Available online at www.sciencedirect.com

ScienceDirect

European Journal of Operational Research 182 (2007) 256–281

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Bayesian demand updating in the lost sales newsvendor problem: A two-moment approximation ☆

Emre Berk [a], Ülkü Gürler [b,*], Richard A. Levine [c]

[a] Faculty of Business Administration, Bilkent University, 06800 Ankara, Turkey
[b] Department of Industrial Engineering, Bilkent University, 06800 Bilkent, Ankara, Turkey
[c] Department of Statistics, San Diego State University, San Diego, USA

Received 14 November 2003; accepted 21 August 2006
Available online 15 November 2006

Abstract

We consider Bayesian updating of demand in a lost sales newsvendor model with censored observations. In a lost sales environment, where the arrival process is not recorded, the exact demand is not observed if it exceeds the beginning stock level, resulting in censored observations. Adopting a Bayesian approach for updating the demand distribution, we develop expressions for the exact posteriors starting with conjugate priors, for negative binomial, gamma, Poisson and normal distributions. Having shown that non-informative priors result in degenerate predictive densities except for negative binomial demand, we propose an approximation within the conjugate family by matching the first two moments of the posterior distribution. The conjugacy property of the priors also ensure analytical tractability and ease of computation in successive updates. In our numerical study, we show that the posteriors and the predictive demand distributions obtained exactly and with the approximation are very close to each other, and that the approximation works very well from both probabilistic and operational perspectives in a sequential updating setting as well.

Keywords: Inventory; Newsboy; Lost sales; Censoring; Bayesian

1. Introduction and literature review

Estimating the demand distribution for a product or service over some time interval is an issue of theoretical interest with practical implications. The impact of demand forecast accuracy on profitability in fashion goods retailing has been well documented by Fisher et al. (1994). Similarly, the success of Wal-Mart, the US-based retail chain, is largely attributed to its superiority in demand and inventory management (*e.g.* Business Week, 2003). Aside from being one of the pioneers in the usage of barcode technology in inventory management, Wal-Mart has also introduced an Internet-based, interactive retailer–manufacturer 'collaborative

☆ This research has been partially sponsored by the grant NSF-INT 0114738.
* Corresponding author. Tel.: +90 312 2902413; fax: +90 312 2664054.
E-mail addresses: eberk@bilkent.edu.tr (E. Berk), ulku@bilkent.edu.tr (Ü. Gürler), ralevine@sciences.sdsu.edu (R.A. Levine).

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2006.08.035

forecasting replenishment' software (CFAR) with the intention of providing more reliable medium-term forecasts and, thereby, improving its stock and sales management (Business Week, 1996). As product lifecycles become shorter and profit margins tighten in the retail industry, the ability to better estimate demand becomes truly a competitive weapon. In this paper, we propose and analyze the efficacy of a Bayesian updating method to estimate demand in the presence of unobservable lost sales.

Depending on whether a seller is able to replenish her stock once or multiple times for a selling season, the season can be viewed as consisting of either a single period or multiple periods, respectively. In a single-period setting, the demand distribution is based only on the prior beliefs about the demand process that will reveal itself within the period. Then, the seller's forecasting ability depends solely on such prior beliefs. In a multiple period setting, however, the seller may also observe the demands that are realized within each period and the estimate for the following periods can be updated using this additional information. The seller's forecasting ability would now depend on how effectively such information is utilized. As long as there are no shortages of goods, the actual sales observed will equal the demands and the seller will have complete data. However, the lost sales environment presents a difficulty in estimation. Once a stockout occurs, if the demands that have not been satisfied are not recorded, there is incomplete data, commonly referred to as *censored* data, about the demand process. Formally, if the demand in a period is denoted by X and the period starts with s units on hand, then at the end of the period, one observes sales as $M = \min(s, X)$, that is, rather than the demand data, the sales data becomes available. The incomplete nature of observed demand introduces more difficulty into the estimation process. The additional effort is worthwhile, however; as our results indicate, ignoring the presence of censoring may result in considerable errors in stock level determination and, thereby, suboptimal management of inventory.

There are two main approaches to demand estimation in the inventory theory literature: the frequentist and the Bayesian approaches. In the frequentist approach, Nahmias (1994) is the first to analyze the censored demand case with normally distributed demand. More recently Agrawal and Smith (1996) investigated the type of the appropriate demand distribution for the sales data of a retailer chain over a number of periods within a sales season, and considered the estimation of negative binomial demand, in the presence of censoring induced by lost sales. Both of these studies consider the stocking levels as given in their respective settings and do not address the optimization problem.

Methods for demand estimation with a Bayesian spirit have been suggested very early on by several authors including Scarf (1959, 1960) and Iglehart (1964). Later, Azoury (1985), Lovejoy (1990), Bradford and Sugrue (1990), Hill (1997, 1999) and Eppen and Iyer (1997) considered various aspects of inventory management with Bayesian updating of demand distributions. However, all of these studies assume fully observed demand, and, therefore, are not applicable in the lost sales environment. Lariviere and Porteus (1999) considered Bayesian updating with partially observed sales but their analysis relies on a special demand distribution which limits its applicability. Recently, Ding et al. (2002) considered the censored demand setting for a perishable product and examined optimal stocking decisions. They showed that, with general continuous demand distributions, stocking levels in the censored case are higher compared to the uncensored case.

When the stocking levels in each period are not given but are instead decision variables, the estimation of demand and the optimization of costs need to be jointly performed over the planning horizon. To this end, Azoury (1985), Eppen and Iyer (1997), and Lariviere and Porteus (1999) use dynamic programming (DP). The DP approach provides the globally optimal solution, albeit at a cost of drastically increasing the scale of the problem when considered with the Bayesian updating procedure. This renders it difficult to numerically elaborate the sensitivity of the problem to the Bayesian aspects such as the choice of the prior and the quality of certain approximations. To overcome this difficulty, Lovejoy (1990) proposed the use of a single-period (newsvendor) approximation and myopic control policies. His findings indicate that such an approximation performs reasonably well and that newsvendor models can be used to study inventory settings with demand updating.

Hence, in this study, we revisit the newsvendor problem in Hill (1997) with the essential difference that lost demands cannot be observed. In the presence of censored observations arising from lost sales, we develop the exact posterior and predictive demand distributions. We consider four distributions commonly employed in inventory theory and supported by empirical analyses (See Agrawal and Smith, 1996 and Nahmias, 1994): Negative binomial, gamma, Poisson and normal. For all of the above, we assume conjugate priors. With the introduction of censoring, the posterior distributions cease to be in the conjugate family after the first censored updating. To ensure analytical tractability in successive estimations, we then propose a conjugacy retaining two-moment approximation and obtain the approximate posterior and corresponding predictive demand distributions. In an earlier work, Berk et al. (2001) suggested the use of such an approximation for posterior generation for poisson and normal demands but did not examine, in detail, estimation of the predictive demand distributions and the resulting costs in an inventory theoretic setting. This is the first work that considers Bayesian updating of negative binomial and gamma demands in an inventory theoretic setting with lost sales. Furthermore, we provide detailed proofs of some of the results and additional results on Poisson and normal demand distributions. Our numerical results indicate that the approximation proposed herein is highly satisfactory from both probabilistic and operational perspectives and that the approximation works well in a sequential update environment as well.

The rest of the paper is organized as follows. In Section 2, we provide the preliminaries of the basic procedure for Bayesian updating of the demand distribution and the cost optimization in the newsvendor problem. In Section 3 we develop the exact and approximate posterior and corresponding predictive demand distributions for the cases considered. We present our numerical study in Section 4 where we investigate the efficacy of the proposed approximation technique and the impact of neglecting the presence of censoring on the inventory control problem. In Section 5, we summarize our findings and discuss future work and in Appendix we provide the proofs of some results.

2. Preliminaries

In this section, we outline the basic steps of Bayesian updating of the demand distribution and optimization of the inventory control problem.

2.1. Bayesian updating of the predictive demand distribution

Consider an inventory system facing random demand in a periodic review setting. Suppose that the demands in successive periods are independent and identically distributed (i.i.d.) random variables denoted by X_i for $i = 1, 2, \ldots$ The distribution of X_i's is characterized by a parameter λ (possibly vector valued), and when the value of λ is given, the common probability density (or probability mass function) of X_i's is given by $f(x|\lambda)$. In accordance with the Bayesian perspective, we treat the unknown parameter λ as a random variable and the initial 'beliefs' about the parameter are expressed by a prior distribution for this random variable.

Suppose the unsatisfied demand in a period is not backlogged and the demand information during a stockout period is lost. In this setting, both the inventory problem under consideration and the Bayesian updating scheme becomes more complicated. The observed sales within a period correspond to the exact demand if there has been enough stock on hand; but, the exact demand of the period is not observed if it exceeds the starting inventory level in the period.

Let us first show how we obtain the posterior for λ at the end of the first period. Suppose at the end of period 1, the observable data is denoted by the random variable $M_1 = \min(s_1, X_1)$ where s_1 is the starting stock level of the first period. Note that M_1 is a truncated random variable and if X_1 is continuous, then M_1 is a mixture variable with both continuous and discrete parts, with a positive probability at $M_1 = s_1$. If X_1 is discrete, then M_1 is again discrete, taking on only the values $0, 1, \ldots, s_1$. Since the exact demand may not be observed, the generation of the posterior distribution for λ is based on M_1. First, we obtain the conditional distribution of the observed data, $f_{M_1}(y \mid \lambda)$ of M_1:

$$f_{M_1}(y|\lambda) = \begin{cases} f(y|\lambda) & \text{if } y < s_1, \\ P(M_1 = s_1) = \int_{s_1} f(u|\lambda)\,\mathrm{d}u & \text{if } y = s_1. \end{cases} \tag{1}$$

The first part above corresponds to the case where exact demand is observed, so that $M_1 = X_1$, and the second part to the censored case where $M_1 = s_1$. For convenience, we assumed above that the demand is continuous; if it is discrete, the integrals must be replaced by summations. Both the marginal distribution of M_1 and the posterior distribution of λ given below are based on the above distribution and due to the complexity induced by lost sales, different expressions are obtained depending on whether the observed data is censored or not. This issue will become more explicit in the next section when we consider specific demand distributions.

The marginal distribution of the observable random variable M_1 is found as

$$f_{M_1}(x) = \int f_{M_1}(x|\lambda)\pi_1(\lambda)\,\mathrm{d}\lambda. \tag{2}$$

Suppose, at the end of the period, $M_1 = y_1$ units of sale are observed. The *posterior* distribution $\pi_2(\lambda|y_1)$ of λ is then based on the distribution of the observed random variable M_1 evaluated at y_1 as follows:

$$\pi_2(\lambda|y_1) = \frac{f_{M_1}(y_1|\lambda)\pi_1(\lambda)}{f_{M_1}(y_1)}. \tag{3}$$

The demand distribution that the inventory system will face in the second period, that is the *predictive* distribution is obtained as

$$f_2(x|y_1) = \int f(x|\lambda)\pi_2(\lambda|y_1)\,\mathrm{d}\lambda. \tag{4}$$

This procedure then continues in the same manner in the successive periods, so that, after observing the data $M_i = \min(X_i, s_i) = y_i$ at the end of period i that starts with inventory level s_i, we obtain the predictive density for the next period as follows: for $i = 2, 3, \ldots$

$$f_{i+1}(x|y_i) = \int f(x|\lambda)\pi_{i+1}(\lambda|y_i)\,\mathrm{d}\lambda, \tag{5}$$

where the π_{i+1} is obtained in a similar way to (3) via (1) and (2) by simply replacing the index 1 with i. Note here that, when finding the predictive distributions we assume that the underlying demand distribution given λ is unchanged and is given by $f(x|\lambda)$, but the prior distribution of λ is replaced by the posterior distributions modified in the successive periods as new demand information for the periods becomes available.

2.2. Two-moment approximation of the posterior distribution

In Bayesian analysis, the prior distribution represents the beliefs of the decision maker about the unknown parameter expressed in a probabilistic statement. This statement may be about an individual's own belief or a consensus statement for a panel of experts. Since the performance of successive updates depends on the prior, its choice is important. If there are no sources of prior information, then, a so-called non-informative prior must be used. A non-informative prior implies that the random variable representing the unknown parameter may take on any value in its domain equally likely. For example, for λ in a Poisson distribution, the non-informative prior would distribute equal chance to the entire interval $(0, \infty)$. Although non-informative priors are intuitively appealing, we establish in the following result that non-informative priors may result in improper demand distributions in certain cases.

Preposition 2.1. *Let X denote the demand in a period with $p(x|\cdot)$ (or, $f(x|\cdot)$) as its probability mass (or, density) function.*

(a) *Suppose $p(x|\lambda)$ is Poisson with mean λ. A non-informative prior for λ results in an improper predictive density for demand.*
(b) *Suppose $f(x|\beta)$ is gamma with scale parameter β and known shape parameter α. A non-informative prior for β results in an improper predictive density for demand.*
(c) *Suppose $f(x|\mu)$ is normal with mean μ and known variance σ^2. A non-informative prior for μ results in an improper predictive density for demand.*

Proof. For the selection and construction of non-informative priors in the proof, we refer to Berger (1985).

(a) The non-informative prior is the uniform density on R^+ (Berger, 1985, p. 82). Hence, when the prior $g(\lambda) = 1$ is used for $\lambda > 0$, it turns out that $f_X(x) = 1$ for $x > 0$, which is not a proper density function since it does not integrate to unity.

(b) Similarly, the non-informative prior for the gamma density is $g(\beta) = 1/\beta$. When it is used, the predictive density is obtained as:

$$f_X(x) = \int_0^\infty \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp\{-\beta x\}\,\mathrm{d}\beta = 1/x \quad \forall x,$$

which does not integrate to unity; hence it is improper.

(c) The non-informative prior for mean μ for the normal density is $g(\mu) = 1$ with $\mu \in R$. Then, the predictive density is given by:

$$f_X(x) = \int \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2}(x-\mu)^2\right\}\mathrm{d}\mu = 1 \quad \forall x,$$

which is improper. □

It is worth noting that the non-informative prior for p in a negative binomial distribution would be a uniform distribution over $[0,1]$ which results in a proper predictive density for demand. Furthermore, this specific distribution can be represented as a Beta distribution with parameters $\alpha = 1$ and $\beta = 1$, which is the so-called *conjugate* prior. In Bayesian analysis, it is common to work with conjugate priors. A conjugate prior for a parameter is a distribution, for which the posterior is also of the same family – *e.g.* a gamma distributed prior resulting in a gamma distributed posterior albeit with different parameters. This conjugacy property is essential for analytical tractability and extremely useful in computation.

With full backordering, the conjugacy property holds since all demands can be observed even though they may not materialize into sales. Furthermore, it has been shown by Azoury (1985) that, with fully observable demands, for commonly encountered demand distributions and conjugate initial prior distributions, the posterior distributions can be expressed through a simple sufficient statistic. Thus, a Bayesian updating procedure is relatively straightforward and computationally easy in the presence of fully observable demand (*i.e.* fully backlogged demand). Systems with lost sales present no difficulty so long as no stockout are experienced. If stockout does not occur in a period and if conjugate priors are used for the unknown parameters, then the conjugacy property holds. However, if a stockout occurs, *i.e.* if the demand observation is censored, then the conjugacy property no longer holds – except for a special distribution (see Lariviere and Porteus, 1999). The resulting posterior distributions are typically too complicated, rendering the updating procedure in successive periods analytically intractable. Although it is possible to use these exact posterior distributions for the optimization of the inventory problem via numerical methods at the expense of computational effort, we propose a simpler method. We propose to use an approximate posterior distribution in the same conjugate family, the parameters of which are selected so as to match the *first two moments* of the exact posterior. In the numerical results we show that this approximation works quite well for the cases considered.

2.3. Optimization of the inventory control problem

We first introduce the newsvendor problem. There is a single selling period of finite duration. The demand during the period is random with a known probability distribution and there is a single replenishment opportunity at the beginning of the period. All of the unsatisfied demand is lost at a unit underage cost of c_u. Likewise, the excess stock at the end of the period, if any, is disposed of at a certain salvage value resulting in a unit overage cost of c_h. All initial ordering costs are ignored. The optimization problem consists of determining the optimal stocking level at the beginning of the period so that expected overage and underage costs are minimized. As such, the newsvendor problem captures the essential trade-off in inventory theory between inventory holding and shortage costs. The setting presented above is the so-called classical newsvendor problem. The classical problem was first introduced by Arrow et al. (1951) and a variety of its extensions have been formulated and investigated since then. (We refer the reader to Khouja (1999) for a comprehensive review of these variants.) The newsvendor problem with Bayesian updating can be formally constructed as follows. Let the demand distribution in period i be given by a known conditional density $f(x|\lambda)$ with an unknown parameter λ, with the posterior found at the end of $(i-1)$th period given by $\pi_i(\lambda|y_{i-1})$. Then the predictive demand density $f_i(x|y_{i-1})$ is as given in (5) with cumulative distribution function $F_i(x|y_{i-1})$. The optimization problem is then stated as

$$\min_{s} \int_{x} [c_u[x-s]^+ + c_h[s-x]^+] f_i(x|y_{i-1})\,\mathrm{d}x. \tag{6}$$

From standard optimization results, the optimal order quantity s^* that solves (6) is the value of s that satisfies

$$F_i(s|y_{i-1}) = \frac{c_u}{c_u + c_h}. \tag{7}$$

In the above, we assumed, for convenience, continuous demand which guarantees that there is a unique s that satisfies the above relation. For discrete demands, s_i^* is the smallest integer s such that $F_i(s|y_{i-1})$ exceeds the RHS of (7). Lowe et al. (1988) showed that the above critical ratio result still holds when the unit underage and overage costs are replaced by their expectations if there is uncertainty about their values. Eq. (7) implies that the demand distribution is the main factor in determining the optimal stocking level and the ensuing costs. Hence, estimation of the distribution of demand during the selling season is critical.

The demand estimation procedure through Bayesian updating described above implicitly assumes multiple periods over which observations are made; therefore, an inventory control problem utilizing this technique makes most sense in a multi-period setting. The solution of multi-period optimization problems suffers from the so-called 'curse of dimensionality'. Fortunately, we know that a periodic review inventory control problem with a planning horizon longer than one period can be reduced to the newsvendor problem under 'myopic' control policies with the appropriate choice of excess and shortage costs incurred at the end of a single period. For certain inventory systems, the myopic policy has already been established to be the optimal policy. That is, it may be sufficient in some cases to analyze the single-period problem to understand the dynamics in a multiple period planning horizon. For instance, the multi-period inventory control problem of perishable goods with a shelflife of one period reduces to a series of independent newsvendor problems. (We refer the interested reader to Zipkin (2000), pp. 378–385 for other applicable cases and a more detailed discussion on myopic policies for inventory systems with linear costs.) Additionally, Lovejoy (1990) showed that myopic policies perform reasonably well *vis a vis* optimal policies when demand is updated in a Bayesian fashion with fully observed demands. (Also see Geunes et al. (2001) for adapting the newsvendor model for infinite-horizon inventory systems.) Therefore, the newsvendor problem provides a versatile inventory theoretic setting to test the efficacy of novel approaches in modeling and analyses. Hence, in this work, we consider the newsvendor setting for the inventory control problem and Bayesian updating for demand estimation over multiple periods. Next, we develop the Bayesian updates for a lost sales newsvendor setting with specific demand distributions.

3. Updating with specific demand distributions

The underlying distribution for demand may vary according to several factors such as product types, marketing strategies, and customer profiles. However from a modeling perspective, negative binomial and Poisson distributions provide families which are sufficiently rich to cover the cases most encountered in practice when demand is considered as a discrete random variable. Furthermore, empirical studies suggest very good fits of these distributions to actual data – Poisson in the case of low coefficient of variation ($\leqslant$1) and negative binomial in the case of high coefficient of variation (>1). We refer the reader to Nahmias and Smith (1994), and Agrawal and Smith (1996) for a discussion of demand distributions of real retail data. Similarly, the normal and gamma families are frequently adopted to describe continuous demand data. (See Nahmias, 1994 for a discussion of empirical issues.) Therefore, in this section, we illustrate the Bayesian demand updating scheme proposed in the previous sections for these commonly encountered demand distributions.

Gamma and negative binomial distributions are characterized by two parameters, one for the shape and the other for the scale. Both of these distributions are right skewed and become more symmetric as the shape parameter increases. Therefore, from a practical point of view, it may not be very unrealistic to assume that a reliable information about the shape of the distribution is available from independent past data. Furthermore, for integer values of the shape parameter, these distributions represent a sum of independent geometrically and exponentially distributed demands, resp. This may correspond to a practically realistic situation. Consider, for instance, a warehouse serving a given number of independently ordering retailers of relatively equal sizes for each of which the demand can be represented by geometric (or exponential) distributions. Then, the total demand that the warehouse faces would be negative binomial (Erlang) with known shape parameters. Similarly, independent daily demand can be aggregated to a negative binomial or gamma weekly demand with known shape parameters. Then, the question of interest becomes estimating the size of demand for each of these retailers or each day. Hence, in the sequel, we make the assumption that the shape parameters of these distributions are known and the Bayesian method is used for dynamically estimating the distribution of the scale variable. We should also mention that, from a technical point, working with only the scale parameter, greatly reduces the analytical tedium and allows for exploiting the conjugacy properties since no bi-variate conjugate prior is known which can be used for estimating both scale and shape parameters. Incidentally, in all of the Bayesian literature with which we are familiar, known shape parameter is the standard assumption. Similarly, for normal demand, we make the standard Bayesian assumption that the variance is known but the mean is uncertain. Such an assumption would again be valid if the variability could be assessed in a fairly precise way from other sources of information. To avoid a cumbersome notation, we drop the index identifying the periods in the following discussions.

3.1. Negative binomial demand

Suppose the demand in a period X has a negative binomial distribution with parameters r and p. Then the probability mass function of $M = \min(s, X)$ given r and p is

$$f_M(y|r,p) = \begin{cases} \binom{y+r-1}{r-1} p^r(1-p)^y & \text{if } y < s, \\ \sum_{x=s}^{\infty} \binom{x+r-1}{r-1} p^r(1-p)^x & \text{if } y = s. \end{cases} \tag{8}$$

Theorem 3.1. *Assume that r is known and p has a beta distribution with parameters α and β. Then*

(a) *The posterior distribution of p is given by*

$$f_p(p|\alpha,\beta,r,M=y) = \begin{cases} \frac{\Gamma(\alpha+\beta+y+r)}{\Gamma(\alpha+r)\Gamma(y+\beta)} p^{\alpha+r-1}(1-p)^{\beta+y-1} & \textit{if } y < s, \\ \frac{p^{\alpha+r-1}(1-p)^{\beta-1}}{\Gamma(\alpha+r)} \dfrac{\sum_{x=s}^{\infty} \binom{x+r-1}{r-1} (1-p)^x}{\sum_{x=s}^{\infty} \binom{x+r-1}{r-1} \frac{\Gamma(x+\beta)}{\Gamma(\beta+\alpha+r+x)}} & \textit{if } y = s. \end{cases}$$

(b) *First two moments of the posterior distribution under censoring are given by*

$$m_1 = E(p|\alpha,\beta,r,M=s) = \frac{\sum_{x=s}^{\infty} \binom{x+r-1}{r-1} \frac{\Gamma(x+\beta)\Gamma(\alpha+r+1)}{\Gamma(\beta+\alpha+r+x+1)}}{\sum_{x=s}^{\infty} \Gamma(a+r) \binom{x+r-1}{r-1} \frac{\Gamma(x+\beta)}{\Gamma(\beta+\alpha+r+x)}}, \tag{9}$$

$$m_2 = E(p^2|\alpha,\beta,r,M=s) = \frac{\sum_{x=s}^{\infty} \binom{x+r-1}{r-1} \frac{\Gamma(x+\beta)\Gamma(\alpha+r+2)}{\Gamma(\beta+\alpha+r+x+2)}}{\sum_{x=s}^{\infty} \Gamma(a+r) \binom{x+r-1}{r-1} \frac{\Gamma(x+\beta)}{\Gamma(\beta+\alpha+r+x)}}. \tag{10}$$

Proof. The joint distribution of p and M is written as

$$f_{p,M}(p,y|\alpha,\beta,r) = \begin{cases} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \binom{y+r-1}{r-1} p^{\alpha+r-1}(1-p)^{y+\beta-1} & \text{if } y < s, \\ \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha+r-1}(1-p)^{\beta-1} \sum_{x=s}^{\infty} \binom{x+r-1}{r-1} (1-p)^x & \text{if } y = s. \end{cases}$$

from which the marginal distribution of M is obtained as

$$f_M(y|\alpha,\beta,r)=\begin{cases}\binom{y+r-1}{r-1}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\frac{\Gamma(\alpha+r)\Gamma(y+\beta)}{\Gamma(\alpha+\beta+y+r)} & \text{if } y<s,\\ \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\sum_{x=s}^{\infty}\frac{\Gamma(\alpha+r)\Gamma(x+\beta)}{\Gamma(\alpha+\beta+k+r)}\binom{x+r-1}{r-1} & \text{if } y=s.\end{cases}$$

Part (a) follows from taking the ratio of the last two expressions. Part (b) is obtained from standard calculations regarding beta functions, the details of which are omitted. □

We observe that if the demand is not censored (*i.e.* $y < s$), the posterior distribution of p is beta with parameters $\alpha^* = \alpha + r$ and $\beta^* = \beta + y$. This follows from the conjugacy property of the beta distribution. However, if the demand is censored, the resulting posterior distribution is outside the conjugate beta family and it is almost impossible to use it for subsequent analytical calculations of predictive distributions. Therefore, we propose a conjugacy retaining approximation below as explained in the previous sections.

3.1.1. Two-moment approximation

If the demand information is censored, we propose to use an approximate conjugate beta distribution for the posterior distribution, with modified parameters. The modification is done so that the first two moments of the approximate distribution matches those of the exact posterior. Recall that for a beta random variable X with parameters α^* and β^*, the first two moments are given as

$$E(X)=\alpha^*/(\alpha^*+\beta^*)\quad\text{and}\quad E(X^2)=\frac{\alpha^*(\alpha^*+1)}{(\alpha^*+\beta^*)(\alpha^*+\beta^*+1)}.$$

Setting $E(X)$ and $E(X^2)$ equal to (9) and (10) respectively and solving for α^* and β^* we obtain

$$\alpha^*=\frac{m_1(m_1-m_2)}{m_2-m_1^2},$$
$$\beta^*=\frac{(1-m_1)(m_1-m_2)}{m_2-m_1^2}.$$

3.1.2. Predictive distribution

Having obtained the posterior distribution either exactly or approximately, the predictive density of demand is found as below according to (5). For $x = 0, 1, \ldots,$

$$f(x|\alpha^*,\beta^*,r)=\binom{x+r-1}{r-1}\frac{\Gamma(\alpha^*+\beta^*)}{\Gamma(\alpha^*)\Gamma(\beta^*)}\frac{\Gamma(\alpha^*+r)\Gamma(x+\beta^*)}{\Gamma(\alpha^*+\beta^*+x+r)}. \tag{11}$$

Now, the optimal stocking level s^* for the inventory problem is found so as to satisfy (7).

3.2. Gamma demand

As in the previous section we observe $M = \min(s, X)$ and now assume that X has a gamma distribution with shape parameter α and scale parameter β. Then the density of the observed M is written as

$$\begin{aligned} f_M(y|\alpha,\beta)&=\frac{\beta^\alpha}{\Gamma(\alpha)}\mathrm{e}^{-\beta y}y^{\alpha-1}\quad \text{if } y<s\\ &=\int_s^\infty\frac{\beta^\alpha}{\Gamma(\alpha)}\mathrm{e}^{-\beta x}x^{\alpha-1}\,\mathrm{d}x\quad\text{if } y=s. \end{aligned}\tag{12}$$

To simplify the notation we present the result below without proof which follows from standard calculations.

Lemma 1. *Let*

$$I(\alpha,\gamma,\tau,s)=\int_s^\infty\frac{x^{\alpha-1}\tau^\gamma}{(\tau+x)^{\alpha+\gamma}}\,\mathrm{d}x.$$

Then

$$I(\alpha,\gamma,\tau,s)=\frac{\Gamma(\alpha)\Gamma(\gamma)}{\Gamma(\alpha+\gamma)}P(B>s/(s+\tau))=\frac{\Gamma(\alpha)\Gamma(\gamma)}{\Gamma(\alpha+\gamma)}\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau)), \tag{13}$$

where B is a beta random variable with parameters α, γ with the tail probability given by $\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))$.

Theorem 3.2. *Suppose demand has a gamma distribution with known shape parameter α and random scale parameter β, where the prior distribution for β is also gamma, with shape and scale parameters γ and τ respectively. Then*

(a) *The posterior distribution of β is given as*:

$$\begin{aligned} f(\beta|M=y,\alpha,\gamma,\tau)&=\frac{\beta^{\alpha+\gamma-1}}{\Gamma(\alpha+\gamma)}\mathrm{e}^{-(\tau+y)\beta}(\tau+y)^{\alpha+\gamma}\quad \text{if } y<s\\ &=\frac{\beta^{\alpha+\gamma-1}}{\Gamma(\gamma)\Gamma(\alpha)}\frac{\mathrm{e}^{-\tau\beta}}{\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}\int_s^\infty \mathrm{e}^{-\beta u}u^{\alpha-1}\,\mathrm{d}u\quad \text{if } y=s. \end{aligned} \tag{14}$$

(b) *The first two moments of the posterior distribution of β under censoring are given by*

$$m_1\equiv E(\beta|M=s,\alpha,\gamma,\tau)=\frac{\gamma}{\tau}\frac{\overline{F}_{B(\alpha,\gamma+1)}(s/(s+\tau))}{\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}, \tag{15}$$

$$m_2\equiv E(\beta^2|M=s,\alpha,\gamma,\tau)=\frac{\gamma(\gamma+1)}{\tau^2}\frac{\overline{F}_{B(\alpha,\gamma+2)}(s/(s+\tau))}{\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}. \tag{16}$$

Proof. Following the lines of proof for negative binomial demand, we obtain the joint distribution of M and β is given by

$$\begin{aligned} f_M(y,\beta|\alpha,\gamma,\tau)&=\frac{\beta^\alpha\tau^\gamma}{\Gamma(\alpha)\Gamma(\gamma)}\mathrm{e}^{-\beta(\tau+y)}y^{\alpha-1}\beta^{\gamma-1}\quad \text{if } y<s\\ &=\int_s^\infty \frac{\beta^\alpha\tau^\gamma}{\Gamma(\alpha)\Gamma(\gamma)}\mathrm{e}^{-\beta(\tau+x)}x^{\alpha-1}\beta^{\gamma-1}\,\mathrm{d}x\quad \text{if } y=s \end{aligned} \tag{17}$$

and the marginal distribution of M as

$$\begin{aligned} f_M(y|\alpha,\gamma,\tau)&=\frac{\Gamma(\alpha+\gamma)}{\Gamma(\alpha)\Gamma(\gamma)}\frac{y^{\alpha-1}\tau^\gamma}{(\tau+y)^{\alpha+\gamma}}\quad \text{if } y<s\\ &=\frac{\Gamma(\alpha+\gamma)}{\Gamma(\alpha)\Gamma(\gamma)}\int_s^\infty \frac{x^{\alpha-1}\tau^\gamma}{(\tau+x)^{\alpha+\gamma}}\,\mathrm{d}x\quad \text{if } y=s. \end{aligned} \tag{18}$$

Part (a) again follows by taking the ratio of the preceding two functions. For Part (b) note that

$$\begin{aligned} m_1&=\int_0^\infty \beta\frac{\beta^{\alpha+\gamma-1}}{\Gamma(\gamma)\Gamma(\alpha)}\frac{\mathrm{e}^{-\tau\beta}}{\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}\int_s^\infty \mathrm{e}^{-\beta u}u^{\alpha-1}\,\mathrm{d}u\,\mathrm{d}\beta\\ &=\frac{\tau^\gamma}{\Gamma(\gamma)\Gamma(\alpha)\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}\int_s^\infty u^{\alpha-1}\int_0^\infty \beta^{\alpha+\gamma}\mathrm{e}^{-(\tau+u)\beta}\,\mathrm{d}\beta\,\mathrm{d}u\\ &=\frac{\Gamma(\alpha+\gamma+1)}{\Gamma(\gamma)\Gamma(\alpha)\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}\int_s^\infty \frac{\tau^\gamma u^{\alpha-1}}{(\tau+u)^{\alpha+\gamma+1}}\,\mathrm{d}u=\frac{\Gamma(\alpha+\gamma+1)}{\Gamma(\gamma)\Gamma(\alpha)\overline{F}_{B(\alpha,\gamma)}(s/(s+\tau))}\frac{1}{\tau}I(\alpha,\gamma+1,\tau,s). \end{aligned}$$

The result follows by reference to Lemma 1. The second moment m_2 is obtained similarly, for which we omit the proof. □

3.2.1. Two-moment approximation

From part (a) of the above theorem we observe that when there is no censoring, the posterior distribution of β is gamma with shape and scale parameters $\gamma^*=\alpha+\gamma$ and $\tau^*=\tau+y$ respectively. If however a stockout occurred within the period, the gamma distribution can not be retained and we make the similar approximation to match the first two moments of the posterior density. Recall that for a gamma random variable with parameters γ and τ, the first moment is given by $m_1 = \gamma/\tau$ and the second one by $m_2 = \gamma(\gamma+1)/\tau^2$. Then for the stockout periods, we again use the gamma family with the following modified parameters:

$$\tau^* = \frac{m_1}{m_2 - m_1^2}, \quad \gamma^* = \tau^* m_1 = \frac{m_1^2}{m_2 - m_1^2}. \tag{19}$$

3.2.2. Predictive distribution

From the above findings the predictive density of demand for the next period is given below. For $x \geqslant 0$

$$f(x|\alpha, \gamma^*, \tau^*) = \frac{\Gamma(\alpha+\gamma^*)}{\Gamma(\alpha)\Gamma(\gamma^*)} \frac{1}{x} \left(\frac{x}{\tau^* + x}\right)^{\alpha} \left(\frac{\tau^*}{\tau^* + x}\right)^{\gamma^*}. \tag{20}$$

3.3. Poisson demand

Poisson and normal distributions are also commonly used for modeling the demand uncertainty. In a preliminary work we suggested (see Berk et al., 2001) the two-moment approximation herein for Poisson and normal demands and presented the expressions for the posteriors starting with respective conjugate priors. The predictive demand distribution and the related proofs however were not given. For the purpose of completeness, we first present below the posterior distributions with their first two moments and then provide the approximate predictive distributions. We also present in the Appendix the detailed proofs not appearing in the earlier work.

Assume now that the demand during a period has Poisson distribution with parameter λ and we consider the gamma distribution as the prior for λ. For ease of presentation of the following results, let $\rho = 1/(\beta+1)$ and

$$A(\rho, s, \alpha) = \sum_{i=s}^{\infty} \frac{\Gamma(\alpha+i)}{i!} \rho^i.$$

We also denote by $A'(\cdot)$ and $A''(\cdot)$ the first two derivatives of $A(\cdot)$ with respect to ρ.

Theorem 3.3. *Suppose demand has a Poisson distribution with parameter λ and λ has a gamma prior with parameters α and β. Then*

(a) *The posterior distribution of λ is given by*

$$h_\lambda(\lambda|M=k) = \begin{cases} \frac{(\beta+1)^{\alpha+k}}{\Gamma(\alpha+k)} \mathrm{e}^{-\lambda(\beta+1)} \lambda^{\alpha+k-1} & \text{if } k < s, \\ \frac{(\beta+1)^{\alpha}}{A(\rho,s,\alpha)} \mathrm{e}^{-\lambda(\beta+1)} \lambda^{\alpha-1} \sum_{i=s}^{\infty} \frac{(\lambda)^i}{i!} & \text{if } k = s. \end{cases}$$

(b) *If the demand is censored, the first two moments of the posterior distribution are*

$$m_1 \equiv E(\lambda|M=s) = \frac{1}{(\beta+1)} \left[\alpha + \rho \frac{A'(\rho,s,\alpha)}{A(\rho,s,\alpha)}\right], \tag{21}$$

$$m_2 = E(\lambda^2|M=s) = (\beta+1)^{-2} \left[\alpha(\alpha+1) + \rho(\alpha+3) \frac{A'(\rho,s,\alpha)}{A(\rho,s,\alpha)} + \rho^2 \frac{A''(\rho,s,\alpha)}{A(\rho,s,\alpha)}\right]. \tag{22}$$

Proof. See Appendix. □

For the evaluation of the above quantities, we use the following computationally handy result, where $B(x, a, b)$ is the cumulative distribution function of a beta random variable with parameters a and b, which is readily available in commercial software and also is well tabulated since it is proportional to incomplete beta function.

Lemma 2

$$A(\rho,s,\alpha)=\frac{\Gamma(\alpha-1)(\beta+1)}{\rho^{\alpha-1}}B(\rho,s,\alpha-1).$$

Proof. Since $A(\rho,s,\alpha)$ appears in the denominator of the conditional density $h(u|M=s)$, it is the scaling constant for the density. Therefore,

$$\begin{aligned}A(\rho,s,\alpha)&=(\beta+1)^{\alpha}\int_0^{\infty}\mathrm{e}^{-\lambda(\beta+1)}\lambda^{\alpha-1}\sum_{i=s}^{\infty}\frac{(\lambda)^i}{i!}\mathrm{d}\lambda=(\beta+1)^{\alpha}\int_0^{\infty}\mathrm{e}^{-\lambda\beta}\lambda^{\alpha-1}\int_0^1\frac{\lambda^{s-1}\mathrm{e}^{-\lambda t}}{\Gamma(s)}t^{s-1}\,\mathrm{d}t\,\mathrm{d}\lambda\\&=(\beta+1)^{\alpha}\int_0^{L}\frac{\Gamma(s+\alpha-1)}{\Gamma(s)}\left(\frac{t}{t+\beta}\right)^{s-1}\left(\frac{1}{t+\beta}\right)^{\alpha}\\&=\frac{(\beta+1)^{\alpha}}{\beta^{\alpha-1}}\int_0^{1/(\beta+1)}\frac{\Gamma(s+\alpha-1)}{\Gamma(s)}u^{s-1}(1-u)^{\alpha-2}\,\mathrm{d}u.\end{aligned}$$

The first equality in the above proof follows from recognizing that the tail sum of a Poisson probability is the c.d.f. of a gamma random variable. The last equality is obtained by making the change of variable $u=t/(t+\beta)$. □

3.3.1. Two-moment approximation

From Theorem 3.3, we observe that if demand is not censored, the posterior distribution of λ is gamma with parameters $\alpha^*=\alpha+k$ and $\beta^*=\beta+1$ as expected from the conjugate property. If however there is censoring, then we approximate the posterior with a gamma distribution with modified parameters $\alpha^*=m_1^2/(m_2-m_1^2)$ and $\beta^*=m_1(m_2-m_1^2)$ as in (19).

3.3.2. Predictive distribution

The predictive distribution of demand corresponding to the distributions found above is given as

$$p(x)=\frac{\Gamma(\alpha^*+x)}{\Gamma(\alpha^*)\Gamma(x+1)}\left(\frac{1}{\beta^*+1}\right)^{x}\left(\frac{\beta^*}{\beta^*+1}\right)^{\alpha^*}\quad x=0,1,\ldots \tag{23}$$

The optimal stocking level s^* for the newsvendor problem can now be found using (7) via the predictive demand distribution (23).

3.4. Normal demand

Suppose, given the mean μ and variance σ^2, the demand during a period (X) has a normal distribution with parameters μ and σ^2. We denote this as $X|\mu,\sigma^2\sim N(\mu,\sigma^2)$. The mean μ is presumed unknown with prior distribution $\mu|\gamma,\tau^2\sim N(\gamma,\tau^2)$ where σ^2 and τ^2 are assumed known as well as the prior mean γ. Assume as before we observe only $M=\min(X,s)$. When the demand is completely observed in that $M=y<s$ (no lost sales), the posterior distribution of μ is easily found from Bayesian calculations with normal priors as

$$\mu|y,\gamma,\sigma,\tau\sim N(\mu^*,\xi^2), \tag{24}$$

where

$$\mu^*=\frac{\gamma\sigma^2+y\tau^2}{\sigma^2+\tau^2}, \tag{25}$$

$$\xi^2=\frac{\tau^2\sigma^2}{\sigma^2+\tau^2}. \tag{26}$$

If however, stockout occurs in a period, then we have the following result regarding the posterior distribution and the respective moments.

Theorem 3.4. *Suppose demand has a normal distribution with mean μ and variance σ^2 and also that μ has a normal prior with parameters γ and τ^2. Then if the observed demand is censored*

(a) *The posterior density of μ is given by*

$$f_\mu(u|M=s) = \frac{\bar{\Phi}\left(\frac{s-\mu}{\sigma}\right)\cdot\frac{1}{\sqrt{2\pi\tau^2}}\exp\left\{\frac{-(u-\gamma)^2}{2\tau^2}\right\}}{\overline{\Phi}\left(\frac{s-\gamma}{\sqrt{\eta}}\right)}. \tag{27}$$

(a) *The first two moments of the density given in part (a) above are*

$$\mu^* = m_1 \equiv E(\mu \mid M=s) = \gamma + \tau^2\lambda(s),$$
$$m_2 \equiv E(\mu^2 \mid M=s) = (1/\eta^2)(\gamma^2\sigma^4 + \eta\tau^2\sigma^2 + 2\tau^2\sigma^2L^2\gamma^2 + \tau^4\eta + \gamma^2\tau^4) + [2\gamma\tau^2\sigma^2 + \tau^4(s+\gamma)]\lambda(s)/\eta,$$

where $\phi(s)$ and $\Phi(s)$ are respectively the density and the cumulative distribution function for standard normal variate and $\lambda(s) = \phi(s)/\bar{\Phi}(s)$ is the corresponding hazard rate.

Proof. See Appendix. □

3.4.1. Two-moment approximation

We observe that the density above corresponds to a weighted normal distribution. The posterior mean is shifted from the prior mean by a constant proportional to the hazard rate and the second moment is also a linear function of the hazard rate where the constant terms are functions of the variance components and the prior mean. Due to the complexity of this exact posterior distribution, we again propose to use a normal distribution for the posterior with the first two moments given by the above theorem. In particular, the mean and the variance of the approximate posterior normal distribution are given by μ^* and $\sigma_*^2 = m_2 - m_1^2$ respectively.

3.4.2. Predictive distribution

The predictive density of demand is obtained as described in Section 2, by integrating the product of the demand and the posterior densities with respect to μ. After some straightforward but somewhat tedious algebra that we skip here, we find the predictive density as normal, with mean μ^* and variance $\sigma^2 + \xi^2$, where $\xi^2 = m_2 - m_1^2$ if there is censoring and otherwise given by the expression (26).

4. Impact of ignoring censoring

We present some theoretical results on the impact of ignoring censoring in the newsvendor setting.

Preposition 4.1

(a) *The stocking level obtained via the predictive distribution computed with censoring ignored is always lower than the optimal for the true distribution for any desired service level.*

(b) *Let X_n denote the true demand identically distributed with X, and $\hat{X}_n$ denote the predictive demand in period $n \geqslant 1$. Also let s be the fixed stocking level for all n and define $\mu_s = E[\min(s, X_n)]$.*
   (i) *Suppose $X \sim NB(r,p)$ where r is known and p has a Beta prior with parameters α and β. Then, $\hat{X}_n$ converges in distribution to $NB(r,\hat{p})$ with $\hat{p} = r/(r+\mu_s)$ as $n \to \infty$.*
   (ii) *Suppose $X \sim Gamma(\alpha,\beta)$ where α is known and β has a gamma prior with parameters γ and τ. Then, $\hat{X}_n$ converges in distribution to $Gamma(\alpha,\hat{\beta})$ with $\hat{\beta} = \alpha/\mu_s$ as $n \to \infty$.*
   (iii) *Suppose $X \sim Poisson(\lambda)$ where λ has a gamma prior with parameters α and β. Then, $\hat{X}_n$ converges in distribution to $Poisson(\hat{\lambda})$ with $\hat{\lambda} = \mu_s$ as $n \to \infty$.*
   (iv) *Suppose $X \sim Normal(\mu,\sigma^2)$ where σ is known and μ has a normal prior with parameters γ and τ^2. Then, $\hat{X}_n$ converges in distribution to $Normal(\hat{\mu}, \sigma^2 + \frac{\tau^2\sigma^2}{\sigma^2+\tau^2})$ with $\hat{\mu} = \mu_s$ as $n \to \infty$.*

Proof

(a) Since we have $M_n = \min(s, X_n)$ as the observation in period n, it follows from Eqs. (1) and (3) that the posterior obtained with censoring ignored is stochastically smaller than the exact posterior for any given prior for that period. From Eq. (5), we also have that the predictive demand is also stochastically smaller when computed with censoring ignored. Hence, the result.

(b) Below let $m_n = \min(s, x_n)$ be the observed value of $M_n = \min(s, X_n)$ in period n. (i) Note that, when censoring is ignored, the posterior distribution of p at nth period is Beta with parameters $\alpha + nr$ and $\beta + \sum_{i=1}^{n} m_i$. The mean of the posterior given by $\left(\frac{\alpha+nr}{\alpha+\beta+nr+\sum_{i=1}^{n} m_i}\right)$ converges to $r/(r+\mu_s)$ and the variance converges to zero as $n \to \infty$. Hence, the posterior distribution converges to a point mass at $p = r/(r+\mu_s)$, which results in the predictive distribution stated above. The other parts follow similarly, since in all cases the variance of the posterior distribution converges to zero; and in (ii), the mean of the posterior distribution in the nth period given by $\left(\frac{\gamma+n\alpha}{\tau+\sum_{i=1}^{n} m_i}\right)$ converges to α/μ_s; in (iii), the expected value of the posterior in the nth period given by $\left(\frac{\alpha+\sum_{i=1}^{n} m_i}{\beta+n}\right)$ converges to μ_s and finally in (*iv*), the posterior mean in period n given by $[\sigma^2/(\sigma^2+\tau^2)]^n \gamma + [\tau^2/(\sigma^2+\tau^2)] \sum_{i=1}^{n} [\sigma^2/(\sigma^2+\tau^2)]^{n-i} m_i$ converges to μ_s. □

The above result shows that if there is a true distribution of demand and censoring is ignored in obtaining the posterior distributions, then the predictive distributions would not converge to the true one as the number of observations tends to infinity, resulting in biased estimation of demand.

Next, we proceed to our numerical analysis.

5. Numerical analysis

In our numerical study, we investigate a number of issues: First, we consider goodness of the proposed two-moment approximation purely from a probabilistic perspective. We compare the approximate posterior density obtained after an observation with the exact posterior and that obtained by ignoring any censoring. We will refer to the last as the naive posterior since it does not use the information about censoring. We illustrate the impact of the approximations on Poisson and negative binomial demand distributions. Second, we look at the impact of using the information about censoring in Bayesian updating from an operational/inventory theoretic perspective. Here, we examine the performance of the approximation in terms of desired and achieved service levels. Third, we address the question of whether or not the performance of the approximation deteriorates in a sequential updating setting. To this end, we conduct a simulation study and investigate the behavior of the predictive demand densities over a number of observation periods. With this simulation study, we report overall performances of the exact, approximate and naive posteriors and an illustrative example of ensuing operational costs.

In our numerical analysis, we used a sample test bed generated from the published statistics of Agrawal and Smith (1996) who analyzed retail sales data at retail outlet level for a particular type of men's slacks at a major retail chain. They investigated goodness of fit of negative binomial and Poisson distributions on the entire sales data, and illustrated, through a simulation study, the impact of censoring on updating with a point-estimate (frequentist) approach. Hence, their numerical study can be viewed as a frequentist version of ours. For illustration, we consider the scenarios that correspond to high-peak and low-off-peak demand patterns in their study (see Table 2 in Agrawal and Smith, 1996). We take the reported distributions as the benchmark, termed 'true', and compare the efficacy of the approximations against this benchmark.

For the negative binomial demand ($X|p \sim NB(r,p)$), the case when $r = r_0$ and $p = p_0$ corresponds to the 'true' (benchmark) distribution. From the reported data, we have $r_0 = 3,4$ and $p_0 = 0.4124, 0.5556$ as the true parameter values. Supposing that p is the parameter unknown to us, we consider a beta prior for p with parameters α and β. These prior parameters are chosen so that the prior is substantially different from the

Table 1
Negative binomial demand

$E_{p_0}(X)$	r_0	p_0	α	β	μ_p	$E_{\mu_p}(X)$
5.7	4	0.4124	0.8	7.2	0.1	36
5.7	4	0.4124	2.4	0.6	0.8	1
5.7	4	0.4124	1.0	1.0	0.5	4
2.4	3	0.5556	0.8	7.2	0.1	27
2.4	3	0.5556	1.0	1.0	0.5	3

$X|p \sim NB(r,p)$, $p \sim Beta(\alpha,\beta)$. $E(X|p) = r(1-p)/p$. Assumed true value of $p = p_0 = \{0.4124, 0.5556\}$, $\mu_p = E(p)$, $E_p(X) = E(X|r,p)$.

assumed 'truth', providing worst case scenarios for the approximations, as well as the uniform prior. (We found that alternative prior parameter specifications result in approximations no worse than those presented here; in most cases the approximations are excellent.) The set of parameters used for negative binomial demand is presented in Table 1. Table 1, μ_p denotes the expected value of the unknown parameter p with the given prior, and $E_{p_0}(X)$ and $E_{\mu_p}(X)$ denote the expected demand in a period with the true value of p and mean of its given prior distribution, respectively.

For the Poisson demand ($X|\lambda \sim Poisson(\lambda)$), we tried to match the mean of the prior distribution and the expected demands given the true mean with those of the negative binomial case. This choice allows us to assess the impact of the change in distributions while the means are kept constant. In particular, for the Poisson distribution, the true values of process parameter, λ_0 are set to 5.7 and 2.4. For the prior distribution of λ, we consider gamma distributions for which the means match the expected demand in the negative binomial demand scenario with success probability set to μ_p in Table 1. The set of parameters used in Poisson demand is given Table 2. Note that in Table 2, $E(\lambda)$ denotes the expected demand in a period with the given prior distribution. As in the negative binomial case, the prior parameters were chosen to reflect the impact of large deviations from the 'truth'. The prior parameters in the first three cases are extreme in that the prior is substantially different from the assumed 'truth', providing worst case scenarios for the approximations. The fourth case demonstrates the performance of the approximation when a reasonable prior distribution is used. We begin the discussion of our findings with the goodness of the posterior approximation. In our numerical study, we test this issue in two ways: (i) comparison of posteriors for specific realization instances, and (ii) comparison of posteriors in an overall sense. We first focus on specific censored observations. Recall that the posterior obtained with our proposed approximation differs from the exact posterior only in cases of censoring; therefore, we construct our experiments as follows. We choose the initial stocking level s at 5, 15, and 30 units and assume that the demand realization during the initial period was such that a stockout has occurred resulting in a censored observation. This construction enables us to observe the effects of our approximation for likely, less likely and highly unlikely cases of demand realizations corresponding to the events that demand in the period exceeds 5, 15 and 30 units, resp.

Figs. 1–4 compare the exact and approximate posterior distributions obtained for negative binomial demand for different priors and initial stocking levels. We also provide in all figures the posterior distributions (dashed lines) that would be obtained, had we ignored the censored nature of the data, *i.e.*, if the censored data were treated as uncensored. As seen from the figures, total ignorance of censoring results in gross errors in the estimation of the posterior distributions. We observe that the approximations are overall quite good. The cases in Figs. 1 and 2 corresponding to high-peak demand exhibit similar poor behavior in the left tail

Table 2
Poisson demand

λ_0	α	β	$E(\lambda)$
5.7	3.24	0.09	36
5.7	0.36	0.36	1
2.4	3.24	1.20	27
2.4	4.14	1.72	2.4

$X|\lambda \sim Poisson(\lambda)$, $\lambda \sim Gamma(\alpha,\beta)$. $E(X|\lambda) = 1/\lambda$. Assumed true value of $\lambda = \lambda_0 = \{5.7, 2.4\}$.

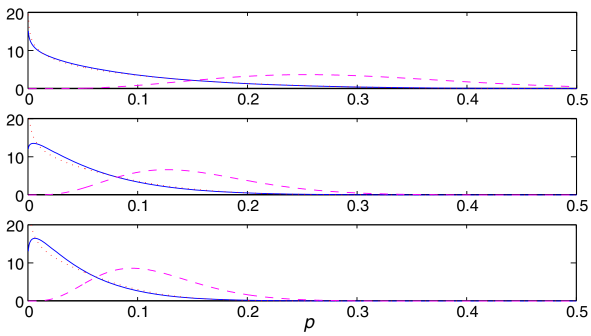

Fig. 1. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim NB(4,p)$, $p \sim Beta(0.8, 7.2)$; $s = 5$, 15, 30 from top to bottom.

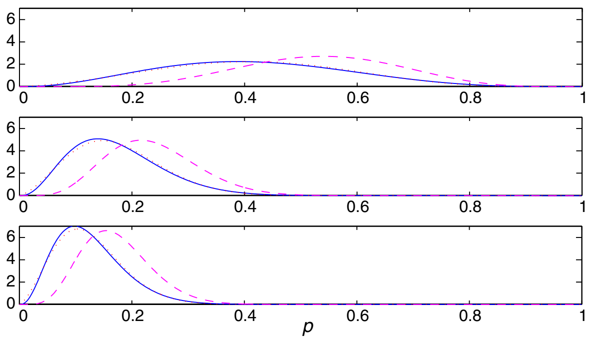

Fig. 2. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim NB(4,p)$, $p \sim Beta(2.4, 0.6)$; $s = 5$, 15, 30 from top to bottom.

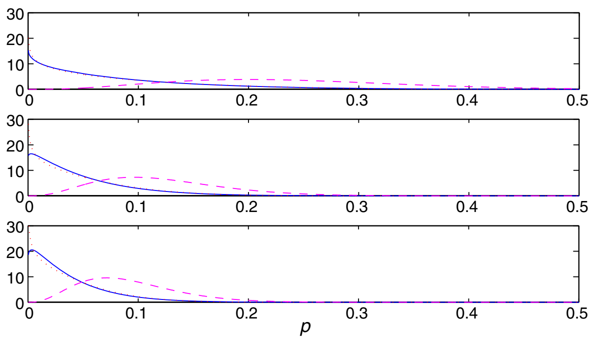

Fig. 3. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim NB(3,p)$, $p \sim Beta(0.8, 7.2)$; $s = 5$, 15, 30 from top to bottom.

indicating that the approximation slightly overestimates the likelihood of the smaller values of the parameter p as s increases. The third and fourth cases corresponding to low-off-peak demand show excellent approximations, although the performance around the mode declines as the skewness of the posterior increases.

Figs. 5–8 compare the exact and approximate posterior distributions for Poisson demand under the same stocking and demand occurrence scenarios. Note that the last parameter set considered for Poisson demand corresponds to the case where the prior information captures the mean of the demand process on average but has some uncertainty about it; that is, we are hitting the true mean of the underlying demand distribution right

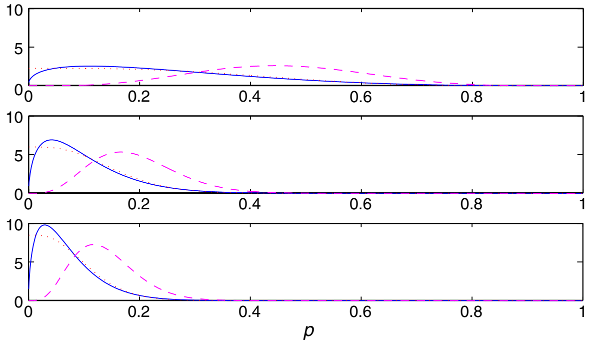

Fig. 4. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim NB(3,p)$, $p \sim Beta(1.0,1.0)$; $s = 5$, 15, 30 from top to bottom.

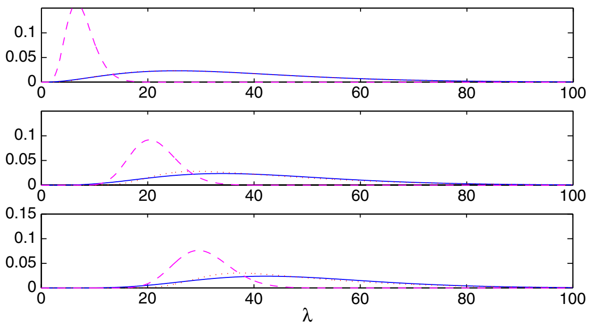

Fig. 5. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim Poisson(\lambda)$, $\lambda \sim Gamma(3.24, 0.09)$; $s = 5$, 15, 30 from top to bottom.

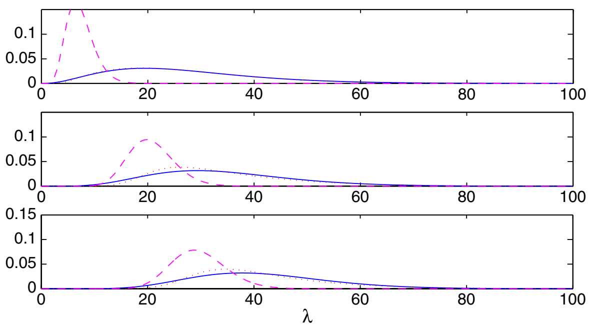

Fig. 6. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim Poisson(\lambda)$, $\lambda \sim Gamma(3.24, 0.12)$; $s = 5$, 15, 30 from top to bottom.

on the mark. The prior with the true mean generates approximations that are indistinguishable from the exact; hence, the impact of the uncertainty around the true mean appears to be minimal.

A general observation in these results is that as the initial stocking level s increases, the approximation worsens. This follows from the design of our experiment since we *assume* a censored observation to compute the posteriors. In reality, a censored observation at large values of demand would be less likely because, as the initial stocking levels increase, the service levels also increase, which in turn implies that a stockout is getting less and less likely. But, in our experiment for illustration purposes, we *deliberately* impose the unlikely event that a stockout (censoring) *has* occurred. To alleviate this problem arising from the artificial construction of

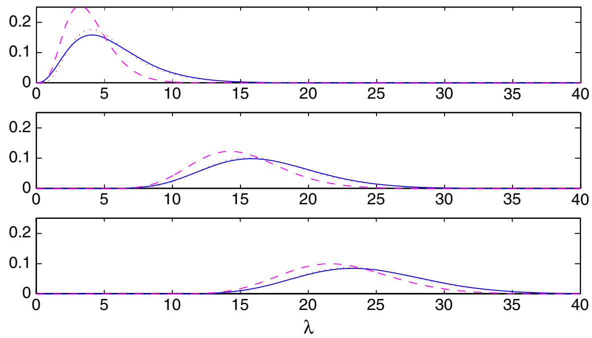

Fig. 7. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim Poisson(\lambda)$, $\lambda \sim Gamma(0.36, 0.36)$; $s = 5$, 15, 30 from top to bottom.

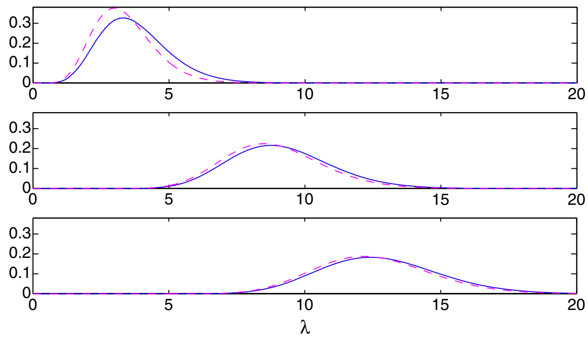

Fig. 8. Exact (solid), approximate (dotted), and exact ignoring censoring (dashed) posterior distributions. $X \sim Poisson(\lambda)$, $\lambda \sim Gamma(4.14, 1.72)$; $s = 5$, 15, 30 from top to bottom.

the experiment, and to obtain an overall goodness measure that takes into account the likelihood of rare events which result in poor posterior calculation, we introduce and define a likelihood-weighted error measure defined as

$$\epsilon = \sum_{x=s}^{\infty} f(x|p) \int_0^1 |\pi(p|x) - \pi_A(p|x)|/\pi(p|x)\,\mathrm{d}p, \tag{28}$$

where $f(x|p)$ is the demand likelihood computed under the true parameter p, and $\pi(p|x)$ and $\pi_A(p|x)$ are the exact and approximate posterior distributions, respectively. Table 3 tabulates ϵ of the posteriors

Table 3
Weighted error, averaged over starting inventory levels $s(\bar{\epsilon})$, for comparing the approximate and exact posterior inference

r	α	β	$\bar{\epsilon}$	ϵ when $s = 5$	ϵ when $s = 30$
$X \sim NB$					
4	2.4	0.6	0.0055	0.0059	0.0064
4	0.8	7.2	0.0057	0.0041	0.030
3	0.8	7.2	0.0016	0.0044	0.020
4	1.0	1.0	0.021	0.023	0.031
3	1.0	1.0	0.0079	0.014	0.019
$X \sim Poisson$					
1	0.36	0.36	0.0026	0.00098	0.000062
36	3.24	0.09	0.00094	0.000011	0.00088
27	3.24	0.12	0.00093	0.000037	0.00098

The weights are the corresponding negative binomial and poisson demand likelihood values for each s.

Table 4
Impact of ignoring censoring of data

s	r	α	β	α^*	β^*	$s^*_{.90}$	$s^*_{.95}$	$s^*_{.99}$	$\bar{\alpha}$	$\bar{\beta}$	$\overline{s_{.90}}$	$\overline{s_{.95}}$	$\overline{s_{.99}}$	$\overline{SL}_{.90}$	$\overline{SL}_{.95}$	$\overline{SL}_{.99}$
5	4	2.4	0.6	3.35	4.71	17	24	47	6.4	5.6	9	12	19	73	82	92
10	4	2.4	0.6	3.51	8.85	29	40	76	6.4	10.6	15	21	33	70	82	92
15	4	2.4	0.6	3.58	13.01	40	56	103	6.4	15.6	23	29	46	73	77	88
20	4	2.4	0.6	3.62	17.17	52	71	131	6.4	20.6	29	37	60	71	80	92
25	4	2.4	0.6	3.65	21.33	64	87	160	6.4	25.6	36	46	74	72	81	93
5	4	1.0	1.0	1.40	4.20	69	121	409	5.0	6.0	13	17	29	51	60	75
10	4	1.0	1.0	1.50	7.50	105	180	565	5.0	11.0	22	29	49	51	60	76
15	4	1.0	1.0	1.55	10.82	143	240	725	5.0	16.0	31	41	70	51	60	77
20	4	1.0	1.0	1.57	14.14	180	300	898	5.0	21.0	41	53	90	51	60	77
25	4	1.0	1.0	1.58	17.47	218	360	1067	5.0	26.0	51	66	110	51	60	77

$X \sim NB(r,p)$ and $p \sim Beta(\alpha,\beta)$.

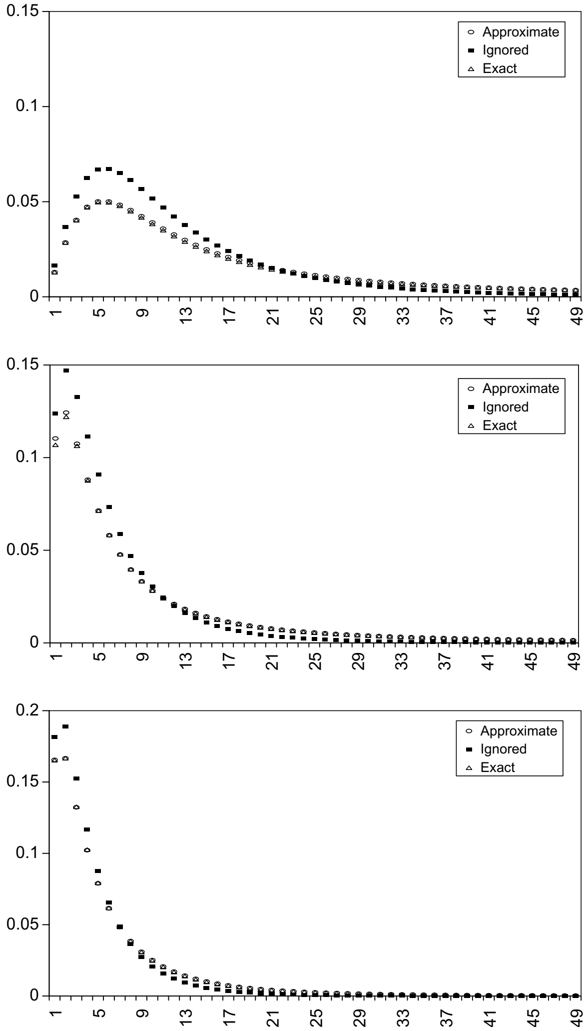

Fig. 9. Comparison of predictive demand p.m.f.s with different initial priors. $s = 6$; $n = 1$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha,\beta)$ with $(\alpha,\beta) \in \{(0.8,7.2),(1,1),(2.4,0.6)\}$.

corresponding to the parameter values used in our test bed. For negative binomial demand, we see that the approximation is much better overall than the best case shown in Fig. 1 (that is, ϵ is smaller than that when $s = 5$ and 30). For the other case, it is slightly worse overall than the worst case shown. For the unreported case of Poisson demand, the same observations hold. Therefore, we can conclude that the two-moment approximation procedure for obtaining the posterior parameters in the presence of censoring provides very good estimates of the underlying distributions when used in single updates.

The next issue we examine is whether or not ignoring the occurrence of censoring impacts the operational decisions of the inventory system, and if so, to what extent? To test this, we first compute the posterior parameters with and without considering the presence of censoring; then, based on the computed posterior parameters, we determine the corresponding optimal stocking levels in the next period under 90%, 95% and 99% service levels.

In Table 4 we report results for the negative binomial demand with different priors under two initial stocking level/demand scenarios in Table 4. The first four columns show the initial stocking levels (s), the known parameter of the demand distribution (r), and the parameters (α, β) of the prior on p; the next four columns show the posterior parameters computed under the consideration of censoring (α^*, β^*) and the corresponding optimal stocking levels $(s^*_{0.90}, s^*_{0.95}, s^*_{0.99})$; and, the last five columns show the posterior parameters $(\bar{\alpha}, \bar{\beta})$ and

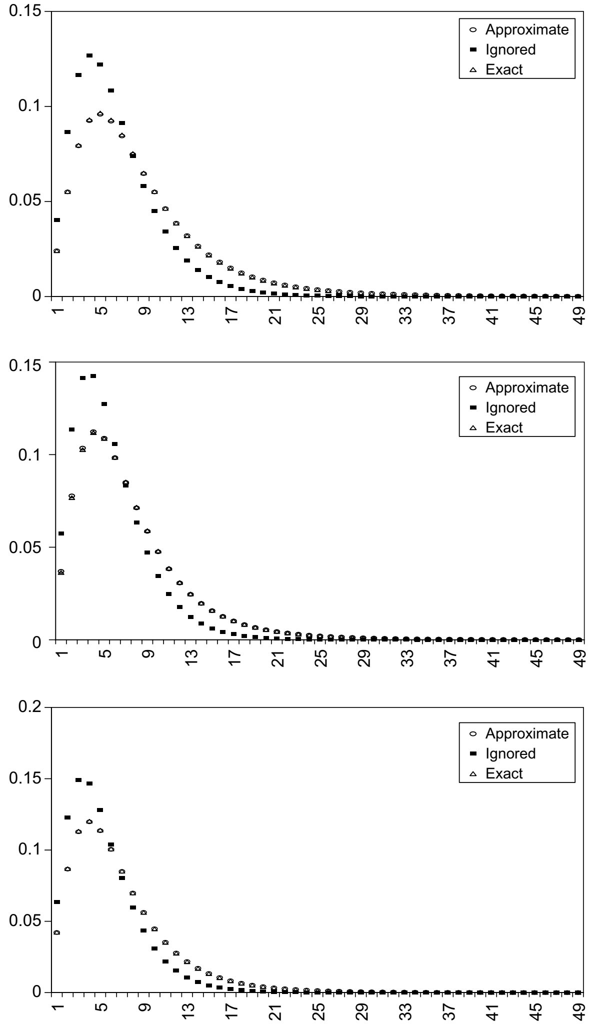

Fig. 10. Comparison of predictive demand p.m.f.s with different initial priors. $s = 6$; $n = 8$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha, \beta)$ with $(\alpha, \beta) \in \{(0.8, 7.2), (1, 1), (2.4, 0.6)\}$ from top to bottom.

corresponding stocking levels $(\bar{s}_{0.90}, \bar{s}_{0.95}, \bar{s}_{0.99})$ computed *as if* the observation was not censored. We note that stocking levels determined for the next period are significantly lower when ignoring censoring as compared to the values computed by taking censoring into account. The operational impact of such underestimations is shown in the last three columns which tabulate the corresponding service levels attained ($\overline{SL}$) if we set the stock quantity at the values specified in the preceding three columns. Observe that these levels are considerably lower than the desired service levels. Therefore, we conclude that ignoring the presence of censoring may result in highly suboptimal control policy parameters.

So far, we have the goodness of the approximation and the impact of information about censoring in single updates only. Next, we address the same issues in a setting of *sequential* updates. The practical setting we envision is as follows. A firm has collected a set of past sales data (observations over a number of periods) with stocking levels over this time horizon. The stocking levels are identical and have been determined on the basis of initial beliefs, as would be the case for new products and/or markets. Now, Bayesian updating is to be used to obtain the predictive demand for the next period of interest (or the rest of the horizon). To study the behavior of the proposed approximation and its benefits in this sequential setting, we performed a simulation study. We used a test horizon consisting of 16 periods, in each of which, demand has a true negative binomial distribution with parameters $r_0 = 4$ and $p_0 = 0.4124$. The initial prior on p is taken to be Beta with

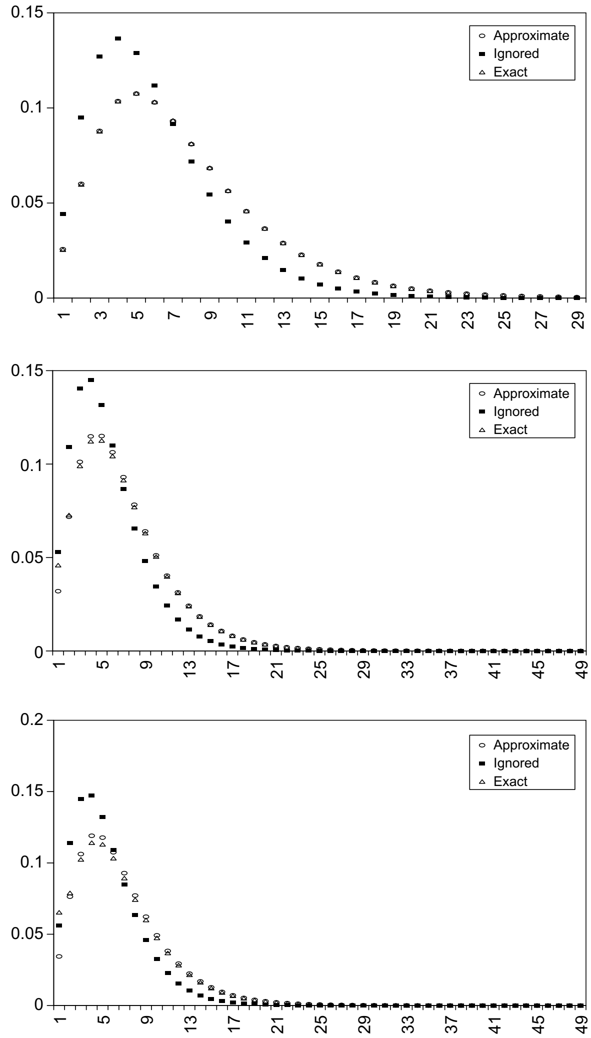

Fig. 11. Comparison of predictive demand p.m.f.s with different initial priors. $s = 6$; $n = 16$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha,\beta)$ with $(\alpha,\beta) \in \{(0.8, 7.2), (1, 1), (2.4, 0.6)\}$ from top to bottom.

$(\alpha, \beta) \in \{(0.8, 7.2), (1.0, 1.0), (2.4, 0.6)\}$. The case (1.0, 1.0) corresponds to the non-informative uniform prior. The stocking levels were fixed for all periods at $s = 2$, 6, 10, ∞. To simulate random observation series, 100 random demand streams were generated over 16 periods from the stationary true demand distribution. Noting that the true demand has a mean of *5.7* and a standard deviation of *3.7*, the imposed stocking levels

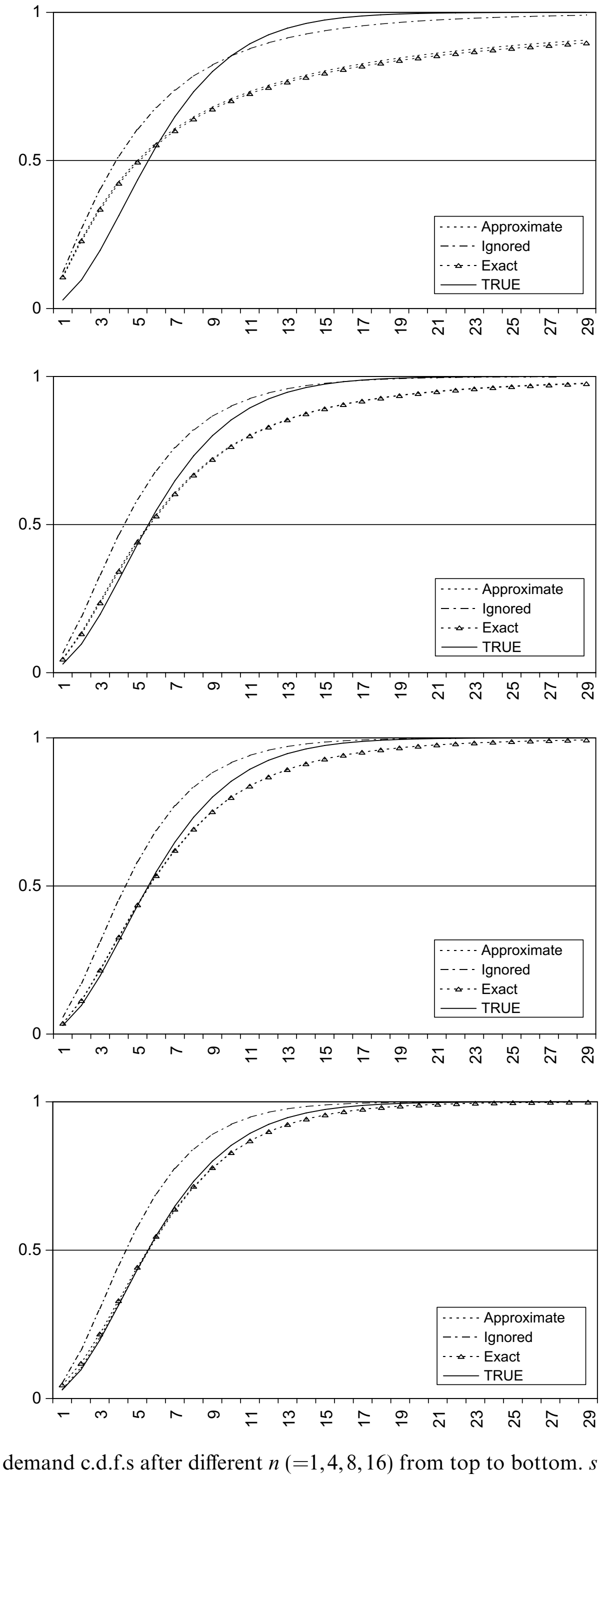

Fig. 12. Comparison of predictive demand c.d.f.s after different n (=1,4,8,16) from top to bottom. $s = 6$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha, \beta)$ with non-informative prior.

correspond to the true demand mean (with a resulting stockout probability of 35%), one standard deviation above and below the true demand mean (with stockout probabilities of 80% and 21%), and the case of no censoring (0% stockout probability). Incidentally, the mean effective demands resulting from the assumed beta-distributed initial priors are 51.4, 11.6 and 1.8. In computing the reported demand densities of this experiment, we did the following. After each period, we calculated the three posteriors – exact, approximate as proposed herein and that computed with any censoring ignored. With these posteriors, we also computed the predictive demand densities. Observe that we end up with 100 predictive densities for each period corresponding to each one of the streams. As the predictive demand density $f_X(x)$ to report in the sequel, we used the arithmetic average of the p.m.f. values at x obtained from the 100 streams for each period. The exact posterior density was obtained numerically as a probability mass function due to its analytical complexity, where we used increments of 0.01 on the unknown parameter p. We address two main issues: (i) the predictive power of the proposed approximation as n, the number of observations in a series of updates increases; and (ii) the impact of ignoring censoring effects in a series of updates.

In Figs. 9–11, we plot the predictive demand densities for different initial priors after one period ($n = 1$), eight periods ($n = 8$) and sixteen periods ($n = 16$) with the stocking level fixed at $s = 6$. We observe that the

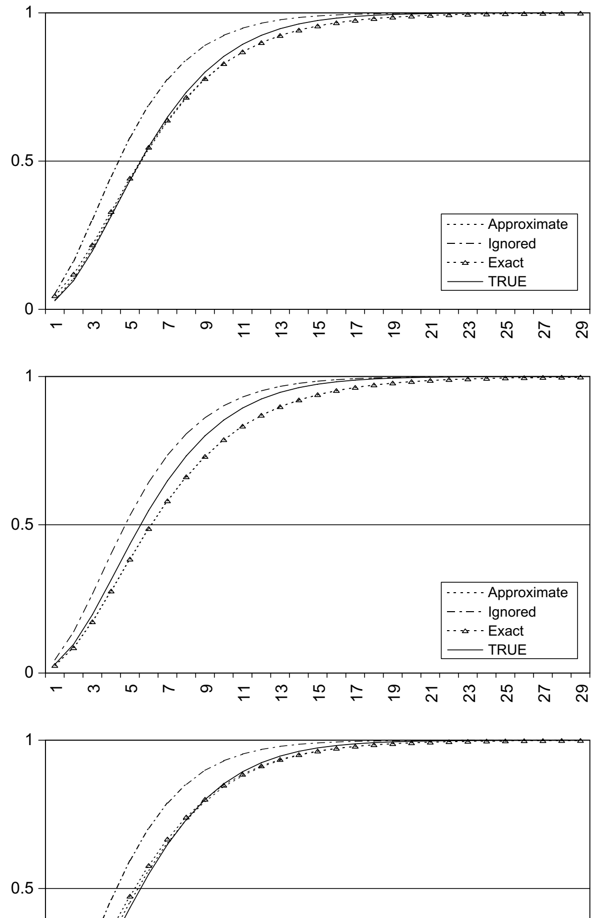

Fig. 13. Comparison of predictive demand c.d.f.s with different initial priors. $s = 6$; $n = 16$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha,\beta)$ with $(\alpha,\beta) \in \{(0.8,7.2),(1,1),(2.4,0.6)\}$ from top to bottom.

proposed approximation yields predictive densities that are very close to those obtained exactly. Noting that each period corresponds to an observation in a series of updates, we conclude that the approximation performs well regardless of the number of consecutive updates. These figures also illustrate that the posterior densities computed when censoring effects are ignored perform badly and always result in underestimation of the true demand, as expected.

To supplement the above findings, we present the c.d.f.s. of the predictive demand for the non-informative prior after $n = 1, 4, 8$ and 16 with $s = 6$ in Fig. 12. Here, we also plot the true demand distribution for comparison. We again conclude that the approximation yields estimates that are very close to those obtained with the exact posteriors. As n increases, the estimated densities approach the true density, as expected. It is important to also observe that the predictive demand obtained by ignoring censoring effects is significantly stochastically smaller than the true density for all n. In Fig. 13, we present the true demand and the three estimated c.d.f.s of the predictive demand after $n = 16$ for all three priors with $s = 6$. We see that the performance of the proposed approximation w.r.t. the exact computation of posteriors does not depend on the choice of the initial prior.

In Fig. 14, we address the impact of stoking level on the estimates. We present the true demand distribution and the three estimated c.d.f.s of the predictive demand after $n = 16$ with the uninformative initial prior for s.

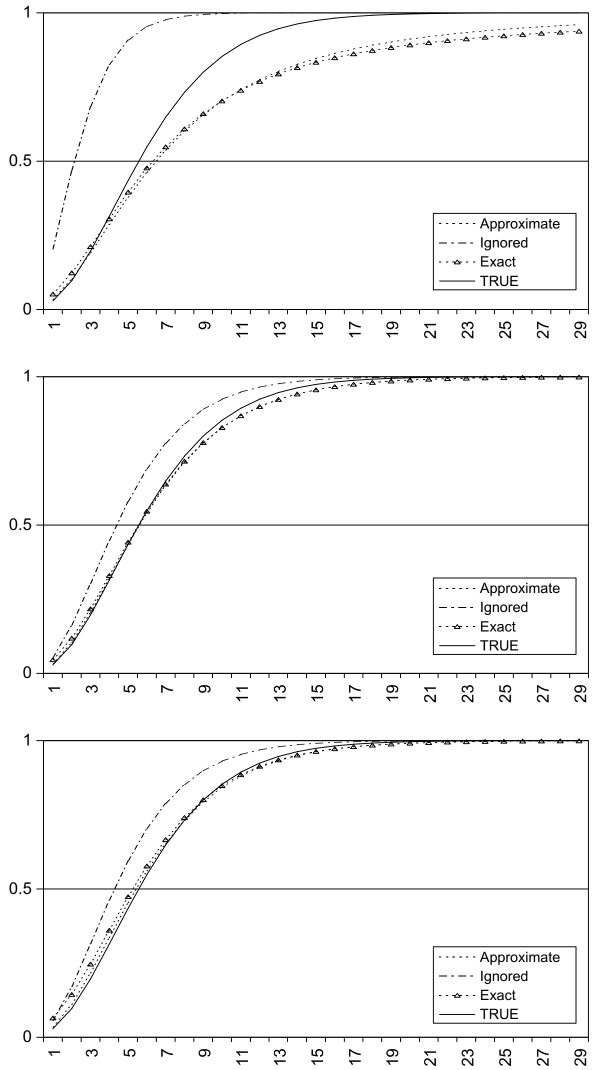

Fig. 14. Comparison of predictive demand c.d.f.s with different stocking levels, $s = 2, 6, 10$ from top to bottom. $n = 16$; $X \sim NB(r,p)$ and $p \sim Beta(\alpha, \beta)$ with non-informative prior.

Table 5
Illustrative example for impact of estimation on computed stocking levels

SL	s^*_{true}	s^*_{exact}	s^*_{approx}	s^*_{ignore}	SL_{exact}	$\text{SL}_{\text{approx}}$	$\text{SL}_{\text{ignore}}$	Δ%
80	8	9	9	7	85.3	85.3	73.2	5.8
90	11	11	11	9	92.4	92.4	85.3	6.7
95	13	14	14	11	97.4	97.4	92.4	3.3
98	15	17	17	13	99.2	99.2	96.3	5.0
99.5	19	22	22	16	99.9	99.9	98.8	3.7

$X \sim NB(4, 0.4124)$; non-informative prior; $s = 6$; $n = 16$.

We see that stocking level has very little impact on the performance of the proposed approximation w.r.t. the exact computation of the predictive demand. However, as expected, ignoring censoring results in highly inaccurate estimation for low s (high ratio of stockout cases).

We should mention that the observations above hold for all of the other cases considered but not reported herein for brevity.

Finally, in order to illustrate the impact of estimation errors on costs, we provide an illustrative example in Table 5 based on Fig. 13 with the non-informative prior and $s = 6$. In this table, we report the stocking levels corresponding to *80*, *90*, *95*, *98* and *99.5* per cent desired service levels with the true demand and the three predictive demand distributions estimated after $n = 16$. (The desired service levels correspond to $C_u/C_h = 4{,}9{,}19{,}49$ and *199*.) One may view this experiment as demand estimation for future periods for a fashion good with past *16* days' sales data. Observe that the stocking levels obtained with censoring ignored are lower than those obtained with the exact (and approximate) determination of posteriors as expected from Proposition 4.1. Taking the unit excess cost as unity, we also report the percentage increase (Δ%) in per period costs as one uses the stocking level computed by ignoring the censoring effects versus the proposed approximation. Due to the discrete nature of demand, Δ% is not monotonic in the desired service level. This small illustrative example shows that the cost impact can be quite high for newly launched products such as textiles where the profit margins are notoriously low (2–3%).

6. Conclusion

In this study, we considered Bayesian updating of demand in the presence of censored observations for the newsvendor problem. Having shown that non-informative priors are improper for some of the demand distributions considered, we developed exact posterior distributions for negative binomial, gamma, Poisson and normal demands starting with conjugate priors. When censoring is introduced, the posterior distributions turn out to be more complicated, usually in the format of a weighted sum of distributions, and no longer retain the conjugacy property for successive updating. Therefore, we proposed a two-moment approximation which retains the conjugate property for the posteriors. This approximation allows for analytical tractability in successive updating. Numerical studies on a test bed designed with actual sales data indicate that the approximation performs extremely well even for sequential updates, and that the adverse impact of neglecting the presence of censoring in updating is significant on the operational performance of the inventory system. A simulation study of showed that the approximation performance does not deteriorate with successive updates. The proposed approximation can also be applied to continuous random variables and settings other than inventory control.

Appendix

Proof of Theorem 3.3. We have the probability mass function of the observed sales given by

$$p(M = k|\lambda) = \begin{cases} \mathrm{e}^{-\lambda}(\lambda)^k/k! & \text{if } k < s, \\ \sum_{i=s}^{\infty} \mathrm{e}^{-\lambda}(\lambda)^i/i! & \text{if } k = s \end{cases}$$

and the joint distribution of λ and M is

$$h_\lambda(\lambda, M=k) = \begin{cases} \frac{\beta^\alpha e^{-\beta\lambda}\lambda^{\alpha-1}}{\Gamma(\alpha)} \cdot \frac{e^{-\lambda}(\lambda)^k}{k!} & \text{if } k < s, \\ \frac{\beta^\alpha e^{-\beta\lambda}\lambda^{\alpha-1}}{\Gamma(\alpha)} \sum_{i=s}^{\infty} \frac{e^{-\lambda}(\lambda)^i}{i!} & \text{if } k = s. \end{cases}$$

From the above, we find the marginal distribution of M as

$$p_M(k) = \begin{cases} \int_0^\infty \frac{\beta^\alpha e^{-\beta\lambda}\lambda^{\alpha-1}}{\Gamma(\alpha)} \cdot \frac{e^{-\lambda}(\lambda)^k}{k!} d\lambda & \text{if } k < s, \\ \int_0^\infty \frac{\beta^\alpha e^{-\beta\lambda}\lambda^{\alpha-1}}{\Gamma(\alpha)} \sum_{i=s}^{\infty} \frac{e^{-\lambda}(\lambda)^i}{i!} d\lambda & \text{if } k = s, \end{cases}$$

which, after some algebra, can be written as

$$p_M(k) = \begin{cases} \frac{\Gamma(\alpha+k)}{\Gamma(\alpha)\Gamma(k)} \left(\frac{1}{\beta+1}\right)^k \left(\frac{\beta}{\beta+1}\right)^\alpha & \text{if } k < s, \\ \sum_{i=s}^{\infty} \frac{\Gamma(\alpha+i)}{\Gamma(\alpha)\Gamma(i)} \left(\frac{L}{\beta+1}\right)^i \left(\frac{\beta}{\beta+1}\right)^\alpha & \text{if } k = s. \end{cases}$$

The first moment of the posterior is given by

$$\begin{aligned} m_1 = E(\lambda|M=s) &= \frac{(\beta+1)^\alpha}{A(\rho,s,\alpha)} \int_0^\infty \lambda e^{-\lambda(\beta+1)}\lambda^{\alpha-1} \sum_{i=s}^{\infty} \frac{(\lambda)^i}{i!} d\lambda = \frac{(\beta+1)^\alpha}{A(\rho,s,\alpha)} \sum_{i=s}^{\infty} \frac{\Gamma(\alpha+i+1)}{i!(\beta+1)^{\alpha+1}} \rho^i \\ &= [(\beta+1)A(\rho,s,\alpha)]^{-1} \sum_{i=s}^{\infty} (\alpha+i) \frac{\Gamma(\alpha+i)}{i!} \rho^i = \frac{1}{(\beta+1)} \left[\alpha + \frac{\sum_{i=s}^{\infty} i \frac{\Gamma(\alpha+i)}{i!} \rho^i}{A(\rho,s,\alpha)} \right]. \end{aligned} \tag{29}$$

Similarly, the second moment is

$$\begin{aligned} m_2 = E(\lambda^2|M=s) &= \frac{(\beta+1)^\alpha}{A(\rho,s,\alpha)} \int_0^\infty \lambda^2 e^{-\lambda(\beta+1)}\lambda^{\alpha-1} \sum_{i=s}^{\infty} \frac{(\lambda)^i}{i!} d\lambda = \frac{(\beta+1)^\alpha}{A(\rho,s,\alpha)} \sum_{i=s}^{\infty} \frac{\Gamma(\alpha+i+2)}{i!(\beta+1)^{\alpha+2}} \rho^i \\ &= (\beta+1)^{-2} \left[\alpha(\alpha+1) + (\alpha+2)\rho \frac{A'(\rho,s,\alpha)}{A(\rho,s,\alpha)} + \frac{\sum_{i=s}^{\infty} i^2 \rho^i \frac{\Gamma(\alpha+i)}{i!}}{A(\rho,s,\alpha)} \right]. \end{aligned} \tag{30}$$

Result follows from noting that

$$\sum_{i=s}^{\infty} i^2 \rho^i \frac{\Gamma(\alpha+i)}{i!} = \sum_{i=s}^{\infty} [i(i-1)+i]\rho^2 \frac{\Gamma(\alpha+i)}{i!} = \rho^2 A''(\rho,s,\alpha) + \rho A'(\rho,s,\alpha). \qquad \square$$

Proof of Theorem 3.4. Let us again denote by ϕ and Φ the standard normal density and cumulative distribution functions respectively, with $\overline{\Phi} = 1 - \Phi$. We first provide the following results, obtained after applying integration by parts

$$\int_s x \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{ -\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 \right\} dz = \sigma[\phi(z_s) + \bar{\Phi}(z_s)[1 + \mu/\sigma] \tag{31}$$

$$\int_s x^2 \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{ -\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 \right\} dz \tag{32}$$

$$= \int_{z_s} \frac{1}{\sqrt{2\pi}} [z^2\sigma + 2\mu\sigma z + (\mu)^2] \exp\left\{ -\frac{1}{2} z^2 \right\} dz = [\sigma + (\mu)^2]\bar{\Phi}(z_s) + \sigma[z_s + 2\mu]\phi(z_s), \tag{33}$$

where $z_s = (s - \mu)/\sigma$. If in a given period lost sales has occurred the posterior distribution function of μ is given by

$$F(u) = P(\mu \leqslant u \mid X > s) = \frac{P(\mu \leqslant u, X > s)}{P(X > s)}, \quad \forall \mu \in R,$$

where the joint probability is given by

$$\begin{aligned} P(\mu \leqslant u, X > s) &= \int_{-\infty}^{u} \int_{s} \frac{1}{2\pi\sigma^2\tau^2} \exp\left\{\frac{-(\mu^2 + \gamma^2 - 2\mu\gamma)}{2\tau^2} + \frac{-(x^2 + \mu^2 - 2\mu x)}{2\sigma^2}\right\} \mathrm{d}x\,\mathrm{d}\mu \\ &= \int_{-\infty}^{u} \bar{\Phi}\left(\frac{s-\mu}{\sigma}\right) \cdot \frac{1}{\sqrt{2\pi\tau^2}} \exp\left\{\frac{-(\mu^2 + \gamma^2 - 2\mu\gamma)}{2\tau^2}\right\} \mathrm{d}\mu. \end{aligned}$$

The posterior density is obtained from taking the derivative of the above expression with respect to u and dividing by one minus the marginal distribution function of X, which has a normal distribution with mean γ and variance η. The moments of the resulting posterior density is obtained referring to the expressions given by (31)–(33).

References

Agrawal, N., Smith, S.A., 1996. Estimating negative binomial demand for retail inventory management with lost sales. Naval Research Logistics 43, 839–861.

Arrow, K.J., Harris, T., Marschak, J., 1951. Optimal inventory policy. Econometrica 19, 250–272.

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

Berger, J.O., 1985. Statistical Decision Theory and Bayesian Analysis, second ed. Springer-Verlag.

Berk, E., Gürler, Ü., Levine, R., 2001. The newsboy problem with Bayesian updating of the demand distribution and censored observations. In: ISBA 2000 6th World Meeting Proceedings – Monographs of Official Statistics: Bayesian Methods, Office for Official Publications of the European Communities, Luxembourg, pp. 21–31.

Bradford, J.W., Sugrue, P.K., 1990. A Bayesian approach to the 2-period style-goods inventory problem with single replenishment and heterogeneous Poisson demands. Journal of the Operations Research Society 41, 211–218.

Business Week, 1996. Clearing the Cob Webs from the Stockroom. October 21, p. 140.

Business Week, 2003. Is Wal-Mart too Powerful. October 6, pp. 46–55.

Ding, X., Puterman, M.L., Bisi, A., 2002. The censored newsvendor and the optimal acquisition of information. Operations Research 50, 517–527.

Eppen, G.D., Iyer, A.V., 1997. Improved fashion buying with Bayesian updates. Operations Research 45, 805–819.

Fisher, M.L., Obermeyer, W.R., Hammond, J.H., Raman, A., 1994. Making supply meet demand in an uncertain world. Harvard Business Review (May–June), 221–240.

Geunes, J.P., Hammond, R.V., Hayya, J.C., 2001. Adapting the newsvendor for infinite-horizon inventory systems. International Journal of Production Economics 72, 237–250.

Hill, R.M., 1997. Applying Bayesian methodology with a uniform prior to the single period inventory model. European Journal of Operations Research 98, 555–562.

Hill, R.M., 1999. Bayesian decision-making in inventory modeling. IMA Journal of Mathematics Applied in Business and Industry 10, 147–163.

Iglehart, D.L., 1964. The dynamic inventory problem with unknown demand distribution. Management Science 10, 429–440.

Khouja, M., 1999. The single-period (News-vendor) problem: Literature review and suggestions for future research. Omega 27, 537–553.

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

Lovejoy, W.S., 1990. Myopic policies for some inventory models with uncertain demand distribution. Management Science 36, 724–738.

Lowe, T.J., Schwarz, L.B., McGavin, E.J., 1988. The determination of optimal base-stock inventory policy when the costs of under- and oversupply are uncertain. Naval Research Logistics 35, 539–554.

Nahmias, S., 1994. Demand estimation in lost sales inventory systems. Naval Research Logistics 41, 739–757.

Nahmias, S., Smith, S., 1994. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. Management Science 40, 582–596.

Scarf, H., 1959. Bayes solutions of the statistical inventory problem. Annals of Mathematical Statistics 30, 490–508.

Scarf, H., 1960. Some remarks on Bayes solutions to inventory problem. Naval Research Logistics Quarterly 7, 591–596.

Zipkin, P., 2000. Foundations of Inventory Management. McGraw-Hill, Singapore.