Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

DEMAND FULFILLMENT PROBABILITY UNDER PARAMETER UNCERTAINTY

Canan G. Corlu

Metropolitan College
Boston University
Boston, MA 02215, USA

Bahar Biller

Software Sciences and Analytics
GE Global Research Center
Niskayuna, NY 12309, USA

Sridhar Tayur

Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA 15213, USA

ABSTRACT

We study a multi-item inventory system with normally distributed demands in the presence of demand parameter uncertainty – the uncertainty that stems from the estimation of the unknown demand parameters necessitated by limited amounts of historical demand data. Using an asymptotic normality approximation, we quantify the variance of the demand fulfillment probability (i.e., the probability that all item demands will be satisfied from stock) that is due to demand parameter uncertainty. We use this quantification to understand the impact of demand parameter uncertainty on the demand fulfillment probability and investigate the sensitivity of the variance of the demand fulfillment probability to selected inventory model parameters.

1 INTRODUCTION

Most of the literature on multi-item inventory control assumes that the demand distributions and the values of their parameters are known with certainty. However, this may not be the case in practice and the unknown demand parameters may have to be estimated from finite historical data. Under the assumption of a stationary demand process, the estimates of the demand parameters often converge to their true counterparts as the length of the historical data stream approaches infinity. However, the inventory manager is rarely fortunate enough to be in an asymptotic situation. Since there is often a short demand history, the asymptotic properties of the demand parameter estimates do not hold and the uncertainty around the demand parameter estimates (i.e., the demand parameter uncertainty) is not factored into the decisions of the inventory manager. This practice leads to poor inventory decisions under high demand parameter uncertainty. Our goal in this paper is to demonstrate the importance of accounting for this uncertainty in the estimation of the demand fulfillment probability (i.e., the probability of satisfying all item demands from stock) in a multi-item inventory setting in the presence of limited historical data.

What is important to recognize in this setting is that the demand fulfillment probability is a random variable because it is a function of the unknown demand parameters estimated from limited demand data. Therefore, we characterize the variance of the demand fulfillment probability due to the use of limited data for demand parameter estimation. To do this, we first capture the demand parameter uncertainty with a Bayesian model. Then, we combine the Bayesian model with a series of asymptotic normal approximations to quantify the variance of the demand fulfillment probability.

It is important to note the simplifying assumptions made in this paper to obtain a closed-form expression for the variance of the demand fulfillment probability due to demand parameter uncertainty. In particular, the product demands are assumed to be independent and normally distributed. This is the key assumption that leads to the derivation of the asymptotic normal approximations for the mean and the variance of the demand fulfillment probability. The motivation behind using the method of asymptotic normality approximation to capture the demand parameter uncertainty is the difficulty in evaluating the exact analytical expressions for the mean and the variance of the demand fulfillment probability in a multi-item inventory system. The difficulty arises from the use of multiple integrals in these expressions. In this paper, however, we avoid the need to evaluate multiple integrals and the computational issues associated with their evaluation in high dimensions. This aspect of our work will be an advantage when the focus switches to solving an multi-item inventory optimization problem under demand parameter uncertainty; we refer the reader to Section 6 for a brief discussion on future work. Furthermore, the numerical study presented in Biller et al. (2016) demonstrates how well our characterization of the demand parameter uncertainty based on the asymptotic normality approximation represents the variance of the demand fulfillment probability computed with respect to the joint posterior density function of the unknown demand parameters. Nevertheless, the relaxation of the assumption of independent and normally distributed demands for non-stationary, multivariate and autocorrelated demand processes will require to resort to numerically intensive methods to capture multivariate demand parameter uncertainty. We refer the reader to Biller and Corlu (2011) for an example method to represent multivariate demand parameter uncertainty in simulations wit stationary input processes.

The severe consequences of ignoring parameter uncertainty in stochastic systems has been discussed by a considerable number of researchers including Cheng and Holland (1997, 1998, 2004), Chick (2001), Barton and Schruben (1993, 2001), Zouaoui and Wilson (2003, 2004), Ng and Chick (2006), Barton (2007), Batarseh and Wang (2008), Barton et al. (2010, 2014), Biller and Corlu (2011), Muñoz et al. (2013), Corlu et al. (2015), and Biller et al. (2016). We refer the reader to Barton (2012) for an excellent review of the literature on the severe consequences of ignoring parameter uncertainty in stochastic system design. The distinguishing feature of our work is the characterization of the variance of the demand fulfillment probability due to the use of limited data for demand parameter estimation in a multi-item inventory system. Assuming a normal demand for each item, we show that with 99.8% probability the demand fulfillment probability deviates from its expected value by up to 40% in the presence of a limited demand history. We study the sensitivity of this quantification to several inventory model parameters including the number of items in the system, the historical data length, and the prior parameters of the Bayesian model.

Under the assumption of independent demands, it can be easily deduced from its functional representation that the demand fulfillment probability is a decreasing function of the number of items in the system (see Section 2). In this case, a natural question to ask is whether the effect of the demand parameter uncertainty disappears in large-scale inventory systems. We find that the answer to this question is no, and the demand fulfillment probability still deviates from its mean by up to 18% (with 99.8% probability) when there are 300 items in the system. We also find that our solutions are robust to the selection of prior values for the demand parameters, which are the key components of the Bayesian demand model.

We organize the remainder of the paper as follows. Section 2 presents our inventory model and a functional form for the demand fulfillment probability. Section 3 provides the Bayesian model to account for the demand parameter uncertainty and quantifies the variance around the demand fulfillment probability due to demand parameter uncertainty. Section 4 provides a simulation algorithm that estimates the mean and the variance of the demand fulfillment probability. Section 5 investigates the sensitivity of the variance of the demand fulfillment probability with respect to several inventory model parameters. Section 6 concludes with a summary of findings.

2 THE INVENTORY MODEL AND THE DEMAND FULFILLMENT PROBABILITY

We consider a P-item inventory setting where each item operates under a periodic-review independent order-up-to policy. We assume deterministic lead times L_i, $i = 1,2,\ldots,P$, and response time window k satisfying $L_i \geq k$ for $i = 1,2,\ldots,P$. The event sequence within a period consists of the inventory review, the ordering decision, the receipt of the replenishment, and the arrival of the item demands. Unsatisfied demands in any period are completely backlogged. The system does not allocate any inventory to meet the demands of a particular period unless all backlogs for the items due to earlier demands are satisfied.

We measure the level of service provided to the customers of our inventory system by the demand fulfillment probability, i.e., the probability of filling all item demands in one period within a time window of k periods. Thus, the focus is on covering the item demands with the current inventories and the replenishments to arrive in the next k periods. To provide a characterization of the demand fulfillment probability, we first define $D_{i,t}$ for the item-i demand. We then assume the item demand random variables $D_{i,t}$, $i = 1,2,\ldots,P$, to be independent both in the same time period t and across time. In addition, we use $A_{i,t}$ for the item-i replenishment arriving in period t and $X_{i,t}$ for the net inventory level at the end of period t. The probability that the period-t demands $D_{i,t}$, $i = 1,2,\ldots,P$, are filled within k periods of arrival is then given by

$$\Pr\left\{X_{i,t} + \sum_{\ell=t+1}^{t+k} A_{i,\ell} \geq 0, \;\; i = 1,2,\ldots,P\right\}.$$

Using the relationship $X_{i,t+k} = X_{i,t} + \sum_{\ell=t+1}^{t+k} A_{i,\ell} - \sum_{\ell=t+1}^{t+k} D_{i,\ell}$, we can write this probability as

$$\Pr\left\{X_{i,t+k} + \sum_{\ell=t+1}^{t+k} D_{i,\ell} \geq 0, \;\; i = 1,2,\ldots,P\right\}.$$

We further define I_i for the item-i inventory target, use the relation $X_{i,t} = I_i - \sum_{\ell=t-L_i}^{t} D_{i,\ell}$ (Hadley and Whitin 1963) to represent $X_{i,t+k} + \sum_{\ell=t+1}^{t+k} D_{i,\ell}$ as $I_i - \sum_{\ell=t-L_i+k}^{t} D_{i,\ell}$, and obtain the following characterization for the demand fulfillment probability (Hausman et al. 1998, Lemma 1):

$$\Pr\left\{\sum_{\ell=t-L_i+k}^{t} D_{i,\ell} \leq I_i, \;\; i = 1,2,\ldots,P\right\}.$$

Under the assumption of independent and normally distributed demands with means μ_i, $i = 1,2,\ldots,P$, and standard deviations σ_i, $i = 1,2,\ldots,P$, that are known with certainty, the demand fulfillment probability takes the following form, where Φ denotes the standard normal cumulative distribution function:

$$\prod_{i=1}^{P} \Phi\left(\frac{I_i - (L_i - k + 1)\mu_i}{\sqrt{L_i - k + 1}\,\sigma_i}\right), \tag{1}$$

Detailed presentation of this inventory model and measure of service can be found in Biller et al. (2016).

However, it is rarely the case in practice that the values of the demand parameters are known with certainty. We consider the situation in which the parameters μ_i, $i = 1,2,\ldots,P$, and σ_i, $i = 1,2,\ldots,P$, in (1) are unknown and they are estimated from historical demand data of finite length, say $d_{i,t}$, $t = 1,2,\ldots,n_i$, $i = 1,2,\ldots,P$, using the method of maximum likelihood estimation. Since the maximum likelihood estimates converge to their true counterparts only as the length of the historical demand data approaches infinity, their use for multi-item inventory control in the presence of limited data, as if they were the true parameter values, ignores the demand parameter uncertainty. Our goal in this paper is to characterize the amount of demand parameter uncertainty in the demand fulfillment probability due to the use of a limited amount of data for demand parameter estimation.

3 BAYESIAN MODEL

The key to achieving our goal is to build a Bayesian model which recognizes the randomness in the unknown demand parameters and propagates this randomness throughout the model to eventually capture it in the demand fulfillment probability. This section describes how to build such a Bayesian model for representing the demand parameter uncertainty in the demand fulfillment probability.

First, we provide the likelihood function of the historical demand data $d_{i,t}, t = 1,2,\ldots,n_i, i = 1,2,\ldots,P$, for normally distributed demands with unknown means μ_i, $i = 1,2,\ldots,P$, and unknown variances σ_i^2, $i = 1,2,\ldots,P$. We use θ_i for the two-dimensional vector of the demand parameters μ_i and σ_i^2; θ for the vector of θ_i, $i = 1,2,\ldots,P$, consisting of the P-item normal demand parameters; and $\mathbf{d}$ for the vector of the historical demand data $d_{i,t}$, $t = 1,2,\ldots,n_i$, $i = 1,2,\ldots,P$:

$$f(\mathbf{d}|\theta) \quad = \quad \prod_{i=1}^{P} \left(2\pi\sigma_i^2\right)^{-n_i/2} \exp\left\{-\frac{1}{2\sigma_i^2}\sum_{t=1}^{n_i}\left(d_{i,t}-\mu_i\right)^2\right\}.$$

Next, we treat the unknown μ_i and σ_i^2 as random variables. We use the normal prior density function with mean $\mu_{i,0}$ and variance $\sigma_i^2/\kappa_{i,0}$ to represent the mean μ_i conditional on the variance σ_i^2; i.e.,

$$\mu_i \mid \sigma_i^2 \propto \left(\sigma_i^2\right)^{-1/2} \exp\left\{-\frac{\kappa_{i,0}}{2\sigma_i^2}\left(\mu_i-\mu_{i,0}\right)^2\right\}.$$

In addition, we choose an inverse gamma prior density function with shape parameter $\nu_{i,0}/2$ and scale parameter $\zeta_{i,0}^2/2$ for modeling the variance σ_i^2; i.e.,

$$\sigma_i^2 \propto \left(\sigma_i^2\right)^{-(\nu_{i,0}+2)/2} \exp\left\{-\frac{\zeta_{i,0}^2}{2\sigma_i^2}\right\}.$$

The hyperparameters $\mu_{i,0}$, $\kappa_{i,0}$, $\nu_{i,0}$, and $\zeta_{i,0}^2$ are, respectively, the mean, shrinkage, degrees of freedom, and scale parameters of the following joint prior density function, which we denote by $\pi_i(\mu_i,\sigma_i^2)$ for μ_i and σ_i^2 (Fraley and Raftery 2007):

$$\pi_i\left(\mu_i,\sigma_i^2\right) \quad \propto \quad \left(\sigma_i^2\right)^{-(\nu_{i,0}+3)/2} \exp\left\{-\frac{1}{2\sigma_i^2}\left(\zeta_{i,0}^2+\kappa_{i,0}\left(\mu_i-\mu_{i,0}\right)^2\right)\right\}.$$

The prior density function $\pi(\theta)$ for the unknown demand parameter vector θ is then given by

$$\pi(\theta) = \prod_{i=1}^{P}\pi_i(\mu_i,\sigma_i^2) \propto \prod_{i=1}^{P}\left(\sigma_i^2\right)^{-(\nu_{i,0}+3)/2} \exp\left\{-\frac{1}{2\sigma_i^2}\left(\zeta_{i,0}^2+\kappa_{i,0}\left(\mu_i-\mu_{i,0}\right)^2\right)\right\}.$$

Using Bayes' rule for combining the likelihood function $f(\mathbf{d}|\theta)$ with the joint prior density function $\pi(\theta)$ leads to a joint posterior density function $h(\theta|\mathbf{d})$:

$$h(\theta|\mathbf{d}) \quad \propto \quad \prod_{i=1}^{P}\left(\sigma_i^2\right)^{-(\nu_{i,0}+n_i+3)/2} \exp\left\{-\frac{1}{2\sigma_i^2}\left(\zeta_{i,0}^2+\kappa_{i,0}\left(\mu_i-\mu_{i,0}\right)^2+\sum_{t=1}^{n_i}\left(d_{i,t}-\mu_i\right)^2\right)\right\}.$$

This density function can be alternatively written as

$$\prod_{i=1}^{P}\left(\sigma_i^2\right)^{-(\nu_{i,n_i}+3)/2} \exp\left\{-\frac{1}{2\sigma_i^2}\left(\zeta_{i,n_i}^2+\kappa_{i,n_i}\left(\mu_i-\mu_{i,n_i}\right)^2\right)\right\}$$

with parameters $\mu_{i,n_i} = (\kappa_{i,0}\mu_{i,0}+\sum_{t=1}^{n_i} d_{i,t})/(\kappa_{i,0}+n_i)$, $\kappa_{i,n_i} = \kappa_{i,0}+n_i$, $\nu_{i,n_i} = \nu_{i,0}+n_i$, and

$$\zeta_{i,n_i}^2 = \zeta_{i,0}^2 + \frac{\kappa_{i,0}n_i}{\kappa_{i,0}+n_i}\left(\sum_{t=1}^{n_i}\frac{d_{i,t}}{n_i}-\mu_{i,0}\right)^2 + \sum_{t=1}^{n_i}\left(d_{i,t}-\sum_{t=1}^{n_i}\frac{d_{i,t}}{n_i}\right)^2.$$

Thus, the joint posterior density function of the demand parameters μ_i and σ_i^2 corresponds to a normal-inverse gamma density function. Furthermore, σ_i^2 has the inverse gamma posterior density function with the shape parameter $\nu_{i,n_i}/2$ and the scale parameter $\zeta_{i,n_i}^2/2$, while μ_i, conditional on the variance σ_i^2, has the normal posterior density function with mean μ_{i,n_i} and variance $\sigma_i^2/\kappa_{i,n_i}$.

Next, we use the joint posterior density function to identify the maximum a posteriori (MAP) estimates of the unknown demand parameters. The MAP estimate $\tilde{\mu}_i$ of the parameter μ_i is μ_{i,n_i}, while $\zeta_{i,n_i}^2/(\nu_{i,n_i}+3)$ is the MAP estimate $\tilde{\sigma}_i^2$ of the parameter σ_i^2 (Biller et al. 2016).

What is important to recognize is that as the length of the historical demand data set increases, the joint posterior density function $h(\theta|\mathbf{d})$ of the normal demand parameter vector $\theta = (\theta_1, \theta_2, \ldots, \theta_P)'$ converges to a $2P$-dimensional normal distribution with a mean vector of $\tilde{\theta} = (\tilde{\theta}_1, \tilde{\theta}_2, \ldots, \tilde{\theta}_P)'$ and a variance-covariance matrix of $\mathbf{G_n}$ with $\mathbf{n} = (n_1, n_2, \ldots, n_P)'$, where

$$\mathbf{G_n} = \begin{pmatrix} \mathbf{G}_{1,n_1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_{2,n_2} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{G}_{P,n_P} \end{pmatrix}$$

with $\mathbf{0}$ as the 2×2 matrix of zeros and $\mathbf{G}_{i,n_i}$ as the 2×2 matrix $[\zeta_{i,n_i}^2/(\kappa_{i,n_i}(\nu_{i,n_i}+3)), 0; 0, 2\zeta_{i,n_i}^4/(\nu_{i,n_i}+3)^3]$ for $i = 1, 2, \ldots, P$ (Biller et al. 2016).

Next, we denote the product in (1) by $y(\mathbf{I}, \theta)$ as a function of the inventory-target vector $\mathbf{I}$ consisting of I_i, $i = 1, 2, \ldots, P$, and the demand parameter vector θ composed of μ_i, $i = 1, 2, \ldots, P$, and σ_i^2, $i = 1, 2, \ldots, P$. Using ϕ for the standard normal probability density function, we quantify the amount of demand parameter uncertainty in $y(\mathbf{I}, \theta)$ due to the use of limited data for normal demand parameter estimation.

Approximation. The demand fulfillment probability $y(\mathbf{I}, \theta)$ is asymptotically normally distributed with a mean of

$$\mathbb{E}(y(\mathbf{I}, \theta)) = \prod_{i=1}^{P} \Phi\left(\frac{(I_i - (L_i - k + 1)\mu_{i,n_i})(\nu_{i,n_i}+3)^{1/2}}{\zeta_{i,n_i}\sqrt{L_i - k + 1}}\right)$$

and a variance of $\mathbb{V}(y(\mathbf{I}, \theta)) = \sum_{i=1}^{P} \mathbb{V}_i(y(\mathbf{I}, \theta))$ with $\mathbb{V}_i(y(\mathbf{I}, \theta))$ as

$$\begin{aligned} = & \left(\frac{L_i - k + 1}{\kappa_{i,n_i}} + \frac{(I_i - (L_i - k + 1)\mu_{i,n_i})^2}{2(L_i - k + 1)\zeta_{i,n_i}^2}\right) \\ \times & \prod_{j=1, j\neq i}^{P} \Phi^2\left(\frac{(I_j - (L_j - k + 1)\mu_{j,n_j})(\nu_{j,n_j}+3)^{1/2}}{\zeta_{j,n_j}\sqrt{L_j - k + 1}}\right) \phi^2\left(\frac{(I_i - (L_i - k + 1)\mu_{i,n_i})(\nu_{i,n_i}+3)^{1/2}}{\zeta_{i,n_i}\sqrt{L_i - k + 1}}\right). \end{aligned}$$

We refer the reader to Biller et al. (2016) for a proof of this approximation.

4 ESTIMATING MEAN AND VARIANCE OF THE DEMAND FULFILLMENT PROBABILITY

We provide a simulation algorithm in Figure 1 that evaluates the asymptotic approximations to the mean and the variance of the demand fulfillment probability. In each macro-replication of the algorithm, we first generate a demand data of length n_i for each product with a given mean and standard deviation. Given the prior density function parameters, we then estimate the posterior density function parameters using the functional forms in Section 3. This follows the estimation of the mean and the variance of the demand fulfillment probability using the asymptotic approximations provided in Section 3. At the end of the macro-replications, we compute the average expected demand fulfillment probability $\bar{\mathbb{E}}$ and the average variance of the demand fulfillment probability $\bar{\mathbb{V}}$. These are the values that will be used in the tables of the next section.

for $r = 1, 2, \ldots, R$ replications **do**

For given inventory targets I_i, $i = 1, 2, \ldots, P$, lead time L_i, $i = 1, 2, \ldots, P$, time window k, and prior density parameters μ_{i0}, ν_{i0}, κ_{i0} and ζ_{i0} for $i = 1, 2, \ldots, P$,:

generate demand data of length n_i for $i = 1, 2, \ldots, P$ with the true parameters;

calculate the posterior density parameters μ_{i,n_i}, ν_{i,n_i}, κ_{i,n_i} and ζ_{i,n_i} for $i = 1, 2, \ldots, P$;

calculate the expectation $\mathbb{E}_r(y(\mathbf{I}, \theta))$ and the variance $\mathbb{V}_r(y(\mathbf{I}, \theta))$;

end for

compute $\bar{\mathbb{E}} := \sum_{r=1}^{R} \mathbb{E}_r(y(\mathbf{I}, \theta))/R$ as an estimate of the mean of the demand fulfillment probability;

compute $\bar{\mathbb{V}} := \sum_{r=1}^{R} \mathbb{V}_r(y(\mathbf{I}, \theta))/R$ as an estimate of the variance of the demand fulfillment probability.

Figure 1: A simulation algorithm for evaluation of the asymptotic approximations to the mean and the variance of the demand fulfillment probability.

5 RESULTS AND INSIGHTS

The objective of this section is to gain insights into the amount of parameter uncertainty in the demand fulfillment probability in the presence of limited data and to study the sensitivity of the quantified demand parameter uncertainty in the demand fulfillment probability to the inventory model parameters. We design our experiments by assuming $n_i \in \{10, 20, 30, 50\}$ and $L_i = k$ for $i = 1, 2, \ldots, P$. We also assume a mean demand of $100 + 20(i-1)$ and a standard deviation of $(0.25 - (0.24/P)(i-1))\mu_i$ for the normally distributed item-i demand. We set the values of the inventory targets, I_i, $i = 1, 2, \ldots, P$ in such a way that they provide a demand fulfillment probability of 80% under the assumption of known demand parameters. We further select the prior density function parameters in such a way that the prior demand parameters deviate by 10% from their (true) counterparts, which are used to generate the limited historical demand data sets of our experiments.

Allowing P, the number of items in the system, to change between 3 and 300, we investigate the variance of the demand fulfillment probability due to demand parameter uncertainty as a function of the number of items in Section 4.1, the length of the historical demand data in Section 4.2, and the prior parameter selection in Section 4.3. The results provided in each of these sections are obtained from a sufficiently large number of macro-replications of the algorithm in Figure 1.

5.1 The Number of Items

Table 1 summarizes our findings on the amount of parameter uncertainty in the fulfillment probability in the form of a percentage of the expected demand fulfillment probability corresponding to the standard deviation of the demand fulfillment probability for normally distributed demands. Therefore, each result reported in Table 1 is obtained from $\sqrt{\bar{\mathbb{V}}}/\bar{\mathbb{E}} \times 100\%$.

When there are 10 historical demand observations available for each item in the three-item inventory setting, we observe the standard deviation of the demand fulfillment probability, due to demand parameter uncertainty, to be 10.07% of its mean (Table 1). It is well known that 99.8% of all the values a normally distributed random variable can take fall within four standard deviations of the mean. Therefore, the following interpretation of the result "10.07%" provides a better illustration of the variability in the demand

Table 1: The percentage of the expected demand fulfillment probability corresponding to the standard deviation of the demand fulfillment probability.

	Length of the Demand History			
P	10	20	30	50
3	**10.07**	7.82	6.66	5.51
5	**9.07**	7.12	6.11	5.00
10	7.89	6.23	5.37	4.45
20	6.84	5.41	4.72	3.83
30	6.38	4.95	4.36	3.51
50	5.78	4.47	3.87	3.13
100	5.13	3.94	3.36	2.67
200	**4.55**	3.52	2.94	2.29
300	**4.42**	3.29	2.72	**2.09**

fulfillment probability due to parameter uncertainty: The demand fulfillment probability deviates from its expected value by up to 40.28% (i.e., 10.07% $\times$ 4) with a probability of 99.8%. For five items (i.e., $n_i = 10$ for $i = 1,2,\ldots,P$ and $P = 5$ in Table 1), the effect of the demand parameter uncertainty on the demand fulfillment probability decreases to 36.28% (i.e., 9.07% $\times$ 4), but still remains significant. This quantification is halved by 200 items in the system; i.e., the demand fulfillment probability deviates from its mean by up to 18.20% (i.e., 4.55% $\times$ 4) with a probability of 99.8%.

It is plausible to expect that the effect of demand parameter uncertainty will increase with the number of items in the system. Contrary to this expectation, the increasing values of P has a decreasing effect on the amount of variability in the demand fulfillment probability. However, when we increase the number of items in the system to 300, we still identify the standard deviation of the demand fulfillment probability to be 4.42% of the expected demand fulfillment probability; i.e., the demand fulfillment probability deviates from its mean by up to 17.68% (i.e., 4.42% $\times$ 4) with a probability of 99.8%. When the length of the demand history increases to 50 observations in the 300-item inventory setting, this particular quantification of the demand parameter uncertainty falls below 10%. In those cases with larger numbers of items, we also find the variance of the demand fulfillment probability due to demand parameter uncertainty to be rather insensitive to the further increase in the number of items in the system. For example, given that there are already 200 items in the inventory system with a length of 10 historical data points for each item, the consideration of an additional 100 items results in a reduction of only 0.13% (i.e., 4.55% $-$ 4.42%) in the effect of the parameter uncertainty on the fulfillment probability.

5.2 The Length of the Demand History

As expected, the variance of the demand fulfillment probability due to demand parameter uncertainty diminishes with the length of the demand history. However, in response to the increasing number of demand observations, the rate of reduction in the values reported in Table 1 is rather slow. For example, in the 10-item setting with historical data of length 10 (Table 1), we achieve a reduction of 21% in our quantification of the demand parameter uncertainty in the demand fulfillment probability with an additional 10 demand observations, 32% reduction with an additional 20 demand observations, and a reduction of approximately 44% with an additional 40 demand observations. Nevertheless, the need for a large historical data set to reduce the variance of the demand fulfillment probability is offset, to an extent, by the increased number of items in the system. Therefore, the answer to the question "what is a large data set?" depends on the value of P. Table 1 shows an up to 18% (i.e., 4.45% $\times$ 4) deviation of the demand fulfillment probability from its mean (with 99.8% probability) for 50 observations in the 10-item inventory setting. We make a similar observation for 30 observations in the 30-item inventory setting and for 20 observations in the 50-item inventory setting.

5.3 The Prior Demand Parameter Selection

Table 2 presents the percentage of the expected demand fulfillment probability corresponding to the standard deviation of the demand fulfillment probability as a function of the deviation of the demand parameters of the prior density function from the true demand parameters; i.e., the values of the demand parameters that are used to generate the historical data sets of our computational study. More precisely, the percentage deviation of priors taking the values of 2%, 5%, 10%, and 20% in Table 2 indicates the discrepancy between the opinions of the experts on the unknown demand parameters and the true distributional characteristics of the demand parameters, which are unknown to the experts. We observe the quantification of the demand parameter uncertainty in the demand fulfillment probability to be insensitive to the selection of the prior parameters for the Bayesian demand model. In comparison to the 2% deviation of the prior demand parameters from the true demand parameters, the 20% deviation has an effect of only 1.5% (i.e., $8.84\% - 7.34\%$) on the percentage standard deviation in the mean demand fulfillment probability when there are 10 observations in the historical data set. When the historical demand data length increases to 20, 30 and 50, this effect falls below 1%, decreasing to 0.76%, 0.46% and 0.09%, respectively.

Table 2: The percentage of the expected demand fulfillment probability corresponding to the standard deviation of the demand fulfillment probability as a function of the deviation of the prior demand parameters from the true parameters in the 10-item inventory setting.

	% Deviation of Priors			
data length	2%	5%	10%	20%
10	**7.34**	7.59	7.89	**8.84**
20	**5.87**	6.03	6.23	**6.63**
30	**5.15**	5.32	5.37	**5.61**
50	**4.42**	4.43	4.45	**4.51**

6 CONCLUSION

We consider a multi-item inventory system with the objective of estimating the demand fulfillment probability when the demand parameters are not known with certainty. There is only a limited amount of historical demand data to provide statistically valid estimates of the demand parameters. However, the parameter estimates converge to their true counterparts only in the presence of a sufficiently large amount of demand data, which is rarely the case in practice. Therefore, the use of the demand parameter estimates for inventory control ignores the demand parameter uncertainty.

Under certain assumptions on the underlying demand process, this paper characterizes the variance of the demand fulfillment probability due to demand parameter uncertainty, and presents insights into the sensitivity of this characterization to the model parameters including the historical data length, the number of items in the system, and the prior demand parameters. In particular, we identify the variance of the demand fulfillment probability to be a decreasing function of the number of items in the system. Nevertheless, we find that with 99.8% probability the demand fulfillment probability to still deviates from its mean by up to 18% when there are 300 items in the system. We also find that our results are robust to the choice of the prior demand parameters of the Bayesian demand model.

A natural setting where the results of this paper can be useful is the budget-constrained inventory systems, where the goal is to identify the optimal inventory targets that maximize the demand fulfillment probability subject to a constraint on the total inventory investment. Such inventory systems have been studied in the literature under the assumption that the demand parameters are known with certainty. Budget-constrained multi-item inventory systems in the presence of limited history of demand is the subject of ongoing work.

REFERENCES

Batarseh, O.G. and Y. Wang. 2008. “Reliable Simulation with Input Uncertainties Using an Interval-Based Approach”. In *Proceedings of the 2008 Winter Simulation Conference*, S. J. Mason, R. R. Hill, L. Monch, O. Rose, T. Jefferson, and J. W. Fowler, eds. 344 – 352. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R. 2007. “Presenting a More Complete Characterization of Uncertainty: Can It Be Done?” In *Proceedings of the INFORMS Simulation Society Research Workshop*. Fontainebleau, France.

Barton, R. R. 2012. “Tutorial: Input Uncertainty in Output Analysis”. In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 162 – 172. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R., B. L. Nelson, and W. Xie. 2010. “A Framework for Input Uncertainty Analysis”. In *Proceedings of the 2010 Winter Simulation Conference*. B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yucesan, eds. 189 – 1198. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R. and L. W. Schruben. 1993. “Uniform and Bootstrap Resampling of Input Distributions”. In *Proceedings of the 1993 Winter Simulation Conference*. G. W. Evans, M. Mollaghasemi, W. E. Biles, and E. C. Russell, eds. 503 – 508. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R. and L. W. Schruben. 2001. “Resampling Methods for Input Modeling”. In *Proceedings of the 2001 Winter Simulation Conference*. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds. 372 – 378. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Barton, R. R., B. L. Nelson, and W. Xie. 2014. “Quantifying Input Uncertainty via Simulation Confidence Intervals”. *Informs Journal on Computing* 26: 74 – 87.

Biller, B. and C. G. Corlu. 2011. “Accounting for Parameter Uncertainty in Large-Scale Stochastic Simulations with Correlated Inputs”. *Operations Research* 59: 661 – 673.

Biller, B., C. G. Corlu, and S. Tayur. 2016. “Demand Fulfillment Probability in a Multi-Item Inventory System with Limited Historical Data”. Working Paper, Metropolitan College, Boston University, Boston, MA.

Corlu, C. G., B. Biller, and S. Tayur. 2015. “Driving Inventory System Simulations with Limited Data: Insights from the Newsvendor Problem”. Working Paper, Metropolitan College, Boston University, Boston, MA.

Cheng, R. C. H. and W. Holland. 1997. “Sensitivity of Computer Simulation Experiments to Errors in Input Data”. *Journal of Statistical Computation and Simulation* 57: 327 – 335.

Cheng, R. C. H. and W. Holland. 1998. “Two-point Methods for Assessing Variability in Simulation Output”. *Journal of Statistical Computation and Simulation* 60: 183 – 205.

Cheng, R. C. H. and W. Holland. 2004. “Calculation of Confidence Intervals for Simulation Output”. *ACM Transactions on Modeling and Computer Simulation* 14: 344 – 362.

Chick, S. E. 2001. “Input Distribution Selection for Simulation Experiments: Accounting for Input Uncertainty”. *Operations Research* 49: 744 – 758.

Fraley, C. and A. E. Raftery. 2007. “Bayesian Regularization for Normal Mixture Estimation and Model-Based Clustering”. *Journal of Classification* 24: 155 – 181.

Hadley, G. and T. M. Whitin. 1963. *Analysis of Inventory Systems*. Prentice Hall, New Jersey.

Hausman, W. H., H. L. Lee, and A. X. Zhang. 1998. “Joint Demand Fulfillment Probability in a Multi-Item Inventory System with Independent Order-up-to Policies”. *European Journal of Operational Research* 109: 646 – 659.

Muñoz, D. F., D. G. Muñoz, and A. Ramirez-López. 2013. “On the Incorporation of Parameter Uncertainty for Inventory Management Using Simulation”. *International Transactions in Operational Research* 20: 493 – 513.

Ng, S. H. and S. E. Chick. 2006. "Reducing Parameter Uncertainty for Stochastic Systems". *ACM Transactions on Modeling and Computer Simulation* 16: 26 – 51.

Zouaoui, F. and J. R. Wilson. 2003. "Accounting for Parameter Uncertainty in Simulation Input Modeling". *IIE Transactions* 35: 781 – 792.

Zouaoui, F. and J. R. Wilson. 2004. "Accounting for Input-Model and Input-Parameter Uncertainties in Simulation". *IIE Transactions* 36: 1135 – 1151.

AUTHOR BIOGRAPHIES

CANAN GUNES CORLU is an Assistant Professor of Administrative Sciences Department of Metropolitan College at Boston University. She holds a Ph.D. in Operations Management from the Tepper School of Business at Carnegie Mellon University. Her research interests are in the areas of design and analysis of stochastic simulations with applications to operations management. Her email address is canan@bu.edu.

BAHAR BILLER is a Research Scientist in the Software Sciences and Analytics organization of the General Electric Global Research Center. Her research lies in the area of quantitative risk management and advances the design and analysis of large-scale business system simulations to aid decision-making under uncertainty. Her email address is bahar.biller@ge.com.

SRIDHAR TAYUR is the Ford Distinguished Research Chair and Professor of Operations Management at Carnegie Mellon University's Tepper School of Business. He has published many scholarly articles in journals such as *Operations Research*, *Management Science*, and *MSOM*, is co-editor of the widely referenced book, *Quantitative Models for Supply Chain Management*, and has served on the editorial boards of *Operations Research*, *Journal of Optimization and Engineering*, *NRLQ*, *MSOM* and *Management Science*. Dr. Tayur also served as the President of Manufacturing and Services Operations Management (MSOM) Society. He is the founder of the software company SmartOps Corporation (2000–) as well as the founder of OrganJet Corporation (2011–). His email address is stayur@andrew.cmu.edu.