M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 8, No. 4, Fall 2006, pp. 351–358
ISSN 1523-4614 | EISSN 1526-5498 | 06 | 0804 | 0351

informs®
DOI 10.1287/msom.1060.0117

A Generalization of the Inventory Pooling Effect to Nonnormal Dependent Demand

Charles J. Corbett, Kumar Rajaram
Anderson School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, California 90095
{charles.corbett@anderson.ucla.edu, kumar.rajaram@anderson.ucla.edu}

Eppen (1979) showed that inventory costs in a centralized system increase with the correlation between multivariate normal product demands. Using multivariate stochastic orders, we generalize this statement to arbitrary distributions. We then describe methods to construct models with arbitrary dependence structure, using the copula of a multivariate distribution to capture the dependence between the components of a random vector. For broad classes of distributions with arbitrary marginals, we confirm that centralization or pooling of inventories is more valuable when demands are less positively dependent.

Key words: inventory control; pooling effect; multivariate dependence; copula
History: Received: August 12, 2004; accepted: June 1, 2006.

1. Introduction

Consider a firm having to determine inventory levels for the same product in many retail locations with stochastic demand. If inventory is centralized, as opposed to being kept at the retail outlets, the demands from all locations are pooled, so the company will face lower aggregate demand uncertainty and hence lower costs. Many variations of this "pooling effect," first analyzed by Eppen (1979) in inventory management, exist. Intuitively, the pooling effect becomes less valuable as demands are more positively dependent, but almost all such analysis to date, including Eppen (1979), has had to focus on the multivariate normal case because of the intractability of dealing with multivariate dependence under nonnormal distributions.

In this paper, we show how Eppen's original results can be generalized to a broad class of nonnormal distributions with arbitrary marginals. Specifically, we formalize the intuitive notion that inventory costs in a centralized system increase as demands are more positively dependent. We also provide examples to illustrate how to construct nonnormal distributions with arbitrary marginals and a wide range of dependence structures, and show how statements about the effect of dependence can still be made in more general contexts.

This paper is organized as follows. In §2, we review relevant literature in the areas of inventory pooling and probability theory. In §3 we formally introduce the inventory pooling problem. In §4 we define the sum-convex order, which we then use to state the more general version of Eppen's (1979) result. Section 5 provides a bivariate and a multivariate application of this generalization, using copulae to model the dependence structure of a multivariate distribution. Section 6 offers conclusions and future research directions.

2. Literature Review

We first summarize relevant literature related to inventory pooling and follow with a short review of some recent work in probability theory. A more comprehensive review is provided in Corbett and Rajaram (2004). A large body of work has grown around various manifestations of Eppen's (1979) notion of pooling of inventories, or Eppen and Schrage's (1981) extension that includes lead times. Federgruen and Zipkin (1984) provide approximations for more general versions of Eppen and Schrage's (1981) model, with finite horizon, other-than-normal demand distributions, and nonidentical retailers. Jönsson and Silver (1987) present an exhaustive study of the impact of changing input parameters on system performance;

Gerchak and Mossman (1992) show how the order quantity and associated costs depend on the randomness parameter in a simple and highly interpretable manner. Erkip et al. (1990) find that high positive correlation among products and successive time periods (around 0.7) results in significantly higher safety stock than the no-correlation case. Alfaro and Corbett (2003) analyze the value of pooling under suboptimal inventory policies and report on numerical and empirical experiments with nonnormal demand data. Dong and Rudi (2004) study the impact of correlation on price interactions under transshipment, and Netessine et al. (2002) study the impact of correlation on flexible service capacity under multivariate normal demand. Van Mieghem and Rudi (2002) further extend the analysis of pooling to "newsvendor networks."

The benefits of delayed product differentiation or postponement are quite similar to those of the pooling effect, referring to multiple products instead of multiple locations (Garg and Lee 1999). Groenevelt and Rudi (2000) and Rudi (2000) have examined the interactions between the optimal inventory policy, the degree of component commonality, demand variability, and correlation under bivariate distributions. Ho and Tang (1998) and the references therein provide further discussion of the pooling effect in the context of product variety.

Most of this literature in inventory pooling, including this paper, focuses on the impact of pooling on expected profits. A related, but usually more intractable problem, concerns the effect of pooling on optimal inventory levels. We do not consider that question here, though some work, including Eppen (1979), Erkip et al. (1990), and Van Mieghem and Rudi (2002) do address that issue under more restrictive distributional assumptions than ours. So far, the work related to pooling of inventories has generally lacked a formal mechanism for assessing the impact of dependence on the value of pooling when demands are nonnormal. Whenever dependence has been explicitly included, it has generally been in the context of bivariate or multivariate normal demands; this paper describes mechanisms that could be used to generalize this work to the nonnormal case.

We refer to work on multivariate stochastic orders and on the copula where appropriate. Recent work on multivariate orders includes Scarsini and Shaked (1996), Shaked and Shanthikumar (1994), Scarsini (1998), Müller and Scarsini (2000, 2001), and Müller and Stoyan (2002); Joe (1997) and Nelsen (1999) provide good overviews of theory and applications of copulae. Clemen and Reilly (1999) introduce the multivariate normal copula and discuss its use in the context of combining expert opinions. The contribution of this paper lies in the combination and application of these recent concepts from probability and statistics to the inventory pooling context.

3. Pooling of Inventories

A well-known problem in inventory theory is to decide how much inventory to carry when faced with uncertain demand; the decision maker has to trade off h, the per unit holding costs of excess inventory against p, the per unit shortage costs of not meeting all demand. For a single product i with stochastic demand x_i and associated cumulative demand distribution $F_i(x_i)$, the decision maker's cost function $C(q_i)$ depends on his or her inventory level q_i as follows: $C(q_i) = E[h(q_i - x_i)^+] + E[p(x_i - q_i)^+]$, where $(z)^+ = \max\{0, z\}$. It is well known that the optimal order quantity is $q_i^* = F_i^{-1}(p/(p+h))$. If the firm sells the product at multiple retail locations, demand is a multivariate random variable $\mathbf{X}$ with corresponding distribution F. It is sufficient to determine the optimal inventory levels for each location independently, regardless of dependence structure, to minimize total expected costs.

However, the firm can also keep a central inventory instead of local inventories (ignoring transportation lead times), hence aggregating demand from multiple locations into a single random variable, allowing it to exploit "statistical economies of scale." This is referred to as the "pooling effect," first characterized by Eppen (1979). For the decentralized case (with subscript D), the problem is to find the vector of local inventories $\bar{q} = (q_1, \ldots, q_N)$ that solves

$$\min_{\bar{q}} E[TC_D(\bar{q})],$$

where

$$\begin{aligned} E[TC_D(\bar{q})] &= E\left[\sum_{i=1}^{N} C_D(q_i)\right] \\ &= \sum_{i=1}^{N} \big(E_i[h(q_i - x_i)^+] + E_i[p(x_i - q_i)^+]\big) \end{aligned}$$

and E_i is the expectation with respect to the ith marginal of the joint demand distribution. For the centralized case this problem is to find a single aggregate inventory q that solves

$$\min_q E[TC_C(q)],$$

where

$$E[TC_C(q)] = E\left[h(q - \sum_{i=1}^{N} x_i)^+\right] + E\left[p(\sum_{i=1}^{N} x_i - q)^+\right].$$

Write $E[TC_D]$ and $E[TC_C]$ for the expected costs with the optimal decentralized and centralized inventory levels.

When needed, we will write $E^X[\cdot]$ to denote the expectation over a random variable $\mathbf{X}$, and $E^n[\cdot]$ for expectations given a normal demand distribution. Let demand follow a normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with correlations ρ_{ij} between all products $i \neq j$. Then, as shown in Eppen (1979), we can represent $E^n[TC_D]$ and $E^n[TC_C]$ as

$$E^n[TC_D] = K \sum_{i=1}^{N} \sigma_i \tag{3.1}$$

$$E^n[TC_C] = K \sqrt{\sum_{i=1}^{N} \sigma_i^2 + 2 \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \sigma_i \sigma_j \rho_{ij}}, \tag{3.2}$$

where, following Eppen's (1979) notation,

$$K = \left(h\Phi^{-1}\left(\frac{p}{p+h}\right) + (p+h)R\left(\frac{p}{p+h}\right)\right),$$

with $\phi(\cdot)$ the probability density function of the standard normal distribution, $\Phi^{-1}(\cdot)$ the inverse of the cumulative distribution function for the standard normal distribution, and $R(u) = \int_u^\infty (w-u)((e^{-w^2}/2)/\sqrt{2\pi})\,dw$ the right-hand unit normal linear-loss integral. We can now examine the effect of dependence on total costs before and after centralization. If $\sigma_i = \sigma$ $\forall\, i$ and $\rho_{ij} = \rho$ $\forall\, i \neq j$, then

$$E^n[TC_C] = K\sigma\sqrt{N + \rho N(N-1)}, \tag{3.3}$$

so the value of pooling $E^n[TC_D] - E^n[TC_C]$ is nonnegative and decreasing in ρ. If $\rho = 1$, then $E^n[TC_D] = E^n[TC_C]$; if $\rho = 0$, then $E^n[TC_C] = E^n[TC_D]/\sqrt{N}$; if $\rho = -1/(N-1)$, then $E^n[TC_C] = 0$. We can summarize this well-known effect of correlation on total costs as follows:

Proposition 1. *When demand follows a normal distribution, total costs after pooling are increasing in all bivariate correlation coefficients ρ_{ij}* (*Eppen* 1979).

Using the concepts gathered in this paper we can state a much more general version of this result, confirm the intuition that the benefits of centralization decrease as the individual demands become more positively dependent, and construct examples of multivariate demand distributions with arbitrary marginals and a broad range of dependence structures. We next describe the multivariate stochastic order that is essential to generalize Proposition 1 to multivariate nonnormal distributions with arbitrary dependence structures.

4. Dependence and the Effect of Pooling

In this section we first introduce the sum-convex order, a simple but general multivariate stochastic order that we will use to compare the effect of pooling under demand portfolios with different dependence structures. We then use this order to generalize Eppen's result to arbitrary distributions. First, recall the well-known univariate convex order. Let $\mathbf{X}$ and $\mathbf{Y}$ be two univariate random variables with distributions F and G, respectively, for which the expectations $E[\psi(\mathbf{X})]$ and $E[\psi(\mathbf{Y})]$ exist for all convex functions ψ. We use $\succ$ to denote weak dominance.

Definition. $\mathbf{X} \succ_{\text{cx}} \mathbf{Y}$ if and only if $E[\psi(\mathbf{X})] \geq E[\psi(\mathbf{Y})]$ for all convex functions ψ: $\Re \to \Re$.

The multivariate version of $\succ_{\text{cx}}$ is easy to define using multivariate convex functions but is often of limited use, as it does not cleanly separate variability and dependence. To overcome this limitation, various researchers have defined a range of new orders for their specific purposes, some of which are briefly defined in the appendix. To study inventory pooling, we follow a similar approach to that which Müller and Scarsini (2001) take and define an order that we call the sum-convex order. Let $\mathbf{X}_i$ denote component i in random vector $\mathbf{X}$.

Definition. Let random vector $\mathbf{X}$ and $\mathbf{Y}$ have dimensions N_X and N_Y, respectively. Then $\mathbf{X}$ dominates $\mathbf{Y}$ in the *sum-convex order*, written as $\mathbf{X} \succ_{\text{scx}} \mathbf{Y}$, if and only if $\sum_{i=1}^{N_X} \mathbf{X}_i \succ_{\text{cx}} \sum_{i=1}^{N_Y} \mathbf{Y}_i$.

One could of course dispense with the sum-convex order by always writing $\sum_{i=1}^{N_X} \mathbf{X}_i \succ_{\mathrm{cx}} \sum_{i=1}^{N_Y} \mathbf{Y}_i$; however, it is notationally convenient to define the sum-convex order separately. Moreover, we can relate it to other multivariate orders that have been defined in the probability literature and show that $\succ_{\mathrm{scx}}$ is strictly weaker (and hence more general) than these existing orders. The ranking of random vectors under $\succ_{\mathrm{scx}}$ depends on the variability of aggregate value, which in turn depends on the variability of the individual components and on the interdependence between them.

Figure 1 places the sum-convex order in the context of several existing multivariate stochastic orders, which are defined in the appendix. More detailed explanations of these orders and the links between them can be found in Corbett and Rajaram (2004), especially in the references therein. Figure 1 shows that the positive linear-convex order (Scarsini 1998), $\succ_{\mathrm{plcx}}$, and the directionally convex order (see, for instance, Müller and Scarsini 2001), $\succ_{\mathrm{dcx}}$, imply $\succ_{\mathrm{scx}}$. Note that these are all partial orderings; i.e., there exist pairs $\mathbf{X}$ and $\mathbf{Y}$ such that neither $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ nor $\mathbf{Y} \succ_{\mathrm{scx}} \mathbf{X}$ holds. However, it is easy to see that both implications are strict; i.e., the $\succ_{\mathrm{scx}}$ order is weaker (and hence allows more general comparisons) than $\succ_{\mathrm{plcx}}$ or $\succ_{\mathrm{dcx}}$. The most obvious case is $\mathbf{X}$ and $\mathbf{Y}$ with different dimensionality, as $\succ_{\mathrm{plcx}}$ and $\succ_{\mathrm{dcx}}$ are then not well defined while $\succ_{\mathrm{scx}}$ is, but examples with $\mathbf{X}$ and $\mathbf{Y}$ with equal dimensionality are also easy to construct. The fact that the sum-convex order allows many more pairs of random vectors to be compared, even with unequal dimensionality, than existing orders, is what makes it useful for our purposes and likely for many other applications in operations research and decision theory.

Figure 1 **Summary of Key Relationships Between Multivariate Orders**

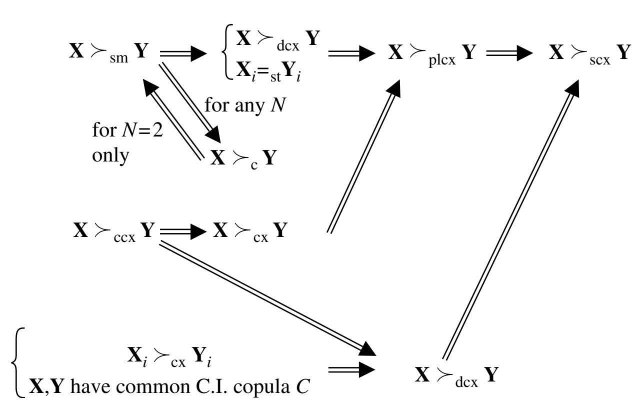

We can now generalize Proposition 1, stating that increased dependence reduces the value of pooling, to much broader classes of distributions:

Proposition 2. *Let* $\mathbf{X}$ *and* $\mathbf{Y}$ *be multivariate random variables with* N_X *and* N_Y *dimensions. Then* $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ *implies* $\min_q E^X[TC_C(q)] \geq \min_q E^Y[TC_C(q)]$; *that is, the cost after pooling is greater under demand* $\mathbf{X}$ *than under* $\mathbf{Y}$.

Proof. It is easy to verify that the centralized objective function $E[TC_C(\mathbf{X}; q)]$ is convex in $\sum_{i=1}^{N_X} \mathbf{X}_i$ for given q, so $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ implies $E^X[TC_C(q)] \geq E^Y[TC_C(q)]$ for any q, so also $\min_q E^X[TC_C(q)] \geq \min_q E^Y[TC_C(q)]$. □

If $\mathbf{X}$ is more positively dependent than $\mathbf{Y}$ under any dependence order that implies $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$, for instance, using the relationships in Figure 1, we can use Proposition 2 to conclude that the costs after pooling are greater under $\mathbf{X}$ than under $\mathbf{Y}$, which generalizes Proposition 1 to multivariate nonnormal distributions with arbitrary dependence structures. Figure 1 summarizes several sufficient conditions for $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ to hold, and hence for costs after pooling to be greater under $\mathbf{X}$ than under $\mathbf{Y}$. Analogously, one can now return to other existing work on pooling of inventories, postponement of differentiation, etc., and verify which of the orders in Figure 1 apply to the objective function considered in that work. This will then show that many of those existing results can also be generalized to nonnormal-dependent distributions. As more orders are defined and more links between them established, more sufficient conditions can be added to the framework in Figure 1.

Proposition 2 immediately raises the question, "When does $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ hold for any given situation?" In the next section, we define broad classes of multivariate distributions and show how Proposition 2 can be used to demonstrate that higher dependence leads to higher costs in a bivariate and a multivariate example. To do so, we need to model dependence in arbitrary nonnormal multivariate distributions, for which we use the copula.

5. Examples: Multivariate Dependence and Pooling

5.1. The Copula

A relatively recent tool for capturing dependence in arbitrary multivariate distributions is the "copula." This will be useful for us in two ways. First, comparing two multivariate random variables with the same marginals is equivalent to comparing their copulae. Second, in constructing stochastic models, the copula allows us to combine arbitrary marginals with an arbitrary dependence structure, rather than limiting us to the few distributions with tractable dependence structures. We illustrate both these uses with examples at the end of this section; first, we introduce the basic ideas behind the copula (see, for instance, Joe 1997).

Definitions. Let $\tilde{F}$ denote the *Fréchet class* given a set of marginal distributions; e.g., $\tilde{F}(F_1, \ldots, F_N)$ is the class of multivariate distributions with given marginals $F_1, \ldots, F_N$.

For any multivariate distribution $F \in \tilde{F}(F_1, \ldots, F_N)$, the *copula associated with* F is a distribution function C: $[0,1]^N \to [0,1]$ that satisfies $F(\mathbf{x}) = C(F_1(x_1), \ldots, F_N(x_N))$, $\mathbf{x} \in \Re^N$. The copula $C(u_1, \ldots, u_N)$ itself is a joint distribution with uniform marginals.

Let $\mathbf{U}$ and $\mathbf{V}$ be multivariate uniform random variables with distributions C_U and C_V, respectively; we will interchangeably write $\mathbf{U} \succ \mathbf{V}$ and $C_U \succ C_V$. Sklar's theorem (see, for instance, Clemen and Reilly 1999) guarantees that a copula always exists:

Sklar's Theorem. *For any multivariate distribution $F \in \tilde{F}(F_1, \ldots, F_N)$, the copula as defined above exists. If the F_i are all continuous, then C is unique; otherwise, C is uniquely determined on $\prod_{i=1}^N \mathrm{Ran}(F_i)$, where $\mathrm{Ran}(F_i)$ is the range of F_i.*

Clemen and Reilly (1999) show that, for differentiable F_i and C, the joint density can be written as $f(x_1, \ldots, x_N) = \prod_{i=1}^N f_i(x_i) c(F_1(x_1), \ldots, F_N(x_N))$, where the $f_i(x_i)$ are the densities of the marginals F_i and $c = \partial^N C / \partial F_1(x_1) \cdots \partial F_N(x_N)$, the copula density. From the definition, it is clear that the copula is entirely general and fully captures the dependence structure inherent in any multivariate distribution F. Using the copula, we can now state results for comparisons between random vectors with equal marginals but different dependence structures.

5.2. Comparing Random Vectors with Equal Marginals but Different Dependence Structures

There is an immediate link between multivariate dependence orders and the copula. Let T: $\mathbf{X} \to T(\mathbf{X})$ be any transform of a multivariate random variable $\mathbf{X}$, where $T(\mathbf{X})$ has the same dimensionality as $\mathbf{X}$ and where each component $T_i(\mathbf{X})$ is an increasing transform of the marginal $\mathbf{X}_i$. A multivariate stochastic order $\succ$ is said to be invariant under increasing transforms if $\mathbf{X} \succ \mathbf{Y}$ implies $T(\mathbf{X}) \succ T(\mathbf{Y})$ for all such T.

Lemma 1. *Let $\mathbf{X}$ and $\mathbf{Y}$ be two multivariate random variables such that $\mathbf{X}_i =_{\mathrm{st}} \mathbf{Y}_i$ $\forall\, i$, with distributions F and G and corresponding copulas C_X and C_Y, respectively. Then for all orders $\succ$ that are invariant under increasing transforms, $\mathbf{X} \succ \mathbf{Y}$ if and only if $C_X \succ C_Y$ (Scarsini and Shaked* 1996, *Remark* 5.6).

Proof. By definition, the random variable $\mathbf{U} = (F_1(\mathbf{X}_1), \ldots, F_N(\mathbf{X}_N))$ has distribution C_X, so that $\mathbf{U}_i = F_i(\mathbf{X}_i)$. Define the inverse $F_i^{-1}(\mathbf{U}_i)$ appropriately to ensure existence. Then $C_X(\mathbf{U}) = (F_1(F_1^{-1}(\mathbf{U}_1)), \ldots, F_N(F_N^{-1}(\mathbf{U}_N)))$ and $F(\mathbf{X}) = C_X(F_1(\mathbf{X}_1), \ldots, F_N(\mathbf{X}_N))$; analogous relations hold between $\mathbf{Y}$, G, $\mathbf{V}$, and C_Y. Both $\mathbf{U}_i = F_i(\mathbf{X}_i)$ and $F_i^{-1}(\mathbf{U}_i)$ are increasing in their respective arguments, so the result follows from the assumption of invariance under increasing transforms of the marginals. $\square$

Any multivariate stochastic order that compares distributions exclusively based on their dependence is called a multivariate positive dependence order (MPDO). Joe (1997) lists nine axioms that any such MPDO must satisfy. The condition of Lemma 1, of invariance under increasing transforms, follows from his Axioms 7 and 8 for MPDOs. This yields:

Corollary 1. *The conditions of Lemma* 1 *are satisfied for all multivariate positive dependence orders as defined by Joe* (1997), *so for all MPDOs one can interchangeably compare the distributions or their copulae.*

The supermodular order (see, e.g., Müller and Scarsini 2000), defined in the appendix, is an MPDO, but orders such as the convex order are *not* dependence orders, and it is easy to find examples in which $C_X \succ_{\mathrm{cx}} C_Y$ but not $\mathbf{X} \succ_{\mathrm{cx}} \mathbf{Y}$. In light of Lemma 1, one can model the dependence structure of a random vector using a copula and use multivariate positive dependence orders to assess the effect of increasing dependence. We illustrate this procedure in §5.3.

5.3. Constructing Multivariate Distributions with Arbitrary Marginals

Here we present two examples of the copula of a multivariate distribution with arbitrary marginals and its relationship to the multivariate orders discussed so far. Let the $\mathbf{X}_i$ follow arbitrary univariate distributions F_i respectively; the F_i need not come from the same family of distributions. To model dependence, we give two examples, the first using bivariate Archimedean copulae, and the second using the multivariate normal copula. For the first, we also use the concordance order (Joe 1997), $\mathbf{X} \succ_{\mathrm{c}} \mathbf{Y}$, which is defined in the appendix, but which essentially means that the elements of $\mathbf{X}$ "move together more closely" than those of $\mathbf{Y}$.

Proposition 3. *Let $\mathbf{X}$ and $\mathbf{Y}$ be arbitrary bivariate random variables with distributions $F, G \in \tilde{F}(F_1, F_2)$ and with copulae C_X and C_Y, respectively. Then $C_X \succ_{\mathrm{c}} C_Y$ implies $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$.*

Proof. Because the concordance ordering is an MPDO (Joe 1997, p. 39), $C_X \succ_{\mathrm{c}} C_Y$ implies that $\mathbf{X} \succ_{\mathrm{c}} \mathbf{Y}$, by Lemma 1. By Theorem 2.5 in Müller and Scarsini (2000, p. 110), we know that $\mathbf{X} \succ_{\mathrm{c}} \mathbf{Y} \Leftrightarrow \mathbf{X} \succ_{\mathrm{sm}} \mathbf{Y}$ for bivariate distributions. It is easy to verify that the function $\psi(\sum_{i=1}^{N} \mathbf{X}_i)$ is supermodular for all convex functions ψ, which then gives $\mathbf{X} \succ_{\mathrm{sm}} \mathbf{Y} \Rightarrow \mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$. □

This means that, presented with any two bivariate demand distributions with pairwise equal (but arbitrary) marginals, the one with the more concordant copula will lead to higher costs in a centralized system. To see how Proposition 3 can be applied, consider the class of bivariate Archimedean copulae, which is broad (Nelsen 1999 lists 22 families on pp. 94–97) and useful for several reasons: They can be constructed easily, a wide variety of families of copulae belong to this class, and they possess many useful properties. We will not define the class in general, but consider, for instance, the following specific family (our example will work with many others):

$$C_\theta(u_1, u_2) = \frac{u_1 u_2}{(1-(1-u_1^\theta)(1-u_2^\theta))^{1/\theta}}, \quad 0 < \theta \le 1. \quad (5.1)$$

The joint distribution is given by $F(X_1, X_2) = C_\theta(F_1(X_1), F_2(X_2))$. If one wishes to combine this copula with marginals that are uniform on $[0,1]$, then the joint distribution itself is also given by $F(X_1, X_2) = (X_1 X_2)/(1-(1-X_1^\theta)(1-X_2^\theta))^{1/\theta}$; moreover, taking the limit gives $\lim_{\theta \downarrow 0} C_\theta(u_1, u_2) = u_1 u_2$, the product copula, so $\mathbf{X}_1$ and $\mathbf{X}_2$ are independent in that case.

Proposition 4. *Let $\mathbf{X}$ and $\mathbf{Y}$ be bivariate random variables with distributions $F, G \in \tilde{F}(F_1, F_2)$ and with copulae $C_{\theta_X}(u_1, u_2)$ and $C_{\theta_Y}(u_1, u_2)$, both from the same family, such that $\partial C_\theta(a,b)/\partial\theta \ge 0$ for all $(a,b) \in [0,1] \times [0,1]$. Then $\theta_X \ge \theta_Y$ implies $\min_q E^X[TC_C(q)] \ge \min_q E^Y[TC_C(q)]$; that is, the cost after pooling is greater under demand $\mathbf{X}$ than under $\mathbf{Y}$.*

Proof. Because $\partial C_\theta(a,b)/\partial\theta \ge 0$, $\theta_X \ge \theta_Y$ also implies that $C_{\theta_X}(u_1, u_2) \ge C_{\theta_Y}(u_1, u_2)$ $\forall u_1, u_2 \in [0,1]$. By the definition of the concordance order (Nelsen 1999, p. 181), $C_{\theta_X}(u_1, u_2) \ge C_{\theta_Y}(u_1, u_2)$ $\forall u_1, u_2 \in [0,1]$ implies that $C_{\theta_X} \succ_{\mathrm{c}} C_{\theta_Y}$. Proposition 3 then gives $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$. By Proposition 2, $\mathbf{X} \succ_{\mathrm{scx}} \mathbf{Y}$ implies that the cost after pooling is greater under $\mathbf{X}$ than under $\mathbf{Y}$. □

The parameter θ in $C_\theta(a,b)$ can be thought of as the counterpart to the correlation coefficient in a multivariate normal distribution, so the condition $\partial C_\theta(a,b)/\partial\theta \ge 0$ implies that for any pair (a,b), the joint probability of falling within $[0,a] \times [0,b]$ is increasing in θ; in other words, the two elements of the joint distribution are more likely to move together as θ increases. Many of the bivariate Archimedean copulae listed in Nelsen (1999), including the family defined in (5.1), satisfy the condition $\partial C_\theta(a,b)/\partial\theta \ge 0$, so this example illustrates how higher dependence leads to higher costs after pooling for bivariate distributions with arbitrary marginals and a range of Archimedean copulae.

For multivariate distributions, the $\succ_{\mathrm{c}}$ order does not imply the $\succ_{\mathrm{sm}}$ order, so the construction above does not work. Joe (1997) discusses properties of a range of multivariate copulae, but these are usually considerably more complex than are likely to be applied in inventory control settings. However, we can still construct a broad class of multivariate random variables using the normal copula, discussed in Clemen and Reilly (1999); this again allows arbitrary marginals but captures dependence exactly as the multivariate normal distribution does, using only pairwise correlations. In other words, the multivariate distribution is fully defined by the marginals F_i and the covariance matrix $\mathbf{\Sigma}$. For modeling multivariate product demands, this dependence structure seems to offer a reasonable compromise, as the information needed to estimate more complex copulae is rarely available, while the combination of a normal copula with arbitrary marginals still offers substantial modeling flexibility. This is because assuming a normal copula is far less restrictive than requiring the entire distribution to be normal, because it allows any shape for the marginals, including nonsymmetric, monotone, and multimodal.

Proposition 5. *Let* $\mathbf{X}$ *and* $\mathbf{Y}$ *be arbitrary multivariate random variables with distributions* $F, G \in \tilde{F}(F_1, \ldots, F_N)$ *and with normal copulae* C_X *and* C_Y*, characterized by covariance matrices* $\boldsymbol{\Sigma}_X$ *and* $\boldsymbol{\Sigma}_Y$ *with elements* $\sigma_{X,ij}$ *and* $\sigma_{Y,ij}$*, respectively. Then* $\sigma_{X,ij} \geq \sigma_{Y,ij}$ $\forall i, j$ *implies* $\min_q E^X[TC_C(q)] \geq \min_q E^Y[TC_C(q)]$*; i.e., the cost after pooling is greater under demand* $\mathbf{X}$ *than under* $\mathbf{Y}$.

Proof. Let $\mathbf{X}$ and $\mathbf{X}'$ be multivariate random variables with copula C_X, and let $\mathbf{Y}$ and $\mathbf{Y}'$ have copula C_Y. Then $\mathbf{X}' \succ_{\text{sm}} \mathbf{Y}' \Leftrightarrow C_X \succ_{\text{sm}} C_Y \Leftrightarrow \mathbf{X} \succ_{\text{sm}} \mathbf{Y}$ because the supermodular order is a dependence order. Now let $\mathbf{X}'$ and $\mathbf{Y}'$ be multivariate normal random variables; then Theorem 4.2 in Müller and Scarsini (2000, p. 117) shows that the conditions in the proposition imply $\mathbf{X}' \succ_{\text{sm}} \mathbf{Y}'$, so also $\mathbf{X} \succ_{\text{sm}} \mathbf{Y}$, from which the rest follows from Figure 1 and Proposition 2. □

Proposition 5 illustrates how higher pairwise dependence leads to higher costs after pooling for multivariate distributions with arbitrary marginals and a normal copula. For instance, one can construct marginal demand distributions for individual products or collections of products at individual locations based on choice models, such as the multinomial logit model or the latest class demand structure; assume a multivariate normal copula for the dependence between product demands across locations; and simply check the pairwise correlation condition to verify the effect of inventory pooling in those settings.

Comparing Proposition 5 with expression (3.2) for normal distributions clearly highlights the trade-off inherent in working with normal distributions. With normal distributions, $\sum_{ij} \sigma_{X,ij} \geq \sum_{ij} \sigma_{Y,ij}$ is sufficient for the costs after pooling to increase; but if only the copula is assumed to be normal but not the marginals, we must have $\sigma_{X,ij} \geq \sigma_{Y,ij}$ for *all* i, j to be able to show that costs after pooling increase. It is possible that tighter conditions than this can be found, though we are not aware of any. In addition, note that Proposition 5 allows any existing result related to pooling with a bivariate normal distribution to be generalized to arbitrary bivariate distributions with normal copula: Inventory costs in a centralized system with arbitrary marginals and a normal copula are increasing in the correlation coefficient ρ, as one would expect.

6. Conclusions

In this paper we have generalized Eppen's (1979) result, on how inventory costs after pooling increase with dependence between the individual demands, to near-arbitrary multivariate-dependent demand distributions, and we have also illustrated how to construct such distributions. In doing so, we have provided a basis to extend the large literature that has sprung from that principle to more general demand distributions. Altogether, this paper shows how one can address problems of pooling of inventories without needing to resort to assumptions of independence or multivariate normality. There are many other potential areas of application of these concepts—in decision theory, risk assessment, reliability, portfolio comparison, and inventory theory. We hope that this paper will stimulate more work in these areas.

Acknowledgments
The authors are grateful to Paul Kleindorfer, John Mamer, Kevin McCardle, Rakesh Sarin, Marco Scarsini, and seminar participants at Carnegie Mellon University, Case Western Reserve University, Northwestern University, Stanford University, University of California, Irvine, University of Texas at Austin, and Washington University in St. Louis for many helpful comments. The authors also thank the senior editor and the referees for their constructive comments.

Appendix. Definitions
Following Müller and Scarsini (2001), define the difference operator $\Delta_i^\varepsilon \psi(\mathbf{X}) := \psi(\mathbf{X} + \varepsilon \boldsymbol{\iota}_i) - \psi(\mathbf{X})$, where $\boldsymbol{\iota}_i$ is the ith unit vector in $\Re^N$ and $\varepsilon > 0$. The orders shown in Figure 1 are defined as follows:

1. A function ψ: $\Re^N \to \Re$ is *supermodular* if $\Delta_i^\varepsilon \Delta_j^\delta \psi(\mathbf{X}) \geq 0$ for all $\mathbf{X} \in \Re^N$, $1 \leq i < j \leq N$, $i \neq j$, and all $\varepsilon, \delta \geq 0$. $\mathbf{X}$ dominates $\mathbf{Y}$ in the *supermodular order*, written as $\mathbf{X} \succ_{\text{sm}} \mathbf{Y}$, if and only if $E[\psi(\mathbf{X})] \geq E[\psi(\mathbf{Y})]$ for all supermodular functions ψ. (See Müller and Scarsini 2001.)

2. A function ψ: $\Re^N \to \Re$ is *directionally convex* if $\Delta_i^\varepsilon \Delta_j^\delta \psi(\mathbf{X}) \geq 0$ for all $\mathbf{X} \in \Re^N$, $1 \leq i, j \leq N$, and all $\varepsilon, \delta \geq 0$. $\mathbf{X}$ dominates $\mathbf{Y}$ in the *directionally convex order*, written as $\mathbf{X} \succ_{\text{dcx}} \mathbf{Y}$, if and only if $E[\psi(\mathbf{X})] \geq E[\psi(\mathbf{Y})]$ for all directionally convex functions ψ. (See Müller and Scarsini 2001.)

3. A function ψ: $\Re^N \to \Re$ is *componentwise convex* if $\Delta_i^\varepsilon \Delta_i^\delta \psi(\mathbf{X}) \geq 0$ for all $\mathbf{X} \in \Re^N$, $1 \leq i \leq N$, and all $\varepsilon, \delta \geq 0$. $\mathbf{X}$ dominates $\mathbf{Y}$ in the *componentwise convex order*, written as $\mathbf{X} \succ_{\text{ccx}} \mathbf{Y}$, if and only if $E[\psi(\mathbf{X})] \geq E[\psi(\mathbf{Y})]$ for all componentwise convex functions ψ. (See Shaked and Shathikumar 1994.)

4. $\mathbf{X}$ dominates $\mathbf{Y}$ in the *positive linear-convex order*, written as $\mathbf{X} \succ_{\text{plcx}} \mathbf{Y}$, if and only if $\mathbf{a}^{\mathrm{T}}\mathbf{X} \succ_{\text{cx}} \mathbf{a}^{\mathrm{T}}\mathbf{Y}$ for all $\mathbf{a}^{\mathrm{T}} \in \Re_+^N$. (See Scarsini 1998.)

5. The *survival function* $\bar{F}$ corresponding to the multivariate distribution function F of a random vector $\mathbf{Z}$ is defined by $\bar{F}(z) = \Pr\{\mathbf{Z}_i > z_i \ \forall\, i = 1, \ldots, N\}$. $\mathbf{X}$ is *more positive lower orthant dependent* than $\mathbf{Y}$, written as $\mathbf{X} \succ_{\text{plo}} \mathbf{Y}$, if and only if $F(z) \geq G(z)\ \forall\, z \in \Re^N$. $\mathbf{X}$ is *more positive upper orthant dependent* than $\mathbf{Y}$, written as $\mathbf{X} \succ_{\text{puo}} \mathbf{Y}$, if and only if $\bar{F}(z) \geq \bar{G}(z)\ \forall\, z \in \Re^N$. (See Joe 1997.)

6. $\mathbf{X}$ is *more concordant* than $\mathbf{Y}$, written as $\mathbf{X} \succ_{\text{c}} \mathbf{Y}$, if and only if $\mathbf{X} \succ_{\text{plo}} \mathbf{Y}$ and $\mathbf{X} \succ_{\text{puo}} \mathbf{Y}$. (See Joe 1997.)

References

Alfaro, J. A., Charles J. Corbett. 2003. The value of SKU rationalization in practice (The pooling effect under suboptimal inventory policies). *Production Oper. Management* **12**(1) 12–29.

Clemen, R. T., T. Reilly. 1999. Correlations and copulas for decision and risk analysis. *Management Sci.* **45**(2) 208–224.

Corbett, C. J., K. Rajaram. 2004. Aggregation of uncertainty and multivariate dependence: The value of pooling of inventories under non-normal dependent demand. Working paper, UCLA Anderson School of Management, Los Angeles, CA.

Dong, L., N. Rudi. 2004. Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Sci.* **50**(5) 645–657.

Eppen, G. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* **25**(5) 498–501.

Eppen, G., L. Schrage. 1981. Centralized ordering policies in a multi-warehouse system with leadtimes and random demand. L. B. Schwarz, ed. *Multi-Level Production/Inventory Systems: Theory and Practice*. North-Holland, New York, 51–67.

Erkip, N., W. H. Hausman, S. Nahmias. 1990. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Sci.* **36**(3) 381–392.

Federgruen, A., P. Zipkin. 1984. Approximations of dynamic, multilocation production and inventory problems. *Management Sci.* **30**(1) 69–84.

Garg, A., H. L. Lee. 1999. Managing product variety: An operations perspective. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, 467–490.

Gerchak, Y., D. Mossman. 1992. On the effect of demand randomness on inventories and costs. *Oper. Res.* **40**(4) 804–807.

Groenevelt, H., N. Rudi. 2000. Product design for component commonality and the effect of demand correlation. Working paper, University of Rochester, Rochester, NY.

Ho, T. H., C. S. Tang. 1998. *Product Variety Management—Research Advances*. Kluwer Academic Publishers.

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

Jönsson, H., E. A. Silver. 1987. Analysis of a two-echelon inventory control system with complete redistribution. *Management Sci.* **33**(2) 215–227.

Müller, A., M. Scarsini. 2000. Some remarks on the supermodular order. *J. Multivariate Analysis* **73** 107–119.

Müller, A., M. Scarsini. 2001. Stochastic comparison of random vectors with a common copula. *Math. Oper. Res.* **26** 723–740.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, New York.

Nelsen, Roger B. 1999. *An Introduction to Copulas*. Springer, New York.

Netessine, S., G. Dobson, R. A. Shumsky. 2002. Flexible service capacity: Optimal investment and the effect of demand correlation. *Oper. Res.* **50**(2) 375–388.

Rudi, N. 2000. Optimal inventory levels in systems with common components. Working paper, University of Rochester, Rochester, NY.

Scarsini, M. 1998. Multivariate convex orderings, dependence, and stochastic equality. *J. Appl. Probab.* **35** 93–103.

Scarsini, M., M. Shaked. 1996. Positive dependence orders: A survey. C. C. Heyde, Y. V. Pohorov, R. Pyke, S. T. Rachev, eds. *Applied Probability*. Vol. 1 of *Athens Conference on Applied Probability and Time Series Analysis*. Springer Verlag, Berlin, Germany, 70–91.

Shaked, M., J. G. Shanthikumar. 1994. *Stochastic Orders and Their Applications*. Academic Press, Inc., Boston, MA.

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor networks: Dynamic inventory management and capacity investment with discretionary pooling. *Manufacturing Service Oper. Management* **4**(4) 313–335.