

**A Simulation-Based Support Tool for Data-Driven Decision Making:
Operational Testing for Dependence Modeling**

Bahar Biller

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

Alp Akcay

Department of Industrial Engineering
Bilkent University
Ankara, 06800, TURKEY

Canan Corlu

Metropolitan College
Boston University
Boston, MA 2215, USA

Sridhar Tayur

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

ABSTRACT

Dependencies occur naturally between input processes of many manufacturing and service applications. When the dependence parameters are known with certainty, the failure to factor the dependencies into decisions is well known to waste significant resources in system management. Our focus is on the case of unknown dependence parameters that must be estimated from finite amounts of historical input data. In this case, the estimates of the unknown dependence parameters are random variables and simulations are designed to account for the dependence parameter uncertainty to better support the data-driven decision making. The premise of our paper is that there are certain cases in which the assumption of an independent input process to minimize the expected cost of input parameter uncertainty becomes preferable to accounting for the dependence parameter uncertainty in the simulation. Therefore, a fundamental question to answer before capturing the dependence parameter uncertainty in a stochastic system simulation is whether there is sufficient statistical evidence to represent the dependence, despite the uncertainty around its estimate, in the presence of limited data. We seek an answer for this question within a data-driven inventory-management context by considering an intermittent demand process with correlated demand size and number of inter-demand periods. We propose two new finite-sample hypothesis tests to serve as the decision support tools determining when to ignore the correlation and when to account for the correlation together with the uncertainty around its estimate. We show that a statistical test accounting for the expected cost of correlation parameter uncertainty tends to reject the independence assumption less frequently than a statistical test which only considers the sampling distribution of the correlation-parameter estimator. The use of these tests is illustrated with examples and insights are provided into operational testing for dependence modeling.

1 INTRODUCTION

An effective management of inventory is crucial for the success of any business system where the goal is to keep stock at a minimum while filling orders on time. In particular, simulation arises as an immediate choice of aid for industry practitioners making inventory decisions under uncertainty. However, when the demand distributions and their parameters are unknown and they must be estimated from finite amounts of

historical demand data, linking historical demand data with simulation design and analysis for inventory management turns into a challenge to be overcome in the search of a successful inventory control policy.

The common practice is to model stochastic demand using a probability distribution and then estimate the parameters of this distribution from the historical data to compute an inventory policy via simulation. However, this inventory policy would be optimal only if the demand parameter values were known with certainty. Therefore, the sequential process of first estimating the demand parameters and then using the demand parameter estimates for inventory-target optimization in a simulation casts doubt on the performance of the resulting inventory policy in the presence of historical demand data of finite length. This is clearly not an issue when there is a large amount of demand data and the demand parameter estimates converge to their true values. Today, however, the product life cycles are getting shorter, leading to the availability of even fewer observations in historical demand data sets. Consequently, the discrepancy between the performance of an optimal but unknown inventory target and the performance of its estimate from the simulation driven by a short demand history arises as a serious, but often ignored, problem in the design and analysis of stochastic inventory systems.

At first glance, our consideration of short demand history in a modern business environment might appear to contradict with the fact that we live in an era of big data. Yet there are still many situations where we unfortunately do not have sufficiently large amounts of data to make decisions with confidence. Brian Lewis, the chief data scientist and co-founder of Fractal Sciences, describes this situation as “big data dreams, small data reality” (Analytics Magazine, January/February 2014). It is this small data reality that our paper addresses when the demand history available for driving the simulation to support inventory decisions is short, and therefore, the inventory targets must be estimated from limited amounts of historical demand data. More importantly, we address this problem when the demand history contains dependent demand data. We present a simulation-based decision support tool to determine when the dependence in the demand process should be taken into consideration for inventory-target estimation in the presence of limited data. We expect this tool to form the building block of a data-driven framework for various streams of fragmented dependent demand data in a fast moving world with an explosive variety of products.

Representing demand uncertainty in a stochastic inventory simulation is often performed by treating item demands as a sequence of independent random variables. Unfortunately, the resulting demand model fails to capture the stochastic properties of demand processes exhibiting marked dependencies that occur naturally in many business systems. For example, the dependencies of demands for different component items may be induced by a bill of materials or product options. In an assemble-to-order system, the demands for component items that are stocked before final order assembly are usually dependent (Hausman et al. 1998). Even when the focus narrows down to the management of inventory for a single item, one of the data challenges to overcome in data-driven inventory management is the temporal dependence in a demand process; i.e., the dependence between consecutive demand realizations (Akcay et al. 2013a). Dependence might also arise between the demand size and the number of inter-demand periods – the number of zero-demand periods between two consecutive positive demands – in an intermittent demand process. In fact, a large percentage of items in industries ranging from aerospace to electronics have intermittent demand histories with positive realizations that only appear at random periods. It is a challenge to manage the inventory of an item with this type of demand process due to the dual sources of variation; i.e., both the demand size and the number of inter-demand periods are uncertain. Furthermore, the demand size and the number of inter-demand periods might be dependent. It is crucial to capture the demand dependencies through accurate multivariate demand modeling to avoid decision-making errors.

However, it is imperative to ask the following question prior to making any decisions to hedge against the dependent demand parameter uncertainty; i.e., the uncertainty around the parameters of the dependent demand distribution obtained from historical data of finite length: Is there sufficient statistical evidence to capture the correlation in the demand process despite the uncertainty around its estimate in the presence of limited demand data? Focusing on an intermittent demand process, seeking an answer to this fundamental question is the primary contribution of our paper. Specifically, we construct two new finite-sample hypothesis

tests to identify the situations in which it may be acceptable to ignore the correlation between the demand size and the number of inter-demand periods, and hence, the dependence parameter uncertainty in a short intermittent demand history, despite the existence of dependence in the demand process. These tests play the role of a new decision support tool for data-driven inventory management. If the test suggests sufficient evidence for the existence of demand dependence, then we incorporate it into the simulation to support the inventory-target estimation. Therefore, the key novelty of our work is to present solutions to guide the data-driven decision making, choosing the domain of application as inventory management where the intermittent demand process exhibits dependence between its size and the number of inter-demand periods.

It is important to note that there are powerful tools in statistics to test the existence of correlation in bivariate normal data (Fosdick and Raftery 2012). Traditionally, the distribution of a test statistic is based on asymptotic theory. In contrast, we generate test-statistic realizations from a limited amount of demand data to build an empirical distribution of the test statistic. We first introduce a statistical test that is solely based on the sampling distribution of the dependence parameter estimator. We then switch our focus to the fact that, as opposed to searching for the statistical significance of a dependence parameter estimate, a decision maker would be more interested in the performance of the decision based on this estimate. This motivates us to introduce an alternative (operational) test building on the difference between the expected cost of an inventory-target estimate and the minimum expected cost. Contrary to our correlation test, this operational test accounts for the costs of overestimation and underestimation of the optimal inventory target, and hence, the asymmetry in the cost function, in its investigation for the existence of demand correlation. We find that the difference between the inventory holding and backlogging costs — which is ignored by the correlation test — plays an important role when the expected cost of parameter uncertainty is the criterion used to decide whether or not to account for the demand correlation in simulation design and analysis.

We show that our operational test — which accounts for the expected cost of correlation parameter uncertainty — tends to reject the independence assumption less frequently than the correlation test — which only considers the sampling distribution of the dependence parameter estimator — especially when the length of the demand history and the strength of the demand dependence are not very high. This can be explained by the additional statistical estimation error around the dependence parameter when the independence assumption is relaxed. To put it another way, the benefit of accounting for the demand correlation is dominated by the additional expected cost associated with the incorrect estimation of the dependence parameter. When the amount of data is limited and the strength of demand dependence is low, the dependence parameter estimator has the highest variance, and therefore the estimated value of the demand dependence is subject to the highest statistical estimation error. This is why the operational test is in favor of the simpler model with no dependence when the demand history is short with low strength of dependence. Therefore, the correlation test is more appropriate if the goal is merely to make an inference about the demand process. The operational test is, on the other hand, more pragmatic to consider when the question is whether the dependence parameter estimate should be used as an input for decision making, as is the case of using simulation to support data-driven inventory management. We conclude with the note that our tests work for any measure of demand dependence. In addition to demand correlation, our tests also work for non-linear dependence measures, which arise as alternatives to the use of correlation for demand modeling when the dependence occurs between very low or very high demand realizations.

The remainder of the paper is organized as follows. Section 2 presents the demand and inventory models. Section 3 develops the hypothesis tests to determine whether to account for the intermittent demand correlation in inventory-target estimation and presents numerical insights into the effect of dependence parameter uncertainty on data-driven decision making. Section 4 concludes with future research directions.

2 BACKGROUND

Section 2.1 presents a copula-based intermittent demand model to capture correlation between a demand size and number of zero-demand periods preceding the demand. Section 2.2 characterizes the optimal policy assuming the complete knowledge of the underlying demand distribution.

2.1 A Copula-Based Intermittent Demand Process

We consider a discrete-time model with inventory review periods that are often shorter than the times between successive demand arrivals. We model the randomness in the demand arrivals by a Bernoulli process; i.e., the number of inter-demand periods, which we denote by $Y \in \{0, 1, 2, \dots\}$, has the geometric distribution with cumulative distribution function (cdf) $G(y; p) = 1 - (1 - p)^{y+1}$. Thus, the probability of observing a positive demand in any period is equal to p . We let the distribution of (positive) demand size, denoted by X , be a member of the location-scale family of distributions with location parameter τ and scale parameter θ . In fact, our copula-based intermittent demand model is flexible enough to combine an arbitrary distribution of the number of inter-demand periods with an arbitrary distribution of the demand size. The motivation behind the choice of the location-scale family of distributions in this paper is to describe the development of our finite-sample hypothesis tests in a way that is independent of the unknown parameters of the demand-size distribution. Hence, we represent the cdf of X , $F(x; \tau, \theta)$ as $F((x - \tau)/\theta; 0, 1)$ in the remainder of the paper, where $F(\cdot; 0, 1)$ is the standardized cdf that does not depend on τ and θ .

Willemain et al. (1994) show the existence of correlation in the historical intermittent shipment data of items ranging from electrical equipment to health products. Eaves (2002) investigates spare parts data from the Royal Air Force and finds statistically significant correlation between demand size and number of inter-demand periods in 18% of more than ten thousand items. More recently, Altay et al. (2012) analyze aircraft service parts data from the U.S. Defense Logistics Agency and identify significant correlation in 35% of the items. We have also encountered correlation in our own experience with a data set from a global luxury manufacturer in consumer-product space: A two-sided t test shows that 292 items (out of 1149) have statistically significant correlation between the demand size and the number of inter-demand periods (Akcay et al. 2013b).

As demonstrated by these real-life examples, the need to address the issue of an unknown correlation between demand size and number of inter-demand periods arises as a challenge in an inventory system simulation with limited intermittent demand history. More specifically, the challenge is to model the distribution of demand size, conditional on the number of periods since the last demand, while retaining the distribution of demand size X from the location-scale family along with the geometric distribution for the number of inter-demand periods Y . We address this issue by constructing the joint distribution of X and Y with a copula, which allows us to model the univariate distributions of X and Y and the dependence structure between them separately. Sklar's theorem (Nelsen 2006) elucidates the role played by a copula in the relation between a joint distribution and its univariate cdfs:

Theorem 1 Given the bivariate cdf $H(x, y)$ for the random variables X and Y with univariate cdfs $F(x; \tau, \theta)$ and $G(y; p)$, there exists a copula \mathcal{C} such that $H(x, y) = \mathcal{C}(F(x; \tau, \theta), G(y; p))$. Conversely, if \mathcal{C} is a copula and $F(x; \tau, \theta)$ and $G(y; p)$ are univariate cdfs, then $H(x, y)$ is a bivariate cdf for random variables X and Y .

More specifically, we construct the bivariate distribution of X and Y using the normal copula, which encodes the dependence precisely the same way a bivariate normal distribution does and, hence, measures the dependence between X and Y with correlation: $H(x, y) = \Phi_2\{\Phi^{-1}(F(x; \tau, \theta)), \Phi^{-1}(G(y; p)); r\}$ where $\Phi_2\{\cdot, \cdot; r\}$ is the bivariate standard normal cdf with correlation r (Nelsen 2006). However, the use of this model for intermittent demand modeling should only follow after the application of our operational test investigating whether there is sufficient statistical evidence to capture this dependence in the presence of limited intermittent demand data.

2.2 Inventory Model

The decision of how much inventory to keep, if any, is made at the beginning of each period and is contingent on the amount of inventory on hand. We consider linear procurement, holding, and backlogging costs c , h , and b per unit, respectively, and the time lag between procurement and delivery is negligible. Holding and backlogging costs are calculated based on the amount of ending inventory in each period. All the backlogged demand is satisfied before the next period starts, and the decision maker cannot dispose any

inventory during the multi-period planning horizon. At the end of the finite planning horizon, the decision maker obtains reimbursement of the procurement cost for each leftover unit and incurs the procurement cost for each backlogged unit. The goal is to find an ordering policy that minimizes the overall expected cost. Therefore, we can characterize the optimal policy when the parameters of our copula-based intermittent demand model (i.e., the demand-size distribution parameters τ and θ , the demand-occurrence probability p , and the normal copula parameter r) are known as follows, where y denotes the number of zero-demand periods since the last demand:

- (i) A base-stock policy is optimal in any time period for all values of y .
- (ii) For the demand-occurrence probability $p > h/(h+b)$, the optimal inventory target q^* is given by $\tau + F^{-1}(\kappa^*; 0, 1) \theta$, where κ^* is the unique value of κ that solves the equation

$$\frac{\Phi_2\{\Phi^{-1}(\kappa), \Phi^{-1}(G(y;p)); r\} - \Phi_2\{\Phi^{-1}(\kappa), \Phi^{-1}(G(y-1;p)); r\}}{(1-p)^y p} = 1 - \frac{h/(h+b)}{p}. \quad (1)$$

- (iii) For $p \leq h/(h+b)$, the optimal inventory target q^* is zero.

We refer the reader to Akcay et al. (2013b) for the derivation of this result. It is important to note that the solution κ^* does not depend on the demand-size parameters due to our copula-based representation of the intermittent demand process. We will make use of this important property in establishing the distribution of the test statistic for the operational test in Section 3.

As outlined above, the decision maker obtains the optimal inventory target q^* for $p > h/(h+b)$ by minimizing the single-period expected cost knowing that the myopic policy is optimal. To implement the policy, however, she obtains the point estimates of the unknown parameters p , τ , θ , and r from the historical demand data, and plugs these estimates into the functional form of q^* which is optimal only when the intermittent demand parameters are completely known. Hence, this well-known practice of inventory-target estimation ignores the uncertainty around the intermittent demand parameter estimates, which are inevitably obtained from a single realization of the demand history.

3 A SIMULATION-BASED DECISION SUPPORT TOOL: OPERATIONAL TESTING

We introduce the correlation test based on the sampling distribution of the correlation parameter estimator in Section 3.1 and the operational test based on the difference between the expected cost of an inventory-target estimate and the minimum expected cost in Section 3.2. We illustrate the use of these tests with examples and provide insights into operational testing for dependence modeling in Section 3.3.

3.1 The Correlation Test

The copula-based representation of intermittent demand allows the marginal demand distribution parameters τ , θ , and p , and the normal copula parameter r to be estimated separately, potentially with different methods. In this section, we first obtain the marginal-distribution parameter estimators $\hat{\tau}$, $\hat{\theta}$, and \hat{p} and transform the historical data $\{(x_t, y_t); t = 1, 2, \dots, n_d\}$ into $\{(u_x^i, u_y^i); i = 1, 2, \dots, n_d\}$ with $u_x^i := F((x_i - \hat{\tau})/\hat{\theta}; 0, 1)$ and $u_y^i := G(y; \hat{p})$. Then we use the maximum likelihood method to estimate the copula parameter r by fitting the normal copula to the transformed data:

$$\hat{r} := \operatorname{argmax}_{r \in (-1, 1)} \sum_{i=1}^{n_d} \log \phi_2\{\Phi^{-1}(u_x^i), \Phi^{-1}(u_y^i); r\}, \quad (2)$$

where $\phi_2\{\cdot, \cdot; r\}$ is the bivariate standard normal probability density function (pdf) with correlation r . The transformation of historical data into a unit hypercube by the parametric estimates of their univariate cdfs is known as the inference functions for margins; we refer the reader to Cherubini et al. (2011 for the details.

Clearly, the estimated value of \hat{r} never takes the exact value of zero. Thus, the immediate task of a decision maker – before accounting for any correlation between demand size and number of inter-demand periods – is to investigate whether the historical data carry sufficient information to reject the independence between the demand size and the number of inter-demand periods. There are powerful tools in statistics to test the existence of correlation in bivariate normal data; see Fosdick and Raftery (2012) for a review.

Traditionally, the distribution of a test statistic is based on asymptotic theory. In contrast, we generate realizations of the test statistic from the limited demand data to build an empirical distribution of the test statistic. The so-called bootstrap hypothesis testing has become increasingly attractive due to advances in computing (MacKinnon 2009). We let the null hypothesis \mathcal{H}_0 be $r = 0$ and the two-tailed alternative hypothesis \mathcal{H}_A be $r \neq 0$. The correlation test compares the observed value of the test statistic \hat{r} with the distribution that it would follow if the null hypothesis were true. The null hypothesis is then rejected if \hat{r} is sufficiently high relative to this distribution. The decision maker performs the test by specifying only the length of the demand history and the number of positive realizations.

Correlation Test Testing for correlation with $\mathcal{H}_0 : r = 0$ and $\mathcal{H}_A : r \neq 0$.

Initialization. Specify the number of bootstrap samples (B), the length of historical data (n), the number of positive realizations (n_d), and the significance level ($1 - \alpha$).

Set $\hat{p} \leftarrow n_d/n$ and $b \leftarrow 1$.

Compute the test statistic \hat{r} as in (2) from the historical data $\{(x_t, y_t); t = 1, 2, \dots, n_d\}$.

while $b \leq B$:

Step 1. Generate the bootstrap sample path:

Set $i \leftarrow 1$ and $t \leftarrow 1$.

while $t \leq n$

Generate $(\mathcal{Z}_{x,i}^b, \mathcal{Z}_{y,i}^b)$ from the bivariate standard normal cdf with correlation $r = 0$.

Obtain $z_i^b = F^{-1}(\Phi(\mathcal{Z}_{x,i}^b); 0, 1)$ and $y_i^b = G^{-1}(\Phi(\mathcal{Z}_{y,i}^b); \hat{p})$.

Set $t \leftarrow t + y_i^b + 1$ and $i \leftarrow i + 1$.

Let n_d^b be the number of nonzero demands in b th bootstrap sample path.

end

Step 2. Compute the test statistic from the bootstrap sample path:

Let \bar{z} and s_z be the sample mean and sample standard deviation of $\{z_1^b, z_2^b, \dots, z_{n_d}^b\}$.

Set $\hat{r}^b \leftarrow \operatorname{argmax}_{r \in (-1, 1)} \sum_{i=1}^{n_d^b} \log \phi_2\{\Phi^{-1}(F((z_i - \bar{z})/s_z; 0, 1)), \Phi^{-1}(G(y_i^b; n_d^b/n)); r\}$.

$b \leftarrow b + 1$

end

Step 3. Construct the critical region to reject the null hypothesis:

Sort \hat{r}^b , $b = 1, 2, \dots, B$ in ascending order $(r_{(1)}, r_{(2)}, \dots, r_{(B)})$.

Reject the null hypothesis \mathcal{H}_0 if $\hat{r} < r_{(\lfloor (\alpha/2)B \rfloor)}$ or $\hat{r} > r_{(\lceil (1-\alpha/2)B \rceil)}$.

We note that the copula-based representation of intermittent demand and the assumption of a location-scale family for the distribution of demand size allow the use of standardized demand samples $\{z_1^b, z_2^b, \dots, z_{n_d}^b\}$ instead of the actual demand samples $\{x_1^b, x_2^b, \dots, x_{n_d}^b\}$, which could be generated only if the true values of τ and θ were known.

3.2 The Operational Test

As opposed to searching for the statistical significance of a normal copula parameter estimate, a decision maker would be more interested in the performance of the decision based on this estimate. In this section, we propose an alternative test — the operational test — considering the sampling distribution of the difference between the expected costs of the optimal inventory target and its estimate. Hence, contrary to

the correlation test, the operational test accounts for the costs of overestimation and underestimation of the optimal inventory target in its investigation for the existence of correlation in the presence of limited historical intermittent demand data.

We start with reducing the expected cost function to an alternative form, which will be used to construct the sampling distribution of the test statistic. First, we let $C(q|y; p, \tau, \theta, r)$ denote the single-period expected cost conditional on the number of periods y since the last demand and represent the expected cost function $C(q|y; p, \tau, \theta, r)$ as

$$\theta C\left(\frac{q-\tau}{\theta}|y; p, 0, 1, r\right) + h\tau(1-p). \tag{3}$$

We refer the reader to Akcay et al. (2013b) for a detailed discussion of this characterization.

The decision maker implements the optimal policy with the threshold $\gamma = h/(h+b)$ as follows: If the estimate of the demand-occurrence probability $\hat{p} > \gamma$, then the inventory target is set to $\hat{\tau} + F^{-1}(\hat{\kappa}_\gamma; 0, 1)\hat{\theta}$; and if $\hat{p} \leq \gamma$, then it is set to zero. We let $\hat{\kappa}_\gamma$ be the value of κ that solves the equation $\Pi(\kappa; \hat{p}, y, \hat{r}) = 1 - \gamma/\hat{p}$, where $\Pi(\kappa; \hat{p}, y, \hat{r})$ is the left-hand side of (1) for notational convenience.

Building on the definitions $\mathcal{U} := (\hat{\tau} - \tau)/\theta$ and $\mathcal{V} := \hat{\theta}/\theta$, the difference between the expected cost of an inventory-target estimate and the minimum expected cost of the optimal inventory target takes the form $\theta\Lambda(\hat{\kappa}_\gamma, \mathcal{U}, \mathcal{V}, y; \gamma, p, r)$, where $\Lambda(\hat{\kappa}_\gamma, \mathcal{U}, \mathcal{V}, y; \gamma, p, r)$ is given by

$$\begin{cases} C(\mathcal{U} + F^{-1}(\hat{\kappa}_\gamma; 0, 1)\mathcal{V}|y; p, 0, 1, r) - C(F^{-1}(1 - \gamma/p; 0, 1)|y; p, 0, 1, r) & \text{for } \hat{p} > \gamma \text{ and } p \leq \gamma, \\ C(\mathcal{U} + F^{-1}(\hat{\kappa}_\gamma; 0, 1)\mathcal{V}|y; p, 0, 1, r) - C(F^{-1}(\kappa^*; 0, 1)|y; p, 0, 1, r) & \text{for } \hat{p} > \gamma \text{ and } p > \gamma, \\ C(F^{-1}(1 - \gamma/p; 0, 1)|y; p, 0, 1, r) - C(F^{-1}(\kappa^*)|y; p, 0, 1, r) & \text{for } \hat{p} \leq \gamma \text{ and } p > \gamma, \\ 0 & \text{for } \hat{p} \leq \gamma \text{ and } p \leq \gamma. \end{cases}$$

We use the nonnegative random variable $\Lambda(\hat{\kappa}_\gamma, \mathcal{U}, \mathcal{V}, y; \gamma, p, r)$ as the test statistic. The operational test compares the observed value of this test statistic to the distribution that it would follow under the null hypothesis of no correlation. The null hypothesis is then rejected when the observed value of the test statistic is sufficiently high relative to this distribution.

Operational Test Testing for correlation with $\mathcal{H}_0 : r = 0$ and $\mathcal{H}_A : r \neq 0$.

Initialization. Specify the number of bootstrap samples (B), the length of historical data (n), the number of positive realizations (n_d), the number of periods since last demand (y), and the significance level ($1 - \alpha$).

Set $\hat{p} \leftarrow n_d/n$ and $b \leftarrow 1$.

Compute the test statistic $\hat{\Lambda} := \Lambda(\hat{\kappa}_\gamma, 0, 1, y; \gamma, \hat{p}, 0)$ from the data set $\{(x_t, y_t); t = 1, 2, \dots, n_d\}$.

while $b \leq B$:

Step 1. Generate the b th bootstrap sample path as in Step 1 of the correlation test.

Step 2. Compute the test statistic from the bootstrap sample path:

Set $\hat{p}^b \leftarrow n_d^b/n$.

Let \bar{z} and s_z be the sample mean and sample standard deviation of $\{z_1^b, z_2^b, \dots, z_{n_d}^b\}$.

Set $\hat{r}^b \leftarrow \operatorname{argmax}_{r \in (-1, 1)} \sum_{i=1}^{n_d^b} \log \phi_2\{F^{-1}(F((z_i - \bar{z})/s_z; 0, 1)), \Phi^{-1}(G(y_i^b; \hat{p}^b)); r\}$.

Set κ_γ^b to the value of κ that solves the equation $\Pi(\kappa; \hat{p}^b, y, \hat{r}^b) = 1 - \gamma/\hat{p}^b$.

Set $u^b \leftarrow \bar{z} - a_1 v^b$ and $v^b \leftarrow s_z/\sqrt{a_2}$.

$\hat{\Lambda}^b \leftarrow \Lambda(\kappa_\gamma^b, u^b, v^b, y; \gamma, \hat{p}^b, 0)$.

$b \leftarrow b + 1$

end

Step 3. Construct the critical region to reject the null hypothesis:

Sort $\hat{\Lambda}^b, b = 1, 2, \dots, B$ in ascending order $(\Lambda_{(1)}, \Lambda_{(2)}, \dots, \Lambda_{(B)})$.

Reject the null hypothesis \mathcal{H}_0 if $\hat{\Lambda} > \Lambda_{(\lceil(1-\alpha)B\rceil)}$.

In the following section, we investigate how the consideration of the expected cost of the correlation parameter uncertainty affects the decision maker in assessing the need for modeling the correlation.

3.3 Results and Insights

We design our experiments by letting the number of independent bootstrap sample paths B take the value of 10,000. We also take the number of positive realizations n_d as 20%, 40%, 60%, and 80% of n , the length of the demand history. Our first numerical finding is that there exists an interval of the correlation between the demand size and the number of inter-demand periods in which the expected cost of parameter uncertainty is smaller if the intermittent demand process is assumed to be independent despite the existence of the correlation. For the significance level of 95% (i.e., $\alpha = 0.05$), we refer to the lower and upper limits of this interval as the 95% confidence bounds and provide the values identified for the correlation test in Table 1. We note that the correlation test rejects the hypothesis that the correlation is zero in favor of the alternative that it is not zero if the copula-parameter estimate is less than 2.5% quantile or greater than 97.5% quantile.

We identify the 95% significance test bounds as $\{-0.89, 0.89\}$ for a demand history of length 30 with 6 ($= (30)(20\%)$) positive demand realizations. That is, the decision maker has sufficient evidence to reject the independence assumption only if the estimated value of the copula parameter is less than -0.89 or greater than 0.89 . On the other hand, the significance test bounds are $\{-0.14, 0.14\}$ for a demand history of length 250 with 200 positive demand realizations. In this case, the estimated value of the copula parameter must be less than -0.14 or greater than 0.14 to justify the modeling of correlation in intermittent demand. The significance bounds approach zero (i.e., the critical region becomes larger) as the length of the demand history and the number of positive realizations increase.

In Table 2, we focus on the fraction of time that the decision maker rejects the null hypothesis at the level $\alpha = 0.05$ based on 500 replications of the correlation test. In each scenario, we let the percentage of positive realizations be 20% (top) or 40% (bottom) of the demand history to focus on highly intermittent demand histories. Clearly, the fraction of rejections approaches one as the length of the demand history and the true value of the copula parameter (i.e., the strength of correlation) increases. For example, the decision maker rejects the null hypothesis of no correlation 5% of the time when the copula parameter is 0.15 in a demand history of length 30 with 12 positive demand realizations. In this case, the fraction of rejections reaches 94% when the copula parameter is as high as 0.90. It is worth noting that the correlation test only considers the sampling distribution of the copula-parameter estimator to decide whether to model the correlation. We next present our findings for the operational test that accounts for the expected cost associated with the overestimation and underestimation of the optimal inventory target from limited historical demand data.

In Table 3 we present the fraction of time in which the decision maker rejects the independence assumption, as a result of using the operational test, at the $\alpha = 0.05$ level when the cost ratio $\gamma = h/(h + b)$ is equal to 0.1. We observe that the fraction of time in which the decision maker rejects the independence assumption is considerably smaller in Table 3 compared to Table 2, especially when the length of the demand history and the strength of the correlation are not very high. This can be explained by the additional statistical

Table 1: The 95% significance bounds in the correlation test to reject $r = 0$ against $r \neq 0$

n	The percentage of positive realizations in the n -period history			
	20%	40%	60%	80%
30	$\{-0.89, 0.89\}$	$\{-0.62, 0.62\}$	$\{-0.48, 0.48\}$	$\{-0.41, 0.41\}$
50	$\{-0.70, 0.70\}$	$\{-0.46, 0.46\}$	$\{-0.36, 0.36\}$	$\{-0.31, 0.31\}$
100	$\{-0.47, 0.47\}$	$\{-0.32, 0.32\}$	$\{-0.26, 0.26\}$	$\{-0.22, 0.22\}$
250	$\{-0.28, 0.28\}$	$\{-0.20, 0.20\}$	$\{-0.16, 0.16\}$	$\{-0.14, 0.14\}$

Table 2: The fraction of time the correlation test rejects $\mathcal{H}_0 : r = 0$ at 95% significance level

The percentage of positive realizations in the n -period history is 20%							
n	n_d	r					
		0.15	0.3	0.45	0.6	0.75	0.90
30	6	3%	3%	5%	12%	17%	38%
50	10	5%	7%	13%	32%	62%	85%
100	20	9%	20%	46%	77%	95%	100%
250	50	20%	53%	90%	99%	100%	100%
The percentage of positive realizations in the n -period history is 40%							
30	12	5%	9%	22%	37%	75%	94%
50	20	7%	20%	41%	71%	94%	100%
100	40	11%	37%	75%	96%	100%	100%
250	100	31%	80%	99%	100%	100%	100%

estimation error around the normal copula parameter when the independence assumption is relaxed. To put it another way, the benefit of accounting for the correlation is dominated by the additional expected cost associated with the incorrect estimation of the normal copula parameter. When the amount of data is limited and the strength of the correlation is low, the copula-parameter estimator has the highest variance, and therefore the estimated value of correlation is subject to the highest statistical estimation error. This is why the operational test, which considers the expected cost of parameter uncertainty, is in favor of the simpler model with no correlation when the intermittent demand history is short with low strength of dependence between demand size and number of inter-demand periods.

We observe that the difference between the inventory holding and backlogging costs, and hence, the asymmetry in the expected cost function – which is ignored by the correlation test – plays an important role when the expected cost of parameter uncertainty is the criterion used to decide whether or not to model the correlation. To illustrate, we present the fraction of time that the decision maker rejects the independence assumption in Table 4 when γ is equal to 0.25 and the percentage of positive realizations in the n -period history is 40%. Clearly, the fraction of time the decision maker rejects the independence assumption is smaller in Table 4 compared to the second half of Table 3. In these tables, the values of n and n_d are both the same, while the difference between the ratio of n_d to n and the value of γ is smaller in Table 4. The decision maker is more likely to set a zero inventory target when the ratio of n_d to n is not sufficiently higher than γ , leading to smaller values for the probability of rejecting the independence assumption in Table 4. Intuitively, the decision maker is less likely to model the correlation since the inventory target is set

Table 3: The fraction of time the operational test rejects $\mathcal{H}_0 : r = 0$ at 95% significance level for $\gamma = 0.1$

The percentage of positive realizations in the n -period history is 20%							
n	n_d	r					
		0.15	0.3	0.45	0.6	0.75	0.90
30	6	0%	0%	0%	0%	0%	2%
50	10	0%	0%	0%	0%	3%	5%
100	20	0%	1%	3%	9%	21%	27%
250	50	3%	7%	26%	33%	46%	69%
The percentage of positive realizations in the n -period history is 40%							
30	12	0%	1%	3%	5%	8%	28%
50	20	0%	3%	8%	13%	35%	45%
100	40	0%	8%	28%	46%	49%	58%
250	100	0%	15%	51%	63%	73%	84%

Table 4: The fraction of time the operational test rejects $\mathcal{H}_0 : r = 0$ at 95% significance level for $\gamma = 0.25$

		The percentage of positive realizations in the n -period history is 40%					
		r					
n	n_d	0.15	0.3	0.45	0.6	0.75	0.90
30	12	0%	0%	0%	1%	3%	5%
50	20	0%	0%	2%	5%	11%	14%
100	40	0%	2%	5%	13%	41%	53%
250	100	1%	10%	35%	69%	72%	81%

to zero anyway. Indeed, we observe that the fraction of time the decision maker rejects the independence assumption in the operational test is almost zero in all the cases when the ratio of n_d to n is 20% with $\gamma = 0.25$ and when the ratio of n_d to n is 20% and 40% with $\gamma = 0.5$.

A natural question to ask is which statistical test to use for investigating the existence of correlation in an intermittent demand history. The correlation test considers the normal copula parameter in isolation, and therefore, it is more appropriate if the goal is merely to make an inference about the intermittent demand process. On the other hand, the operational test incorporates the expected cost of incorrectly estimating the copula parameter into the decision of modeling the correlation. Thus, the operational test is more pragmatic to consider when the question is whether the copula parameter estimate should be used as an input in decision making, as is the case in this paper.

4 CONCLUSION

When we have the full knowledge of the dependence parameters of an input distribution function, it is well known that accounting for the dependence in the input process can provide significant benefits. In the absence of this information, however, the benefit of accounting for the correlation may be dominated by the additional expected cost associated with the incorrect estimation of the correlation parameter. This insight into the price of correlation parameter uncertainty led us to ask whether we should always model the correlation despite the statistical estimation error around the copula parameter. We answer this fundamental question before attempting to account for the correlation in simulation design and analysis with a focus on data-driven inventory management for a product with an intermittent demand process. We present an operational test that is in favor of the simpler model with no dependence when the demand history is short with low strength of dependence. As we consider settings with many input processes, this simplicity may potentially lead to sparse correlation matrices for system inputs, leading to the identification of clusters of input processes, where only the input processes in a cluster are considered to be dependent under parameter uncertainty. In certain cases, this simplicity may also allow us to ensure computational efficiency by the cluster-based decomposition of large-scale simulation optimization problems with dependent inputs into a series of small-scale simulation optimization problems. This is the subject of ongoing work.

REFERENCES

- Akcay, A., B. Biller, S. Tayur. 2013a. Data-driven newsvendor: A probabilistic guarantee for near optimality when demand is temporally dependent. Working Paper 2013 - E11, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA.
- Akcay, A., B. Biller, and S. Tayur. 2013b. Managing inventory with limited history of intermittent demand. Working Paper 2013 - E6, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA.
- Altay, N., L.A. Litteral, F. Rudisill. 2012. Effects of correlation on intermittent demand forecasting and stock control. *International Journal of Production Economics*. 135: 275 – 283.
- Cherubini, U., S. Mulinacci, F. Gobbi, S. Romagnoli. 2011. *Dynamic Copula Methods in Finance*. Wiley.
- Eaves, A.H.C. 2002. Forecasting for the ordering and stock-holding of consumable spare parts. *Ph.D. thesis*, Lancaster University.

- Fosdick, B.K., A.E. Raftery. 2012. Estimating the correlation in bivariate normal data with known variances and small sample sizes. *The American Statistician* 66(1):34–41.
- 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.
- MacKinnon, J.G. 2009. Bootstrap hypothesis testing. *Handbook of Computational Econometrics*. Chichester, UK: Wiley 183–213.
- Nelsen, R.B. 2006. *An Introduction to Copulas*. 2nd edition, Springer, New York.
- Willemain, T.R., C.N. Smart, J.H. Shockor, P.A. DeSautels. 1994. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston’s method. *International Journal of Forecasting*. 10: 529 – 538.

AUTHOR BIOGRAPHIES

BAHAR BILLER is an Associate Professor of Operations Management and Manufacturing at Carnegie Mellon University. Her primary research interest lies in the area of computer simulation experiments for stochastic systems and more specifically, in multivariate input modeling for dependent input processes with applications to operations management. Her email address is billerb@andrew.cmu.edu.

ALP AKCAY is an Assistant Professor of Industrial Engineering at Bilkent University. He received his Ph.D. in Operations Management and Manufacturing from Tepper School of Business at Carnegie Mellon University. His research interests include the design and analysis of stochastic system simulations and data-driven decision making under uncertainty with applications in operations management. His email address is alp.akcay@bilkent.edu.tr.

CANAN GUNES CORLU is an Assistant Professor of Administrative Sciences at Boston University. Her research interests include the design of large-scale simulations with applications to inventory management, and the applications of Operations Research techniques (e.g., vehicle routing, inventory control) tailored to the problems of non-profit sectors. Her email address is canan@bu.edu.

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 publications in journals such as *Operations Research*, *Management Science*, and *MSOM Journal*, 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 Journal* and *Management Science*. Dr. Tayur also served as 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.