Accounting for Parameter Uncertainty in Large-Scale Stochastic Simulations with Correlated Inputs

Bahar Biller
Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, billerb@andrew.cmu.edu

Canan G. Corlu
Department of Industrial Engineering, Bilkent University, 06800 Bilkent, Ankara, Turkey, canan.corlu@bilkent.edu.tr

This paper considers large-scale stochastic simulations with correlated inputs having normal-to-anything (NORTA) distributions with arbitrary continuous marginal distributions. Examples of correlated inputs include processing times of workpieces across several workcenters in manufacturing facilities and product demands and exchange rates in global supply chains. Our goal is to obtain mean performance measures and confidence intervals for simulations with such correlated inputs by accounting for the uncertainty around the NORTA distribution parameters estimated from finite historical input data. This type of uncertainty is known as the parameter uncertainty in the discrete-event stochastic simulation literature. We demonstrate how to capture parameter uncertainty with a Bayesian model that uses Sklar's marginal-copula representation and Cooke's copula-vine specification for sampling the parameters of the NORTA distribution. The development of such a Bayesian model well suited for handling many correlated inputs is the primary contribution of this paper. We incorporate the Bayesian model into the simulation replication algorithm for the joint representation of stochastic uncertainty and parameter uncertainty in the mean performance estimate and the confidence interval. We show that our model improves both the consistency of the mean line-item fill-rate estimates and the coverage of the confidence intervals in multiproduct inventory simulations with correlated demands.

Subject classifications: Bayesian; correlation; design of experiments; sampling; statistical analysis.
Area of review: Simulation.
History: Received April 2008; revisions received November 2008, August 2009, January 2010, March 2010; accepted April 2010.

1. Introduction

In recent years, large-scale discrete-event stochastic simulation has become a tool that is used routinely for the design and analysis of manufacturing and service systems. Two important components of the large-scale stochastic simulation are multivariate input modeling and output analysis. Multivariate input modeling is the estimation of an appropriate multivariate probability distribution that characterizes the stochastic behavior of the system inputs. Output analysis is the study of the simulation output data to estimate the distributional properties (e.g., mean, probability, or quantile) of the performance measure.

In this paper, we are interested in the case where the objective of the output analysis is to predict a mean performance measure and a confidence interval. There are three main sources of uncertainty to account for in the output analysis: stochastic uncertainty, model uncertainty, and parameter uncertainty. Stochastic uncertainty arises from the dependence of the output on the simulation's random input streams (Helton 1997). Model uncertainty arises due to the uncertainty around the selection of an appropriate family of distributions for the system inputs, whereas parameter uncertainty arises due to the uncertainty around

the parameter values of a given probability distribution (Raftery et al. 1996).

The goal of this paper is to account for the stochastic uncertainty and the parameter uncertainty in the estimation of the mean performance measure and the confidence interval of the stochastic simulation with correlated inputs. Accounting for parameter uncertainty in a stochastic simulation is not common practice. The simulation often starts with fitting a probability distribution to the historical input data of finite length. Although the parameters of the fitted distribution are shown to have asymptotical properties (e.g., consistency and normality) for the number of historical data points approaching infinity, the simulation is driven with the probability distribution estimated from the finite historical input data. The output data obtained from the simulation are analyzed for predicting the mean performance measure and constructing the confidence interval. This practice of using the estimated probability distribution for driving the simulation ignores both the model uncertainty and the parameter uncertainty, and accounts only for the stochastic uncertainty in the output analysis. Consequently, the simulation analyst obtains not only an inconsistent estimate for the mean performance measure, but also an inconsistent coverage for the confidence interval.

The problem of accounting for model uncertainty and/or parameter uncertainty in stochastic simulations has been studied by a number of researchers, including Cheng and Holland (1997, 1998, 2004), Chick (1997, 1999, 2001), Barton and Schruben (2001), and Zouaoui and Wilson (2003, 2004). Cheng and Holland (1997) made the first attempt to show the dependence of the simulation output on stochastic and parameter uncertainties. They continued the study of this problem in Cheng and Holland (1998 and 2004). However, using frequentist techniques to represent parameter uncertainty did not allow the incorporation of any relevant information other than the historical input data into the simulation output. Furthermore, the problem of accounting for model uncertainty remained unsolved. As a result of following the Bayesian Model Averaging (BMA) approach, Chick (1999, 2001) captured not only stochastic uncertainty and parameter uncertainty, but also model uncertainty. Although the BMA approach had been used in a number of different settings to capture model uncertainty and/or parameter uncertainty (e.g., Draper 1995, Hoeting et al. 1999, and George 1999), it was Chick (1997, 1999, 2001) who outlined the basic methodology for implementing the BMA approach in discrete-event stochastic simulations.

Chick's formulation led to the simulation replication algorithm that allowed the simulation analyst to capture both model uncertainty and parameter uncertainty by sampling input distributions and their parameters from Bayesian posterior density functions before each replication. This made it possible to drive the simulation without using a single distribution and a single set of parameter values. Consequently, the use of the simulation replication algorithm improved the consistency of the mean performance estimate and the coverage of the confidence interval. However, the simulation replication algorithm did not provide separate estimates for the variances due to different sources of uncertainty. The Bayesian simulation replication algorithm introduced by Zouaoui and Wilson (2003) provided a solution for the problem of decomposing the variance in the simulation output data into variances due to stochastic and parameter uncertainties. They extended their framework to also account for model uncertainty in Zouaoui and Wilson (2004).

The focus in these papers has been on discrete-event stochastic simulations with independent inputs. However, building large-scale simulations may require the development of multivariate input models. Examples of multivariate inputs include the processing times of a workpiece across several workcenters in a manufacturing facility (Xu 1999) and the product demands and exchange rates in a global supply chain (Kouvelis and Su 2007). In this paper, we focus on simulations with multivariate inputs that have stochastic dependencies among them, and we describe how to account for both stochastic uncertainty and parameter uncertainty in their output analyses. Although the BMA approach formulated in Chick (2001) can accommodate the

stochastic dependencies among the inputs, the simulation analyst still needs a Bayesian model that works for simulations with many correlated inputs. Also, it is not clear how to sample the parameters of the joint distribution of the correlated inputs before each replication so that the Bayesian model is easily incorporated into the simulation replication algorithm. The development of a Bayesian model, which overcomes these challenges and leads to a fast sampling algorithm well suited for handling a large number of correlated inputs, is the primary contribution of our paper to the discrete-event stochastic simulation literature.

We represent multivariate inputs using random vector $\mathbf{X} = (X_1, X_2, \ldots, X_k)'$ denoting a collection of k correlated components, each of which is a real-valued input random variable. We characterize the joint stochastic behavior of these correlated inputs with the flexible normal-to-anything (NORTA) distribution of Cario and Nelson (1997). Thus, we construct our k-dimensional random vector $\mathbf{X} = (X_1, X_2, \ldots, X_k)'$ by first taking Z_i as the ith component of a k-dimensional standard normal random vector $\mathbf{Z} = (Z_1, Z_2, \ldots, Z_k)'$ with positive definite $k \times k$ correlation matrix $\mathbf{\Sigma}_k$. Then we obtain $X_i = F_i^{-1}(\Phi(Z_i); \mathbf{\Psi}_i)$ for $i = 1, 2, \ldots, k$, where F_i is the arbitrary continuous marginal cumulative distribution function (cdf) of component i with parameter vector $\mathbf{\Psi}_i$ and Φ is the cdf of the standard normal random variable. A random vector constructed in this way is said to have a k-dimensional NORTA distribution. From this point on, we call the parameters of this distribution, $\mathbf{\Psi}_i$, $i = 1, 2, \ldots, k$, and $\mathbf{\Sigma}_k$, the NORTA parameters.

Considering a stochastic simulation with inputs having a k-dimensional NORTA distribution, we demonstrate how to account for parameter uncertainty (i.e., the uncertainty around the values of the NORTA parameters estimated from finite historical input data) in the mean performance estimate and the confidence interval. More specifically, we develop a Bayesian model that samples NORTA parameters from their Bayesian posterior density functions before each replication of the simulation replication algorithm. This allows us to drive the stochastic simulation without using a single set of NORTA parameters. However, it is a challenging task to develop such a Bayesian model as the number of NORTA parameters needed to sample increases very quickly with k, the number of components. For example, if each of the k components are exponentially distributed, we need to sample k different parameters for the component marginal distributions and $2k$ different parameters for the gamma distributed components. In addition to the parameters of the component marginal distributions, we need to sample $k(k-1)/2$ correlations of $\mathbf{\Sigma}_k$ in a way that the resulting correlation matrix is positive definite.

We overcome the challenge of sampling a large number of NORTA parameters using Sklar's marginal-copula representation together with Cooke's copula-vine specification. Sklar's marginal-copula representation allows us to write the joint posterior density function as the multiplication of the marginal posterior density functions and the posterior

normal copula density function. Thus, we separate the problem of sampling the parameters of the component marginals from the problem of sampling the dependence parameters of the NORTA distribution. Furthermore, Cooke's copula-vine specification enables us to represent the k-dimensional posterior normal copula density function as the product of $k(k-1)/2$ two-dimensional posterior normal copula density functions. Therefore, we do not need to satisfy any algebraic constraints for positive definiteness.

We organize the remainder of the paper as follows. In §2, we describe the NORTA distribution and present a copula-based representation for its joint density function. We use this representation in §3 for developing a Bayesian model that samples NORTA parameters. In §4, we describe how to incorporate the Bayesian model into the simulation replication algorithm for estimating the mean performance measure and the confidence interval that accounts for both stochastic uncertainty and parameter uncertainty. In §5, we show that our model allows the simulation analyst to improve both the consistency of the mean line-item fill-rate estimates and the coverage of the confidence intervals in multiproduct inventory simulations with correlated demands. We conclude with a summary of results in §6. For clarity in the presentation of the results, we moved the implementation details of sampling NORTA parameters to an electronic companion that is available as part of the electronic version at http://or.journal.informs.org/.

2. The NORTA Distribution and Its Copula-Based Representation

We introduce the k-dimensional NORTA distribution in §2.1 and provide a copula-based representation for its joint density function in §2.2.

2.1. The k-Dimensional NORTA Distribution

We characterize the joint stochastic behavior of correlated inputs X_i, $i = 1, 2, \ldots, k$ using the NORTA distribution of Cario and Nelson (1997). The central idea is to transform a standard multivariate normal random vector into the random vector referred to as having a NORTA distribution. Specifically, we let

$$\mathbf{X} = \left(F_1^{-1}(\Phi(Z_1); \boldsymbol{\Psi}_1), F_2^{-1}(\Phi(Z_2); \boldsymbol{\Psi}_2), \ldots, \right.$$
$$\left. F_k^{-1}(\Phi(Z_k); \boldsymbol{\Psi}_k) \right)',$$

where $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$ are arbitrary marginal cdfs with parameter vectors $\boldsymbol{\Psi}_i$, $i = 1, 2, \ldots, k$ and the base vector $\mathbf{Z} = (Z_1, Z_2, \ldots, Z_k)'$ is a k-dimensional standard normal random vector with $k \times k$ positive definite correlation matrix $\boldsymbol{\Sigma}_k \equiv [\rho(i, j); i, j = 1, 2, \ldots, k]$. In this characterization, $\rho(i, j)$ is the Pearson product-moment correlation between Z_i (i.e., $\Phi^{-1}[F_i(X_i; \boldsymbol{\Psi}_i)]$) and Z_j (i.e., $\Phi^{-1}[F_j(X_j; \boldsymbol{\Psi}_j)]$), whereas the transformation $X_i = F_i^{-1}(\Phi(Z_i); \boldsymbol{\Psi}_i)$ ensures that X_i has cdf $F_i(\cdot; \boldsymbol{\Psi}_i)$. A more detailed description of the NORTA distribution is available in Biller and Ghosh (2006).

2.2. A Copula-Based Representation for the NORTA Distribution

First, we present a brief review of the copula theory in §2.2.1. Then, we derive a copula-based representation for the NORTA distribution using Sklar's marginal-copula representation in §2.2.2 and Cooke's copula-vine specification in §2.2.3.

2.2.1. Copula Theory. We begin this section with the definition of a k-dimensional copula (Nelsen 2006, Definition 2.10.6). The first condition of this definition provides the lower bound on the distribution function and insures that the marginal distributions of the copula are uniform, whereas the second condition insures that the probability of observing a point in a k-box is nonnegative:

DEFINITION 1. A k-dimensional copula is a function C_k: $[0, 1]^k \to [0, 1]$ with the following properties: (1) For every $\mathbf{u} = (u_1, u_2, \ldots, u_k)$ in $[0, 1]^k$, $C_k(\mathbf{u}) = 0$ if at least one coordinate of $\mathbf{u}$ is 0, and if all coordinates of $\mathbf{u}$ are 1 except u_l, then $C_k(\mathbf{u}) = u_l$ for $l = 1, 2, \ldots, k$. (2) For every $\mathbf{a} = (a_1, a_2, \ldots, a_k)$ and $\mathbf{b} = (b_1, b_2, \ldots, b_k)$ in $[0, 1]^k$ such that $\mathbf{a} \leqslant \mathbf{b}$, i.e., $a_l \leqslant b_l$, $l = 1, 2, \ldots, k$, and for every k-box $[\mathbf{a}, \mathbf{b}]$, i.e., $[a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_k, b_k]$, the C_k-volume given by $\Delta_{\mathbf{a}}^{\mathbf{b}} C_k(\mathbf{t}) = \Delta_{a_k}^{b_k} \Delta_{a_{k-1}}^{b_{k-1}} \cdots \Delta_{a_2}^{b_2} \Delta_{a_1}^{b_1} C_k(\mathbf{t})$ with $\Delta_{a_l}^{b_l} C_k(\mathbf{t}) = C_k(t_1, \ldots, t_{l-1}, b_l, t_{l+1}, \ldots, t_k) - C_k(t_1, \ldots, t_{l-1}, a_l, t_{l+1}, \ldots, t_k)$ is $\geqslant 0$.

The use of a copula for understanding the joint distribution of a random vector has been studied extensively for the last two decades (Schweizer 1991, Joe 1997, Nelsen 2006). In this paper, we restrict our attention to Sklar's theorem, which describes how to extract the dependence structure of a random vector from its joint distribution with arbitrary continuous marginal distributions (Nelsen 2006, Theorem 2.10.9):

THEOREM 1. *Let H_k be a k-dimensional distribution function with continuous marginal cdfs $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$. Then there exists a k-dimensional unique copula C_k such that for all x_i, $i = 1, 2, \ldots, k$ in $\Re$,*

$$H_k(x_1, x_2, \ldots, x_k)$$
$$= C_k(F_1(x_1; \boldsymbol{\Psi}_1), F_2(x_2; \boldsymbol{\Psi}_2), \ldots, F_k(x_k; \boldsymbol{\Psi}_k)). \quad (1)$$

Conversely, if C_k is a k-dimensional copula and $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$ are continuous distribution functions with parameter vectors $\boldsymbol{\Psi}_i$, $i = 1, 2, \ldots, k$, then the function H_k defined in (1) is a k-dimensional distribution function with marginal cdfs $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$.

The major implication of this theorem is that copula C_k is the joint distribution function of $U_i \equiv F_i(X_i; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$, where U_i, $i = 1, 2, \ldots, k$ are the probability integral transforms of X_i, $i = 1, 2, \ldots, k$. Thus, each of the random variables U_i, $i = 1, 2, \ldots, k$ follows a uniform distribution in $[0, 1]$, regardless of the distributions of the input random variables X_i, $i = 1, 2, \ldots, k$. More importantly,

C_k can be interpreted as the dependence structure of the joint cdf H_k and written as $C_k(u_1, u_2, \ldots, u_k) = H_k(F_1^{-1}(u_1; \boldsymbol{\Psi}_1), F_2^{-1}(u_2; \boldsymbol{\Psi}_2), \ldots, F_k^{-1}(u_k; \boldsymbol{\Psi}_k))$, where $F_i^{-1}(\cdot; \boldsymbol{\Psi}_i)$ is the generalized inverse of marginal cdf $F_i(\cdot; \boldsymbol{\Psi}_i)$ (Nelsen 2006, Corollary 2.10.10).

Another important implication of the representation in (1) is that a joint probability density function (pdf) can be written as a product of marginal pdfs and copula density function, which encodes all of the information about the stochastic dependencies among the components. More specifically, for differentiable marginal cdfs $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$ and differentiable copula C_k, the k-dimensional pdf, which is denoted by h_k below, can be written as

$$
\begin{aligned}
h_k(x_1, x_2, \ldots, x_k) \\
= \prod_{i=1}^{k} f_i(x_i; \boldsymbol{\Psi}_i) \\
\times c_k(F_1(x_1; \boldsymbol{\Psi}_1), F_2(x_2; \boldsymbol{\Psi}_2), \ldots, F_k(x_k; \boldsymbol{\Psi}_k)).
\end{aligned}
$$

In this representation, $f_i(\cdot; \boldsymbol{\Psi}_i)$ is the pdf of X_i, i.e., $f_i(x; \boldsymbol{\Psi}_i) \equiv \partial F_i(x; \boldsymbol{\Psi}_i)/\partial x$, and c_k is the k-dimensional copula density function given by $\partial^k C_k(u_1, u_2, \ldots, u_k)/(\partial u_1 \partial u_2 \ldots \partial u_k)$. This copula density function takes the value of 1 when X_i, $i = 1, 2, \ldots, k$ are independent and therefore, the joint density function reduces to the product of only the marginal pdfs.

2.2.2. Sklar's Marginal-Copula Representation and NORTA. The use of Sklar's theorem for representing a k-dimensional NORTA distribution shows that the k-dimensional random vector with the NORTA distribution and the k-dimensional normal random vector share the same copula:

COROLLARY 1. *Let* $\mathbf{X} = (X_1, X_2, \ldots, X_k)'$ *correspond to a k-dimensional random vector with a NORTA distribution characterized by continuous marginal cdfs $F_i(\cdot; \boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$ and positive definite correlation matrix $\boldsymbol{\Sigma}_k$. Then there exists a k-dimensional unique normal copula that represents the dependence structure of $\mathbf{X}$.*

The key to the proof of this corollary is that the joint cdf H_k of X_i, $i = 1, 2, \ldots, k$ is given by

$$
\begin{aligned}
H_k(x_1, x_2, \ldots, x_k) \\
= \Phi_k(\Phi^{-1}[F_1(x_1; \boldsymbol{\Psi}_1)], \Phi^{-1}[F_2(x_2; \boldsymbol{\Psi}_2)], \ldots, \\
\Phi^{-1}[F_k(x_k; \boldsymbol{\Psi}_k)]; \boldsymbol{\Sigma}_k),
\end{aligned}
$$

where Φ^{-1} is the functional inverse of Φ and $\Phi_k(\cdot; \boldsymbol{\Sigma}_k)$ is the k-dimensional standard normal cdf with correlation matrix $\boldsymbol{\Sigma}_k$. Because the normal copula is the dependence function implicitly assumed whenever the multivariate normal distribution is used, the dependence structure of a k-dimensional NORTA distribution is represented by a k-dimensional normal copula.

The joint pdf of the k-dimensional NORTA distribution, h_k, can now be written as the multiplication of the k component marginal density functions and the k-dimensional normal copula density function; i.e.,

$$
\begin{aligned}
h_k(x_1, x_2, \ldots, x_k) \\
= \prod_{i=1}^{k} f_i(x_i; \boldsymbol{\Psi}_i) \times \phi_k(\Phi^{-1}[F_1(x_1; \boldsymbol{\Psi}_1)], \Phi^{-1}[F_2(x_2; \boldsymbol{\Psi}_2)], \\
\ldots, \Phi^{-1}[F_k(x_k; \boldsymbol{\Psi}_k)]; \boldsymbol{\Sigma}_k). \quad (2)
\end{aligned}
$$

The normal copula density function ϕ_k is further given by

$$
\begin{aligned}
\phi_k\big(\Phi^{-1}[u_1], \Phi^{-1}[u_2], \ldots, \Phi^{-1}[u_k]; \boldsymbol{\Sigma}_k\big) \\
\equiv \frac{\partial^k \Phi_k(\Phi^{-1}[u_1], \Phi^{-1}[u_2], \ldots, \Phi^{-1}[u_k]; \boldsymbol{\Sigma}_k)}{\partial u_1 \partial u_2 \ldots \partial u_k}, \\
\equiv |\boldsymbol{\Sigma}_k|^{-1/2} \exp\left\{-\frac{1}{2}\mathbf{s}'(\boldsymbol{\Sigma}_k^{-1} - \mathbf{I}_k)\mathbf{s}\right\}, \quad (3)
\end{aligned}
$$

where $u_i \equiv F_i(x_i; \boldsymbol{\Psi}_i)$ for $i = 1, 2, \ldots, k$, $\mathbf{s} \equiv (\Phi^{-1}[u_1], \Phi^{-1}[u_2], \ldots, \Phi^{-1}[u_k])'$, and $\mathbf{I}_k$ is the k-dimensional identity matrix. Thus, the copula density function ϕ_k captures all of the information about the dependence structure of $\mathbf{X}$ using correlations $\rho(i, j)$, $i, j = 1, 2, \ldots, k$.

2.2.3. Cooke's Copula-Vine Specification and NORTA. A vine is a graphical model for constructing high-dimensional joint distributions using a series of two-dimensional (conditional) distributions. It was introduced in Cooke (1997), studied extensively in Bedford and Cooke (2001, 2002) and Kurowicka and Cooke (2003), and described comprehensively in Kurowicka and Cooke (2006). In this paper, we use a vine for representing NORTA's normal copula density function in (3). More specifically, we represent the joint distribution of the base random variables of the k-dimensional NORTA distribution (i.e., $Z_i \equiv \Phi^{-1}[F_i(X_i; \boldsymbol{\Psi}_i)]$, $i = 1, 2, \ldots, k$) with a regular vine defined as follows (Kurowicka and Cooke 2006, Definition 4.4).

DEFINITION 2. $\mathcal{V}$ is a vine on k elements under the following conditions: (1) $\mathcal{V} = (\mathcal{T}_1, \mathcal{T}_2, \ldots, \mathcal{T}_{k-1})$. (2) $\mathcal{T}_1$ is a connected tree with nodes $\mathcal{N}_1 = \{1, 2, \ldots, k\}$ and edges $\mathcal{E}_1$; for $i = 2, 3, \ldots, k - 1$, $\mathcal{T}_i$ is a connected tree with nodes $\mathcal{N}_i = \mathcal{E}_{i-1}$. $\mathcal{V}$ is a regular vine on k elements if additionally the following condition is satisfied: (3) For $i = 2, 3, \ldots, k - 1$, if $\{a, b\} \in \mathcal{E}_i$, then $\#a\Delta b = 2$, where Δ denotes the symmetric difference. In other words, if a and b are nodes of $\mathcal{T}_i$ connected by an edge in $\mathcal{T}_i$, where $a = \{a_1, a_2\}$ and $b = \{b_1, b_2\}$, then exactly one of the a_i equals one of the b_i.

No unique regular vine exists for representing the dependence structure of the NORTA distribution. Figure 1 provides examples of three different regular vines, each of which can be used for representing the dependence structure of the 5-dimensional NORTA distribution. The first

Figure 1.   Three different regular vine specifications for the five-dimensional NORTA distribution.
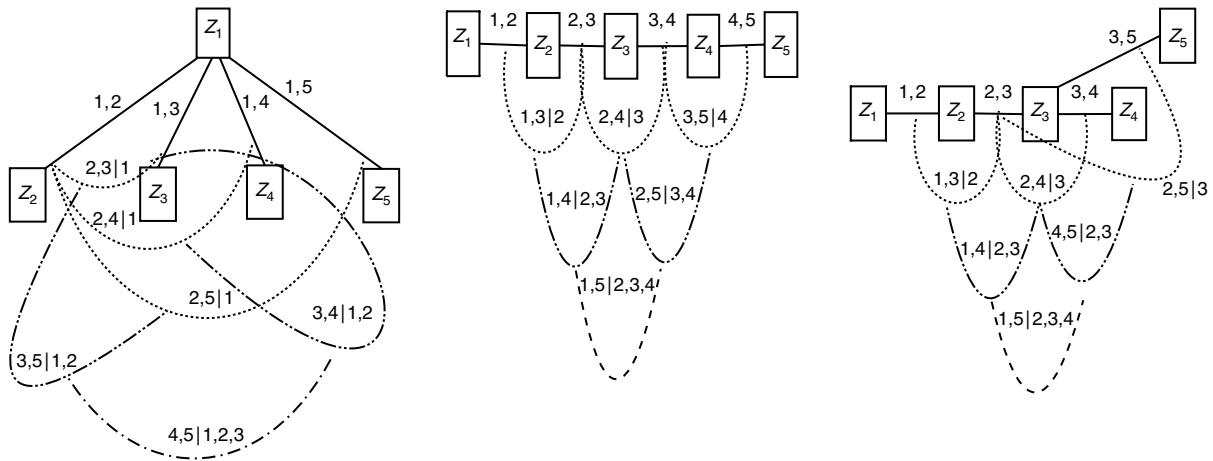
two vines are known as the $\mathscr{C}$-vine and the $\mathscr{D}$-vine, respectively, and described in Kurowicka and Cooke (2006), whereas the last one is introduced in Bedford and Cooke (2002). Our solution approach works for any of these vines as well as any regular vine constructed as described in Definition 2. Because using different regular vine specifications leads to different sampling algorithms for NORTA's dependence parameters, we use a $\mathscr{C}$-vine for describing our solution approach in the remainder of the paper due to the ease of its implementation (Kurowicka and Cooke 2006, §6.4.2).

We represent NORTA's k-dimensional normal copula density function in (3) with a $\mathscr{C}$-vine on $\Phi^{-1}[F_i(X_i; \Psi_i)]$, $i = 1, 2, \ldots, k$. This vine has the following unconditional and conditional correlations assigned to its edges: $\rho(1, i)$, $i = 2, 3, \ldots, k$ and $\rho(j-1, i; 1, 2, \ldots, j-2)$, $i = j$, $j+1, \ldots, k$, $j = 3, 4, \ldots, k$. Whereas $\rho(1, i)$ is the (unconditional) correlation between random variables $\Phi^{-1}[F_1(X_1; \Psi_1)]$ and $\Phi^{-1}[F_i(X_i; \Psi_i)]$, $\rho(j-1, i; 1, 2, \ldots, j-2)$ is the (conditional) correlation between conditional random variables $\Phi^{-1}[F_{j-1}(X_{j-1}; \Psi_{j-1})] \mid \Phi^{-1} \cdot [F_l(X_l; \Psi_l)]$, $l = 1, 2, \ldots, j-2$ and $\Phi^{-1}[F_i(X_i; \Psi_i)] \mid \Phi^{-1} \cdot [F_l(X_l; \Psi_l)]$, $l = 1, 2, \ldots, j-2$. Because the dependence structure of the NORTA distribution is represented by a normal copula (Corollary 1), the conditional correlation $\rho(j-1, i; 1, 2, \ldots, j-2)$ is also the partial correlation (i.e., the correlation between the orthogonal projections of $\Phi^{-1}[F_{j-1}(X_{j-1}; \Psi_{j-1})]$ and $\Phi^{-1}[F_i(X_i; \Psi_i)]$ on the plane orthogonal to the space spanned by $\Phi^{-1}[F_l(X_l; \Psi_l)]$, $l = 1, 2, \ldots, j-2$) (Morales et al. 2006). Recursive formulas exist that allow the identification of the partial correlations from the correlations of Σ_k (Yule and Kendall 1965).

All of the (partial) correlations in the $\mathscr{C}$-vine specification of the k-dimensional NORTA distribution are algebraically independent. Therefore, they do not need to satisfy any algebraic constraints for positive definiteness. Furthermore, the resulting copula-vine specification uniquely determines the correlation matrix Σ_k:

COROLLARY 2. *For the $\mathscr{C}$-vine on $\Phi^{-1}[F_i(X_i; \Psi_i)]$, $i = 1, 2, \ldots, k$, there is a one-to-one correspondence between the set of $k \times k$ positive definite correlation matrices of the form Σ_k and the set of correlations $\rho(1, i)$, $i = 2, 3, \ldots, k$ and partial correlations $\rho(j-1, i; 1, 2, \ldots, j-2)$, $i = j$, $j+1, \ldots, k$, $j = 3, 4, \ldots, k$ of the k-dimensional NORTA distribution.*

The proof of this corollary is from the application of Theorem 4.4 of Kurowicka and Cooke (2006) to the copula-vine specification of the NORTA distribution. Therefore, all assignments of the numbers between -1 and 1 to the edges of the $\mathscr{C}$-vine (and to the edges of any arbitrary regular vine) are consistent in the sense that there is a NORTA distribution realizing these (partial) correlations.

We are now ready to replace $h_k(x_1, x_2, \ldots, x_k)$, the joint pdf of the k-dimensional NORTA distribution in (2), with the product of the k component marginal density functions and the $k(k-1)/2$ two-dimensional normal copula density functions, each of which is associated with an edge of the $\mathscr{C}$-vine:

$$= \prod_{i=1}^{k} f_i(x_i; \Psi_i)$$

$$\times \prod_{i=2}^{k} \phi_2\big(F_1(x_1; \Psi_1), F_i(x_i; \Psi_i); \Sigma_2(1, i)\big)$$

$$\times \prod_{j=3}^{k} \prod_{i=j}^{k} \phi_2\Big(\Phi_{j-1 \mid 1, 2, \ldots, j-2}\big(\Phi^{-1}[F_{j-1}(x_{j-1}; \Psi_{j-1})]\,\big|$$

$$\Phi^{-1}[F_l(x_i; \Psi_l)], \, l = 1, 2, \ldots, j-2\big),$$

$$\cdot \Phi_{i \mid 1, 2, \ldots, j-2}\big(\Phi^{-1}[F_i(x_i; \Psi_i)] \,\big|\, \Phi^{-1}[F_l(x_l; \Psi_l)],$$

$$l = 1, 2, \ldots, j-2);$$

$$\cdot \Sigma_2(j-1, i; 1, 2, \ldots, j-2)\Big). \quad (4)$$

In this representation, $\Sigma_2(1, i)$ is the two-dimensional correlation matrix with correlation $\rho(1, i)$ as its (1, 2)th entry

(i.e., $\mathbf{\Sigma}_2(1, i) \equiv [1\rho(1, i); \rho(1, i)1]$). Similarly, $\mathbf{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2) \equiv [1\rho(j-1, i; 1, 2, \ldots, j-2); \rho(j-1, i; 1, 2, \ldots, j-2)1]$. Additionally, $\Phi_{s\,|\,1,2,\ldots,j-2}$ is the marginal cdf of the conditional random variable $\Phi^{-1}[F_s(X_s; \mathbf{\Psi}_s)] \mid \Phi^{-1}[F_l(X_l; \mathbf{\Psi}_l)]$, $l = 1, 2, \ldots, j-2$ with mean $\mu_{s\,|\,1,2,\ldots,j-2}$ and variance $\sigma^2_{s\,|\,1,2,\ldots,j-2}$ that can be obtained using Theorem 3.3.4 of Tong (1990) with the recursive formulas of Yule and Kendall (1965). Appendix B describes how to do this for a 5-dimensional NORTA distribution.

3. Bayesian Model for Sampling NORTA Parameters

The key to the development of our Bayesian model is to separate the sampling of the component parameter vectors from the sampling of the (partial) correlations using Sklar's marginal-copula representation and Cooke's copula-vine specification. Therefore, in §3.1 we first focus on the ith component of the NORTA vector and describe how to sample parameter vector $\mathbf{\Psi}_i$. Then we discuss the sampling of correlation $\rho(1, i)$ and partial correlation $\rho(j-1, i; 1, 2, \ldots, j-2)$ in §3.2. Finally, in §3.3 we describe how to sample all parameters of the k-dimensional NORTA distribution (i.e., $\mathbf{\Psi}_i$, $i = 1, 2, \ldots, k$, $\rho(1, i)$, $i = 2, 3, \ldots, k$, and $\rho(j-1, i; 1, 2, \ldots, j-2)$, $i = j, j+1, \ldots, k$, $j = 3, 4, \ldots, k$) using the sampling algorithms of §§3.1 and 3.2.

3.1. Sampling the Parameters of the Component Marginal Distributions

Well-established literature exists on Bayesian probability theory for sampling the parameters of the standard families of distributions (Gelman et al. 2000, Carlin and Louis 2000). Assuming the availability of the historical data of finite length, this section describes how to use this literature to sample the parameters of the ith component having the exponential distribution or the gamma distribution of the standard input-modeling packages.

First, we choose the distribution of the ith component X_i as exponential with scale parameter β_i; i.e., $f_i(x_i; \beta_i) = \beta_i^{-1} \exp(-x_i\beta_i^{-1})$. Therefore, we can write the likelihood function $\prod_{t=1}^n f_i(x_{i,t}; \beta_i)$, which describes the joint pdf of the historical data $x_{i,t}$, $t = 1, 2, \ldots, n$ of length n, as $\beta_i^{-n} \exp(-\beta_i^{-1} \sum_{t=1}^n x_{i,t})$. The next step is to construct a prior density function $\pi_i(\beta_i)$ on scale parameter β_i. We do this with the conjugate,[1] inverse gamma prior with shape parameter θ_i and scale parameter ν_i; i.e., $\pi_i(\beta_i) = \nu_i^{\theta_i} \Gamma^{-1}(\theta_i) \beta_i^{-(\theta_i+1)} \exp(-\nu_i\beta_i^{-1})$. Finally, we denote the vector of the historical data available for the ith component with $\mathbf{x}_i$ and combine the prior density function with the likelihood function using Bayes' rule to obtain the posterior density function $p_i(\beta_i \mid \mathbf{x}_i)$ of parameter β_i (Bernardo and Smith 1994):

$$p_i(\beta_i \mid \mathbf{x}_i) \propto \pi_i(\beta_i) \prod_{t=1}^n f_i(x_{i,t}; \beta_i)$$

$$\propto \beta_i^{-(n+\theta_i+1)} \exp\left\{-\frac{\nu_i + \sum_{t=1}^n x_{i,t}}{\beta_i}\right\}. \tag{5}$$

Thus, representing parameter uncertainty for component X_i reduces to the sampling of β_i^{-1} from a gamma distribution with shape parameter $n + \theta_i$ and scale parameter $(\nu_i + \sum_{t=1}^n x_{i,t})^{-1}$, for which an efficient algorithm is available in Appendix A.1. The appendix is part of the online version that can be found at http://or.journal.informs.org/.

Next, we consider a gamma component with shape parameter α_i and scale parameter β_i; i.e., $f_i(x_i; \alpha_i, \beta_i) = x_i^{\alpha_i-1} \Gamma^{-1}(\alpha_i) \beta_i^{-\alpha_i} \exp(-x_i\beta_i^{-1})$. We use Bayes' rule for combining Jeffreys' prior density function $\pi_i(\alpha_i, \beta_i) \propto \beta_i^{-1}$ with the likelihood function of $\mathbf{x}_i$ and obtain the following joint posterior density function for the parameters of the gamma component (Son and Oh 2006):

$$p_i(\alpha_i, \beta_i \mid \mathbf{x}_i) \propto \pi_i(\alpha_i, \beta_i) \prod_{t=1}^n f_i(x_{i,t}; \alpha_i, \beta_i)$$

$$= \frac{\beta_i^{-\alpha_i n-1}}{[\Gamma(\alpha_i)]^n} \left(\prod_{t=1}^n x_{i,t}^{\alpha_i-1}\right) \exp\left\{-\frac{\sum_{t=1}^n x_{i,t}}{\beta_i}\right\}.$$

Parameters α_i and β_i can be sampled from this joint posterior density function using the Markov Chain Monte Carlo (MCMC) method. The idea behind the MCMC method is to simulate a random walk in the space of (α_i, β_i) that converges to the joint posterior density function $p_i(\alpha_i, \beta_i \mid \mathbf{x}_i)$ (Gilks et al. 1996). A widely used MCMC method is the Gibbs sampler algorithm (Geman and Geman 1984, Gelfand and Smith 1990). We describe the implementation of this algorithm for the parameters of the gamma distribution in Appendix A.2.

3.2. Sampling the (Partial) Correlations

First, we describe the sampling of correlation $\rho(1, i)$. Because the focus is on the correlation between random variables $\Phi^{-1}[F_1(X_1; \mathbf{\Psi}_1)]$ and $\Phi^{-1}[F_i(X_i; \mathbf{\Psi}_i)]$, we provide an explicit representation of their joint density function $\phi_2(F_1(x_1; \mathbf{\Psi}_1), F_i(x_i; \mathbf{\Psi}_i); \mathbf{\Sigma}_2(1, i))$:

$$|\mathbf{\Sigma}_2(1, i)|^{-1/2} \exp\left\{-\frac{1}{2}\begin{pmatrix}\Phi^{-1}[F_1(X_1; \mathbf{\Psi}_1)] \\ \Phi^{-1}[F_i(X_i; \mathbf{\Psi}_i)]\end{pmatrix}' \right.$$
$$\left. (\mathbf{\Sigma}_2^{-1}(1, i) - \mathbf{I}_2)\begin{pmatrix}\Phi^{-1}[F_1(X_1; \mathbf{\Psi}_1)] \\ \Phi^{-1}[F_i(X_i; \mathbf{\Psi}_i)]\end{pmatrix}\right\}.$$

Thus, defining

$$\mathbf{S}_2(1, i \mid \mathbf{\Psi}_1, \mathbf{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i) \equiv \sum_{t=1}^n \begin{pmatrix}\Phi^{-1}[F_1(x_{1,t}; \mathbf{\Psi}_1)] \\ \Phi^{-1}[F_i(x_{i,t}; \mathbf{\Psi}_i)]\end{pmatrix}$$
$$\cdot \begin{pmatrix}\Phi^{-1}[F_1(x_{1,t}; \mathbf{\Psi}_1)] \\ \Phi^{-1}[F_i(x_{i,t}; \mathbf{\Psi}_i)]\end{pmatrix}'$$

and using the trace operator (tr), we can represent the likelihood function associated with the dependence structure (i.e.,

$\prod_{t=1}^{n} \phi_2(F_1(x_{1,t}; \boldsymbol{\Psi}_1), F_i(x_{i,t}; \boldsymbol{\Psi}_i); \boldsymbol{\Sigma}_2(1, i)))$ as follows:

$$|\boldsymbol{\Sigma}_2(1, i)|^{-n/2} \exp\{\text{tr}(-\tfrac{1}{2}\mathbf{S}_2(1, i \mid \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i) \cdot (\boldsymbol{\Sigma}_2^{-1}(1, i) - \mathbf{I}_2))\}. \quad (6)$$

The form of this likelihood function suggests the use of the inverse Wishart density function as a conjugate prior for $\boldsymbol{\Sigma}_2(1, i)$ (Rossi et al. 2006). Therefore, we follow Jeffreys' invariance principle (Kass and Wasserman 1996) and choose prior density function $\pi(\boldsymbol{\Sigma}_2(1, i)) \propto |\boldsymbol{\Sigma}_2(1, i)|^{-3/2}$ for correlation matrix $\boldsymbol{\Sigma}_2(1, i)$. The right-hand side of this representation corresponds to the inverse Wishart density function with zero degrees of freedom and serves as a diffuse prior density function. Furthermore, it coincides with the beta prior density function of Barnard et al. (2000) with $\nu = 0$.

Combining the prior density function $\pi(\boldsymbol{\Sigma}_2(1, i))$ with the likelihood function in (6), using Bayes' rule leads to the following conditional posterior copula density function:

$$p(\boldsymbol{\Sigma}_2(1, i) \mid \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i)$$
$$\propto |\boldsymbol{\Sigma}_2(1, i)|^{-(n+3)/2} \exp\{\text{tr}(-\tfrac{1}{2}\mathbf{S}_2(1, i \mid \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i) \cdot \boldsymbol{\Sigma}_2^{-1}(1, i))\}.$$

Thus, sampling $\rho(1, i)$ reduces to the sampling of the correlation matrix $\boldsymbol{\Sigma}_2(1, i)$ from the inverse Wishart density function with parameters n and $\mathbf{S}_2(1, i \mid \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i)$. An algorithm for sampling $\boldsymbol{\Sigma}_2(1, i)$ is provided in Appendix A.3.

Next, we consider the partial correlation $\rho(j - 1, i; 1, 2, \ldots, j - 2)$ between conditional random variables $\Phi^{-1}[F_{j-1}(X_{j-1}; \boldsymbol{\Psi}_{j-1})] \mid \Phi^{-1}[F_l(X_l; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j - 2$ and $\Phi^{-1}[F_i(X_i; \boldsymbol{\Psi}_i)] \mid \Phi^{-1}[F_l(X_l; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j - 2$. Similarly, we choose the conjugate, inverse Wishart prior density function

$$\pi(\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2))$$
$$\propto \boldsymbol{\Sigma}_2 |(j - 1, i; 1, 2, \ldots, j - 2)|^{-3/2}$$

for partial correlation matrix $\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2)$. This leads to the conditional posterior copula density function $p(\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2) \mid \boldsymbol{\Lambda}_j, \mathbf{x})$ of the form

$$\propto |\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2)|^{-(n+3)/2}$$
$$\times \exp\{\text{tr}(-\tfrac{1}{2}\mathbf{S}_2(j - 1, i; 1, 2, \ldots, j - 2 \mid \boldsymbol{\Lambda}_j, \mathbf{x}) \cdot \boldsymbol{\Sigma}_2^{-1}(j - 1, i; 1, 2, \ldots, j - 2))\},$$

where $\mathbf{x}$ is the kn-dimensional vector of the historical data; $\boldsymbol{\Lambda}_j$ is the vector of NORTA parameters $\boldsymbol{\Psi}_m$, $m = 1, 2, \ldots, k$, $\rho(1, m)$, $m = 2, 3, \ldots, k$, and $\rho(l - 1, m; 1, 2, \ldots, l - 2)$, $m = l, l + 1, \ldots, k$, $l = 3, 4, \ldots, j - 1$; and $\mathbf{S}_2(j - 1, i; 1, 2, \ldots, j - 2 \mid \boldsymbol{\Lambda}_j, \mathbf{x})$ is the two-dimensional

matrix defined by

$$\sum_{t=1}^{n} \begin{pmatrix} \dfrac{\Phi^{-1}[F_{j-1}(x_{j-1,t}; \boldsymbol{\Psi}_{j-1})] - \mu_{j-1|1,2,\ldots,j-2}}{\sigma_{j-1|1,2,\ldots,j-2}} \\[1em] \dfrac{\Phi^{-1}[F_i(x_{i,t}; \boldsymbol{\Psi}_i)] - \mu_{i|1,2,\ldots,j-2}}{\sigma_{i|1,2,\ldots,j-2}} \end{pmatrix}$$
$$\cdot \begin{pmatrix} \dfrac{\Phi^{-1}[F_{j-1}(x_{j-1,t}; \boldsymbol{\Psi}_{j-1})] - \mu_{j-1|1,2,\ldots,j-2}}{\sigma_{j-1|1,2,\ldots,j-2}} \\[1em] \dfrac{\Phi^{-1}[F_i(x_{i,t}; \boldsymbol{\Psi}_i)] - \mu_{i|1,2,\ldots,j-2}}{\sigma_{i|1,2,\ldots,j-2}} \end{pmatrix}'.$$

Therefore, sampling $\rho(j - 1, i; 1, 2, \ldots, j - 2)$ reduces to the sampling of the partial correlation matrix $\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2)$ from the inverse Wishart density function with parameters n and $\mathbf{S}_2(j - 1, i; 1, 2, \ldots, j - 2 \mid \boldsymbol{\Lambda}_j, \mathbf{x})$.

An alternative to the use of conjugate, inverse Wishart priors for the (partial) correlations of the copula-vine specification is to use noninformative priors (Liechty et al. 2004). These priors include Jeffreys' prior (Jeffreys 1961), log-matrix prior (Leonard and Hsu 1992), reference prior (Berger and Sun 2008), uniform shrinkage prior (Daniels 1999), and uniform prior (Barnard et al. 2000). Barnard et al. (2000) further note that $\pi(\rho) = (1 - \rho^2)^{(\nu-3)/2}$, which is a beta density function with shape parameters $(\nu - 1)/2$ as well as a uniform density function for $\nu = 3$, can be chosen as a prior density function for $\rho \in [-1, 1]$. Selecting any of these prior density functions for the (partial) correlation requires the use of the MCMC method. Therefore, more computational effort is needed than that required by the sampling of the (partial) correlation from the inverse Wishart density function.

3.3. Sampling All NORTA Parameters

Motivated by the decomposition of the joint pdf in (4) into separate terms associated with the component marginal distributions and the (partial) correlations, we independently choose a prior density function for each of the NORTA parameters. Specifically, we select prior density functions $\pi_i(\boldsymbol{\Psi}_i)$, $i = 1, 2, \ldots, k$ for component parameter vectors $\boldsymbol{\Psi}_i$, $i = 1, 2, \ldots, k$, utilizing the well-established Bayesian literature on standard families of distribution as described in §3.1. Assuming the probabilistic independence of the (partial) correlations, we choose prior density functions $\pi(\boldsymbol{\Sigma}_2(1, i))$, $i = 2, 3, \ldots, k$ for correlation matrices $\boldsymbol{\Sigma}_2(1, i)$, $i = 2, 3, \ldots, k$, and $\pi(\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2))$, $i = j, j + 1, \ldots, k$, $j = 3, 4, \ldots, k$ for partial correlation matrices $\boldsymbol{\Sigma}_2(j - 1, i; 1, 2, \ldots, j - 2)$, $i = j, j + 1, \ldots, k$, $j = 3, 4, \ldots, k$. The selection of the prior density functions for the (partial) correlation matrices is discussed in §3.2.

Because different regular vine specifications are characterized by different (partial) correlations, the probabilistic independence of the partial correlations of the $\mathscr{C}$-vine does not imply the probabilistic independence of the partial correlations of any other regular vine. Therefore, if

the simulation analyst chooses to represent the dependence structure of the NORTA distribution with a different regular vine, then she must assume the probabilistic independence of the partial correlations of that vine. This assumption not only provides flexibility in choosing prior density functions for NORTA's dependence parameters, but also allows the use of the existing literature on Bayesian inference for correlation matrices without being challenged by the high-dimensional nature of the large-scale stochastic simulations.

Assuming the availability of the k-dimensional historical data of length n (i.e., $x_{i,t}$, $i = 1, 2, \ldots, k$, $t = 1, 2, \ldots, n$), we use Bayes' rule for combining the prior density functions with the likelihood function $\prod_{t=1}^{n} h_k(x_{1,t}, x_{2,t}, \ldots, x_{k,t})$, i.e.,

$$\prod_{t=1}^{n} \prod_{i=1}^{k} f_i(x_{i,t}; \boldsymbol{\Psi}_i)$$

$$\times \prod_{t=1}^{n} \prod_{i=2}^{k} \phi_2\big(F_1(x_{1,t}; \boldsymbol{\Psi}_1), F_i(x_{i,t}; \boldsymbol{\Psi}_i); \boldsymbol{\Sigma}_2(1,i)\big)$$

$$\times \prod_{t=1}^{n} \prod_{j=3}^{k} \prod_{i=j}^{k} \phi_2\Big(\Phi_{j-1|1,2,\ldots,j-2}\big(\Phi^{-1}[F_{j-1}(x_{j-1,t}; \boldsymbol{\Psi}_{j-1})]\big|$$

$$\Phi^{-1}[F_l(x_{l,t}; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j-2\big),$$

$$\Phi_{i|1,2,\ldots,j-2}\big(\Phi^{-1}[F_i(x_{i,t}; \boldsymbol{\Psi}_i)]\big|$$

$$\Phi^{-1}[F_l(x_{l,t}; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j-2\big);$$

$$\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)\Big),$$

and obtain the following joint posterior density function:

$$\propto \prod_{i=1}^{k} \overbrace{\left[\pi_i(\boldsymbol{\Psi}_i) \prod_{t=1}^{n} f_i(x_{i,t}; \boldsymbol{\Psi}_i) \right]}^{p_i(\boldsymbol{\Psi}_i | \mathbf{x}_i)}$$

$$\times \prod_{i=2}^{k} \overbrace{\left[\pi(\boldsymbol{\Sigma}_2(1,i)) \prod_{t=1}^{n} \phi_2\big(F_1(x_{1,t}; \boldsymbol{\Psi}_1), F_i(x_{i,t}; \boldsymbol{\Psi}_i); \boldsymbol{\Sigma}_2(1,i)\big) \right]}^{p(\boldsymbol{\Sigma}_2(1,i) | \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i)}$$

$$\times \prod_{j=3}^{k} \prod_{i=j}^{k} \overbrace{\left[\pi(\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)) \right.}^{p(\boldsymbol{\Sigma}_2(j-1,i; 1,2,\ldots,j-2) | \boldsymbol{\Lambda}_j, \mathbf{x})}$$

$$\cdot \prod_{t=1}^{n} \phi_2\Big(\Phi_{j-1|1,2,\ldots,j-2}\big(\Phi^{-1}[F_{j-1}(x_{j-1,t}; \boldsymbol{\Psi}_{j-1})]\big|$$

$$\Phi^{-1}[F_l(x_{l,t}; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j-2\big),$$

$$\Phi_{i|1,2,\ldots,j-2}\big(\Phi^{-1}[F_i(x_{i,t}; \boldsymbol{\Psi}_i)]\big|$$

$$\Phi^{-1}[F_l(x_{l,t}; \boldsymbol{\Psi}_l)], l = 1, 2, \ldots, j-2\big);$$

$$\left. \boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)\big)\right]. \quad (7)$$

Thus, the joint posterior density function of the NORTA parameters is the product of the k marginal posterior density

functions $p_i(\boldsymbol{\Psi}_i | \mathbf{x}_i)$, $i = 1, 2, \ldots, k$; the two-dimensional posterior copula density functions $p(\boldsymbol{\Sigma}_2(1,i) | \boldsymbol{\Psi}_1, \boldsymbol{\Psi}_i, \mathbf{x}_1, \mathbf{x}_i)$, $i = 2, 3, \ldots, k$ associated with the first tree of the $\mathscr{C}$-vine; and the two-dimensional posterior copula density functions $p(\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2) | \boldsymbol{\Lambda}_j, \mathbf{x})$, $i = j$, $j+1, \ldots, k$ associated with the $(j-1)$th tree of the $\mathscr{C}$-vine for $j = 3, 4, \ldots, k$. Appendix B provides an explicit representation of this posterior density function for the 5-dimensional NORTA random vector with exponentially distributed components.

The form of the joint posterior density function in (7) allows us to develop a fast, $(k(k+1)/2)$-stage algorithm for sampling the NORTA parameters. We provide the resulting algorithm in Figure 2. In the first k stages, we sample parameter vectors $\boldsymbol{\Psi}_i$, $i = 1, 2, \ldots, k$ (i.e., $\tilde{\boldsymbol{\Psi}}_i$, $i = 1, 2, \ldots, k$) from posterior density functions $p_i(\boldsymbol{\Psi}_i | \mathbf{x}_i)$, $i = 1, 2, \ldots, k$, as described in §3.1 for exponentially and gamma distributed components. This allows us to account for the uncertainty around the parameters of the marginal distributions of the components. The next $k-1$ stages sample $\boldsymbol{\Sigma}_2(1, i)$, $i = 2, 3, \ldots, k$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2(1, i)$, $i = 2, 3, \ldots, k$) from posterior density functions $p(\boldsymbol{\Sigma}_2(1, i) | \tilde{\boldsymbol{\Psi}}_1, \tilde{\boldsymbol{\Psi}}_i, \mathbf{x}_1, \mathbf{x}_i)$, $i = 2, 3, \ldots, k$ using $\tilde{\boldsymbol{\Psi}}_i$, $i = 1, 2, \ldots, k$ obtained in the first k stages of the algorithm. We sample each of these $k-1$ correlation matrices as described in §3.2 and set the $(1,2)$th entry of $\tilde{\boldsymbol{\Sigma}}_2(1, i)$ to $\tilde{\rho}(1, i)$ for $i = 2, 3, \ldots, k$. This allows us to account for the uncertainty around the (unconditional) correlations of the $\mathscr{C}$-vine specification. In the remaining stages, we capture the uncertainty around the partial correlations. To do this, we first construct vector $\tilde{\boldsymbol{\Lambda}}_j$ with the sampled NORTA parameters; i.e., $\tilde{\boldsymbol{\Psi}}_m$, $m = 1, 2, \ldots, k$, $\tilde{\rho}(1, m)$, $m = 2, 3, \ldots, k$, and $\tilde{\rho}(l-1, m; 1, 2, \ldots, l-2)$, $m = l, l+1, \ldots, k$, $l = 3, 4, \ldots, j-1$. Then, we characterize the conditional normal cdfs $\Phi_{i|1, 2, \ldots, j-2}$ and $\Phi_{j-1|1, 2, \ldots, j-2}$ by obtaining their means $\mu_{i|1, 2, \ldots, j-2}$ and $\mu_{j-1|1, 2, \ldots, j-2}$ and variances $\sigma_{i|1, 2, \ldots, j-2}^2$ and $\sigma_{j-1|1, 2, \ldots, j-2}^2$ from $\tilde{\boldsymbol{\Lambda}}_j$ via Theorem 3.3.4 of Tong (1990) and the recursive formulas of Yule and Kendall (1965). We use the resulting cdfs for determining the posterior density function $p(\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2) | \tilde{\boldsymbol{\Lambda}}_j, \mathbf{x})$. We sample $\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2(j-1, i; 1, 2, \ldots, j-2)$) from this posterior density function as described in §3.2 and set the $(1,2)$th entry of $\tilde{\boldsymbol{\Sigma}}_2(j-1, i; 1, 2, \ldots, j-2)$ to $\tilde{\rho}(j-1, i; 1, 2, \ldots, j-2)$. Repeating this for $i = j$, $j+1, \ldots, k$ and $j = 3, 4, \ldots, k$ allows us to account for the uncertainty around the partial correlations of the $\mathscr{C}$-vine specification. Appendix B provides a detailed implementation of this NORTA parameter-sampling algorithm for the 5-dimensional NORTA distribution with exponentially distributed components.

4. Estimation of the Mean Performance Measure and the Confidence Interval

In this section, we let Y be the performance measure whose mean is relevant to the decision-making process. Our goal

Figure 2.   An algorithm that samples NORTA parameters for capturing parameter uncertainty in a stochastic simulation with k correlated inputs having a k-dimensional NORTA distribution.

> **for** $i = 1, 2, \ldots, k$ **do**
>   sample the ith component parameter vector $\boldsymbol{\Psi}_i$ (i.e., $\tilde{\boldsymbol{\Psi}}_i$) from $p_i(\boldsymbol{\Psi}_i \mid \mathbf{x}_i)$
> **end loop**
> **for** $i = 2, 3, \ldots, k$ **do**
>   sample correlation matrix $\boldsymbol{\Sigma}_2(1, i)$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2(1, i)$) from $p(\boldsymbol{\Sigma}_2(1, i) \mid \tilde{\boldsymbol{\Psi}}_1, \tilde{\boldsymbol{\Psi}}_i, \mathbf{x}_1, \mathbf{x}_i)$
> **end loop**
> **for** $j = 3, 4, \ldots, k$ **do**
>   construct vector $\tilde{\boldsymbol{\Lambda}}_j$ consisting of the sampled NORTA parameters
>   **for** $i = j, j+1, \ldots, k$ **do**
>   (a) compute the means and variances of cdfs $\Phi_{i \mid 1, 2, \ldots, j-2}$ and $\Phi_{j-1 \mid 1, 2, \ldots, j-2}$ using $\tilde{\boldsymbol{\Lambda}}_j$,
>   Theorem 3.3.4 of Tong (1990), and recursive formulas of Yule and Kendall (1965),
>   and insert them into the density function $p(\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2) \mid \tilde{\boldsymbol{\Lambda}}_j, \mathbf{x})$
>   (b) sample partial correlation matrix $\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2(j-1, i; 1, 2, \ldots,$
> $j-2)$) from the posterior density function $p(\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2) \mid \tilde{\boldsymbol{\Lambda}}_j, \mathbf{x})$
>   **end loop**
> **end loop**

is to incorporate the Bayesian model of §3 into the simulation replication algorithm of Chick (2001) and generate a point estimate and a confidence interval of $\mathrm{E}_{Y \mid \mathbf{x}}(Y \mid \mathbf{x})$ (i.e., the posterior mean of the output random variables Y_r, $r = 1, 2, \ldots, R$ of a stochastic simulation with R replications, given the historical input data vector $\mathbf{x}$ and the prior information about the NORTA parameters).

We present the resulting simulation replication algorithm in Figure 3. Step 1 uses our Bayesian model to sample NORTA parameters for each of the R replications of Step 2. We denote the parameters sampled in the rth replication by $\tilde{\boldsymbol{\Psi}}_i^r$, $i = 1, 2, \ldots, k$, $\tilde{\rho}^r(1, i)$, $i = 2, 3, \ldots, k$, and $\tilde{\rho}^r(j-1, i; 1, 2, \ldots, j-2)$, $i = j, j+1, \ldots, k$, $j = 3, 4, \ldots, k$. Step 2 captures both stochastic uncertainty and parameter uncertainty by generating k correlated input variates $\tilde{x}_i^r$, $i = 1, 2, \ldots, k$ from the k-dimensional NORTA distribution with parameters $\tilde{\boldsymbol{\Psi}}_i^r$, $i = 1, 2, \ldots, k$, $\tilde{\rho}^r(1, i)$, $i = 2, 3, \ldots, k$, and $\tilde{\rho}^r(j-1, i; 1, 2, \ldots, j-2)$, $i = j, j+1, \ldots, k$, $j = 3, 4, \ldots, k$ for $r = 1, 2, \ldots, R$. Thus, the key difference of this algorithm from the one presented in Chick (2001) is its first step, where the NORTA parameters are sampled from their Bayesian posterior density functions.

The proper implementation of the simulation replication algorithm requires the consideration of two important issues: the independent sampling of the NORTA parameters for each of the R replications and the analysis of the simulation output data y_r, $r = 1, 2, \ldots, R$ for estimating a point estimate and a confidence interval of $\mathrm{E}_{Y \mid \mathbf{x}}(Y \mid \mathbf{x})$. The first issue arises when the prior density functions chosen for the component parameter vectors and/or the (partial) correlations are not conjugate. In this case, we sample the NORTA parameters using the Gibbs sampler algorithm. When this algorithm is used for generating a distribution parameter from its posterior density function, there often appear autocorrelations between the values sampled within the chain. There might also appear cross correlations between different parameters sampled in different chains. However, the

simulation replication algorithm requires the independent sampling of the NORTA parameters for each of its replications. Therefore, it is important to implement the Gibbs sampler algorithm in a way that it provides statistically independent values of a NORTA parameter for each of the R replications. Appendix C describes how to do this using the method of batching.

There are also cases in which the sampling of the NORTA parameters does not require the use of an MCMC method and thus, the independent sampling of the NORTA parameters is easy to achieve. One such case occurs when the components of the NORTA vector are exponentially distributed and conjugate inverse gamma density functions are chosen as the priors for the scale parameters of the components, whereas conjugate inverse Wishart prior density functions are used for the (partial) correlations. Therefore, we can easily generate R independent sets of NORTA parameters using well-known algorithms for sampling from gamma and Wishart density functions (Appendix A).

Despite obtaining statistically independent sets of NORTA parameters in the first step of the simulation replication algorithm, this is an approximation to the independent sampling of the NORTA parameters when the prior density functions chosen for the component parameter vectors and/or the (partial) correlations are not conjugate. The failure to independently sample NORTA parameters often leads to dependent simulation output data. Therefore, the second issue that might arise in the implementation of the simulation replication algorithm is related to the lack of independence in the output data y_r, $r = 1, 2, \ldots, R$. Law (2007) provides an excellent overview of the methods that have been proposed for the analysis of dependent simulation output data. In Appendix D, we describe how to use the method of batching for analyzing the (dependent) output data y_r, $r = 1, 2, \ldots, R$ to obtain a point estimate and a confidence interval of $\mathrm{E}_{Y \mid \mathbf{x}}(Y \mid \mathbf{x})$.

Figure 3. The simulation replication algorithm that captures stochastic uncertainty and parameter uncertainty, and generates a point estimate and a confidence interval of $\mathrm{E}_{Y|\mathbf{x}}(Y \mid \mathbf{x})$.

Step 1
for $i = 1, 2, \ldots, k$ **do**
    sample R independent variates of component parameter vector $\boldsymbol{\Psi}_i$ (i.e., $\tilde{\boldsymbol{\Psi}}_i^r$, $r = 1, 2, \ldots, R$)
    from $p_i(\boldsymbol{\Psi}_i \mid \mathbf{x}_i)$, independent of the parameters generated for other components
end loop
for $i = 2, 3, \ldots, k$ **do**
  for $r = 1, 2, \ldots, R$ replications **do**
    sample $\boldsymbol{\Sigma}_2(1, i)$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2^r(1, i)$) from $p(\boldsymbol{\Sigma}_2(1, i) \mid \tilde{\boldsymbol{\Psi}}_1^r, \tilde{\boldsymbol{\Psi}}_i^r, \mathbf{x}_1, \mathbf{x}_i)$, independent of the
    correlation matrices generated in other replications, and set $\tilde{\boldsymbol{\Sigma}}_2^r(1, i)[1, 2]$ to $\tilde{\rho}^r(1, i)$
  end loop
end loop
for $j = 3, 4, \ldots, k$ **do**
  for $i = j, j+1, \ldots, k$ **do**
    for $r = 1, 2, \ldots, R$ replications **do**
      sample $\boldsymbol{\Sigma}_2(j-1, i; 1, 2, \ldots, j-2)$ (i.e., $\tilde{\boldsymbol{\Sigma}}_2^r(j-1, i; 1, 2, \ldots, j-2)$) from $p(\boldsymbol{\Sigma}_2(j-1,$
      $i; 1, 2, \ldots, j-2) \mid \tilde{\boldsymbol{\Lambda}}_j^r, \mathbf{x})$, independent of the partial correlation matrices of other
      replications, and set $\tilde{\boldsymbol{\Sigma}}_2^r(j-1, i; 1, 2, \ldots, j-2)[1, 2]$ to $\tilde{\rho}^r(j-1, i; 1, 2, \ldots, j-2)$
    end loop
  end loop
end loop
Step 2
for $r = 1, 2, \ldots, R$ replications **do**
    sample input random variates (i.e., $\tilde{x}_i^r$, $i = 1, 2, \ldots, k$) given NORTA parameters of the
    rth replication (i.e., $\tilde{\boldsymbol{\Psi}}_i^r$, $i = 1, 2, \ldots, k$, $\tilde{\rho}^r(1, i)$, $i = 2, 3, \ldots, k$, $\tilde{\rho}^r(j-1, i; 1, 2, \ldots, j-2)$,
    $i = j, j+1, \ldots, k$, $j = 3, 4, \ldots, k$) and calculate the output y_r as a function of the input
    random variates $\tilde{x}_i^r$, $i = 1, 2, \ldots, k$
end loop
analyze the simulation output data y_r, $r = 1, 2, \ldots, R$ and generate a point estimate and
a confidence interval of $\mathrm{E}_{Y|\mathbf{x}}(Y \mid \mathbf{x})$

5. An Inventory Simulation Example

This section performs a numerical study demonstrating the importance of the joint representation of stochastic and parameter uncertainties in the estimation of the mean line-item fill rates[2] and the confidence intervals of multiproduct inventory simulations with correlated demands. We refer the reader to §5.1 for the experimental design and §5.2 for the results.

5.1. Experimental Design

We consider a periodic-review inventory system with $k \geq 1$ different products whose demands follow a k-dimensional NORTA distribution. We assume the following properties for the true NORTA distribution: (i) The ith product demand has an exponential distribution with a mean of $10(k+1-i)/k$ units for $i = 1, 2, \ldots, k$. (ii) Each (partial) correlation in the $\mathscr{C}$-vine specification of the k-dimensional NORTA demand distribution is 0.30. We let the number of different products, k, take the values of 1, 2, 3, 5, and 10 and manage the inventories of the products with the base-stock policy assuming zero ordering cost and zero lead time. More specifically, we identify the base-stock levels I_i, $i = 1, 2, \ldots, k$ via the use of the single-product models, each of which has a nonstockout probability of 0.90; i.e., $I_i \equiv F_i^{-1}(0.90; 10(k+1-i)/k)$ for $i = 1, 2, \ldots, k$. This

results in a true mean line-item fill rate of 0.90 in each of the k-product inventory simulations.

We assume the availability of the historical demand data of length 100 generated from the true NORTA distribution. We let Y be the line-item fill rate whose mean is relevant to the inventory manager and use $\mathbf{x}$ for denoting the vector of the historical demand data. We implement the simulation replication algorithm as described in §4 for generating a point estimate and a 95% confidence interval of $\mathrm{E}_{Y|\mathbf{x}}[Y \mid \mathbf{x}]$. Our goal is to compare the performances of the point estimates and the confidence intervals obtained from the implementation of our approach to those obtained from stochastic simulations that consider only stochastic uncertainty. We assess the performance of the point estimate using the mean absolute percentage error (MAPE) and the mean square error (MSE), while we evaluate the performance of the confidence interval using the average confidence-interval half-width (HW) and the average coverage probability (CP) (Zouaoui and Wilson 2003).

5.2. Results

Table 1 presents the results obtained when the view of a frequentist is taken and the exponentially distributed product demands are assumed to be independent. Table 2 uses the Bayesian model and presents the results obtained assuming independent product demands, despite the increasing strength of dependence with the number of products.

Table 1. The results obtained via frequentist approach assuming independent demands.

	Mean fill rate		95% confidence interval	
k	MAPE (%)	MSE	HW	CP (%)
$R = 1{,}000$				
1	1.64	3.44×10^{-4}	4.12×10^{-2}	79.90
2	2.01	6.60×10^{-4}	2.93×10^{-2}	73.60
3	2.88	1.21×10^{-3}	2.44×10^{-2}	70.90
5	4.45	2.56×10^{-3}	1.19×10^{-2}	67.10
10	8.98	5.98×10^{-3}	1.35×10^{-2}	49.40
$R = 5{,}000$				
1	1.53	3.40×10^{-4}	1.34×10^{-2}	59.50
2	1.94	6.30×10^{-3}	9.43×10^{-3}	55.20
3	2.60	1.12×10^{-3}	7.78×10^{-3}	52.30
5	4.15	2.13×10^{-3}	6.05×10^{-3}	44.80
10	8.41	5.73×10^{-3}	4.27×10^{-3}	31.70
$R = 10{,}000$				
1	1.52	3.31×10^{-4}	4.26×10^{-3}	57.10
2	1.73	3.61×10^{-4}	3.02×10^{-3}	53.20
3	2.37	8.91×10^{-4}	2.47×10^{-3}	48.70
5	4.12	2.09×10^{-3}	1.91×10^{-3}	41.90
10	8.36	5.29×10^{-3}	1.32×10^{-3}	30.20

Finally, Table 3 presents the results obtained when both the view of the Bayesian is taken and the stochastic dependencies among the product demands are considered. Each of these tables reports the results for three different values of R (i.e., run length) of the simulation replication algorithm: 1,000, 5,000, and 10,000. As the run length increases, we observe that the average confidence-interval half-width approaches zero, whereas the accuracy in the estimation of the mean absolute percentage error and the

Table 2. The results obtained via Bayesian approach assuming independent demands.

	Mean fill rate		95% confidence interval	
k	MAPE (%)	MSE	HW	CP (%)
$R = 1{,}000$				
1	1.48	3.06×10^{-4}	4.86×10^{-2}	88.35
2	1.77	5.70×10^{-4}	3.05×10^{-2}	83.12
3	2.12	9.70×10^{-4}	2.97×10^{-2}	79.88
5	3.24	1.85×10^{-3}	1.82×10^{-2}	76.10
10	5.63	3.47×10^{-3}	2.01×10^{-2}	68.24
$R = 5{,}000$				
1	1.38	2.90×10^{-4}	3.26×10^{-2}	87.29
2	1.72	5.30×10^{-4}	2.41×10^{-2}	82.18
3	1.94	6.00×10^{-4}	9.03×10^{-3}	78.57
5	3.08	1.52×10^{-3}	6.82×10^{-3}	73.36
10	5.34	2.99×10^{-3}	5.94×10^{-3}	66.87
$R = 10{,}000$				
1	1.35	2.73×10^{-4}	5.12×10^{-3}	82.83
2	1.68	3.46×10^{-4}	4.02×10^{-3}	79.52
3	1.88	5.87×10^{-4}	2.98×10^{-3}	78.04
5	3.01	1.50×10^{-3}	1.62×10^{-3}	72.96
10	5.14	2.82×10^{-3}	2.62×10^{-3}	64.24

Table 3. The results obtained via Bayesian approach assuming correlated demands.

	Mean fill rate		95% confidence interval	
k	MAPE (%)	MSE	HW	CP (%)
$R = 1{,}000$				
2	1.32	3.27×10^{-4}	5.20×10^{-2}	88.76
3	2.08	6.54×10^{-4}	3.98×10^{-2}	84.02
5	2.67	1.17×10^{-3}	3.01×10^{-2}	81.45
10	4.34	2.32×10^{-3}	2.62×10^{-2}	80.00
$R = 5{,}000$				
2	1.26	3.01×10^{-4}	3.65×10^{-2}	87.93
3	1.83	5.78×10^{-4}	9.87×10^{-3}	84.00
5	2.19	6.67×10^{-4}	8.05×10^{-3}	80.43
10	3.92	2.07×10^{-3}	6.13×10^{-3}	78.12
$R = 10{,}000$				
2	1.23	2.92×10^{-4}	7.02×10^{-3}	84.72
3	1.82	5.76×10^{-4}	4.72×10^{-3}	82.61
5	2.12	6.54×10^{-4}	3.61×10^{-3}	79.28
10	3.41	2.02×10^{-3}	3.06×10^{-3}	77.83

mean square error increases. We also observe that the use of our model allows the simulation analyst to obtain point estimates with lower mean absolute percentage errors and confidence intervals with higher coverage than those of the stochastic simulations that account only for stochastic uncertainty.

Specifically, the comparison of the results tabulated in Table 1 to those in Table 2 and Table 3 shows that the point estimator accuracy for the Bayesian approach is better than the point estimator accuracy for the frequentist approach. The mean absolute percentage error is 1.23% in the 2-product setting, whereas it is 2.12% in the 5-product setting for the Bayesian approach with a run length of 10,000 replications (Table 3). On the other hand, for the frequentist approach the mean absolute percentage errors are 1.73% and 4.12% in the 2-product and 5-product settings (Table 1). Although the average confidence-interval half-widths are tighter than their Bayesian counterparts, the frequentist approach delivers decreasing coverage probabilities with increasing number of products. Because the confidence intervals of the frequentist approach are centered on biased estimates of the mean line-item fill rate, the confidence-interval coverage eventually drops to zero as the number of products increases.

On the other hand, the confidence intervals based on the Bayesian approach, even under the assumption of independent product demands, show much higher coverage probabilities as they account for the uncertainty around the parameters of the component marginal distributions as well as the stochastic uncertainty. We find that the mean absolute percentage error is 3.01% and the coverage probability is 72.96% in the 5-product setting under the assumption of independent demands (Table 2). However, the mean absolute percentage error increases to 5.14% and the coverage probability decreases to 64.24% in the 10-product setting,

while accounting for the correlations among the product demands results in a mean absolute percentage error of 3.41% and a coverage probability of 77.83%. Thus, the consideration of the demand correlations further improves both the mean absolute percentage error of the point estimates and the coverage probability of the confidence intervals, especially as the number of products increases.

6. Conclusion

We consider a large-scale stochastic simulation whose correlated inputs have a NORTA distribution with arbitrary continuous marginal distributions. We investigate how to account for stochastic and parameter uncertainties in the estimation of the mean performance measure and the confidence interval of this simulation. Utilizing Sklar's marginal-copula representation together with Cooke's copula-vine specification, we develop a Bayesian model for the fast sampling of the parameters of the NORTA distribution. The development of such a Bayesian model, which enables simulation analysts to capture parameter uncertainty in stochastic simulations with correlated inputs, is the primary contribution of this paper to the discrete-event stochastic simulation literature. We incorporate the Bayesian model into the simulation replication algorithm for the joint representation of stochastic uncertainty and parameter uncertainty in the computation of the mean performance estimate and the confidence interval.

We demonstrate the effectiveness of the Bayesian model in decreasing the mean absolute percentage error of the mean line-item fill-rate estimate and increasing the coverage of the confidence interval in a multiproduct inventory simulation with correlated stochastic demands. It may be possible to improve the performance of the mean line-item fill-rate estimate and the confidence interval further with additional multivariate demand data. However, the collection of data for many inputs of a large-scale stochastic simulation might be a challenging task, in which case it becomes necessary to focus the data-collection effort on the inputs with significant impact on the performance. A way of achieving this is to decompose the total variation of the simulation output into distinct terms representing stochastic uncertainty and parameter uncertainty. Such decomposition of the output variance in simulations with correlated inputs is the subject of ongoing work.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Endnotes

1. A prior density function is said to be conjugate to a likelihood function if the resulting posterior density function is in the same family of distributions as the prior density function (Bernardo and Smith 1994).

2. The line-item fill rate compares the number of different products shipped complete to the number of different products demanded. The use of the line-item fill rate, which is joint across products, is common in settings where the demands for the items can be correlated as they are frequently used in sets.

Acknowledgments

References

Barnard, J., R. McCulloch, X. Meng. 2000. Modeling covariance matrices in terms of standard deviations and correlations with applications to shrinkage. *Statistica Sinica* **10**(4) 1281–1311.

Barton, R. R., L. W. Schruben. 2001. Resampling methods for input modeling. B. A. Peters, J. S. Smith, D. J. Medeiros, M. W. Rohrer, eds. *Proc. 2001 Winter Simulation Conf.*, IEEE, Piscataway, NJ, 372–378.

Bedford, T., R. M. Cooke. 2001. *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, Cambridge, UK.

Bedford, T., R. M. Cooke. 2002. Vines—A new graphical model for dependent random variables. *Ann. Statist.* **30**(4) 1031–1068.

Berger, J. O., D. Sun. 2008. Objective priors for the bivariate normal model. *Ann. Statist.* **36**(2) 963–982.

Bernardo, J. M., A. F. M. Smith. 1994. *Bayesian Theory*. John Wiley & Sons, Inc., New York.

Biller, B., S. Ghosh. 2006. Multivariate input processes. B. L. Nelson, S. G. Henderson, eds. *Simulation*. Handbooks in Operations Research and Management Science. Elsevier Science, Amsterdam.

Cario, M. C., B. L. Nelson. 1997. Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Working paper, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.

Carlin, B. P., T. A. Louis. 2000. *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman & Hall, London.

Cheng, R. C. H., W. Holland. 1997. Sensitivity of computer simulation experiments to errors in input data. *J. Statist. Comput. Simulation* **57**(1–4) 327–335.

Cheng, R. C. H., W. Holland. 1998. Two-point methods for assessing variability in simulation output. *J. Statist. Comput. Simulation* **60**(3) 183–205.

Cheng, R. C. H., W. Holland. 2004. Calculation of confidence intervals for simulation output. *ACM Trans. Modeling Comput. Simulation* **14**(4) 344–362.

Chick, S. E. 1997. Bayesian analysis for simulation input and output. S. Andradottir, K. J. Healy, D. H. Withers, B. L. Nelson, eds. *Proc. 1997 Winter Simulation Conf.*, IEEE, Piscataway, NJ, 253–260.

Chick, S. E. 1999. Steps to implement Bayesian input distribution selection. P. A. Farrington, H. B. Nembhard, D. T. Sturrock, G. W. Evans, eds. *Proc. 1999 Winter Simulation Conf.*, IEEE, Piscataway, NJ, 317–324.

Chick, S. E. 2001. Input distribution selection for simulation experiments: Accounting for input uncertainty. *Oper. Res.* **49**(5) 744–758.

Cooke, R. M. 1997. Markov and entropy properties of tree- and vine-dependent variables. *Proc. ASA Section Bayesian Statist. Sci.*, American Statistical Association, Alexandria, VA, 166–175.

Daniels, M. J. 1999. A prior for the variance in hierarchical models. *Canadian J. Statist.* **27**(3) 567–578.

Draper, D. 1995. Assessment of propagation of model uncertainty (with discussion). *J. Roy. Statist. Soc. Ser. B* **56**(1) 501–514.

Gelfand, A., A. Smith. 1990. Sampling based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85**(410) 398–409.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Rubin. 2000. *Bayesian Data Analysis*. Chapman & Hall, London.

Geman, S., A. Geman. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intelligence* **6**(6) 721–740.

George, E. I. 1999. Bayesian model selection. S. Kotz, C. Read, D. Banks, eds. *Encyclopedia of Statistical Sciences*, Vol. 3. Wiley, New York, 39–46.

Gilks, W. R., S. Richardson, D. J. Spiegelhalter. 1996. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, London.

Helton, J. C. 1997. Uncertainty and sensitivity analysis in the presence of stochastic and subjective uncertainty. *J. Statist. Comput. Simulation* **57**(1–4) 3–76.

Hoeting, J. A., D. Madigan, A. E. Raftery, C. T. Volinsky. 1999. Bayesian model averaging: A tutorial. Technical Report 9814, Department of Statistics, Colorado State University, Fort Collins, CO.

Jeffreys, H. 1961. *Theory of Probability*. Oxford University Press, Oxford, UK.

Joe, H. 1997. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.

Kass, R. E., L. Wasserman. 1996. The selection of prior distributions by formal rules. *J. Amer. Statist. Assoc.* **91**(435) 1343–1370.

Kouvelis, P., P. Su. 2007. *The Structure of Global Supply Chains*. Research monograph for Foundations and Trends® in Technology, Information, and Operations Management. Now Publishers, Boston.

Kurowicka, D., R. Cooke. 2003. A parameterization of positive definite matrices in terms of partial correlation vines. *Linear Algebra Appl.* **372**(1) 225–251.

Kurowicka, D., R. Cooke. 2006. *Uncertainty Analysis with High Dimensional Dependence Modeling*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ.

Law, A. M. 2007. *Simulation Modeling and Analysis*, 4th ed. McGraw-Hill, New York.

Leonard, T., J. S. Hsu. 1992. Bayesian inference for a covariance matrix. *Ann. Statist.* **20**(4) 1669–1696.

Liechty, J. C., M. W. Liechty, P. Muller. 2004. Bayesian correlation estimation. *Biometrika* **91**(1) 1–14.

Morales, O., D. Kurowicka, A. Roelen. 2006. Elicitation procedures for conditional and unconditional rank correlations. *Resources for the Future, Expert Judgement Policy Symposium and Technical Workshop*, Washington, DC, March 13–14.

Nelsen, R. B. 2006. *An Introduction to Copulas*. Springer-Verlag, New York.

Raftery, A. E., D. Madigan, C. T. Volinsky. 1996. Accounting for model uncertainty in survival analysis improves predictive performance (with discussion). J. M. Bernardo, J. O. Berger, A. P. David, A. F. M. Smith, eds. *Bayesian Statistics 5*. Oxford University Press, Oxford, UK, 323–349.

Rossi, P. E., G. M. Allenby, R. McCulloch. 2006. *Bayesian Statistics and Marketing*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ.

Schweizer, B. 1991. Thirty years of copulas. G. Dall'Aglio, S. Kotz, G. Salinetti, eds. *Advances in Probability Distributions with Given Marginals*. Kluwer, Dordrecht, The Netherlands, 13–50.

Son, Y. S., M. Oh. 2006. Bayesian estimation of the two-parameter gamma distribution. *Comm. Statist. - Simulation Comput.* **35**(2) 285–293.

Tong, Y. L. 1990. *The Multivariate Normal Distribution*. Springer-Verlag, New York.

Xu, S. H. 1999. Structural analysis of a queueing system with multiclasses of correlated arrivals and blocking. *Oper. Res.* **47**(2) 264–276.

Yule, G., M. G. Kendall. 1965. *An Introduction to the Theory of Statistics*. Charles Griffin & Company, Belmont, CA.

Zouaoui, F., J. R. Wilson. 2003. Accounting for parameter uncertainty in simulation input modeling. *IIE Trans.* **35**(9) 781–792.

Zouaoui, F., J. R. Wilson. 2004. Accounting for input-model and input-parameter uncertainties in simulation. *IIE Trans.* **36**(11) 1135–1151.