

NEAR OPTIMALITY GUARANTEES FOR DATA-DRIVEN NEWSVENDOR WITH TEMPORALLY DEPENDENT DEMAND: A MONTE CARLO APPROACH

Alp Akcay

Department of Industrial Engineering
Bilkent University
06800, Ankara, TURKEY

Bahar Biller
Sridhar Tayur

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

ABSTRACT

We consider a newsvendor problem with stationary and temporally dependent demand in the absence of complete information about the demand process. The objective is to compute a probabilistic guarantee such that the expected cost of an inventory-target estimate is arbitrarily close to the expected cost of the optimal critical-fractile solution. We do this by sampling dependent uniform random variates matching the underlying dependence structure of the demand process – rather than sampling the actual demand which requires the specification of a marginal distribution function – and by approximating a lower bound on the probability of the so-called near optimality. Our analysis sheds light on the role of temporal dependence in the resulting probabilistic guarantee, which has been only investigated for independent and identically distributed demand in the inventory management literature.

1 INTRODUCTION

An important problem in inventory management is to set inventory targets in the absence of complete information about the demand generating process. In this paper, we address this problem in a newsvendor setting for stationary and temporally dependent demand. It is well known that the *critical fractile solution* is optimal when the distribution of the demand, conditional on past demand realizations, is known. In practice, however, this distribution is unknown and must be estimated using only a finite (and sometimes, very limited) amount of historical demand data. Consequently, the expected cost associated with an *estimate* of the critical fractile solution can be far from the minimum expected cost.

The dependence of demand on its past realizations is not uncommon in practice. Erkip, Hausman, and Nahmias (1990) find that the autocorrelation in monthly demand can reach up to 0.70 in a consumer products company. Lee, Padmanabhan, and Whang (1997) report dependence between demand realizations over time – especially positive autocorrelation – in high-tech and grocery industries. Lee, So, and Tang (2000) show that the sales data of 91% of the items in a supermarket have autocorrelations between 0.26 and 0.89. Similarly, Hosoda et al. (2008) analyze the sales data of soft drink products at a grocery retailer and find that the autocorrelation in weekly demand varies between 0.77 and 0.83.

A widely used method to model autocorrelation is to construct the temporally dependent process via classical time series. For example, a linear autoregressive (AR) process with a normally distributed random shock is often used for demand modeling in inventory management and supply chain analysis; e.g., Lee, So, and Tang (2000), Luong (2007), Chen and Lee (2012). In this case, an AR process can be expressed as a linear combination of independent and normally distributed random shocks, implying that the marginal demand distribution is normal. However, a normal distribution often falls short of an adequate representation of the demand distribution, leading to inaccurate prediction models and poor operational performance.

For autocorrelated demand, as in this paper, a considerable amount of effort has been devoted to modeling time series with exponential, gamma, geometric, or general discrete marginal distributions; e.g.,

Block, Langberg, and Stoffer (1990), Tiku, Wong, Vaughan, and Bian (2002), Akkaya and Tiku (2005), Gouieroux and Jasiak (2006), and Jose, Lishamol, and Sreekumar (2008). Nevertheless, these models often allow only limited control of the dependence structure for a given marginal demand distribution. In addition, there is no one-size-fit-all solution for modeling and estimation of time-series processes. We overcome these challenges by (i) modeling the temporal dependence in a transformed demand process which has a standard uniform marginal distribution and matches the underlying dependence structure of the actual demand process, and (ii) capturing the demand's distributional shape with the empirical demand distribution function. To be specific, we model the temporal dependence in $\{F(X_t); t = 1, 2, \dots\}$ rather than modeling the temporal dependence in the actual demand process $\{X_t; t = 1, 2, \dots\}$, where F is the marginal cumulative distribution function (cdf) of the demand process. We do this by using a *copula* which allows a decision maker to avoid any restrictive assumption on the functional form of F . To estimate the critical fractile solution, the decision maker first obtains the empirical demand distribution function from the temporally dependent demand data and then uses it in lieu of the true marginal demand distribution while estimating the copula parameters that characterize the temporal dependence. Finally, an estimate of the critical fractile solution is obtained as a function of the empirical demand distribution and the estimated values of the copula parameters.

A natural question to ask is as follows. *How good is an inventory target estimated in this way?* We define the *goodness* of an inventory-target estimate as in Levi, Roundy, and Shmoys (2007) by using the notion of ε -optimality; i.e., an inventory-target estimate is ε -optimal if its expected cost is at most $1 + \varepsilon$ of the minimum but unknown expected cost. Clearly, the ε -optimality of an inventory-target estimator is a random event because the inventory-target estimator is a random variable as a result of being a function of the historical data randomly generated by the true demand process. In this paper, we consider all the possible realizations of a demand history, and hence, all the realizations of an inventory-target estimator in a frequentist framework. Accordingly, we obtain a lower bound on the probability of ε -optimality when the inventory target is a function of the empirical demand distribution and copula-parameter estimators. This lower bound, which we refer to as *a probabilistic guarantee for near optimality*, serves as a level of confidence for the decision maker to ensure the ε -optimality of an inventory target obtained from a limited amount of demand data.

We propose a sampling-based method to compute this probabilistic guarantee independent of any specific marginal demand distribution and for an arbitrary choice of copula. In particular, our method builds on the idea of sampling dependent uniform random variates matching the underlying dependence structure of the demand process rather than sampling the actual demand which requires the specification of the marginal distribution. We believe that our copula-based demand model and the sampling-based method of computing a probabilistic guarantee for near optimality have application areas not only in inventory management but essentially in any decision problem in which an overage and underage trade-off exists and the random realizations observed in consecutive time periods depend on each other.

The paper is organized as follows. Section 2 presents our copula-based demand model and the solution of the newsvendor problem with complete information about the temporally dependent demand process. Section 3 provides a two-step estimation method consistent with the marginal-copula representation of demand. Section 4 characterizes the ε -optimality for temporally dependent demand and presents our sampling-based method to compute a probabilistic guarantee for ε -optimality. We conclude with a summary of the paper in Section 5.

2 MODELING FRAMEWORK

Section 2.1 introduces our copula-based demand model. Section 2.2 presents an overview of the newsvendor model and characterizes the optimal critical-fractile solution when the marginal demand distribution and the temporal dependence structure are known.

2.1 Demand Process

We let the demand $\{X_t; t = 1, 2, \dots\}$ be a stationary first-order Markov process; i.e.,

$$\mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t, X_{t-1} = x_{t-1}, \dots, X_1 = x_1) = \mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t).$$

In contrast to using a conditional probability distribution directly, we consider an alternative approach based on *copulas* to model the first-order temporal dependence. We do this by constructing the joint distribution of transformed demand random variables $U_t = F(X_t)$ and $U_{t+1} = F(X_{t+1})$ by a bivariate copula $C(\cdot, \cdot; \theta)$ with $\theta \in \Theta$, where the parameter space Θ is a compact subset of \mathbb{R}^d and d is the number of parameters characterizing the underlying dependence structure.

Definition 1 (Nelsen 2006) A bivariate copula is a function $C : [0, 1] \times [0, 1] \rightarrow [0, 1]$; i.e., a mapping of the unit square into the unit interval with the following properties:

- (i) For every (u_t, u_{t+1}) in $[0, 1] \times [0, 1]$, $C(u_t, 0; \theta) = C(0, u_{t+1}; \theta) = 0$, and $C(u_t, 1; \theta) = u_t$ and $C(1, u_{t+1}; \theta) = u_{t+1}$;
- (ii) For every $u'_t, u''_t, u'_{t+1}, u''_{t+1}$ in $[0, 1]$ such that $u'_t \leq u''_t$ and $u'_{t+1} \leq u''_{t+1}$, $C(u''_t, u''_{t+1}; \theta) - C(u'_t, u''_{t+1}; \theta) - C(u''_t, u'_{t+1}; \theta) + C(u'_t, u'_{t+1}; \theta) \geq 0$.

The first condition provides the lower bound on the distribution function and ensures that a bivariate copula is a two-dimensional distribution function with standard uniform marginal distributions. The second condition guarantees the probability of observing a point in $[0, 1] \times [0, 1]$ to be nonnegative. The use of copulas for representing the joint distribution of a random vector has been studied extensively for the past two decades. We refer the reader to Joe (1997) and Nelsen (2006) for the widely known properties of the copulas. The use of copulas for modeling the temporal dependence of time series data has recently gained attention; see Chen and Fan (2006), Beare (2010), and Patton (2012) as the example studies.

Sklar's theorem (Sklar 1959) and the stationarity of the demand process allow us to construct the distribution of demand random variables X_t and X_{t+1} with a marginal distribution function F and a copula $C(\cdot, \cdot; \theta)$:

Theorem 1 (Sklar's Theorem) Let H be a bivariate distribution function with the continuous marginal cdf F . Then, there exists a bivariate unique copula C such that $H(x_t, x_{t+1}) = C(F(x_t), F(x_{t+1}); \theta)$ for all $(x_t, x_{t+1}) \in \mathbb{R}^2$. Conversely, if C is a bivariate copula and the marginal cdf F is continuous, then the function $C(F(x_t), F(x_{t+1}); \theta)$ is a bivariate distribution function with marginal cdf F .

The advantage of using copula for demand modeling is that we have the freedom to specify the marginal demand distribution and the dependence structure separately. That is, we can choose any arbitrary continuous demand distribution, link the consecutive demand random variables with a copula, and obtain a legitimate bivariate distribution to characterize the first-order time series. We only impose the following technical assumption on the dependence structure:

Assumption 1 The copula $C(u_t, u_{t+1}; \theta)$ is absolutely continuous with respect to Lebesgue measure on $[0, 1] \times [0, 1]$, and is neither the Fréchet-Hoeffding upper bound (i.e., $C(u_t, u_{t+1}; \theta) \neq \min(u_t, u_{t+1})$) nor the Fréchet-Hoeffding lower bound (i.e., $C(u_t, u_{t+1}; \theta) \neq \max(u_t + u_{t+1} - 1, 0)$).

Assumption 1 is standard in the context of dependence modeling to rule out the deterministic cases of $X_t = X_{t-1}$ for the upper bound and $X_t = F^{-1}(1 - F(X_{t-1}))$ for the lower bound. We now illustrate our method of demand modeling in Example 1 by using the *normal copula*, which encodes the dependence precisely the same way a bivariate normal distribution does; we refer the reader to Joe (1997) for other copulas that enable us to capture any form of dependence structure.

Example 1 (Normal Copula). Let Φ be the standard normal cdf and $\Phi_2(\cdot, \cdot; \theta)$ be the standard normal bivariate cdf with correlation $\theta \in (-1, 1)$. The normal copula is defined as

$$C(u_t, u_{t+1}; \theta) = \Phi_2(\Phi^{-1}(u_t), \Phi^{-1}(u_{t+1}); \theta).$$

As a result of representing a standard normal random variable Z_t as $\Phi^{-1}(U_t)$ from a standard uniform random variable $U_t = F(X_t)$, it can be easily seen that normal copula constructs a valid joint distribution for demand random variables X_t and X_{t+1} :

$$\begin{aligned} \Phi_2(\Phi^{-1}(u_t), \Phi^{-1}(u_{t+1}); \theta) &= \mathbb{P}(Z_t \leq \Phi^{-1}(u_t), Z_{t+1} \leq \Phi^{-1}(u_{t+1})) \\ &= \mathbb{P}(\Phi^{-1}(U_t) \leq \Phi^{-1}(u_t), \Phi^{-1}(U_{t+1}) \leq \Phi^{-1}(u_{t+1})) \\ &= \mathbb{P}(U_t \leq u_t, U_{t+1} \leq u_{t+1}) \\ &= \mathbb{P}(F(X_t) \leq F(x_t), F(X_{t+1}) \leq F(x_{t+1})) \\ &= \mathbb{P}(X_t \leq x_t, X_{t+1} \leq x_{t+1}) \end{aligned}$$

The last equation follows because an inequality still holds after applying a monotonically increasing function to both sides of the inequality.

It is worth noting that the dependence structure of a normal copula implies that the demand process $\{X_t; t = 1, 2, \dots\}$ accepts an Autoregressive-To-Anything (ARTA) process representation of Cario and Nelson (1996). In Example 1, the ARTA process first takes an AR(1) model $Z_t = \theta Z_{t-1} + Y_t$ with standard normal Z_t and normally distributed independent error term Y_t with mean zero and variance $1 - \theta^2$, and then obtains the demand random variable X_t via the transformation $X_t = F^{-1}(\Phi(Z_t))$. Clearly, the use of normal copula together with normal marginal distribution further reduces the demand process to a classical AR model.

There is a one-to-one mapping between the normal copula parameter θ – which is the autocorrelation coefficient in the AR(1) model described above – and the autocorrelation of the demand process $\{X_t; t = 1, 2, \dots\}$. We refer the reader to Cario and Nelson (1996) for details of the so-called correlation matching problem. In this paper, we focus on estimating the copula parameter directly.

2.2 Newsvendor Model

The decision maker aims to set the correct number of units in stock to meet the unknown demand in period $n + 1$. The period starts with zero inventory on hand. Ordering too few incurs a shortage cost of b per unit short, while ordering too many incurs a holding cost of h per unit over. The goal is to minimize the sum of expected shortage and holding costs *conditional* on the most recent demand realization since the demand is known to be a first-order Markov process.

The resulting objective function is convex and minimized by the critical fractile solution (Porteus 2002). The critical fractile solution requires the knowledge of the distribution of demand X_{n+1} conditional on $X_n = x_n$, which can be derived directly from the copula as follows:

$$\begin{aligned} \mathbb{P}(X_{n+1} \leq x_{n+1} | X_n = x_n) &= \mathbb{P}(U_{n+1} \leq u_{n+1} | U_n = u_n) \\ &= \lim_{\delta \rightarrow 0} \frac{\mathbb{P}(U_{n+1} \leq u_{n+1}, U_n \in (u_n - \delta, u_n + \delta))}{\mathbb{P}(U_n \in (u_n - \delta, u_n + \delta))} \\ &= \lim_{\delta \rightarrow 0} \frac{C(u_n + \delta, u_{n+1}; \theta) - C(u_n - \delta, u_{n+1}; \theta)}{2\delta} \\ &= \left. \frac{\partial C(u_n, F(x_{n+1}); \theta)}{\partial u_n} \right|_{u_n = F(x_n)}. \end{aligned}$$

We denote this conditional distribution by $C_{2|1}(F(\cdot) | F(x_n); \theta)$. The critical fractile solution q^* is then the value of q that solves the first-order condition

$$C_{2|1}(F(q) | F(x_n); \theta) - \frac{b}{h + b} = 0, \tag{1}$$

and it can be written as $q^* = F^{-1}(C_{2|1}^{-1}(b/(h + b) | F(x_n); \theta))$. We next discuss a two-step method to estimate the unknown marginal distribution function F and the d -dimensional copula-parameter vector θ .

3 MODEL ESTIMATION

The marginal-copula representation of the demand process allows the marginal distribution function and copula parameters to be estimated separately, potentially with different methods. In this study, we focus on the case where the decision maker estimates the marginal distribution function with the empirical demand distribution. Subsequently, the copula parameters are estimated via the maximum likelihood method by treating the empirical demand distribution as the true marginal distribution. This method is known as *semiparametric estimation* as well as *canonical maximum likelihood estimation* (Patton 2012).

The joint density function of the demand random variables X_t and X_{t+1} is given by

$$\frac{\partial^2 C(F(x_t), F(x_{t+1}); \theta)}{\partial x_t \partial x_{t+1}} = f(x_t) f(x_{t+1}) c(F(x_t), F(x_{t+1}); \theta),$$

where $f(\cdot)$ is the marginal probability density function (pdf) and $c(u, v; \theta) := \partial^2 C(u, v; \theta) / (\partial u \partial v)$ is the copula density function. Therefore, the density function of X_{t+1} conditional on $X_t = x_t$, which we denote with $f_{2|1}(x_{t+1}|x_t)$, reduces to $f(x_{t+1})c(F(x_t), F(x_{t+1}); \theta)$. The likelihood function of the historical demand data $\{x_t; t = 1, 2, \dots, n\}$ then takes the form

$$f(x_1) \prod_{t=1}^{n-1} f_{2|1}(x_{t+1}|x_t) = \prod_{t=1}^n f(x_t) \prod_{t=1}^{n-1} c(F(x_t), F(x_{t+1}); \theta). \quad (2)$$

It is worth noting that the copula density function of independent random variables takes the value of one and the likelihood function in (2) reduces to the likelihood function $\prod_{t=1}^n f(x_t)$ of independent and identically distributed data.

In the first step of semiparametric estimation, the decision maker estimates the marginal distribution function $F(x)$ by using the empirical demand distribution function $F_n(x) := (1/n) \sum_{t=1}^n \mathbb{1}(X_t \leq x)$, where $\mathbb{1}(\cdot)$ is the indicator function. In the second step, the copula-parameter vector θ is estimated by maximizing the log-likelihood function after replacing the unknown F by F_n and ignoring the terms that do not depend on θ :

$$\hat{\theta} = \operatorname{argmax}_{\theta \in \Theta} \sum_{t=1}^{n-1} \log c(F_n(x_t), F_n(x_{t+1}); \theta).$$

Finally, the decision maker estimates the critical fractile solution q^* by replacing the marginal distribution function F by F_n and the copula parameters θ by $\hat{\theta}$ in the functional form of q^* . The inventory-target estimator, which is a function of the demand random variables $\{X_t; t = 1, 2, \dots, n\}$, then takes the form

$$\hat{Q} = F_n^{-1} \left(C_{2|1}^{-1} \left(\frac{b}{h+b} | F_n(X_n); \hat{\theta} \right) \right), \quad (3)$$

where

$$F_n^{-1}(\tau) = \min_{j=1,2,\dots,n} \left\{ X_j : \frac{1}{n} \sum_{t=1}^n \mathbb{1}(X_t \leq X_j) \geq \tau \right\}.$$

The consistency of the semiparametric copula-parameter estimator $\hat{\theta}$ (Chen and Fan 2006) and the convergence of the empirical demand distribution to the true marginal distribution imply that, when there is a large number of demand observations, the inventory-target estimator \hat{Q} approaches the critical fractile solution with full knowledge of copula parameters and marginal distribution function. In practice, a decision maker is rarely fortunate enough to observe a large number of demand observations. In contrast, the demand history can be very short, casting doubt on the performance of the inventory-target estimator \hat{Q} in minimizing the sum of expected inventory holding and shortage costs. In this paper, we focus on this expected cost associated with \hat{Q} for *finite* number of demand observations, which is a more relevant measure for the decision maker than the asymptotic properties of the estimators F_n and $\hat{\theta}$.

4 ε -OPTIMALITY IN INVENTORY-TARGET ESTIMATION

We let \hat{q} denote the realization of the inventory-target estimator \hat{Q} from the historical demand data $\{x_t; t = 1, 2, \dots, n\}$ of length n . The expected cost associated with \hat{q} is given by

$$L(\hat{q}|x_n) := h \int_{-\infty}^{\hat{q}} (\hat{q} - x_{n+1}) f_{2|1}(x_{n+1}|x_n) dx_{n+1} + b \int_{\hat{q}}^{\infty} (x_{n+1} - \hat{q}) f_{2|1}(x_{n+1}|x_n) dx_{n+1}.$$

The decision maker does not necessarily achieve the minimum expected cost $L(q^*|x_n)$ by using the inventory-target estimate \hat{q} especially when n is small. In this study, we are interested in finding a lower bound to the probability of the difference $L(\hat{q}|x_n) - L(q^*|x_n)$ not exceeding a certain threshold for *all* possible realizations of the demand history. To this end, we let $0 < \varepsilon \leq 1$ and define the ε -optimality in inventory-target estimation as follows:

Definition 2 The inventory-target estimate \hat{q} is ε -optimal if its expected cost is at most $1 + \varepsilon$ of the minimum expected cost; i.e., $L(\hat{q}|x_n) \leq (1 + \varepsilon)L(q^*|x_n)$.

Considering all possible realizations of the demand observation in period n , we let Q^* denote the critical fractile solution which has the functional form $F^{-1}(C_{2|1}^{-1}(b/(h+b)|F(X_n); \theta))$. We aim to provide a lower bound to the probability of the event

$$\left[L(\hat{Q}|X_n) \leq (1 + \varepsilon)L(Q^*|X_n) \right]$$

to measure the quality of the inventory-target estimator \hat{Q} in a frequentist framework. If this probability is sufficiently large, then \hat{Q} can be used confidently even though it is estimated from a finite number of demand observations. This is because the expected cost of its realization (before the realization of demand in period $n + 1$) cannot be more than $1 + \varepsilon$ of the minimum expected cost with high probability for *any* historical demand data.

We organize the remainder of this section as follows. Section 4.1 generalizes the characterization of ε -optimality for temporally dependent demand. Section 4.2 proposes a sampling-based method to compute a lower bound to the probability of ε -optimality. Section 4.3 provides insights shedding light on the role of temporal dependence in inventory-target estimation.

4.1 The ε -Optimality for Temporally Dependent Demand Data

The ε -optimality in inventory-target estimation is introduced by Levi, Roundy, and Shmoys (2007) under the assumption of *independent* and identically distributed demand data. We start with generalizing the concept of ε -optimality for stationary and temporally dependent demand data.

Proposition 1. Let $C_{2|1}(F(\hat{q})|F(x_n); \theta) \geq b/(h+b) - \alpha$ and $C_{2|1}(F(\hat{q})|F(x_n); \theta) \leq b/(h+b) + \alpha$. The expected cost of the inventory-target estimate \hat{q} is at most $1 + \varepsilon$ of the expected cost of the optimal inventory target q^* when α is equal to $(\varepsilon hb/(h+b))/(h+b + \varepsilon \max(b, h))$.

The proof of Proposition 1 is available in Akçay, Biller, and Tayur (2013). In Proposition 1, we characterize the value of α so that an inventory-target estimate \hat{q} is *guaranteed* to be ε -optimal. Therefore, if the value of α is taken as characterized in Proposition 1, the probability of the event

$$Y(\hat{Q}, X_n) := \left[C_{2|1}(F(\hat{Q})|F(X_n); \theta) \geq \frac{b}{h+b} - \alpha \right] \cap \left[C_{2|1}(F(\hat{Q})|F(X_n); \theta) \leq \frac{b}{h+b} + \alpha \right]$$

gives a *lower bound* to the probability of the ε -optimality of the inventory-target estimator \hat{Q} . We denote the probability of the event $Y(\hat{Q}, X_n)$ with δ for ease in exposition.

Remark 1 Levi, Roundy, and Shmoys (2007) set α to $(\varepsilon/3) \min(b, h)/(h+b)$ when the historical demand data is independent and identically distributed. The lower bound on the probability of ε -optimality is higher when α is set as in Proposition 1, leading to a tighter bound. This is because $(\varepsilon hb/(h+b))/(h+b + \varepsilon \max(b, h))$ is always greater than $(\varepsilon/3) \min(b, h)/(h+b)$.

4.2 Identifying the Probabilistic Guarantee for ε -Optimality under General Copula

The objective of this section is to provide an arbitrarily close approximation to the lower bound δ on the probability of ε -optimality; i.e.,

$$\mathbb{P}(L(\hat{Q}|X_n) \leq (1 + \varepsilon)L(Q^*|X_n)) \geq \delta. \tag{4}$$

Clearly, the probability of the random event $Y(\hat{Q}, X_n)$ – which gives the lower bound δ – is a function of the unknown marginal distribution function F , while our goal is to find a probabilistic guarantee on ε -optimality for any marginal distribution. The marginal-copula representation of the demand process allows us to achieve this goal by waiving the requirement to know the true marginal distribution function of the demand process.

We let $\{U_t; t = 1, 2, \dots, n\}$ correspond to a series of *dependent* standard uniform random variables represented by the copula function $C(\cdot, \cdot; \theta)$, and let $G_n(u) := (1/n)\sum_{t=1}^n \mathbb{1}(U_t \leq u)$ be the empirical distribution function built from these uniform random variables. Then, we define a random variable U^* as follows:

$$U^* = \min_{j=1,2,\dots,n} \left\{ U_j : G_n(U_j) \geq C_{2|1}^{-1} \left(\frac{b}{h+b} | G_n(U_n); \hat{\theta} \right) \right\}. \tag{5}$$

We also define the event $\tilde{Y}(U^*, U_n)$ as

$$\left[C_{2|1}(U^*|U_n; \theta) \geq \frac{b}{h+b} - \alpha \right] \cap \left[C_{2|1}(U^*|U_n; \theta) \leq \frac{b}{h+b} + \alpha \right],$$

which will lead us to the key result to identify the lower bound δ on the probability of ε -optimality for any marginal demand distribution:

Proposition 2. *The probability of the event $Y(\hat{Q}, X_n)$ is equal to the probability of the event $\tilde{Y}(U^*, U_n)$, where U^* is defined as in (5) from a series of dependent standard uniform random variables $\{U_t; t = 1, 2, \dots, n\}$ generated from the copula $C(\cdot, \cdot; \theta)$.*

Proof. It follows from the probability integral transformation and the definition of a copula that the series of standard uniform random variables $\{U_t; t = 1, 2, \dots, n\}$ is equivalent to $\{F(X_t); t = 1, 2, \dots, n\}$. We consider the realizations $\{u_t; t = 1, 2, \dots, n\}$ and $\{x_t; t = 1, 2, \dots, n\}$. Since only the ordinal relation matters in building an empirical distribution function, the values of $F_n(x_t)$ and $G_n(u_t)$ are the same for $t = 1, 2, \dots, n$. Therefore, $C_{2|1}^{-1}(b/(h+b)|F_n(x_n); \hat{\theta})$ is equal to $C_{2|1}^{-1}(b/(h+b)|G_n(x_n); \hat{\theta})$. The result follows because $u_t = F(x_t)$, $t = 1, 2, \dots, n$ and

$$\hat{q} = \min_{j=1,2,\dots,n} \left\{ x_j : F_n(x_j) \geq C_{2|1}^{-1}(b/(h+b)|F_n(x_n); \hat{\theta}) \right\}$$

imply the equivalence between

$$\min_{j=1,2,\dots,n} \left\{ u_j : G_n(u_j) \geq C_{2|1}^{-1}(b/(h+b)|G_n(u_n); \hat{\theta}) \right\}$$

and $F(\hat{q})$ for any realizations of $\{U_t; t = 1, 2, \dots, n\}$ and $\{X_t; t = 1, 2, \dots, n\}$. \square

Proposition 2 plays a critical role in identifying the lower bound δ because it allows us to focus on the transformed demand random variables $\{U_t; t = 1, 2, \dots, n\}$ instead of the actual demand process $\{X_t; t = 1, 2, \dots, n\}$. We next present an algorithm based on this property to approximate the lower bound δ on the ε -optimality probability by sampling dependent uniform random variates rather than sampling the actual demand which requires the knowledge of the marginal distribution.

Algorithm 1. Computation of the lower bound δ on the probability of ε -optimality

Initialization. Specify an accuracy parameter $\gamma > 0$ and a confidence parameter $\beta \in (0, 1)$.

Let $M = \lceil \log(2/\beta)/(2\gamma^2) \rceil$ and $m = 1$.

while $m \leq M$:

Generate dependent standard uniform variates (u_1, u_2, \dots, u_n) from the copula $C(\cdot, \cdot; \theta)$.

$\hat{\theta} \leftarrow \operatorname{argmax} \{ \sum_{t=1}^{n-1} \log c(G_n(u_t), G_n(u_{t+1}); \theta) : \theta \in \Theta \}$.

$u^* \leftarrow \min_{j=1,2,\dots,n} \{ u_j : G_n(u_j) \geq C_{2|1}^{-1}(b/(h+b) | G_n(u_n); \hat{\theta}) \}$.

$B_m \leftarrow \begin{cases} 1 & \text{if } |C_{2|1}(u^* | u_n; \theta) - b/(h+b)| \leq (\varepsilon hb/(h+b))/(h+b+\varepsilon \max(h,b)), \\ 0 & \text{otherwise.} \end{cases}$

$m \leftarrow m + 1$.

end

Set $\bar{B} := M^{-1} \sum_{m=1}^M B_m$.

Return $\bar{B} - \gamma$ as the value of δ .

Remark 2 The decision maker may question the existence of temporal dependence in the demand process and choose to set the inventory target as if the demand realizations were independent and identically distributed. In this case, the algorithm provides the lower bound on the probability of ε -optimality by choosing u^* as $\min_{j=1,2,\dots,n} \{ u_j : G_n(u_j) \geq b/(h+b) \}$ for each sample path.

The theoretical support for the algorithm above rests upon Hoeffding's inequality for bounded random variables (Hoeffding 1963):

$$\mathbb{P}(|\bar{B} - \mathbb{E}(B)| > \gamma) \leq 2 \exp(-2M\gamma^2).$$

In this representation, $\mathbb{E}(B)$ is the unknown mean of the independently sampled indicator random variables $\{B_m; m = 1, 2, \dots, M\}$, and it corresponds to the true value of the lower bound δ on the probability of ε -optimality. Hoeffding's inequality immediately leads to the guarantee $\mathbb{P}(\mathbb{E}(B) \in [\bar{B} - \gamma, \bar{B} + \gamma]) \geq 1 - \beta$. The algorithm provides an arbitrarily close approximation to δ because the values of γ and β can be chosen arbitrarily smaller to make $\bar{B} - \gamma$ close enough to the true value of the δ with high confidence. To sum up, given the confidence level $1 - \beta$, the expected cost of any realized value of the inventory-target estimator \hat{Q} is at most $1 + \varepsilon$ of the minimum expected cost with probability at least $\bar{B} - \gamma$.

4.3 Results

In this section, we implement the algorithm in Section 4.2 to investigate how the lower bound δ for the probability of ε -optimality is affected by the temporal dependence in the demand process, the length of the demand history, and the cost parameters. We assume that the decision maker does not know the marginal demand distribution and the normal-copula parameter θ is the measure of temporal dependence in the demand process. Letting φ denote the critical fractile $b/(h+b)$ for ease in presentation, we set the values of γ and β to 0.001 and 0.05, respectively.

We let the normal-copula parameter θ take values between -0.9 and 0.9 to represent a wide range of negative and positive autocorrelations in the demand process. Figure 1 plots the lower bound δ on the probability of 0.25-optimality as a function of the length of the demand history n . For all values of n , we observe that the lower bound δ takes its maximum value when there is no autocorrelation in the demand process; i.e., when θ is equal to zero. Furthermore, we see that the lower bound δ decreases as the absolute value of the copula parameter θ increases. That is, the probabilistic guarantee to assure that the expected cost of an inventory-target estimate is at most 1.25 of the optimal expected cost decreases with the increasing strength of the autocorrelation in the demand process.

The decrease in δ can be explained by the slower convergence of the empirical demand distribution to the true marginal distribution in the presence of temporal dependence compared to the case with no temporal dependence. Intuitively, it takes more time for the demand history to discard the initial condition (i.e.,

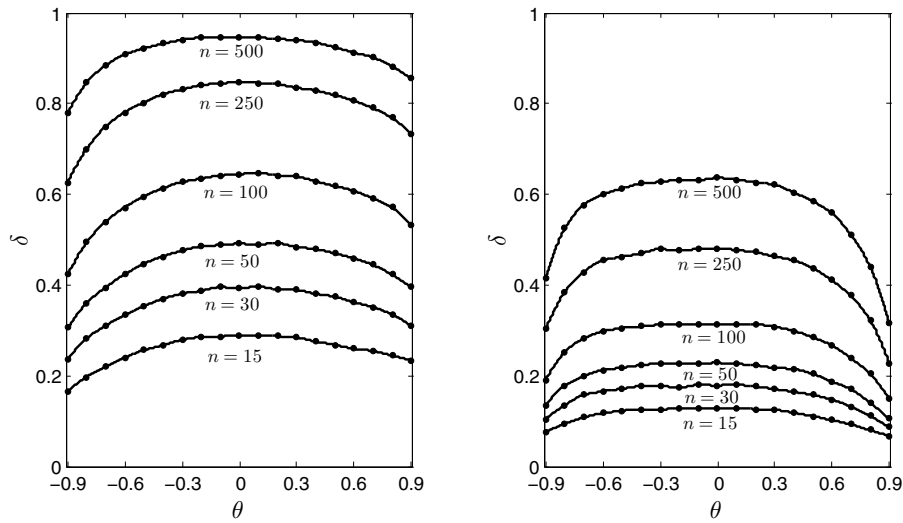


Figure 1: $\varepsilon = 0.25$ and $\varphi = 0.5$ (left); $\varepsilon = 0.25$ and $\varphi = 0.95$ (right).

the demand realizations at the beginning of the process) which controls the later demand realizations, and thus, delays the convergence of the empirical demand distribution when the strength of the autocorrelation is high.

Table 1 and Table 2 provide the number of demand observations necessary to achieve an ε -optimality guarantee of 50%. They verify the need for a larger number of demand observations with increasing strength of autocorrelation. For example, Table 1 shows that the expected cost of the inventory-target estimator in (3) is at most 1.5 of the optimal expected cost with a probability of at least 50% if the historical demand data includes the past 45 observations of a highly negatively correlated demand process (i.e., $\theta = -0.9$). The required number of demand observations decreases to 14 when there is no temporal dependence (i.e., $\theta = 0$), and then increases to 30 for a highly positively correlated demand process (i.e., $\theta = 0.9$). Tables 1 and 2 also shed light on how fast the number of demand observations necessary to achieve a certain level of probabilistic guarantee increases as ε decreases.

Table 1: The number of demand observations necessary to achieve $\delta = 0.5$ with $\varphi = 0.5$.

ε	θ						
	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
0.50	45	23	17	14	17	20	30
0.25	145	73	57	52	53	62	87
0.10	749	374	309	290	287	326	401

Table 2: The number of demand observations necessary to achieve $\delta = 0.5$ with $\varphi = 0.95$.

ε	θ						
	-0.9	-0.6	-0.3	0	0.3	0.6	0.9
0.50	265	108	100	96	102	136	475
0.25	778	322	280	274	299	387	1384
0.10	3812	1581	1390	1370	1442	1892	6738

5 CONCLUSION

We consider the problem of estimating the critical-fractile solution in a newsvendor problem when the demand depends on its past realizations. As opposed to widely-used linear time-series models with normally distributed random shocks, we introduce a copula-based demand model that allows us to represent the stationary and temporally dependent demand process with any marginal demand distribution and an arbitrary dependence structure. Consequently, the decision maker estimates the marginal demand distribution and the copula parameters that characterize the temporal dependence separately without making any assumptions on the parametric form of the marginal demand distribution. The objective of our study is to identify a probabilistic guarantee for the ε -optimality of an estimate of the critical-fractile solution, which is obtained from the empirical demand distribution and the estimates of the copula parameters.

We propose a sampling-based method to compute this probabilistic guarantee when the marginal distribution is unknown to the decision maker without any restrictions on the functional form of copula and the cost parameters. Our method builds on the idea of sampling dependent uniform random variates matching the underlying dependence structure of the demand process rather than sampling the actual demand which requires the specification of the marginal demand distribution. Our numerical analysis shows that the probabilistic guarantee takes its maximum value when there is no temporal dependence, and it decreases as the strength of the demand autocorrelation increases.

REFERENCES

- Akçay, A., B. Biller, and S. Tayur. 2013. "Data-Driven Newsvendor: A Probabilistic Guarantee for Near Optimality when Demand is Temporally Dependent". Working Paper, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA.
- Akkaya, A., and M. Tiku. 2005. "Time series AR(1) model for short-tailed distributions". *Statistics* 39 (2): 117–132.
- Beare, B. 2010. "Copulas and temporal dependence". *Econometrica* 78 (1): 395–410.
- Block, H., N. Langberg, and D. Stoffer. 1990. "Time series models for non-Gaussian processes". *Lecture Notes-Monograph Series*:69–83.
- Cario, M., and B. Nelson. 1996. "Autoregressive to anything: Time-series input processes for simulation". *Operations Research Letters* 19 (2): 51–58.
- Chen, L., and H. Lee. 2012. "Bullwhip effect measurement and its implications". *Operations Research* 60 (4): 771–784.
- Chen, X., and Y. Fan. 2006. "Estimation of copula-based semiparametric time series models". *Journal of Econometrics* 130 (2): 307–335.
- Erkip, N., W. Hausman, and S. Nahmias. 1990. "Optimal Centralized Ordering Policies in Multi-Echelon Inventory Systems with Correlated Demands". *Management Science* 36:381–392.
- Gourieroux, C., and J. Jasiak. 2006. "Autoregressive gamma processes". *Journal of Forecasting* 25 (2): 129–152.
- Hoeffding, W. 1963. "Probability inequalities for sums of bounded random variables". *Journal of the American Statistical Association* 58 (301): 13–30.
- Hosoda, T., M. Naim, S. Disney, and A. Potter. 2008. "Is there a benefit to sharing market sales information? Linking theory and practice". *Computers & Industrial Engineering* 54 (2): 315–326.
- Joe, H. 1997. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.
- Jose, K., T. Lishamol, and J. Sreekumar. 2008. "Autoregressive processes with normal-Laplace marginals". *Statistics & Probability Letters* 78 (15): 2456–2462.
- Lee, H., V. Padmanabhan, and S. Whang. 1997. "Information Distortion in a Supply Chain: The Bullwhip Effect". *Management Science* 43:546–558.
- Lee, H., K. So, and C. Tang. 2000. "The Value of Information Sharing in a Two-Level Supply Chain". *Management Science* 46:626–643.

- Levi, R., R. Roundy, and D. Shmoys. 2007. "Provably near-optimal sampling based policies for stochastic inventory control models". *Mathematics of Operations Research* 32 (4): 821–839.
- Luong, H. 2007. "Measure of bullwhip effect in supply chains with autoregressive demand process". *European Journal of Operational Research* 180:1086–1097.
- Nelsen, R. 2006. *An Introduction to Copulas*. 2nd edition, Springer, New York.
- Patton, A. 2012. "A review of copula models for economic time series". *Journal of Multivariate Analysis* 110:4–18.
- Porteus, E. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Sklar, A. 1959. "Fonctions de repartition à n dimensions et leurs marges". *Publications de l'Institut de Statistique de L'Universit de Paris* 8:229–231.
- Tiku, M., W.-K. Wong, D. Vaughan, and G. Bian. 2002. "Time Series Models in Non-Normal Situations: Symmetric Innovations". *Journal of Time Series Analysis* 21 (5): 571–596.

AUTHOR BIOGRAPHIES

ALP AKÇAY 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.akçay@bilkent.edu.tr.

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 and financial time-series modeling. Her email address is billerb@andrew.cmu.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.