Using simulation to solve the newsvendor problem under parameter uncertainty

David F. Muñoz and David G. Muñoz

***Abstract*—** We discuss the formulation and solution to the newsvendor problem under a Bayesian framework that allows us to incorporate the uncertainty in the parameters of demand modeling (introduced in the process of parameter estimation). We present an example with an analytical solution and use this example to show that a classical approach (without parameter uncertainty) tends to overestimate the expected benefit. Furthermore, we conduct experiments that confirm our results and illustrate the estimation of the optimal order size using stochastic simulation, method that is suggested when model complexity does not allow us to obtain an analytical solution.

***Keywords*—** Bayesian forecasting, inventory management, newsvendor problem, parameter uncertainty, stochastic simulation.

I. Introduction

LET D represent the demand (during the sales period) of a seasonal item. If $w \geq 0$ denotes the loss for every unsold unit at the end of the period, and $u \geq 0$ denotes the profit for every unit sold during the period, the total profit for an order size of Q units is given by

$$b_D(Q) = \begin{cases} uD - w(Q-D), & D < Q, \\ uQ, & D \geq Q. \end{cases} \quad (1)$$

The most common approach (see, e.g, [1]) to find the optimal order size Q_C^* consists in defining a density function $f(y|\theta)$ for the demand D (the analysis is similar for the discrete case), where θ is the parameter vector and, assuming θ is known, we define the expected profit as

$$B_C(Q|\theta) \overset{def}{=} E[b_D(Q)|\Theta=\theta] = u\int_0^Q yf(y|\theta)dy - w\int_0^Q (Q-y)f(y|\theta)dy + uQ\int_Q^\infty f(y|\theta)dy. \quad (2)$$

This function has a derivative when $Q > 0$ so that (by imposing first-order optimality conditions) the optimal order size Q_C^* that maximizes $B_C(Q|\theta)$ satisfies

$$F_C\left(Q_C^*|\theta\right) \overset{def}{=} \int_0^{Q_C^*} f(y|\theta)dy = \frac{u}{u+w}. \quad (3)$$

This research was supported by Asociación Mexicana de Cultura A.C., and the National Council for Science and Technology (CONACYT) of Mexico..

David F. Muñoz is Professor of the Department of Industrial & Operations Engineering at the Instituto Tecnológico Autónomo de México, 01080 Ciudad de México, (phone: 52-55-56284118; e-mail: davidm@ itam.mx).

David G. Muñoz is Director of Yield Analytics at AOL Advertising, Palo Alto, CA 94306 USA (e-mail: davidmm82@gmail. com).

Note that $F_C(y|\theta)$ is the cumulative distribution function (cdf) of the demand given $[\Theta=\theta]$. Furthermore, if $f(x|\theta) > 0$ is continuous in a neighbourhood of Q_C^*, condition (3) is sufficient for finding the optimal value Q_C^*. We also mention, for the sake of clarity, that the formulation presented in [1] is an equivalent formulation where the authors minimize the expected value of $-b_D(Q) - uD$, so that (3) follows from [1], and we are using index C to remark that this is a classical approach described in most textbooks.

In practice, the value of θ is estimated from a data set $x = (x_1, \ldots, x_n)$ using, for example, the (maximum likelihood) estimator that maximizes a likelihood function $L(\theta|x)$. The most common approach for finding the optimal order size consists in setting $\theta = \hat{\theta}$ in (3), where $\hat{\theta} = \hat{\theta}(x)$ is a point estimator. While this procedure is found extensively in Operations Management textbooks, it has the downside of assuming that the point estimator equals the parameter. Thus, in this article we discuss a Bayesian approach to the newsvendor problem (i.e., finding the optimal order size) incorporating uncertainty (introduced by the estimation process) in the parameter vector.

Bayesian methods to incorporate parameter uncertainty for inventory management have been proposed since the pioneer work of Scarf [2], where the author discusses the optimality of a Bayesian updating rule for inventory management. Also, the incorporation of parameter uncertainty using Bayesian methods is proposed in [3], where the author shows how to compute a reorder point by modeling the demand as a multinomial distribution. Bayesian methods have been extensively applied to inventory management in order to propose updating rules for the optimal inventory policy based on new information on the product's demand, see e.g., [4]-[9], and the references therein. However, the use of simulation techniques to estimate performance measures for inventory management is not considered in these articles and the related literature.

When a random sample from demand D is available, the solution of (3) can be approximated by the corresponding sample quantile and (as shown in [10] and [11]) the quality of this approximation improves as the sample size increases. However, our methodology can be applied to a more general situation, since a random sample from future demand D is not necessarily required. Our approach is similar to the one proposed in [9], where the authors illustrate how it can be applied using data on sales only (without information on lost sales) or data on sales occurrences (as in the example shown below in this article). Extensions to the newsvendor problem have been also proposed, for example, uncertainty on price and order size are considered in [12], and the effect of risk aversion in the optimal order size is analyzed in [13].

The following section describes the theory behind the proposed Bayesian approach, which allows the incorporation of parametric uncertainty in the newsvendor problem. Afterwards, in the subsequent section, we illustrate how to estimate the optimal order size using simulation by means of a simple example. This example showcases how to estimate the optimal order size in more complicated problems. In the same section, we present a comparison of the results obtained from applying a classical approach versus the results obtained using the Bayesian approach. Finally, in the last section, we present our conclusions and recommendations.

II. Problem Formulation

Under a Bayesian framework, the parameter vector is a random variable that has a prior density function , thus the posterior density function (given a data set) is given by

$$p(\theta|x) = \frac{p(\theta)L(x|\theta)}{\int_{S_0} p(\theta)L(x|\theta)d\theta}, \tag{4}$$

where $x \in \Re^d$, $\theta \in S_0$ and $L(x|\theta)$ is the likelihood function. From (4) and following the same notation as in (3), the cdf of the demand (given $[X = x]$) is given by

$$\begin{aligned} F_B(y|x) &= E[E[F_C(y|\Theta)]|X = x] \\ &= \int_{S_0} F_C(y|\theta)p(\theta|x)d\theta, \end{aligned} \tag{5}$$

for $y \geq 0$, where $F_C(y|\theta)$ and $p(\theta|x)$ are defined in (3) and (4), respectively. Similarly, from (1) we obtain the expected profit (given $[X = x]$) as

$$\begin{aligned} B_B(Q|x) &\overset{def}{=} u\int_0^Q y dF_B(y|x) \\ &- w\int_0^Q (Q-y)dF_B(y|x) + uQ\int_Q^\infty dF_B(y|x), \end{aligned} \tag{6}$$

where $F_B(y|x)$ is the cdf defined in (5). This shows that $B_B(Q|x)$ has a similar form to $B_C(Q|x)$ defined in (2). Consequently, the optimal order size Q_B^* considering parametric uncertainty satisfies

$$F_B\left(Q_B^*|x\right) = \frac{u}{u+w}, \tag{7}$$

where $F_B(y|x)$ is defined in (5). It is worth mentioning that our problem formulation is different from the one proposed in [9], where the authors propose a dynamic program to solve the newsvendor problem in a multi-period setting in order to show the advantages of using a more efficient updating rule. The main difference is that we consider a single-period expected profit that is explicitly dependent on the available data set *x*. This formulation allowed us to obtain a simple solution in the form of (7).

It is important to point out that for the case where demand is discrete, taking values $d_1 < d_2 < \dots$, the function $F_B(y|x)$ is not continuous, and equation (7) might not have a solution, in which case we must find the value of d_k that satisfies:

$$P[D \leq d_k|X = x] \leq \frac{u}{u+w} \leq P[D \leq d_{k+1}|X = x], \tag{8}$$

in order to evaluate $B_B(d_k|x)$ and $B_B(d_{k+1}|x)$, where:

$$\begin{aligned} B_B(Q|x) &= u\sum_{j\leq Q} jP[D = j|X = x] \\ &- w\sum_{j\leq Q}(Q-j)P[D = j|X = x] + uQP[D > Q|X = x]. \end{aligned} \tag{9}$$

If $B_B(d_k|x) \geq B_B(d_{k+1}|x)$, the optimal order size will be given by $Q_B^* = d_k$, otherwise it will be given by $Q_B^* = d_{k+1}$. Note that in the discrete case (8) is equivalent to (7), in the sense that neither equation considers fixed ordering costs. If there is an initial inventory of Q_0, we should not order when $Q_0 \geq Q_B^*$, otherwise we should order $Q_B^* - Q_0$ units only if $B_B\left(Q_B^*|x\right) - B_B(Q_0|x) > C_0$, where C_0 is the fixed ordering cost.

III. An Illustrative Example

In this section we illustrate the application of the proposed methodology through a similar model to the one presented in [14] for the forecast of an item of intermittent demand. We know that the demand for service parts follows a Poisson process, though there exists uncertainty in the arrival rate Θ_0. Thus, given $[\Theta_0 = \theta_0]$, the times between customers' arrivals are i.i.d. according to the exponential density function:

$$f(y|\theta_0) = \begin{cases} \theta_0 e^{-\theta_0 y}, & y > 0, \\ 0, & \text{otherwise}, \end{cases}$$

where $\theta_0 \in (0,\infty)$. Every client can order j units of an item with probability P_j, $j = 1,\dots,q$, $q \geq 2$. Let $\Theta_1 = (\mathrm{P}_1,\dots,\mathrm{P}_{q-1})$ and $\sum_{j=1}^q \mathrm{P}_j = 1$, then $\Theta = (\Theta_0, \Theta_1)$ denotes the vector of parameters and the parameter space is given by $S_0 = (0,\infty) \otimes S_{01}$, where

$$S_{01} = \left\{(\rho_1,\dots,\rho_{q-1}) : \sum_{j=1}^{q-1}\rho_j \leq 1; \rho_j \geq 0,\ j = 1,\dots,q-1\right\}$$

The total demand for a period of length T is given by

$$D = \begin{cases} \sum_{i=1}^{N(T)} U_i, & N(T) > 0, \\ 0, & \text{otherwise}, \end{cases} \tag{10}$$

where $N(s)$ is the number of clients that arrived in the interval $[0,s]$, $s \geq 0$, and $U_1, U_2 \dots$ are the individual item demands (we assume they are conditionally independent with respect to Θ). The information on Θ consists of (i.i.d.) observations $v = (v_1,\dots,v_n)$, $u = (u_1,\dots,u_n)$ of past clients, where v_i is the time between the arrival of client i and previous client ($i-1$), and u_i is the number of items ordered by client i. The likelihood functions for v and u are given by

$$L(v|\theta_0) = \theta_0^n e^{-\theta_0 \sum_{i=1}^{n} v_i}, \quad (11)$$

and

$$L(u|\theta_1) = \left(1 - \sum_{j=1}^{q-1} \rho_j\right)^{c_q} \prod_{j=1}^{q-1} \rho_j^{c_j}, \quad (12)$$

respectively, where $\theta_1 = (\rho_1, \ldots, \rho_{q-1})$, and $c_j = \sum_{i=1}^{n} I[u_i = j]$ is the number of clients that ordered j items.

From an objective point of view, we can assume a non-informative prior density function for Θ, using Jeffrey's prior density. In the case of the exponential model, Jeffrey's prior density (see, for instance, [15]), is given by $p(\theta_0) = \theta_0^{-1}$, $\theta_0 > 0$. By taking $\theta = \theta_0$ and $x = v$ in (4), from (11) we have

$$p(\theta_0|v) = \frac{\theta_0^{n-1} \left(\sum_{i=1}^{n} v_i\right)^n e^{-\theta_0 \sum_{i=1}^{n} v_i}}{(n-1)!}, \quad (13)$$

which corresponds to a $\text{Gamma}\left(n, \sum_{i=1}^{n} v_i\right)$ distribution, where $\text{Gamma}(\beta_1, \beta_2)$ denotes a gamma distribution with expectation $\beta_1 \beta_2^{-1}$. Similarly, Jeffrey's prior density for the multinomial model (see, for instance, [16]) corresponds to a Dirichlet distribution with density function:

$$p(\theta_1) = \frac{\left(1 - \sum_{j=1}^{q-1} \rho_j\right)^{-1/2} \prod_{j=1}^{q-1} \rho_j^{-1/2}}{B(1/2, \ldots, 1/2)},$$

where $B(a_1, \ldots, a_q) = \prod_{j=1}^{q} \Gamma(a_j) / \Gamma\left(\sum_{j=1}^{q} a_j\right)$, for $a_1, \ldots, a_q > 0$, thus, it follows from (4) and (12) that

$$p(\theta_1|u) = \frac{\left(1 - \sum_{j=1}^{q-1} \rho_j\right)^{c_q - 1/2} \prod_{j=1}^{q-1} \rho_j^{c_j - 1/2}}{B(c_1 + 1/2, \ldots, c_q + 1/2)}, \quad (14)$$

which corresponds to a Dirichlet distribution with parameter $(c_1 + 1/2, \ldots, c_q + 1/2)$. Let $x_i = (v_i, u_i)$, $i = 1, \ldots, n$, $x = (x_1, \ldots, x_n)$, $\theta = (\theta_0, \theta_1)$, and assuming independence, the posterior density is given by $p(\theta|x) = p(\theta_0|v) p(\theta_1|u)$, where $p(\theta_0|v)$ and $p(\theta_1|u)$ are defined in (13) and (14), respectively.

Note that, in this example, we can obtain a closed-form expression for the point estimate of the demand $\mu = E[D|X = x]$, since, from (13) and (14) we have $E[\Theta_0|V = v] = n\left(\sum_{i=1}^{n} v_i\right)^{-1}$, and

$$E\left[\Theta_{1j}|U = u\right] = c^{-1}\left(c_j + 1/2\right),$$

where $c = \sum_{j=1}^{n}\left(c_j + 1/2\right) = n + q/2$), so that following (10) we have

$$\begin{aligned}
\mu &= E\left[E\left[D|\Theta, X = x\right] X = x\right] \\
&= E\left[E\left[N(T)|\Theta\right] E\left[U_1|\Theta\right] X = x\right] \\
&= TE\left[\Theta_0 \sum_{j=1}^{q} j\Theta_{1j} | X = x\right] \\
&= TE\left[\Theta_0|V = v\right] \sum_{j=1}^{q} jE\left[\Theta_{1j}|U = u\right] \\
&= Tn\left(\sum_{j=1}^{n} v_i\right)^{-1} (n + q/2)^{-1} \sum_{j=1}^{q} j\left(c_j + 1/2\right).
\end{aligned}$$

The above expression allows us to calculate a forecast for demand D based on a data set x. In this case, however, it is not easy to obtain a closed-form expression for the cdf and the optimal order size. Thus, we can apply the posterior sampling (PS) algorithm described in section 3.1 of [17] in order to calculate, via simulation, the corresponding optimal order size, given a service level $\alpha = u/(u + w)$. It is worth mentioning that the main idea behind the PS algorithm is to generate simulated observations for the demand by first sampling a parameter value θ from the posterior distribution density $p(\theta|x)$, and the sampling a demand observation from the forecasting model for the demand D (conditional on $\Theta = \theta$).

For the case when $q = 1$ (i.e., every client orders just one unit), the model is simpler and it is not necessary to turn to simulation in order to find the optimal order size. In this case, we can ignore θ_1 and the values u_i (since they are always equal to 1). Let $x = v$, $\theta = \theta_0$, we have that

$$P[D = j|\Theta = \theta] = P[N(T) = j\|\Theta = \theta] = e^{-\theta T} (\theta T)^j \big/ j!,$$

and considering (13), it can be proven that

$$P[D = j|X = x] = \binom{n + j - 1}{j} \left(\frac{\sum_{i=1}^{n} x_i}{T + \sum_{i=1}^{n} x_i}\right)^n \left(\frac{T}{T + \sum_{i=1}^{n} x_i}\right)^j, \quad (15)$$

for $j = 0, 1, \ldots$, which corresponds to a negative binomial distribution. Using (8), (9) and (15), we can determine the optimal order size Q_B^* for this particular case, without resorting to the PS algorithm or simulation.

IV. Experimental Results

In order to illustrate the validity of the PS algorithm and, in particular, how it can be applied in order to determine the optimal order size, we will use the example from the previous section that has a closed-form expression (15) for the posterior distribution of the demand. First of all, we should point out that we considered the values of $T = 15$, $n = 20$, $\sum_{i=1}^{n} x_i = 10$, $u = 9$, $w = 1$. With this data, the optimal service level is $\alpha = u/(u + w) = 0.9$. After applying the Bayesian approach described in equations (8) and (9), and the posterior distribution defined in (14), we obtained an optimal order size of $Q_B^* = 41$, then, by following (6), we have an expected profit of $B_B(Q_B^*|x) = 253.38$, and a service level of $F_B(Q_B^*|x) = 0.901$ (slightly higher than 0.9).

With the objective of comparing the results obtained through the classical approach, notice that, from (10), we can find that the cdf $F_C(y|\theta)$ defined in (3) corresponds to a Poisson distribution with mean θT. On the other hand, the maximum likelihood estimator of θ is $\hat{\theta} = 2$, thus, when applying the classical approach with conditions similar to (8) and (9), we obtained $Q_C^* = 37$, reporting an expected profit of $B_C(Q_C^*|x) = 260.05 > B_B(Q_B^*|x)$ from (2), and a service level, from the posterior distribution in (15), of

$$F_B\left(Q_C^*|x\right) = 0.803 < F_B\left(Q_B^*|x\right)$$

These results suggest that under the classical approach, expected profit is overestimated, and results in a more conservative service level when compared to the Bayesian approach, confirming the intuition that parametric uncertainty proposes a posterior distribution for the demand $F_B(y|x)$ with greater dispersion than the classical approach distribution $F_C(y|x)$. In the following section, we present empirical results that confirm these observations. Subsequently, we will also illustrate how we can estimate the optimal order size when it is not possible (or is extremely complicated) to find a closed-form solution.

A. *Empirical Comparison between the Classical and Bayesian Approaches*

In our first experiment, we assumed an arrival rate for clients of $\theta = 2$ and generated $m = 1000$ samples of arrival times, each of size $n = 20$. For every sample, we calculated $\sum_{i=1}^{n} x_i$ and the optimal order size (under both the classical and Bayesian approaches) with the data from the previous section ($T = 15$, $n = 20$, $u = 9$, $w = 1$). For every sample, we calculated the difference in expected profit between both approaches, $B_C(Q_C^*|x) - B_B(Q_B^*|x)$, and the service level for the (classical) optimal order size: $F_B\left(Q_C^*|x\right)$.

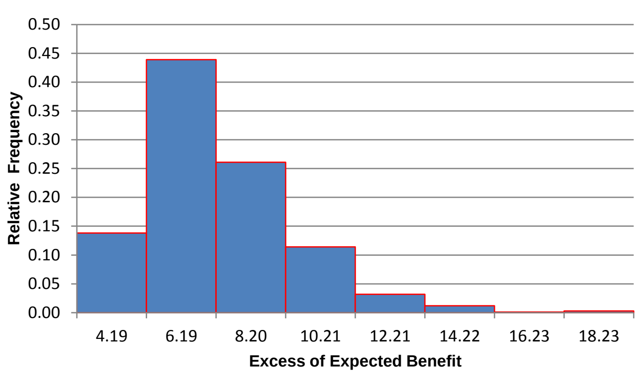

Fig. 1: Histogram of the difference $B_C(Q_C^*|x) - B_B(Q_B^*|x)$ based on 1000 replications of the estimation experiments under the classical and Bayesian approaches.

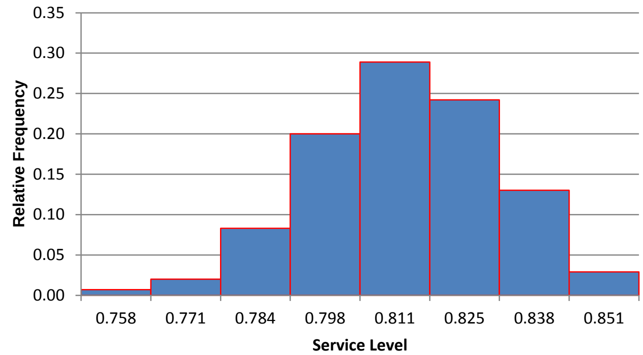

Fig. 2: Histogram of the actual service level $F_B\left(Q_C^*|x\right)$ for the optimal order size under the classical approach based on 1000 replications of the estimation experiments under the classical and Bayesian approaches with *n* = 50.

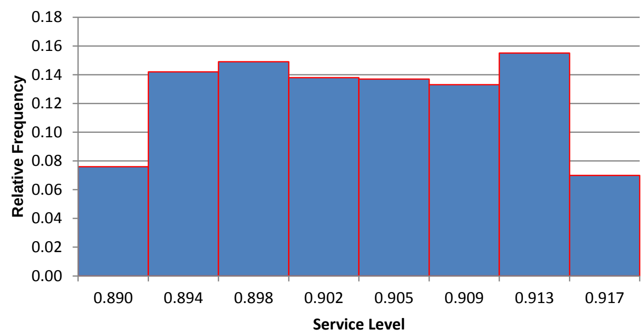

Fig. 3: Histogram of the actual service level $F_B\left(Q_C^*|x\right)$ for the optimal order size under the classical approach based on 1000 replications of the estimation experiments under the classical and Bayesian approaches with *n* = 300.

Table I: Excess of benefits and service levels for the classical order size under different sample sizes.

n	Excess of Benefit		Service Level	
	Mean	St. Dev.	Mean	St. Dev.
5	25.95	18.02	0.732	0.032
10	13.61	6.13	0.770	0.025
20	7.23	2.16	0.813	0.018
50	3.10	0.59	0.861	0.011
100	1.61	0.22	0.885	0.009
150	1.08	0.12	0.894	0.008
200	0.82	0.08	0.899	0.008
250	0.66	0.06	0.901	0.008
300	0.55	0.05	0.903	0.008

Based on Fig. 1, notice that the classical approach has overestimated the expected profit in all replications of the experiment. Similarly, based on Fig. 2, notice that the classical approach has provided a more conservative service level in every replication of the experiment. Despite the results of Fig.2, there should be a positive probability (very small in this case) of having Q_C^* large enough to provide a service level larger than 0.9. In Fig. 3 we report the results of the same experiments of Fig. 2 with a sample size of n = 300, and we illustrate that a service level larger than 0.9 can actually be obtained.

In order to remark that the results provided by the proposed Bayesian approach are consistent with the results of a classical approach, we replicated our previous experiments under different values for the sample size n (of times between arrivals), and the main results are provided in Table I and, as can be seen from this table, the excess of benefit tends to zero as the sample size increases, and the service level tend to the optimal 0.9, illustrating that both methods tend to provide a similar result as the sample size increases (and parameter uncertainty disappears).

B. Analysis of Some Empirical Results

In this Section we study the possibility that the results of Fig. 1 can be generalized for different parameter values.

Let us suppose that $\bar{x}$ is fixed (after simulating a random sample of size n from an exponential distribution with mean $1/\theta$), and let Y and Z be random variables distributed as Poisson with expectation $\lambda = T/\bar{x}$, and negative binomial with parameters n and $p = \bar{x}/(T+\bar{x})$, respectively. Under these definitions, we can easily verify that the classical solution Q_C^* to the newsvendor problem maximizes $E[b_Y(Q)]$, and the Bayesian solution Q_B^* maximizes $E[b_Z(Q)]$, where $b_D(Q)$ is defined in (1).

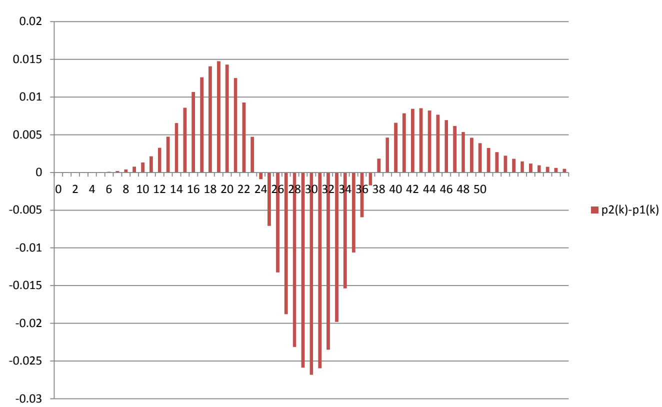

Fig. 4: Graph of $p_2(k) - p_1(k)$ for k = 0, 1,….

In Fig. 4 we show a graph of $p_2(k) - p_1(k)$ under the same parameters of Fig. 1 and $\bar{x} = 0.5$, where $p_1(k) = P[Y = k]$ and $p_2(k) = P[Z = k]$, k = 0, 1,…. In this figure we illustrate that this function satisfies the following property.

Property 1. For $n \geq 1$, $\bar{x} > 0$ and $T > 0$ given, and $g(y)$ as in (16), there exist $0 \leq k_1 < k_2 < \infty$ such that $g(y) > 0$ for $0 \leq y \leq k_1$, $y \geq k_2$; and $g(y) \leq 0$ for $k_1 < y < k_2$.

We will show that Property 1 is satisfied for Y and Z, for which we define

$$g(y) \overset{def}{=} \ln[p_2(y)/p_1(y)] = \ln[\Gamma(n+y)] - \ln[\Gamma(n)] + n\ln(p) + y\ln(1-p) + \lambda - y\ln(\lambda), \quad (16)$$

where $\lambda = T/\bar{x}$ and $p = \bar{x}/(\bar{T}+\bar{x})$. Note that, from $e^{\bar{T}/\bar{x}} > 1 + \bar{T}/\bar{x}$ we easily obtain $g(0) > 0$, so that there exists $k_1 = 0$ such that $g(y) > 0$ for $0 \leq y \leq k_1$. On the other hand, $\lim_{y\to\infty} g(y) = \infty$, so that there exists k_2 such that $g(y) > 0$ and $y \geq k_2$. Finally, $g'(y) = \Psi(n+y) - \ln(a)$, where $a = n/p$ and $\Psi(y)$ is the well-known digamma function, which shows that $g(y)$ decreases for $y < y_0$ and increases for $y > y_0$, where y_0 is the only solution of $\Psi(n+y) - \ln(a) = 0$. Property 1 then follows from the Intermediate Value Theorem by using $\sum_{k=0}^{\infty} p_1(k) = \sum_{k=0}^{\infty} p_2(k) = 1$.

We include the next definition for completeness, although it is well-known. The next Proposition follows from Property 1, Theorem 2 of [18] and $E[Y] = E[Z] = T/\bar{x}$.

Definition 1. We say that a random variable Y has second-order stochastic dominance on a random variable Z if and only if $E[U(Y)] \geq E[U(Z)]$ for any concave function U.

Proposition 1. Let $n \geq 1$, $\bar{x} > 0$ and $T > 0$ be given. If Y, Z are random variables distributed as Poisson with expectation $\lambda = T/\bar{x}$ and negative binomial with parameters n and $p = \bar{x}/(\bar{T}+\bar{x})$, respectively, then Y has second-order stochastic dominance on Z.

Note that the function $b_D(Q)$ is concave, so that $E[b_Y(Q_C^*)] \geq E[b_Z(Q_B^*)]$ follows from Proposition 1, and from $P[\bar{X} > 0] = 1$, where $\bar{X}$ is the mean of a random sample of size n from an exponential distribution with expectation $1/\theta$, we have the following Corollary.

Corollary 1. Let $n \geq 1$, $\bar{x} > 0$ and $T > 0$ be given, then $P[B_C(Q_C^*|x) - B_B(Q_B^*|x) \geq 0] = 1$, where $B_C(Q_C^*|x)$ and $B_B(Q_B^*|x)$ correspond to the example defined in (15).

C. Estimation of the Optimal Order Size using Simulation

With the objective of illustrating how to calculate the optimal order size when the complexity of the model does not allow the calculation of a closed form expression for the solution, in this section we show the use of the PS algorithm to find the optimal order size using simulation.

In order to apply the PS algorithm in our example, we once again use the data with $T = 15$, $n = 20$, $\sum_{i=1}^{n} x_i = 10$, $u = 9$, $w = 1$. Using these settings, we know that the optimal order size is $Q_B^* = 41$, with an expected profit of $B_B(Q_B^*|x) = 253.38$. Based on the algorithm described in

Figure 2 of [17] the PS algorithm consists in simulating m observations $w_1,\ldots,w_m$ of the demand. Each observation w_i is obtained by first simulating the value of the parameter via the posterior distribution $p(\theta|x)$, and then simulating w_i using the forecast model (given the parameter value), which in our case corresponds to model (10).

Table II: Results after applying the PS algorithm for m = 100 and m = 1000.

m	Optimal Order Size		Estimation of the Expected Profit		
			Point	Lower Bound	Upper Bound
100	d_k	42	256.00	240.65	271.35
	d_{k+1}	44	256.20	240.14	272.26
1000	d_k	40	252.37	247.89	256.85
	d_{k+1}	41	252.49	247.90	257.08

A In the case where the demand allows for a density function, the optimal order size is obtained by setting the service level to $\alpha = u/(u+w)$ and applying a valid method for quantile estimation. Nonetheless, for the discrete case, it is convenient to apply the method described by equations (8) and (9), replacing the cdf $F_B(y|x)$ for the empirical distribution of the observations $w_1,\ldots,w_m$.

From Table II, notice that for m = 1000 observations, the PS algorithm provides an optimal order size of 41, and estimates an expected profit between 247.9 and 257.08, which covers the actual value (253.38). For m = 100, the number of observations is insufficient for obtaining an optimal order size (surprisingly, we saw no observation with a value of 43). For values of m > 1000, the PS algorithm should still provide an optimal order size of 41, with a better estimate of the expected profit.

V. Conclusion

The results obtained by experimenting with the proposed approach show that the classical approach tends to overestimate the expected profit when compared to the Bayesian approach. On the other hand, as the number of real data observations increases, the results with both methods tend to coincide.

Based on the obtained results, we recommend applying the proposed Bayesian approach when the number of observations is small since in this case, the uncertainty in the parameters is significant. On the other hand, if we use stochastic simulation in order to estimate the optimal order size, we have to consider a large enough number of simulated observations in order to obtain an adequate precision.

References

[1] S. Nahmias, S. and T. L. Olsen, *Production and Operations Analysis*, Long Grove:Waveland Press, 2015.

[2] H. Scarf, "Bayes solutions of the statistical inventory problem," *The Annals of Mathematical Statistics*, vol. 30, no. 2, pp. 490-508, 1959.

[3] E. A. Silver, "Bayesian determination of the reorder point of a slow moving item," *Operations Research*, vol. 13, no. 6, pp. 989-997, 1965.

[4] K. S. Azouri, "Bayes solution to dynamic inventory models under unknown demand distribution," *Management Science*, vol. 31, no. 9, pp. 1150-1160, 1985.

[5] G. D. Eppen, and A. V. Iyer, "Improved fashion buying with Bayesian updates," *Operations Research*, vol. 45, no. 6, pp. 805-819, 1997.

[6] M. A. Lariviere, and E. L. Porteus, "Stalking information: Bayesian inventory management with unobserved lost sales," *Management Science*, vol. 45, no. 3, pp. 343-363, 1999.

[7] L. Chen, and E. L. Plambeck, "Dynamic inventory management with learning about the demand distribution and substitution probability," *Manufacturing Service and Operations Management*, vol. 10, no. 2, pp. 236-256, 2008.

[8] L. Chen, "Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales," *Operations Research*, vol. 58, no. 2, pp. 396-413, 2000.

[9] A. Jain, N. Rudi, and T. Wang, "Demand estimation and ordering under censoring: stocking-out timing is (almost) all you need," *Operations Research*, vol. 61, no. 1, pp. 134-150, 2015.

[10] R. Levi, R. Roundy, and D. B. Shmoys, "Provably near-optimal sampling based policies for stochastic inventory control models, *Mathematics of Operations Research*, vol. 32, no. 4, pp. 821–839, 2007.

[11] R. Levi, G. Perakis, and J. Uichanco, "The data-driven newsvendor problem: new bounds and insights," *Operations Research*, vol. 63, no. 6, pp. 1294-1306, 2015.

[12] V. Agrawal, and S. Seshadri, "Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem," *Manufacturing & Service Operations Management*, vol. 2, no. 4, pp. 410-423, 2000.

[13] B. Keren, and J. S. Pliskin, "A benchmark solution for the risk-averse newsvendor problem," *European Journal of Operational Research*, vol. 174, no. 3, pp. 1643-1650, 2006.

[14] D. F. Muñoz, and D. F. Muñoz, "Bayesian forecasting of spare parts using simulation," in *Service Parts Management: Demand Forecasting and Inventory Control*, edited by N. Altay y L. A. Litteral, pp. 105-123. New York:Springer, 2011.

[15] J. M. Bernardo, and A. F. M. Smith, *Bayesian Theory*, Chichester:Wiley, 2009.

[16] J. O. Berger, and J. M. Bernardo, "Ordered group reference priors with application to the multinomial problem," *Biometrika*, vol. 79, no. 1, pp. 25-37, 1992.

[17] D. F. Muñoz, D. G. Muñoz, and A. Ramírez-López, "On the incorporation of parameter uncertainty for inventory management using simulation, *International Transactions in Operational Research* vol. 20, no. 4, pp. 493-513, 2013.

[18] M. Rothschild, and J. E. Stiglitz, "Increasing risk: I. A definition," *Journal of Economic Theory*, vol. 2, no. 3, pp. 225-243, 1970.