that the optimal order quantities are smaller in general, the marginal value of learning based on timing is higher than learning based on a stock-out event. Then the firm's optimal order quantity with timing observations is larger than that with event observations.

For commonly used Poisson and normal distributions with unknown mean, we show that the information contained in the timing of sales occurrences (i.e., observation type (ii)) is fully captured by two numbers: observed sales and the timing of stock-out (if any). So in the event of stock-out, the single number that indicates the time of the demand for the last unit sold is sufficient. In such cases, the use of stock-out timing leads to remarkably simple and intuitive sufficient statistics reflective of learning from available historical data. We also relax the implicit assumption underlying observation type (ii)—namely, that the status of inventory is continuously visible—by analyzing “checkpoint models” in which observations are made only at predetermined times (referred to as “checkpoints”)

within each period. For this we consider two types of observations at each checkpoint: the inventory level and the binary variable indicating whether a stock-out has occurred.

Our numerical study indicates that, on average, the use of timing observations eliminates 76.1% of the loss in expected profit—with respect to the benchmark type (iii)—that results from using only stock-out event observations. To evaluate the relative importance of demand estimation and subsequent optimization, we consider the *myopic timing* model, under which the demand distribution is updated in Bayesian fashion using timing observations, but ordering decisions are made myopically (i.e., single-period optimal newsvendor quantities are ordered). In essence, the myopic timing model ignores the effect of the current period's inventory decision (and resulting learning) on the expected profit earned after current period. This simplification leads to a practically appealing solution with computational complexity linear in the problem size. We find that this model also performs well in terms of expected profit: it eliminates 75.0% of the loss in expected profit (i.e., only 1 percentage point less than under nonmyopic optimization) when only stock-out event observations are used. This result indicates that using rich data to learn about demand makes it largely unnecessary to account for learning when establishing optimal order quantities.

1.1. Relevant Literature

There is a vast literature in operations management that addresses estimation of demand using sales data that might be censored owing to stock-outs. This paper belongs to the substream of research on *Bayesian inventory models*, which focuses on simultaneous demand estimation and inventory optimization while using a Bayesian framework to estimate demand. Early papers on Bayesian inventory management (e.g., Scarf 1959, Iglehart 1964, Azoury 1985) assume that unmet demand is both observed and backordered. This assumption makes it easier to define the conjugate priors, which greatly simplify the dynamic updating procedure. However, when actual demand is unknown because of stock-outs and observations are censored, the problem becomes difficult to solve. Conrad (1976) was the first to highlight this issue and noted that, when observations are censored, inventory levels in the current period affects the demand estimate for the next period. Furthermore, with censored data the sufficient statistics (which capture the cumulative learning from historical data) tend to explode in dimension for general distributions. To tackle this problem, Braden and Freimer (1991) propose a family of distributions with a conjugate prior for censored observations. This family, the so called newsvendor distributions, includes exponential and Weibull distributions, but not the Poisson and normal distributions that are more commonly used for modeling demand. Lariviere and Porteus (1999), and more recently Bisi et al. (2011), adopt the newsvendor distributions in their nonperishable inventory models and thereby further reduce the state space. Three

notable papers—Ding et al. (2002), Lu et al. (2008), and Bensoussan et al. (2009)—consider the interaction between inventory decisions and demand learning for a perishable inventory model. They show that a firm should order more than the myopically optimal quantity (i.e., over-order) so as to compensate for loss of learning from censoring. Chen and Plambeck (2008) generalize the over-ordering result to cases in which demand is partially satisfied by a substitute product in the event of a stock-out. Given the complexity of solving inventory models with Bayesian demand updates under censoring, Chen (2010) develops bounds on costs and optimal order quantities. Ulu et al. (2012) employ a Bayesian framework to discern customer preference while simultaneously optimizing assortment; here censoring arises from the firm offering a limited assortment of products (and not from stock-outs). Other papers propose using nonparametric, data-driven approaches to solve the problem of simultaneous demand estimation and inventory optimization. Along these lines, Burnetas and Smith (2000), Godfrey and Powell (2001), Huh and Rusmevichientong (2009), and Huh et al. (2009 and 2011) develop adaptive inventory policies based on historical observations. Besbes and Muharremoglu (2013) take a nonparametric approach to analyze the effect—measured in terms of decision makers “regret”—of demand censoring.

There is a large body of literature in the retailing and revenue management settings that addresses the correction of censoring-induced errors (on demand estimation) via procedures to “uncensor” or “unconstrain” demand data. In the context of inventory management, Wecker (1978) illustrates the adverse effect of stock-outs on demand forecast accuracy; Nahmias (1994) and Agrawal and Smith (1996) analyze the problem of uncensoring normal and negative binomial demand, respectively. Lau and Lau (1996) detail a procedure—for obtaining the demand distribution from censored data—that combines a nonparametric “product limit” method with extrapolation of hourly sales. In the revenue management setting, Queenan et al. (2007) propose an unconstraining method that employs double exponential smoothing to estimate lost sales. Several papers have considered the problem of estimating demand or a consumer choice model in multiproduct settings, where stock-outs result in demand censoring as well as product substitution. The expectation-maximization (EM) algorithm outlined in Dempster et al. (1977) has been widely adopted. Anupindi et al. (1998) consider sales data from a periodic review system and propose using the EM algorithm along with the *time of stock-out* as missing data in order to estimate both demand rates and substitution probabilities. Talluri and van Ryzin (2004) and Vulcano et al. (2010) illustrate how the EM algorithm can be used to estimate general choice models from sales data with unobserved no-purchase incidences. Kök and Fisher (2007) develop an EM-based method for estimating demand and substitution probabilities within a hierarchical model of consumer choice in a retail setting. Vulcano et al. (2012) integrate the EM algorithm with a multinomial logit choice

model to estimate the arrival rate and customers' preferences. van Ryzin and Vulcano (2014) carry out similar estimation using a more general rank-based, nonparametric choice model. Papers by Stefanescu (2009) and Conlon and Mortimer (2013) also use the EM algorithm to correct for demand censoring in their proposed demand estimation methods. In contrast to the studies cited previously, Bruno and Vilcassim (2008) estimate demand using the "random utility maximization" model with time-varying product availability. Musalem et al. (2010) propose a method for inferring product availability from sales data; they use it together with a Markov chain Monte Carlo method to estimate a choice model as well as lost sales. Farias et al. (2013) is another notable recent work to consider estimation of customer choice model from limited data.

1.2. Contribution and Outline

Our paper makes two main contributions to the literature. First, we propose Bayesian updating based on the timing sales occurrences. We show that for commonly used demand distributions, this proposed method is equivalent to updating demand based on timing of stock-outs. In these cases, the proposed method preserves Bayesian conjugacy and hence achieves the same simplicity and tractability obtained with the newsvendor distribution (proposed by Braden and Freimer 1991). Second, we analyze the effect of the proposed demand learning method on optimal ordering decisions and on expected profits. Specifically, we show that even though with the proposed method leads to over-ordering, in most cases the extent of over-ordering is significantly less than when only stock-out event observations are used. Employing our proposed method also eliminates much of the loss in expected profit that results from using only stock-out event observations.

In §2, we describe our most general model setup, which includes observations of sales occurrences and minimal assumptions about the demand process; here we require only that the demand be modeled as a stochastic process with independent and stationary increments. All the discrete renewal processes and continuous-time processes (such as Brownian motion) satisfy this requirement. Under this general setup we compare optimal profits and stocking levels among the three types of observations. In §3 we consider observations only on stock-out timing. For the special cases of Poisson and normal demand, we show that observing stock-out timing is equivalent to observing all sales occurrences. In §4, we study checkpoint models in which inventory status is observed at predetermined times in each period, and in §5 we give a numerical illustration of the models discussed. Section 6 concludes the paper.

2. Main Model with Observations of Sales Occurrences

The main goal of this paper is to analyze the value for an inventory system of information about demand arrival timing when knowledge about demand is incomplete. Toward

this end, we consider a parsimonious model of a newsvendor selling a perishable product over a finite horizon of N periods. At the beginning of each period, the newsvendor determines stocking level to satisfy the uncertain demand for the period. We assume that, in each period, inventory is delivered immediately after the stocking decision, unmet demand is lost, and any inventory left over (after satisfying demand in the period) has zero salvage value. The unit cost c and revenue r are both stationary.

Unlike most papers on the periodic review inventory models, in which uncertain demand over a period is modeled as a random variable that is realized in a single instant, we "open up" the period and view demand arrivals as a stochastic process that transpires over a span of time. More specifically, each period in our model consists of a fixed duration of unit length from $t = 0$ to $t = 1$. We model cumulative demand arrivals during a period as a stochastic process $\{D(t), 0 \leq t \leq 1\}$ with unknown parameter(s) θ . Conditional on θ , the demand process has independent and stationary increments. For ease of exposition, we focus on discrete demand processes. We suppress time notation in $D(t)$ when $t = 1$ and use subscripts to denote period-specific variables and functions. Cumulative demand $D(t)$ arriving over duration t has probability mass function (pmf) $f(\cdot | t, \theta)$, cumulative distribution function (cdf) $F(\cdot | t, \theta)$, and complementary cdf $\bar{F}(\cdot | t, \theta)$. The arrival epoch of the i th demand within a period is denoted t^i ; the corresponding interarrival time is $\tau^i = t^i - t^{i-1}$ (let $t^0 = 0$), with probability density function (pdf) $\psi(\tau | \theta)$, cdf $\Psi(\tau | \theta)$, and finite moments. Knowledge about parameter θ is described by prior distribution $\pi(\theta)$. For a period with given prior π , the newsvendor with stocking level y earns an expected single-period profit of

$$W(y | \pi) = \mathbb{E}_{D|\pi}[r \min\{D, y\}] - cy. \quad (1)$$

At the end of the period, the newsvendor makes an observation $O(y)$ that may (as a function of the "learning model" to be specified) depend on the realization of demand D , the inventory level y , and/or the demand arrival timing $\{t^1, t^2, \dots, t^D\}$ (or, equivalently, the interarrival times $\{\tau^1, \tau^2, \dots, \tau^D\}$). Let $l(O(y) | \theta)$ denote the likelihood of observing $O(y)$ for a given θ . The newsvendor updates her knowledge about θ by combining the prior knowledge π with observation $O(y)$ in Bayesian fashion:

$$\pi'(\theta) = \pi(\theta) \oplus O(y) = \frac{l(O(y) | \theta) \cdot \pi(\theta)}{\int_{\Theta} l(O(y) | \theta') \cdot \pi(\theta') d\theta'}. \quad (2)$$

Let $V_n(\pi_n)$ denote the optimal expected profit of the newsvendor from the beginning of period n until the end of the horizon—given current knowledge π_n about demand parameter θ . Without loss of generality, we assume that $V_{N+1}(\cdot) = 0$. Determining the optimal stocking level can then be formulated as the following dynamic program:

$$V_n(\pi_n) = \max_{y_n \geq 0} \{W(y_n | \pi_n) + \mathbb{E}[V_{n+1}(\pi_n \oplus O_n(y_n))]\}, \quad n = 1, \dots, N. \quad (3)$$

Next we introduce three learning models and present the main proposition of this paper. For each learning model, we specify the observation $O(y)$ and its likelihood function; the posterior distribution of θ can be obtained by applying Equation (2).

(i) *Observation of sales quantity and stock-out event* ($O^E(y)$). This is the most commonly studied scenario in the literature: in each period, the newsvendor observes only the realized total sales and whether a stock-out has occurred. Demand exceeding the available inventory is lost and unobserved. The observation for a period is $O^E(y) = \{s, e\}$, where $s = \min\{D, y\}$ is the sales and $e = \mathbf{1}_{\{D \geq y\}}$ is an indicator variable for stock-out events. Two outcomes are possible in a period: if there is no stock-out ($e = 0$) then the newsvendor observes the true demand as the sales ($D = s$), the likelihood of which is simply $f(s | 1, \theta)$; if there is a stock-out ($e = 1$) then all the inventory are sold ($s = y$), and the true demand is even larger ($D \geq y$), which yields likelihood $\bar{F}(y - 1 | 1, \theta)$. The likelihood of $O^E(y)$ can then be written as

$$l(O^E(y) | \theta) = \begin{cases} f(s | 1, \theta) & \text{if } e = 0, s < y; \\ \bar{F}(y - 1 | 1, \theta) & \text{if } e = 1, s = y. \end{cases} \quad (4)$$

(ii) *Observation of sales occurrences* ($O^T(y)$). This is the new learning model that we propose. Here, in addition to the sales quantity and stock-out event information, the newsvendor has access to the arrival timing of all the observed sales, $\vec{t}^s = (t^1, t^2, \dots, t^s)$. The observation for a period is $O^T(y) = \{s, e, \vec{t}^s\}$ with likelihood

$$l(O^T(y) | \theta) = \begin{cases} \prod_{i=1}^s \psi(\tau^i | \theta) \bar{\Psi}(1 - t^s | \theta) & \text{if } e = 0, s < y; \\ \prod_{i=1}^y \psi(\tau^i | \theta) & \text{if } e = 1, s = y. \end{cases} \quad (5)$$

Recall that the interarrival times are $\tau^i = t^i - t^{i-1}$. If there is no stock-out then demand is equal to sales ($D = s$), and the newsvendor observes interarrival times $\{\tau^i\}_{i=1}^s$ and one duration of length $1 - t^s$ without any arrivals; in this case the likelihood function is $\prod_{i=1}^s \psi(\tau^i | \theta) \bar{\Psi}(1 - t^s | \theta)$. If there is a stock-out then only the first y interarrival times are observed and no additional observations are made after the last unit of inventory is sold. Hence the likelihood function in this case is $\prod_{i=1}^y \psi(\tau^i | \theta)$.

(iii) *Full observation of all demand occurrences* (O^F). This represents an ideal case where the newsvendor observes all demand occurrences during a period, including sales and lost sales, as well as their arrival timing. The observation for a period is $O^F = \{D, \vec{t}^D\}$, where $\vec{t}^D = (t^1, t^2, \dots, t^D)$, regardless of whether a stock-out occurs. The likelihood function is then

$$l(O^F | \theta) = \prod_{i=1}^D \psi(\tau^i | \theta) \bar{\Psi}(1 - t^D | \theta). \quad (6)$$

Note that observation O^F and its likelihood are independent of the inventory level y .

Hereafter we use the terms “event,” “timing,” and “full” in reference to observations $O^E(y)$, $O^T(y)$, and O^F (respectively). The superscripts E, T, and F denote functions and variables relevant to models with (respectively) event, timing, and full observations. More specifically, $V_n^E(\cdot)$, $V_n^T(\cdot)$, and $V_n^F(\cdot)$ denote expected profits to go— $y_n^E(\cdot)$, $y_n^T(\cdot)$, and $y_n^F(\cdot)$ denote optimal inventory levels—in period n under the three types of observations. Finally, $y^M(\cdot)$ denotes the myopically optimal inventory level that maximizes the newsvendor’s single-period expected profit in Equation (1).

Proposition 1 orders the learning models just described in terms of their optimal profits and optimal ordering decisions.

PROPOSITION 1. *For any given π_n at the beginning of period n , the following statements hold:*

- (a) $V_n^E(\pi_n) \leq V_n^T(\pi_n) \leq V_n^F(\pi_n)$;
- (b) $y_n^E(\pi_n) \geq y^M(\pi_n)$, $y_n^T(\pi_n) \geq y^M(\pi_n)$, and $y_n^F(\pi_n) = y^M(\pi_n)$;
- (c) $y_n^E(\pi_n)$ and $y_n^T(\pi_n)$ have finite upper bounds.

We prove Proposition 1 by employing results from the statistics literature on comparison of experiments (for a detailed overview see Ginebra 2007). The resulting proof is shorter and more intuitive than corresponding proofs in the literature. Here, we sketch the proof, relegating its detailed mathematical version to appendix. First, we formally define three statistical experiments—namely, $\mathcal{E}^E(y)$, $\mathcal{E}^T(y)$, and $\mathcal{E}^F$ —corresponding to the observations (respectively) $O^E(y)$, $O^T(y)$, and O^F . Experiments $\mathcal{E}^E(y)$ and $\mathcal{E}^T(y)$ each have one parameter, the inventory level y , whereas experiment $\mathcal{E}^F$ has none. Using the probability distribution of the outcomes of these experiments, we compare them in terms of their *statistical sufficiency*: one experiment is said to be statistically sufficient for another experiment if the outcome of the latter can be obtained by a stochastic transformation of the former’s outcome or, in other words, if the former experiment yields a dominating set of observations. To prove part (a), we establish the following sufficiency ordering: for any given inventory level y , experiment $\mathcal{E}^F$ is sufficient for $\mathcal{E}^T(y)$ and, $\mathcal{E}^T(y)$ is sufficient for $\mathcal{E}^E(y)$. Furthermore, given the same prior knowledge π_n , optimal future profits are increasing in the informativeness of the observations. Thus, the sufficiency ordering between experiments leads directly to part (a).


To prove the first two inequalities of part (b), we show that experiment $\mathcal{E}^E(y)$ (respectively, $\mathcal{E}^T(y)$) with an inventory parameter y is sufficient for experiment $\mathcal{E}^E(y')$ (respectively, $\mathcal{E}^T(y')$) with a smaller inventory parameter $y' < y$. This sufficiency ordering implies increasing informational value as an additional unit of inventory is held: a higher inventory level y will lead to less censoring of demand and hence to a statistically dominating observation, which translates into superior knowledge about θ . In order to obtain

that superior knowledge, the newsvendor sets a higher inventory level than the myopic one, y^M . The third equality, $y_n^F(\pi_n) = y^M(\pi_n)$, follows because the informativeness of experiment $\mathcal{E}^F$ with full observations does not depend on the inventory level y .

An intuitive way to understand the over-ordering result is by considering trade-off, between exploration and exploitation, that is inherent in the inventory decision for any period. In our models with censored observations, the inventory decision plays two roles: it determines the profit in the current period $W(y_n | \pi_n)$ (i.e., exploitation), and it affects the informativeness of the end-of-period observation, which is the basis for improving future profits $\mathbb{E} V_{n+1}(\pi_n \oplus O_n(y_n))$ (i.e., exploration). On the one hand, if exploitation were the only concern (or, when exploration is unaffected by the inventory decision, as under full observation), then the optimal inventory level is simply the single-period newsvendor quantity $y^M(\pi_n)$. On the other hand, if exploration were the only concern then a maximal inventory would be maintained because future profit increases with each marginal unit of inventory (although the marginal increase diminishes). Given event and timing observations, the optimal inventory decision trades off exploitation for exploration and takes a value no smaller than $y^M(\pi_n)$. We illustrate this trade-off in Figure 2, which is based on one of the scenarios in our numerical study in §5. For the case when there are observations on both sales timing and stock-out events, the marginal value of exploration $\Delta_V(y) = \mathbb{E} V_{n+1}(\pi_n \oplus O_n(y+1)) - \mathbb{E} V_{n+1}(\pi_n \oplus O_n(y))$ and the marginal loss from exploitation $\Delta_W(y) = -(W(y+1 | \pi_n) - W(y | \pi_n))$ are plotted against inventory decision y . In both cases, the optimal inventory decision is to order the smallest inventory level at which the marginal loss $\Delta_W(y)$ exceeds the marginal gain $\Delta_V(y)$ (i.e., the level at which an additional unit would reduce expected profit). The myopically optimal inventory decision ignores exploration and so is the smallest inventory level at which $\Delta_W(y)$ exceeds 0 (point $\textcircled{M}$ in Figure 2). With full observation, the same information is gathered at all inventory levels and the marginal value of exploration is therefore 0; thus the optimal inventory decision (point $\textcircled{F}$ in Figure 2) is the same as the myopically optimal inventory decision. With event and timing observations, the marginal values of exploration ($\Delta_V^E(y)$ and $\Delta_V^T(y)$, respectively) are positive for all y , and the value of $\Delta_W(y)$ is greater than both of them only when the inventory level exceeds the myopic optimum. The inventory levels with event and timing observations are indicated by points $\textcircled{E}$ and $\textcircled{T}$, respectively, in Figure 2.

It is noteworthy in Figure 2 that the plots for marginal values of exploration with event and timing observations ($\Delta_V^E(y)$ and $\Delta_V^T(y)$, respectively) have markedly different shapes, which can be explained as follows. Under *event* observation, the marginal value of exploration is proportional to the reduction in probability of stock-out. For values y well below (respectively, above) expected demand in

Figure 2. Marginal values of exploration, marginal loss of exploitation, and optimal ordering decisions.



a period, the probability of stock-out is high (respectively, low); in both cases, the marginal value of learning is relatively insensitive to the inventory decision y . Yet for a range of intermediate y , this marginal value of learning is significant; these dynamics explain the unimodal shape in Figure 2. In contrast, the marginal value of exploration under *timing* observations is the expected value of information (given a stock-out) multiplied by the probability of a stock-out. More simply, the first term can be viewed as the value of observing the timing of y units of demand (which would have a decreasing marginal value). That value is multiplied by the second term to reflect that the former yields informational value only when there is a stock-out, the probability of which is decreasing in y . This dynamic suggests that the marginal value of learning from timing observations is decreasing in the inventory level, as illustrated in Figure 2. The implications are that, if the myopically optimal inventory level y^M is large, then $y^T \leq y^E$; in other words, less over-stocking occurs with timing observations. However, one cannot rule out the possibility of $y^T > y^E$ at the other extreme (i.e., for small values of y^M); the existence of such cases is confirmed by the numerical study in §5.

3. Observations of Stock-Out Timing

The general model setup in §2 allowed us to demonstrate the informational benefit of “opening up” a period and viewing uncertain demand as a stochastic process rather than a random variable. However, that setup requires information on all sales occurrences in order to update demand information. We shall now discuss some special cases of demand processes under which the information requirement can be significantly simplified. We focus on Poisson and normal distributions, which are the two most commonly used in newsvendor models. With only stock-out event observations, Bayesian updating of Poisson and normal demand with unknown mean is notoriously complex because neither distribution has conjugate priors under censored observations (Braden and Freimer 1991). Yet with

timing observations, under which demand is modeled in terms of corresponding stochastic processes, conjugate priors do exist. Furthermore, the information requirement for Bayesian updating can be significantly simplified.

We define stock-out timing observation $O^{\text{ST}}(y) = \{s, e, t^s\}$, as a special case of the timing observation $O^{\text{T}}(y)$, in which the newsvendor observes only sales quantity and timing of stock-out (if one occurs). We use superscript ST to denote functions and variables for model with stock-out timing observation. Given observation $O^{\text{ST}}(y) = \{s, e, t^s\}$ retains only part of information contained in $O^{\text{T}}(y)$, but carries more information than in event observation $O^{\text{E}}(y)$, we have the following result:

COROLLARY 1. For any given π_n at period n , $V_n^{\text{E}}(\pi_n) \leq V_n^{\text{ST}}(\pi_n) \leq V_n^{\text{T}}(\pi_n)$.

Although the above inequalities hold for general demand processes specified in §2, for special cases of Poisson or normal demand processes we show that the stock-out timing observation $O^{\text{ST}}(y)$ is sufficient for timing observation $O^{\text{T}}(y)$ (i.e., the second inequality becomes an equality). In the following, we first illustrate this for Poisson demand and then formalize the finding in Proposition 2 (the proof and an illustration for normal demand are provided in appendix).

Consider a Poisson demand process with unknown rate λ . The cumulative demand $D(t)$ during time t is a Poisson random variable with mean λt , and the interarrival times τ are exponentially distributed with rate λ (mean $1/\lambda$). The conjugate prior for the unknown parameter λ is gamma distribution with shape and rate parameters α and β , respectively. Formally, we have

$$f(d | t, \lambda) = \frac{(\lambda t)^d e^{-\lambda t}}{d!}, \quad \psi(\tau | \lambda) = \lambda e^{-\lambda \tau},$$

$$\pi(\lambda) = \frac{\beta^\alpha \lambda^{\alpha-1} e^{-\lambda \beta}}{\Gamma(\alpha)}.$$

With inventory level y , the newsvendor observes $O^{\text{T}}(y) = \{s, e, \vec{t}^s\}$ and updates her demand information in Bayesian fashion using the likelihood function given in (5), which yields the following posterior distribution for λ (see appendix for the derivation):

$$\pi'(\lambda | O^{\text{T}}(y)) = \begin{cases} \text{Gamma}(\alpha + s, \beta + 1) & \text{if } e = 0, s = 0, \dots, y - 1; \\ \text{Gamma}(\alpha + s, \beta + t^s) & \text{if } e = 1, s = y, 0 < t^s \leq 1. \end{cases} \quad (7)$$

If no stock-out occurs ($e = 0$), then the distribution of λ is updated in the standard way using the exact observations (Raiffa and Schlaifer 2000); the result is a gamma posterior with rate and shape parameters $\alpha + s$ and $\beta + 1$. Notably, the Poisson–gamma conjugacy is preserved even when stock-out does occur ($e = 1$), in which case we obtain a gamma

posterior with rate and shape parameters $\alpha + s$ and $\beta + t^s$. In either case, the timing observations of all sales occurrences (i.e., all elements of vector $\vec{t}^s$) are no longer needed: in the case of no stock-out, information on sales s (which equals exact demand D) is sufficient; and in the case of stock-out, the last element of $\vec{t}^s$, namely, t^s (or equivalently, t^y), which is the time of sale of the last unit of inventory and the occurrence of stock-out—is sufficient. In a multi-period setting, on repeatedly applying Bayesian updating as outlined above, it follows that two numbers, namely, *cumulative sales quantity* and *cumulative sales duration*, constitute a sufficient statistic for updating demand. More specifically, given historical data—over a set $\mathbf{N}$ of past periods—consisting of $\{y_i\}_{i \in \mathbf{N}}$ stocking decisions, $\{d_i\}_{i \in \mathbf{N}_1}$ exact demand observations over periods in the set $\mathbf{N}_1 \subseteq \mathbf{N}$ without stock-out events, and $\{t_i\}_{i \in \mathbf{N} \setminus \mathbf{N}_1}$ stock-out timing observations, only two numbers $\mathcal{S} = \sum_{i \in \mathbf{N}_1} d_i + \sum_{j \in \mathbf{N} \setminus \mathbf{N}_1} y_j$ and $\mathcal{T} = |\mathbf{N}_1| + \sum_{j \in \mathbf{N} \setminus \mathbf{N}_1} t_j$ are needed for updating demand. Further, a new observation can be incorporated by adding observed sales quantity to $\mathcal{S}$ and sales duration to $\mathcal{T}$.

Our next proposition summarizes the findings for these special cases.

PROPOSITION 2. For a Poisson demand process with unknown rate λ or a normal demand process with unknown mean μ , with timing observations $O^{\text{T}}(y)$:

- (a) there exists a conjugate prior and a fixed dimensional sufficient statistic;
- (b) the sufficient statistic comprises cumulative sales quantity and cumulative sales duration;
- (c) stock-out timing observation $O^{\text{ST}}(y)$ is sufficient, i.e., $V_n^{\text{ST}}(\pi_n) = V_n^{\text{T}}(\pi_n)$.

As mentioned previously, no conjugate priors exist for Poisson and normal demand processes with stock-out event observations. Therefore, the Bayesian updating in (2) does not result in a posterior distribution that is the same irrespective of whether a stock-out event took place. So, to solve dynamic program (3), one must keep track of the entire history of observations, the dimension of which naturally grows as more observations are added. With timing observations, a conjugate prior exists even in case of stock-outs, and beliefs about the unknown parameter (in both cases, the demand mean) follow a common distribution. More importantly, as noted in part (b) of Proposition 2, the whole history of observations is encapsulated in a two-dimensional sufficient statistic. As a result, the corresponding dynamic program remains two-dimensional. In addition to significantly simplifying the analysis and computation of Bayesian inventory models, this finding has a remarkably intuitive appeal. A manager who seeks to update demand information on a regular basis need only track cumulative sales and the duration over which those sales were made.

A natural question that arises here is whether stock-out timing is sufficient for other demand distributions and, if not, whether there exist similar simple statistics. With timing observations, the Bayesian inference is directed toward

learning about the underlying distribution of interarrival times. The likelihood in (5) shows that, when a stock-out occurs, the newsvendor observes y exact observations of interarrivals ($\{\tau^i\}_{i=1}^y$); when there is no stock-out, the observation consists of s exact observations of interarrivals ($\{\tau^i\}_{i=1}^s$) and one censored observation of the $(s+1)$ th interarrival time ($\tau^{s+1} > 1 - \sum_{i=1}^s \tau^i$). This censored observation leads to technical difficulty in obtaining sufficient statistics. We shed light on how this difficulty can be resolved for two families of interarrival time distributions: (i) the newsvendor family and (ii) the exponential family. Note that (ii) is a subset of (iii), and the exponential (i.e., Poisson demand) considered in Proposition 2 belongs to both. For other distributions in the newsvendor family, part (a) of Proposition 2 continues to hold: there exist a conjugate prior and a fixed-dimensional sufficient statistic (Braden and Freimer 1991), but the latter is a more complex function of observed interarrival times than a simple sum, which is the case of exponential distribution. Expanding further to the whole exponential family, there is not much that can be done as is. However, if it is reasonable to suppose that the additional information carried in the censored observation of the $(s+1)$ th interarrival time is negligible, an approximation can be obtained simply by omitting the censored observation. This way, the existence of a conjugate prior and a fixed-dimensional sufficient statistic can be generalized to the whole exponential family of distributions (see Braden and Freimer 1991, §2).

4. Observations at Discrete Time Points: Checkpoint Models

The timing observation proposed in our paper requires the newsvendor to have continuous visibility of demand arrivals (or, equivalently, of inventory levels) over a period. In this section, we relax that requirement and consider a case in which the newsvendor can make observations only at predetermined time points in each period. Let $\mathbf{M} = \{p^1, \dots, p^M\}$ denote a set of M time points in a period, each referred to as a “checkpoint.” The checkpoints are ordered by their indices (i.e., $p^m < p^{m+1}$), where $p^0 = 0$ and $p^M = 1$ denote (respectively) a period’s beginning and end. We consider two levels of checkpoint observations: exact inventory levels and stock-out indicator (i.e., whether or not a stock-out has occurred by a checkpoint).

Observations of Inventory Levels. First consider the inventory checkpoint model IC[M] in which the newsvendor observes inventory levels at all checkpoints in the set $\mathbf{M}$. Let $I^0 = y$ be the initial inventory and I^m , $m = 1, \dots, M$, denote the inventory level at the m th checkpoint (which is at time p^m in a period). Then, in the inventory checkpoint model IC[M], the newsvendor’s observation for a period is the vector containing inventory level information at checkpoints; that is, $O^{\text{IC[M]}}(y) = \bar{I} = \{I^0, \dots, I^M\}$. Note that the information on both sales and the stock-out indicator during a period can be deduced from $\bar{I}$. Define

$\tilde{m} = \max\{m: I^m > 0, m = 0, 1, \dots, M\}$, which is the last checkpoint at which the newsvendor observes a positive inventory level. Then, with prior $\pi(\theta)$, the likelihood of the observation $O^{\text{IC[M]}}(y)$ is

$$l(O^{\text{IC[M]}}(y) | \theta) = \begin{cases} \prod_{m=1}^M f(I^{m-1} - I^m | p^m - p^{m-1}, \theta) & \text{if } \tilde{m} = M; \\ \prod_{m=1}^{\tilde{m}} f(I^{m-1} - I^m | p^m - p^{m-1}, \theta) \cdot \bar{F}(I^{\tilde{m}} - 1 | p^{\tilde{m}+1} - p^{\tilde{m}}, \theta) & \text{if } 0 < \tilde{m} < M; \\ \bar{F}(y - 1 | p^1, \theta) & \text{if } \tilde{m} = 0. \end{cases} \quad (8)$$

The case $\tilde{m} = M$ implies that no stock-out occurs during the period. In this case, the newsvendor obtains, at each checkpoint m , the exact demand arrivals in the preceding time interval $(p^{m-1}, p^m]$ of length $p^m - p^{m-1}$ by tracking the change in inventory level $I^{m-1} - I^m$ (i.e., the newsvendor observes exact demand over all M time intervals). Thus, the likelihood of this observation over the period is the product of the likelihoods of exact demand observations at all M intervals. In the second, where $0 < \tilde{m} < M$, a stock-out takes place between checkpoints $p^{\tilde{m}}$ and $p^{\tilde{m}+1}$. In this case, exact demand arrivals are observed for time intervals only up to the checkpoint $\tilde{m}$; for the time interval $(p^{\tilde{m}}, p^{\tilde{m}+1}]$ immediately following that checkpoint, the newsvendor can infer only that demand exceeded $I^{\tilde{m}}$ (i.e., the demand observation in this interval is censored), for which the likelihood is $1 - F(I^{\tilde{m}} - 1 | p^{\tilde{m}+1} - p^{\tilde{m}}, \theta)$. Multiplying the likelihoods of all $\tilde{m} + 1$ observations then gives the expression for likelihood of this observation in (8). In the third case, $\tilde{m} = 0$, a stock-out occurs in the interval $(0, p^1]$ and so the newsvendor obtains only a censored observation of demand during this interval.

Observations of Stock-out Status. Next, consider the stock-out checkpoint model SC[M], in which the newsvendor collects information on whether a stock-out has occurred at each checkpoint in $\mathbf{M}$. In other words, the observation made at each checkpoint m is of the stock-out indicator variable $\varepsilon^m = \mathbf{1}_{\{I^m > 0\}}$. Since this variable is binary and nondecreasing in nature, it follows that the information contained in the vector $\{\varepsilon^0, \varepsilon^1, \dots, \varepsilon^M\}$ is sufficiently captured by the simple statistic $\tilde{m}$, which is the last checkpoint with positive inventory. Thus the observation made in a period is $O^{\text{SC[M]}}(y) = \{s, \tilde{m}\}$, and its likelihood is given by

$$l(O^{\text{SC[M]}}(y) | \theta) = \begin{cases} f(s | 1, \theta) & \text{if } \tilde{m} = M, s < y; \\ \sum_{d_{\tilde{m}}=0}^{y-1} f(d_{\tilde{m}} | p^{\tilde{m}}, \theta) \cdot \bar{F}(y - d_{\tilde{m}} - 1 | p^{\tilde{m}+1} - p^{\tilde{m}}, \theta) & \text{if } 0 < \tilde{m} < M; \\ \bar{F}(y - 1 | p^1, \theta) & \text{if } \tilde{m} = 0. \end{cases} \quad (9)$$

This likelihood is obtained by unconditioning the likelihood for the inventory checkpoint model given in (8) over intermediate inventory levels $\{I^m\}_{m \leq \bar{m}}$, which is equivalent to taking expectation over demand occurred in the preceding checkout intervals.

Our next proposition gives the ordering of expected profits under discrete observations as the set of checkpoints $\mathbf{M}$ is refined; and it also extends Proposition 1 to the two checkpoint models. The proof of this result is much like that of Proposition 1 and so is omitted.

PROPOSITION 3. *For any given π_n at the beginning of period n , the following statements hold:*

- (a) $V_n^E(\pi_n) = V_n^{SC[\{1\}]}(\pi_n) = V_n^{IC[\{1\}]}(\pi_n)$;
- (b) *for $\mathbf{M} \supseteq \mathbf{M}'$, $V_n^{IC[\mathbf{M}]}(\pi_n) \geq V_n^{IC[\mathbf{M}']}(\pi_n)$ and $V_n^{SC[\mathbf{M}]}(\pi_n) \geq V_n^{SC[\mathbf{M}']}(\pi_n)$;*
- (c) $V_n^E(\pi_n) \leq V_n^{SC[\mathbf{M}]}(\pi_n) \leq V_n^{IC[\mathbf{M}]}(\pi_n) \leq V_n^T(\pi_n) \leq V_n^F(\pi_n)$;
- (d) $y_n^{IC[\mathbf{M}]}(\pi_n) \geq y^M(\pi_n)$.

Part (a) indicates that with a single checkpoint (i.e., $\mathbf{M} = \{1\}$), both checkpoint models reduce to the model with event observations. Part (b) establishes that, under both checkpoint models, expected profit increases as more checkpoints are added, which makes sense because more checkpoints result in improved demand information. For the same set of checkpoints, the observations obtained in the inventory checkpoint model carry more information than those in the stock-out checkpoint model; hence the optimal profit of the former is larger (this is the second inequality in part (c)). Both checkpoint models yield more information than the model with event observations but less than the model with timing observations. These two facts together imply the first and the third inequalities in part (c). Part (d) states that the newsvendor over-orders in the inventory checkpoint model. This claim follows because, in that model, a higher level of initial inventory results in better information.

COROLLARY 2. *Suppose the checkpoints are evenly spaced on the interval $[0, 1]$, then*

- (a) $\lim_{M \rightarrow \infty} V_n^{SC[\mathbf{M}]}(\pi_n) = V_n^{ST}(\pi_n)$ and $\lim_{M \rightarrow \infty} V_n^{IC[\mathbf{M}]}(\pi_n) = V_n^T(\pi_n)$;
- (b) *for a Poisson demand process with unknown rate λ or a normal demand process with unknown mean μ ,*
 $\lim_{M \rightarrow \infty} V_n^{SC[\mathbf{M}]}(\pi_n) = \lim_{M \rightarrow \infty} V_n^{IC[\mathbf{M}]}(\pi_n) = V_n^{ST}(\pi_n) = V_n^T(\pi_n)$.

As the number of checkpoints increases to infinity, the observations of inventory at discrete checkpoints converge to the observation of sale occurrences during a period (i.e., timing observation). In contrast, the observations of stock-out status at discrete checkpoints converge to the observation of the stock-out timing. Combining the results in Proposition 3(a) and Corollary 2, we can see the checkpoint models essentially bridge the event and timing models, and gradual transition from one to the other can be achieved by adding/removing more checkpoints.

5. Numerical Illustration


In this section we illustrate numerically the models studied in §§2–4. We shall consider a Poisson demand process with unknown arrival rate λ , where the prior distribution of λ is gamma with shape and rate parameters (α, β) . We choose values of the model parameters to span a wide spectrum of practical scenarios. The reciprocal of the rate parameter $1/\beta$ (also known as the scale parameter), which captures the degree of uncertainty about λ , takes values $1/\beta \in \{0.5, 1, 2, 4, 8, 16\}$. The shape parameter α takes values such that the expected arrival rate $\mathbb{E}[\lambda] = \alpha/\beta \in \{10, 20, 30, 40, 50\}$. The selling price r is fixed at 2, and the unit cost c varies from 0.2 to 1.8 in increments of 0.1. Thus the newsvendor ratio, $(r - c)/r$, takes 17 different values ranging from 0.1 to 0.9 in increments of 0.05. To conduct an exact comparison of all the discussed models (i.e., those with event, timing, and full observations as well as the inventory and stock-out checkpoint models), we implement the models for each of these 510 scenarios and solve their corresponding dynamic programs specified in (3) for $N = 4$ periods via backward induction. In each recursion, the optimal inventory level is obtained by a linear search from its lower bound—the myopic inventory level—to the upper bound given in Proposition 1(c). More details on the issues related to computation are provided in the online supplement (available as supplemental material at <http://dx.doi.org/10.1287/opre.2014.1326>). We also simulate the myopic policies for a longer horizon $N = 100$ and illustrate their usability in practical settings.

In describing our numerical results, we consider the most widely studied event observations as our benchmark case with censored demand and consider the full observations as the ideal case. Thus, $V_1^F - V_1^E$ is the benchmark loss in expected profit because of censoring, and $V_1^X - V_1^E$ is the amount of loss recovered using observation type $X \in \{T, IC[\mathbf{M}], SC[\mathbf{M}]\}$. We focus on the relative performance measure $\eta^X = [(V_1^X - V_1^E)/(V_1^F - V_1^E)]\%$, which represents the percentage recovery in the loss of expected profit by using observation type X .

5.1. Performance of Timing Observations

We evaluate the performance of our model with timing observations by comparing its optimal order quantity and optimal expected profit with those of models with full and event observations. In Proposition 1 we established that, compared with the case of full observations, the newsvendor over-orders in the case of either event or timing observations. It is worthwhile to compare the magnitude of over-ordering in these latter two models. In Figure 3(a) we plot the first-period optimal inventory decisions y_1^E , y_1^T , and y_1^F against the newsvendor ratio for a set of scenarios with $\alpha = 10/16$ and $\beta = 1/16$ (i.e., $\mathbb{E}[\lambda] = 10$). The order quantities y_1^E , y_1^T , and y_1^F are scale dependent and so, to obtain a better representation of the magnitude of over-ordering, we evaluate the service level measure of the targeted in-stock

Figure 3. Comparison of optimal inventory decisions under the model with event, timing, and full observations.




probability for the first period: $SL(y) = F(y | 1, \pi)$. Let $\delta^E = SL(y_1^E) - SL(y_1^F)$ and $\delta^T = SL(y_1^T) - SL(y_1^F)$; then δ^E and δ^T capture the increase in in-stock probability due to over-ordering. Figure 3(b) uses solid lines to plot the average values of δ^E and δ^T as functions of the newsvendor ratio $(r - c)/r$; the figure also has two dotted lines for the 90th percentile of δ^E and δ^T . Comparing δ^E and δ^T in Figure 3(b) reveals that, on average, the use of timing observations results in significantly less over-ordering than does the use of only event observations. This result is expected because the learning mechanism with timing observations is more efficient and therefore lessens the need to compensate by over-ordering. In Figure 3 we can also observe that the magnitude of over-ordering is much smaller for higher newsvendor ratios. This finding is intuitive: a higher newsvendor ratio implies a lower probability of stock-out, so less over-ordering is required in order to compensate for the loss of observations due to censoring.

A natural question that arises is whether the optimal inventory level with event observations, y^E , is *never* less than that level with timing observations, y^T . Our numerical investigation yields a negative answer. In particular, if the newsvendor ratio is 0.1, then Figure 3(a) shows that $y^T > y^E$. Still more counterexamples can be found when both β and the newsvendor ratio are low; such values describe scenarios characterized by high demand uncertainty and a low profit margin. Under these conditions, a myopic newsvendor is better-off not ordering at all because ordering a positive quantity leads to a negative expected profit in the current period—although a newsvendor more concerned with future than current profits might nonetheless order simply for the sake of learning about demand. There are cases in which the gain in future profits dominates the immediate loss when learning occurs via timing but *not* via event observations, resulting in $y^T > y^E$.

Figure 4 illustrates the comparison of expected profits V_1^E , V_1^T , and V_1^F with event, timing, and full observations. In Figure 4(a), optimal expected profits V_1^E , V_1^T , and V_1^F are plotted against the newsvendor ratio for the set of scenarios identical to those in Figure 3(a): $\alpha = 10/16$, $\beta = 1/16$, $\mathbb{E}[\lambda] = 10$. Evaluating performance in terms of η^T , the percentage recovery of loss with timing observations, in most cases we find that the timing observations recover a major portion of the loss in profit. Across the 510 examples in our study, the mean and median of η^T are 76.1% and 77.7%, respectively. Moreover, in 90% of the cases η^T exceeds 62.7% and in 75% of the cases it exceeds 71.2%. Figure 4(b) displays the mean (solid line) as well as the 50% (dark gray area) and 90% (light gray area) intervals of η^T grouped by newsvendor ratio. This graph illustrates that the use of timing observations enables the newsvendor to recover a large portion of the loss in expected profit that is due to demand censoring. Moreover, this loss recovery of loss via timing observations is robust across the whole range of newsvendor ratios.

We also consider the so-called *myopic timing* model, wherein demand is updated via the timing observations; but order quantities in each period are determined by optimizing single-period expected profits (i.e., a single-period newsvendor quantity based on updated demand information is ordered in each period). In other words, under this model the order quantity decisions ignore the effect of learning on expected profits in subsequent periods. Let $\hat{V}_1^T$ denote the expected profit of the myopic timing model and let $\hat{\eta}^T = [(\hat{V}_1^T - V_1^E)/(V_1^F - V_1^E)]\%$ denote its performance in terms of percentage recovery of loss in expected profit. In our numerical examples, we find that the value of $\hat{\eta}^T$ is close to that of η^T : its average and median values are 75.0% and 77.2%, respectively. Furthermore, in 90% of the cases $\hat{\eta}^T$ exceeds 60.9% and in 75% of the cases it exceeds

Figure 4. Comparison of optimal profits of models under event, timing, and full observations.

70.3%. Figure 4(b) also plots (using a dashed line) average values of $\hat{\eta}^T$ by newsvendor ratio. The implication is that using timing observations for learning alone usually achieves most of the benefit and that accounting for the optimization adds relatively less value. There are few cases where the value of $\hat{\eta}^T$ deviates significantly from η^T : when both demand uncertainty is high (large $1/\beta$) and the relative profit margin is small (i.e., a low newsvendor ratio).

5.2. Performance of Checkpoint Models

We consider a series of checkpoint models with checkpoint sets $\mathbf{M}_k, k = 0, 1, \dots, 5$, where the set $\mathbf{M}_k$ contains $M = 2^k$ evenly spaced checkpoints in a period. For example $\mathbf{M}_0 = \{1\}$, $\mathbf{M}_1 = \{0.5, 1\}$, $\mathbf{M}_2 = \{0.25, 0.5, 0.75, 1\}$, and so forth. Such a power-of-two structure assures a strictly increasing set of checkpoints as k increases (i.e., $\mathbf{M}_k \subset \mathbf{M}_{k+1}$) and, together with Proposition 3, implies that the expected profits of both the inventory and stock-out checkpoint models are increasing in k . In Figure 5(a), we plot the percentage recovery in loss of expected profit under our two checkpoint models, $\eta^{IC[\mathbf{M}_k]}$ and $\eta^{SC[\mathbf{M}_k]}$ (in solid lines), and their myopic counterpart, $\hat{\eta}^{IC[\mathbf{M}_k]}$ and $\hat{\eta}^{SC[\mathbf{M}_k]}$ (in dashed lines), against the number of checkpoints in $\mathbf{M}_k$. The parameter values in this figure are $r = 2$, $c = 1$, $\alpha = 10$, and $\beta = 1$. The figure illustrates the transition from the case $M = 1$, where both checkpoint models reduce to the event model (Proposition 3(a)), to the cases with larger M , where both models converge to the timing model (Corollary 2). A similar observation can be made on the myopic models, too. The plot also suggests that, for each checkpoint model, convergence takes place in a concave fashion: moving from one to two checkpoints results in a large increment in expected profit, but further increases in the number of checkpoints yield smaller profit increments. This implies

that a newsvendor who cannot observe inventory continuously to make timing observations can still achieve much of the possible benefit by making a few discrete observations within a period.

To evaluate the overall performance of the checkpoint models, in Figure 5(b) we include all the models and plot their average loss recovery η across all 510 scenarios. So in addition to the timing and myopic timing models investigated previously in Figure 4(b), we also include stock-out and inventory checkpoint models with 2 and 4 checkpoints. Computational constraints induce us to focus on the myopic checkpoint models. Just as in the myopic timing model, demand learning occurs in response to the respective checkpoint observations but inventory optimization is performed myopically (i.e., without considering future periods). It is evident from the figure that a significant portion of the loss due to censored observations can be recovered even with only a few checkpoints and a myopic inventory policy.

5.3. Simulation of Myopic Policies

The analysis above, which is based on solving the exact dynamic programs, provides accurate performance estimates of the models discussed. For real-world applications, what is also important is to have a policy that is easy to calculate and implement. In the following, we focus on the myopic policies, which has been shown to perform reasonably well in the 4-period scenarios, and conduct a simulation study of a practical sized problem with $N = 100$ periods. Note that in each period, the computations required by the myopic timing model consist of (i) updating posterior by incorporating the new observation into the two-dimensional sufficient statistic according to (7) and (ii) solving for the single-period newsvendor quantity under the new belief. Both computations consume constant

Figure 5. Illustration of the checkpoint models.





Figure 6. Simulated profit gaps of the myopic policies.



time, hence the overall computational complexity is *linear* in the number of periods N . However, for the myopic event model, because of lack of conjugate priors and fixed-dimensional sufficient statistics, the whole history needs to be tracked and the posterior is updated by evaluating (2) by brute force every period.

We simulate the models with full (F), timing (T), and event (E) observations, together with a benchmark model with perfect knowledge about the underlying demand distribution (P). All models start with the same prior belief $\lambda \sim \text{Gamma}(\alpha, \beta)$. In each simulation instance, a sample path of the 100-period demand realization, including individual customer arrivals within each period, is first simulated.

In each period, full, event, and timing observations are constructed based on the sample path, beliefs are updated accordingly, myopically optimal inventory levels are calculated, and profits are accounted for. We simulate a total of 1,000,000 such instances, and obtain average profit of each model in each period, W_n^X for $X \in \{P, F, T, E\}$ and $n = 1, \dots, 100$.

We define profit gap in period n as percentage profit difference from the perfect model, $\hat{\xi}_n^X = (W_n^P - W_n^X)/W_n^P$, and plot them as a function of period n (on a logarithm scale) in Figure 6. Figures 6(a) and 6(b) are two scenarios with the same prior (α, β) as in Figures 3(a) and 4(a) but different unit cost c . The figures illustrate as knowledge accumu-

lates, how the single-period newsvendor profits under the full, event, and timing model improve over time and eventually converge to the profit under perfect knowledge.

6. Concluding Remarks and Extensions

Understanding and estimating the demand process are among the most important challenges in the field of operations management. In industry, there is an increasing focus on managing stock-outs that stems, in part, from improved access to quality data through the widespread adaptation of ERP systems. Clearly, the most difficult (and least salient) cases of estimating demand arise when demand is not fully observed, a situation that occurs when a product is out of stock and hence demand is censored by insufficient inventory.

This paper makes an essential contribution to understanding and managing demand estimation in the presence of demand censoring. In particular, we consider the timing of sales occurrences and show that the use of such *timing* observations can mitigate the information loss because of demand censoring (as compared with using only data about the stock-out *event*). This finding has important managerial implications. First and foremost, it suggests that, by using readily available information on the timing of sales observations, managers can recover losses of information and profit that are due to censoring. Second, if managers use such timing observations to update demand but still order the single-period newsvendor quantity in each period (i.e., if they ignore the implications of the current period's learning on future profits and hence do not over-order), then they will still gain most of the benefits of ordering optimally. In other words, improved estimation largely obviates the need for multiperiod optimization. Finally, we show that the aforementioned insights from timing observations are also applicable to settings where the newsvendor may not be able to track inventory and stock-out status continuously but can nonetheless make observations at discrete checkpoints within a time period.

The results of our parsimonious model suggest several follow-up studies. A natural extension would be to consider a multiproduct setting that involves uncertainties with respect to the arrival times and basket compositions. A retailer could combine timing data from multiple sources (traffic counters, POS information, etc.) to improve its knowledge about the patterns of customer arrival at the store and about their purchase decisions once they are in the store. When one or more products stock out, timing observations enable the retailer to observe the purchase decisions of customers who face a different product assortment. Thus, such a model would improve not only inventory decision but also assortment decisions. Another extension would be to consider multiple stores and allow learning about demand over multiple locations using larger data sets. Our checkpoint models illustrate that much of the increase in profit can be achieved by observing inventory/stock-out

status just once within a period. It would be worthwhile to examine the selection of when that observation should occur during a period in order to gain the maximum benefit.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/opre.2014.1326>.

Acknowledgments

The authors thank the editor, the area editor, the associate editor, and two anonymous referees for their constructive feedback, which significantly improved the paper. The work of the third author is supported in part by the National University of Singapore Academic Research Fund [R-314-000-084-112 and R-314-000-098-112].

Appendix

Proof of Proposition 1

The proof of Proposition 1 uses results from the literature on comparison of experiments (see Ginebra 2007). In the following, we first describe our three observations as outcomes of three statistical experiments; we then establish the sufficiency ordering of these experiments and prove Proposition 1 with the ordering so obtained.

DEFINITION 1 (GINEBRA 2007). A *statistical experiment* $\mathcal{E} = \{(X, S_X); (P_\theta, \Theta)\}$ ($\mathcal{E} = \{X; P_\theta\}$ for short) yields an observation on a random variable X defined on S_X , with an unknown probability distribution that is known to be in the family $(P_\theta, \theta \in \Theta)$.

The three observation types in §2 can be described as outcomes of statistical experiments defined as follows.

(i) *Observation of sales quantity and stock-out event* ($O^E(y)$). This is an experiment right-censored by inventory level y . Denote the experiment by $\mathcal{E}^E(y) = \{O^E(y); P_\theta^E\}$, where $O^E(y) = \{s, e\}$, $s = \min\{D, y\}$, and $e = \mathbf{1}_{\{D \geq y\}}$. The joint distribution P_θ^E of observation $\{s, e\}$ is given by (4).

(ii) *Observation of sales occurrences* ($O^T(y)$). We have an experiment $\mathcal{E}^T(y) = \{O^T(y); P_\theta^T\}$, where $O^T(y) = \{s, e, \tilde{r}^s\}$ for $s = \min\{D, y\}$, $e = \mathbf{1}_{\{D \geq y\}}$, and $\tilde{r}^s = (t^1, t^2, \dots, t^s)$. We can also derive $\tilde{r}^s = (\tau^1, \tau^2, \dots, \tau^s)$ as $\tau^i = t^i - t^{i-1}$. The joint distribution is given by (5).

(iii) *Full observation of all demand occurrences* (O^F). The experiment with full observations is $\mathcal{E}^F = \{O^F; P_\theta^F\}$, where $O^F = \{D, \tilde{r}^D\}$ and $\tilde{r}^D = (t^1, t^2, \dots, t^D)$. The joint distribution is given by (6).

Next we define and establish sufficiency ordering among these three experiments.

DEFINITION 2 (GINEBRA 2007). Experiment $\mathcal{E} = \{X; P_\theta\}$ is *sufficient* for experiment $\mathcal{F} = \{Y; Q_\theta\}$ if there is a stochastic transformation of X to a random variable $\mathcal{S}(X)$ such that $\mathcal{S}(X)$ and Y have identical distributions under each $\theta \in \Theta$.

Lemma 1 establishes ordering of the two $\mathcal{E}^E(y)$ experiments with different initial inventory levels y .

LEMMA 1. Let $\mathcal{E}^E(y) = \{O^E(y); P_\theta^E\}$ and $\mathcal{E}^E(y') = \{O^E(y'); Q_\theta^E\}$ be two experiments with event observations. If $y \geq y'$, then $\mathcal{E}^E(y)$ is sufficient for $\mathcal{E}^E(y')$.

This lemma follows directly from Theorem 3.1 of Goel (1988). The theorem states that, if only sales quantity is observed, then a larger initial inventory leads to less censoring and hence to more information. We can similarly show that, for timing observations, the amount of information is also increasing in the initial inventory level.

LEMMA 2. Let $\mathcal{E}^T(y) = \{O^T(y); P_\theta^T\}$ and $\mathcal{E}^T(y') = \{O^T(y'); Q_\theta^T\}$ be experiments with timing observations. If $y \geq y'$, then $\mathcal{E}^T(y)$ is sufficient for $\mathcal{E}^T(y')$.

PROOF. Experiment $\mathcal{E}^T(y)$ has the observation $O^T(y)$ and experiment $\mathcal{E}^T(y')$ has the observation $O^T(y')$. We will construct a stochastic transformation from $O^T(y) = \{s, e, \tau^s\}$ to $O^{T*} = (s^*, e^*, \tau^*)$ such that O^{T*} has the same distribution as $O^T(y')$. Consider the following transformation.

1. If $e = 1$, $s = y$, and $\tau^s = (\tau^1, \tau^2, \dots, \tau^y)$, then let $e^* = 1$, $s^* = y'$, and $\tau^* = (\tau^1, \tau^2, \dots, \tau^{y'})$.
2. If $e = 0$, $s = d$, and $\tau^s = (\tau^1, \tau^2, \dots, \tau^d)$, then
 - (a) if $y' \leq d < y$, let $e^* = 1$, $s^* = y'$, and $\tau^* = (\tau^1, \tau^2, \dots, \tau^{y'})$;
 - (b) if $0 \leq d < y'$, let $e^* = 0$, $s^* = d$, and $\tau^* = (\tau^1, \tau^2, \dots, \tau^d)$.

It can be readily verified that O^{T*} has the same distribution as $O^T(y')$, which completes the proof. $\square$

LEMMA 3. For any given inventory level y , the experiment $\mathcal{E}^F$ is sufficient for $\mathcal{E}^T(y)$ and experiment $\mathcal{E}^T(y)$ is sufficient for $\mathcal{E}^E(y)$.

PROOF. Experiment $\mathcal{E}^F$ can be viewed as an experiment $\mathcal{E}^T(y)$ with an infinite censoring level $y \rightarrow \infty$ or $E^T(\infty)$. Then Lemma 2 implies that $\mathcal{E}^F$ is sufficient for $\mathcal{E}^T(y)$. With the same inventory level y , experiments $\mathcal{E}^T(y)$ and $\mathcal{E}^E(y)$ differ only with respect to whether or not the arrival timing is observed; the observations in $\mathcal{E}^T(y)$ are a superset of those in $\mathcal{E}^E(y)$. By a randomization argument similar to the one used in the proof of Lemma 2, we can show that $\mathcal{E}^T(y)$ is sufficient for $\mathcal{E}^E(y)$. $\square$

Finally, we combine the sufficiency ordering results obtained so far with the following result, which is restatement of Proposition 3.2 in Ginebra (2007).

LEMMA 4 (GINEBRA 2007). Experiment $\mathcal{E}$ is sufficient for experiment $\mathcal{F}$ if and only if, for every decision problem involving θ , the Bayes risk for $\mathcal{E}$ does not exceed the Bayes risk for $\mathcal{F}$ —that is, iff $r_\pi(\mathcal{E}) \leq r_\pi(\mathcal{F})$ for every $\pi(\theta)$.

PROOF OF PROPOSITION 1(B). The equality $y_n^F(\pi_n) = y_n^M(\pi_n)$ follows directly from the dynamic program. Here we will show $y_n^T(\pi_n) \geq y_n^M(\pi_n)$; a similar proof can be obtained for $y_n^E(\pi_n) \geq y_n^M(\pi_n)$. Since $y_n^M(\pi_n)$ maximizes $W(y_n | \pi_n)$ and $y_n^T(\pi_n)$ maximizes $W(y_n | \pi_n) + \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y_n))]$, it suffices to show that $\mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y_n))]$ is increasing in y_n . The dynamic program can be written as

$$V_n^T(\pi_n) = \max_{y_n \geq 0} \mathbb{E}[W_n(y_n, D_n | \pi_n)], \quad n = 1, \dots, N;$$

$$W_n(y_n, d_n | \pi_n) = \mathcal{W}(y_n, d_n) + V_{n+1}^T(\pi_n \oplus O_n^T(y_n)),$$

$$n = 1, \dots, N.$$

Here $\mathcal{W}(y, d) = r \min\{d, y\} - cy$ (i.e., $W(y | \pi) = \mathbb{E}_{D|\pi}[\mathcal{W}(y, D)]$), and $W_n(\cdot)$ is an auxiliary function for the profit-to-go in period n with a given inventory level and demand realization. In any period n , a newsvendor faces the following

statistical decision problem. For given knowledge π_n and an experiment $\mathcal{E}^T(y_n)$, the newsvendor incurs an experimental cost $W(y_n | \pi_n)$, observes $O_n^T(y_n)$, updates her knowledge to π_{n+1} , and optimizes the action y_{n+1} so as to maximize her future payoff $\mathbb{E}[W_{n+1}(y_{n+1}, D_{n+1} | \pi_{n+1})]$. The corresponding Bayes risk is

$$\begin{aligned} r_{\pi_n}(\mathcal{E}^T(y_n)) &= W_{n+1}^* - \mathbb{E} \left[\min_{y_{n+1} \geq 0} \mathbb{E}[W_{n+1}(y_{n+1}, D_{n+1} | \pi_{n+1})] \right] \\ &= W_{n+1}^* - \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y_n))], \end{aligned}$$

where W_{n+1}^* denotes the expected profit-to-go for the ideal case in which there is no uncertainty about θ . Consider $y'_n \leq y_n$; by Lemma 2, we know that experiment $\mathcal{E}^T(y_n)$ is sufficient for $\mathcal{E}^T(y'_n)$. Hence, by Lemma 4, $r_{\pi_n}(\mathcal{E}^T(y_n)) \leq r_{\pi_n}(\mathcal{E}^T(y'_n))$. This inequality implies that $\mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y_n))] \geq \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y'_n))]$, completing the proof. $\square$

PROOF OF PROPOSITION 1(A). For the same sequence of inventory levels $y_n, y_{n+1}, \dots, y_N$, from Lemma 3 it follows that, in all periods from n to N , experiment $\mathcal{E}^F$ is sufficient for $\mathcal{E}^T(y_n)$ and $\mathcal{E}^T(y_n)$ is sufficient for $\mathcal{E}^E(y_n)$. Thus, by Lemma 4, the expected profits-to-go for the three scenarios under the same sequence of inventory levels follow the same order. Because this order holds for any sequence of inventory levels, the optimal expected profits satisfy $V_n^E(\pi_n) \leq V_n^T(\pi_n) \leq V_n^F(\pi_n)$. $\square$

PROOF OF PROPOSITION 1(C). Consider the model with timing observation $O^T(y)$. In any period n with prior π_n , an inventory level $y \geq y^M(\pi_n)$ cannot be optimal if

$$\begin{aligned} &\mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y))] - \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y-1))] \\ &< W(y-1 | \pi_n) - W(y | \pi_n); \end{aligned}$$

in other words, the marginal benefit is dominated by the marginal cost. Note that the marginal cost on the right-hand side is increasing (by concavity of the single-period profit function) and converges to c as y approaches infinity. The left-hand side need not be monotone, but it is bounded from above by $\mathbb{E}[V_{n+1}^T(\pi_n \oplus O^F)] - \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y-1))]$ because it is always better to replace a timing observation $O^T(y)$ with full observation O^F . The upper bound is decreasing in y and converges to 0, as already proved. So there must exist a $\bar{y}$ such that

$$\begin{aligned} &\mathbb{E}[V_{n+1}^T(\pi_n \oplus O^F)] - \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(\bar{y}))] \\ &< W(\bar{y} | \pi_n) - W(\bar{y} + 1 | \pi_n). \end{aligned}$$

By monotonicity, it then follows that, for all $y \geq \bar{y} + 1$,

$$\begin{aligned} &\mathbb{E}[V_{n+1}^T(\pi_n \oplus O^F)] - \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y))] \\ &< W(y | \pi_n) - W(y + 1 | \pi_n); \end{aligned}$$

this inequality implies that y cannot be optimal. Let $G_n^T(y | \pi_n) = W(y | \pi_n) + \mathbb{E}[V_{n+1}^T(\pi_n \oplus O_n^T(y))]$ and let its upper bound $\bar{G}_n^T(y | \pi_n) = W(y | \pi_n) + \mathbb{E}[V_{n+1}^T(\pi_n \oplus O^F)]$. Then the upper bound of the optimal inventory level $\bar{y}^T(\pi_n)$ is the smallest y satisfying

$$y \geq y^M(\pi_n) \quad \text{and} \quad G_n^T(y | \pi_n) > \bar{G}_n^T(y + 1 | \pi_n).$$

A similar bound can be derived for the event model; we omit the details. $\square$

Proof of Proposition 2

We prove part (a) and (b) for Poisson and normal demand below. Part (c) follows directly from these.

Poisson Demand. Recall that the demand arrival process $\{D(t), 0 \leq t \leq 1\}$ is a Poisson process with unknown rate λ and that $\lambda \sim \text{Gamma}(\alpha, \beta)$. The timing observation $O^T(y) = \{s, e, \bar{r}^s\}$ has a likelihood specified in (5). When there is no stock-out ($e = 0, s < y$), the newsvendor observes s interarrival times $(\tau^1, \dots, \tau^s)$ and the last time interval $1 - t^s$ during which no arrival is observed. We can update the posterior by applying the Bayes formula (2):

$$\begin{aligned}\pi'(\lambda | O^T(y)) &\propto l(O^T(y) | \lambda) \cdot \pi(\lambda) \\ &= \prod_{i=1}^s \psi(\tau^i | \lambda) \bar{\Psi}(1 - t^s | \lambda) \pi(\lambda) \\ &\propto \prod_{i=1}^s \lambda e^{-\lambda \tau_i} e^{-\lambda(1-t^s)} \cdot \lambda^{\alpha-1} e^{-\lambda\beta} \\ &\propto \lambda^s e^{-\lambda(\sum_{i=1}^s \tau_i + 1 - t^s)} \lambda^{\alpha-1} e^{-\lambda\beta} \propto \lambda^{\alpha+s-1} e^{-\lambda(\beta+1)},\end{aligned}$$

which is essentially the kernel of $\text{gamma}(\alpha + s, \beta + 1)$. If a stock-out does occur ($e = 1, s = y$), then

$$\begin{aligned}\pi'(\lambda | O^T(y)) &\propto l(O^T(y) | \lambda) \cdot \pi(\lambda) = \prod_{i=1}^s \psi(\tau^i | \lambda) \cdot \pi(\lambda) \\ &\propto \prod_{i=1}^s \lambda e^{-\lambda \tau_i} \cdot \lambda^{\alpha-1} e^{-\lambda\beta} \propto \lambda^s e^{-\lambda(\sum_{i=1}^s \tau_i)} \cdot \lambda^{\alpha-1} e^{-\lambda\beta} \\ &\propto \lambda^{\alpha+s-1} e^{-\lambda(\beta+t^s)},\end{aligned}$$

which is $\text{gamma}(\alpha + s, \beta + t^s)$.

Normal demand. Suppose the demand in a period follows a normal distribution with mean μ and precision ν (the reciprocal of variance). We can open up the period and model demand arrival as a Brownian motion with drift $D(t) = \mu t + W(t)/\sqrt{\nu}$, where $W(t)$ is a standard Brownian motion. Demand up to any time t is a normal random variable with mean μt and precision ν/t (thus demand for the entire period $D \sim \text{Normal}(\mu, \nu)$). The precision ν is known; the mean μ is unknown and is the parameter of interest. The newsvendor employs a normal conjugate prior for μ ; that is, $\mu \sim \text{Normal}(m, \nu_0)$. Formally,

$$\begin{aligned}f(d | t, \mu) &= \sqrt{\frac{\nu}{2\pi t}} \exp\left\{-\frac{\nu(d - \mu t)^2}{2t}\right\}, \\ \pi(\mu) &= \sqrt{\frac{\nu_0}{2\pi}} \exp\left\{-\frac{\nu_0(\mu - m)^2}{2}\right\}.\end{aligned}$$

We remark that, since Brownian motion is a continuous process, our original definition of timing observation $O^T(y) = \{s, e, \bar{r}^s\}$ no longer applies. Now $O^T(y)$ represents the entire sample path of the Brownian motion $\{D(t)\}$ until it hits y . Because of censoring, y can be viewed as an absorbing state: once the Brownian motion hits y , it stays there forever and so no further observations can be made.

Since Brownian motions are infinitely divisible, the observation of a sample path is infinite; therefore, learning from such observations is a concept with only theoretical meaning. The following lemma proves that all we need in order to learn the unknown drift is the last effective observation on the sample path.

LEMMA 5. For a Brownian motion $D(t)$ defined as before, if there are two observations $D(t_1) = x_1$ and $D(t_2) = x_2$ made at time $t_1 \leq t_2$, then the posterior distribution obtained by using both observations, $\pi'(\mu | D(t_1) = x_1, D(t_2) = x_2)$, is the same as one obtained by using only the last observation, $\pi'(\mu | D(t_2) = x_2)$.

PROOF. Suppose we use both observations, $O = \{D(t_1) = x_1, D(t_2) = x_2\}$. The likelihood is then

$$\begin{aligned}l(O | \mu) &= \Pr(D(t_1) = x_1, D(t_2) = x_2 | \mu) \\ &= \Pr(D(t_1) = x_1 | \mu) \cdot \Pr(D(t_2 - t_1) = x_2 - x_1 | \mu) \\ &\propto \exp\left\{-\frac{\nu(x_1 - t_1\mu)^2}{2t_1}\right\} \\ &\quad \cdot \exp\left\{-\frac{\nu((x_2 - x_1) - (t_2 - t_1)\mu)^2}{2(t_2 - t_1)}\right\}.\end{aligned}$$

Recall that the normal prior $\pi(\mu) \propto \exp\{-\nu_0(\mu - m)^2/2\}$. We have the posterior

$$\begin{aligned}\pi'(\mu | O) &\propto l(O | \mu) \cdot \pi(\mu) \\ &\propto \exp\left\{-\frac{\nu(x_1 - t_1\mu)^2}{2t_1}\right\} \cdot \exp\left\{-\frac{\nu((x_2 - x_1) - (t_2 - t_1)\mu)^2}{2(t_2 - t_1)}\right\} \\ &\quad \cdot \exp\left\{-\frac{\nu_0(\mu - m)^2}{2}\right\} \\ &\propto \exp\left\{-\frac{(\nu_0 + t_2\nu)\mu^2 - 2(m\nu_0 + x_2\nu)\mu}{2}\right\} \\ &= \exp\left\{-\frac{(\nu_0 + t_2\nu)(\mu - (m\nu_0 + x_2\nu)/(\nu_0 + t_2\nu))^2}{2}\right\}.\end{aligned}$$

This is essentially the kernel of a normal distribution with mean $(m\nu_0 + x_2\nu)/(\nu_0 + t_2\nu)$ and precision $\nu_0 + t_2\nu$, and it is exactly the same as the posterior derived when we use only the last observation $O' = \{D(t_2) = x_2\}$. $\square$

Lemma 5 concludes that, for the purpose of learning the unknown drift, observing the sample path can be simplified to $O^{\text{ST}}(y) = \{s, e, t^s\}$. Here t^s is the time at which the Brownian motion first hits the absorbing state y , which follows an inverse Gaussian distribution (Pettit and Young 1999). The likelihood of $O^T(y)$ is thus

$$l(O^T(y) | \mu) = \begin{cases} f(s | 1, \mu) & \text{if } e = 0, s = 0, \dots, y-1; \\ g(t^s | s, \mu) & \text{if } e = 1, s = y, 0 < t^s < 1. \end{cases}$$

This expression, $g(t | s, \mu) = \sqrt{\nu s^2/(2\pi t^3)} \exp\{-\nu(s - \mu t)^2/(2t)\}$ is the pdf of the inverse Gaussian distribution.

The Bayesian updating proceeds as follows. If no stock-out occurs ($e = 0$) then the newsvendor observes the true demand $D = s$, after which her knowledge about μ can be updated in the standard way (Raiffa and Schlaifer 2000) for normal random variables with normal priors: $\mu' \sim \text{Normal}(m', \nu')$; here the posterior precision $\nu' = \nu_0 + \nu$ increases by ν , and the posterior mean $m' = (m\nu_0 + s\nu)/(\nu_0 + \nu)$ is the precision-weighted average of prior mean m and the observed demand s . If a stock-out does occur ($e = 1, s = y$) then the newsvendor observes $D(t^s) = s$. As already shown in the proof of Lemma 5, the normal-normal conjugacy can still be preserved: the posterior is of the same form,

$\mu' \sim \text{Normal}(m', \nu')$; the posterior mean is still the precision-weighted average of the prior mean and observed sales; but now the increase in posterior precision $\nu' = \nu_0 + t^s \nu$ is limited by t^s , the duration of sales observations. The resulting posterior is

$$\pi'(\mu | O^T(y)) = \begin{cases} \text{Normal}\left(\frac{m\nu_0 + s\nu}{\nu_0 + \nu}, \nu_0 + \nu\right) & \text{if } e = 0, s = 0, \dots, y-1; \\ \text{Normal}\left(\frac{m\nu_0 + s\nu}{\nu_0 + t^s \nu}, \nu_0 + t^s \nu\right) & \text{if } e = 1, s = y, 0 < t^s < 1. \end{cases}$$

Proof of Corollary 2

Part (a) requires us to show that the likelihood for observations of inventory status $l(O^{\text{IC}[M]}(y) | \theta)$ converges to $l(O^T(y) | \theta)$, and the likelihood for observations of stock-out status $l(O^{\text{SC}[M]}(y) | \theta)$ converges to $l(O^{\text{ST}}(y) | \theta)$, as $M \rightarrow \infty$. The proof for part (b) follows immediately from Proposition 2. To prove part (a), consider a realization of sales occurrences $\tilde{t}^s = (t^1, t^2, \dots, t^s)$ without a stock-out event. Consider an evenly spaced set of checkpoints $\mathbf{M}$. For sales observation i , which takes place at time t^i , let $\underline{p}^i = \max\{p^k: p^k < t^i\}$ and $\bar{p}^i = \min\{p^k: p^k > t^i\}$. Then the likelihood function for inventory checkpoint model can be written as

$$\frac{1}{\prod_{i=1}^s (\bar{p}^i - \underline{p}^i)} \int_{t^s=\underline{p}^s}^{\bar{p}^s} \dots \int_{t^1=\underline{p}^1}^{\bar{p}^1} \left(\prod_{i=1}^s \psi(u^i - u^{i-1} | \theta) \cdot \bar{\Psi}(1 - u^s | \theta) \right) du^1 \dots du^s.$$

In this expression variable u^i represents the timing of i th sales that takes place in the interval $(p^i, p^{i+1}]$. For a given set of checkpoints, the denominator is a constant, therefore does not affect demand updating process (i.e., it cancels out in updating process in (2)). As $M \rightarrow \infty$, the duration of each of these intervals approaches 0, but it always contains the timings of sales occurrences by definition. Given this and continuity of the integrand, it follows from the fundamental theorem of calculus that

$$\begin{aligned} \lim_{M \rightarrow \infty} \frac{1}{\prod_{i=1}^s (\bar{p}^i - \underline{p}^i)} \int_{t^s=\underline{p}^s}^{\bar{p}^s} \dots \int_{t^1=\underline{p}^1}^{\bar{p}^1} \left(\prod_{i=1}^s \psi(u^i - u^{i-1} | \theta) \cdot \bar{\Psi}(1 - u^s | \theta) \right) du^1 \dots du^s \\ = \prod_{i=1}^s \psi(t^i - t^{i-1} | \theta) \bar{\Psi}(1 - t^s | \theta). \end{aligned}$$

The proof for the case with stock-out taking place ($e = 1$) follows similarly.

Let $g(t | y, \theta)$ denote the distribution of time t during which cumulative demand y is observed. Then the likelihood for the observation of stock-out timing is $g(t | y, \theta)$ when a stock-out occurs, and it is $f(y | 1, \theta)$ when no stock-out occurs. For the latter case, the likelihood is the same as that for model SC[M]. For the case of a stock-out, we know that $t^y \in (p^{\bar{m}}, p^{\bar{m}+1}]$, thus likelihood for SC[M] can be expressed as

$$\frac{1}{p^{\bar{m}+1} - p^{\bar{m}}} \int_{p^{\bar{m}}}^{p^{\bar{m}+1}} g(t | y, \theta) dt.$$

Taking limit $M \rightarrow \infty$, while preserving definition of $\bar{m}$, it follows that

$$\lim_{M \rightarrow \infty} \frac{1}{p^{\bar{m}+1} - p^{\bar{m}}} \int_{p^{\bar{m}}}^{p^{\bar{m}+1}} g(t | y, \theta) dt = g(t^y | y, \theta).$$

This completes proof of part (a).

References

- Agrawal N, Smith S (1996) Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.* 43(6):839–861.
- Anupindi R, Dada M, Gupta S (1998) Estimation of consumer demand with stock-out based substitution: An application to vending machine products. *Marketing Sci.* 17(4):406–423.
- Azoury KS (1985) Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* 31(9):1150–1160.
- Bensoussan A, Çakanyildirim M, Sethi SP (2009) Technical note—a note on “the censored newsvendor and the optimal acquisition of information.” *Oper. Res.* 57(3):791–794.
- Besbes O, Muharremoglu A (2013) On implications of demand censoring in the newsvendor problem. *Management Sci.* 59(6):1407–1424.
- Bisi A, Dada M, Tokdar S (2011) A censored-data multiperiod inventory problem with newsvendor demand distributions. *Manufacturing Service Oper. Management* 13(4):525–533.
- Braden DJ, Freimer M (1991) Informational dynamics of censored observations. *Management Sci.* 37(11):1390–1404.
- Bruno HA, Vilcassim NJ (2008) Structural demand estimation with varying product availability. *Marketing Sci.* 27(6):1126–1131.
- Burnetas AN, Smith CE (2000) Adaptive ordering and pricing for perishable products. *Oper. Res.* 48(3):436–443.
- Chen L (2010) Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales. *Oper. Res.* 58(2):396–413.
- Chen L, Plambeck EL (2008) Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing Service Oper. Management* 10(2):236–256.
- Conlon C, Mortimer J (2013) Demand estimation under incomplete product availability. *Amer. Econom. J.-Microeconomics* 5(4):1–30.
- Conrad SA (1976) Sales data and the estimation of demand. *Oper. Res. Quart.* 27(1):123–127.
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B (Methodological)* 39(1):1–38.
- Ding X, Puterman ML, Bisi A (2002) The censored newsvendor and the optimal acquisition of information. *Oper. Res.* 50(3):517–527.
- Farias V, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Sci.* 59(2):305–322.
- Ginebra J (2007) On the measure of the information in a statistical experiment. *Bayesian Anal.* 2(1):167–211.
- Godfrey GA, Powell WB (2001) An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Sci.* 47(8):1101–1112.
- Goel PK (1988) Comparison of experiments and information in censored data. Gupta S, Berger J, eds. *Statistical Decision Theory and Related Topics IV*, Vol. 2 (Springer-Verlag, New York), 335–349.
- Huh WT, Rusmevichientong P (2009) A nonparametric asymptotic analysis of inventory planning with censored demand. *Math. Oper. Res.* 34(1):103–123.
- Huh WT, Janakiraman G, Muckstadt JA, Rusmevichientong P (2009) An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Math. Oper. Res.* 34(2):397–416.
- Huh WT, Levi R, Rusmevichientong P, Orlin JB (2011) Adaptive data-driven inventory control with censored demand based on Kaplan-Meier estimator. *Oper. Res.* 59(4):929–941.
- Iglehart DL (1964) The dynamic inventory problem with unknown demand distribution. *Management Sci.* 10(3):429–440.

- Kök AG, Fisher ML (2007) Demand estimation and assortment optimization under substitution: Methodology and application. *Oper. Res.* 55(6):1001–1021.
- Lariviere MA, Porteus EL (1999) Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* 45(3):346–363.
- Lau H, Lau AH (1996) Estimating the demand distributions of single-period items having frequent stockouts. *Eur. J. Oper. Res.* 92(2):254–265.
- Lu X, Song JS, Zhu K (2008) Analysis of perishable-inventory systems with censored demand data. *Oper. Res.* 56(4):1034–1038.
- Musalem A, Olivares M, Bradlow E, Terwiesch C, Corsten D (2010) Structural estimation of the effect of out-of-stocks. *Management Sci.* 56(7):1180–1197.
- Nahmias S (1994) Demand estimation in lost sales inventory systems. *Naval Res. Logist.* 41(6):739–757.
- Pettit LI, Young KDS (1999) Bayesian analysis of inverse Gaussian lifetimes with measures of degradation. *J. Statist. Comput. Simulation* 63(3):217–234.
- Queenan CC, Ferguson M, Higbie J, Kapoor R (2007) A comparison of unconstraining methods to improve revenue management systems. *Production Oper. Management* 16(6):729–746.
- Raiffa H, Schlaifer R (2000) *Applied Statistical Decision Theory* (Wiley-Interscience, New York).
- Scarf HE (1959) Bayes solutions of the statistical inventory problem. *Ann. Math. Statist.* 30(2):490–508.
- Stefanescu C (2009) Multivariate Customer demand: Modeling and estimation from censored sales. Working paper, London Business School, <http://ssrn.com/abstract=1334353>.
- Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1):15–33.
- Ulu C, Honhon D, Alptekinoglu A (2012) Learning consumer tastes through dynamic assortments. *Oper. Res.* 60(4):833–849.
- van Ryzin G, Vulcano G (2014) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Sci.*, ePub ahead of print October 29. <http://dx.doi.org/10.1287/mnsc.2014.2040>.
- Vulcano G, van Ryzin G, Chaar W (2010) Choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing Service Oper. Management* 12(3):371–392.
- Vulcano G, van Ryzin G, Ratliff R (2012) Estimating primary demand for substitutable products from sales transaction data. *Oper. Res.* 60(2):313–334.
- Wecker WE (1978) Predicting demand from sales data in the presence of stockouts. *Management Sci.* 24(10):1043–1054.

Aditya Jain is an assistant professor of operations management at Indian School of Business, Hyderabad, India. His research is primarily focused on supply chain management, retailing, and healthcare operations.

Nils Rudi is a professor of technology and operations management at INSEAD. His research focuses on analytics of supply chains, retail, behavioral issues, and sports.

Tong Wang is an assistant professor at the Department of Decision Sciences, NUS Business School, National University of Singapore. He is interested in the broad area of joint estimation and optimization with a focus on demand estimation and inventory or pricing optimization.