

This file was downloaded from BI Open Archive, the institutional repository (open access) at BI Norwegian Business School <http://brage.bibsys.no/bi>.

It contains the accepted and peer reviewed manuscript to the article cited below. It may contain minor differences from the journal's pdf version.

Grønneberg, S., & Foldnes, N. (2017). Covariance model simulation using regular vines. *Psychometrika*, 82(4), 1035-1051 DOI: <https://doi.org/10.1007/s11336-017-9569-6>

Copyright policy of Springer, the publisher of this journal:

"Authors may self-archive the author's accepted manuscript of their articles on their own websites. Authors may also deposit this version of the article in any repository, provided it is only made publicly available 12 months after official publication or later. He/ she may not use the publisher's version (the final article), which is posted on SpringerLink and other Springer websites, for the purpose of self-archiving or deposit..."

<http://www.springer.com/gp/open-access/authors-rights/self-archiving-policy/2124>

COVARIANCE MODEL SIMULATION USING REGULAR VINES (ACCEPTED IN PSYCHOMETRIKA)

STEFFEN GRØNNEBERG AND NJÅL FOLDNES

ABSTRACT. We propose a new and flexible simulation method for non-normal data with user-specified marginal distributions, covariance matrix and certain bivariate dependencies. The VITA (VIne To Anything) method is based on regular vines and generalises the NORTA (NORmal To Anything) method. Fundamental theoretical properties of the VITA method are deduced. Two illustrations demonstrate the flexibility and usefulness of VITA in the context of structural equation models. R code for the implementation is provided.

1. INTRODUCTION

Multivariate data simulation is the main tool in evaluating the finite-sample performance of estimation methods and goodness-of-fit measures in structural equation modeling (SEM). This paper proposes a new method for simulating non-normal random vectors with a pre-specified covariance matrix and pre-specified univariate marginal distributions. The method is very flexible and yields a large class of multivariate distributions, and is based on the statistical concept of a regular vine (R-vine). R-vines combine certain bivariate distributions that jointly define a multivariate distribution. In the proposed method, we search for bivariate distributions so that the resulting R-vine has a given covariance matrix and given univariate marginals. We refer to the new method as VITA (VIne To Anything). VITA supports highly non-normal features, such as tail dependence and deviations from symmetry not supported by elliptical distributions. In practice, this means that VITA may be used in simulation studies where a high level of control of multivariate non-normality is needed. We also illustrate how the flexibility of VITA allows a user to simulate data that resemble a given real-world sample. A VITA implementation conducted in the R (R Core Team, 2015) environment, and relying on the VineCopula package (Schepsmeier et al., 2015), is included in the supplementary material.

The present article is structured as follows. After an overview of related literature, we proceed to Section 2, where we briefly review copula and R-vine theory, and introduce the VITA algorithm. We also state a theorem implying that VITA is well-behaved, with a proof provided in the appendix. Section 3 contains two

illustrative examples of the VITA method. The first example, in Section 3.1, is a five-dimensional factor model where we specify normal marginals and interpolate between a fully multivariate normal distribution and a regular vine constructed by the VITA method. We illustrate that the normal-theory based likelihood ratio (NT-LR) statistic is not robust to multivariate non-normality, even when the marginals are exactly normal. In Subsection 3.2 we identify a regular vine that emulates the well-known Holzinger-Swineford nine-dimensional dataset. We use non-normal marginals and consider two sets of bivariate copulas, one containing only normal copulas, and one containing a mixture of non-normal copulas determined by the data. We illustrate that the two resulting distributions, although sharing the same marginals and covariance matrix, differ with respect to the performance of the NT-LR statistic. In these two illustrations we employ a new and simple interpolation idea that allows us to investigate the transition from established covariance simulation techniques to VITA. The concluding section contains a discussion of VITA.

1.1. An overview of related methods. A well-known method for simulating data with pre-specified covariance matrix and univariate marginals is the NORTA (NORmal To Anything) method of Cario & Nelson (1997) and Li & Hammond (1975). This method is popular in the time series and computer science literature, but has to our knowledge not been employed much in the SEM literature (a notable exception is the PhD thesis of Boomsma (1983)). In SEM the technique of Vale & Maurelli (1983) is the most widely used non-normal simulation method. The Vale-Maurelli (VM) approach is implemented in software packages like Mplus (Muthén & Muthén, 2012), EQS (Bentler, 2006) and lavaan (Rosseel, 2012), and Monte Carlo studies have employed the VM technique for data generation over a span of several decades (e.g., Bentler & Tong, 2013; Fouladi, 2000; Curran et al., 1996). The VM technique only allows the specification of univariate moments, in addition to the covariance matrix, while NORTA supports near arbitrary marginals. However, the NORTA method has a Gaussian/normal¹ dependence structure, formalized by the simulated random vector’s copula. This limits its applicability for simulation studies focusing on the finite sample performance of statistical methods for non-normal data. The situation is similar for the VM method. Foldnes & Grønneberg (2014) showed theoretically that the VM method produces data with either an exact or near-exact Gaussian copula, and illustrated by Monte Carlo that this might result in overly optimistic evaluations of the performance of normal-theory based maximum likelihood estimation in SEM under non-normal data conditions. Hence, the evaluation of the finite sample behaviour of statistical techniques using VM or NORTA may be skewed in favour of methods that work well with data having a Gaussian dependence structure. Recently Foldnes & Olsson (2016) proposed a

¹We use the terms “Gaussian” and “normal” interchangeably.

simulation method with non-Gaussian dependence, but with little control over the copula and the marginals beyond matching moments.

A simple copula-based approach to non-normal data simulation has recently been proposed by Mair et al. (2012). Their two-stage method allows simulation from a large class of distributions with non-Gaussian copulas. The basic idea is to first freely choose a multivariate copula and univariate marginals. The resulting distribution will not match the target covariance matrix, and this is remedied by applying a linear transformation to the simulated data. This assures that the simulated data match the pre-specified covariance matrix. However, it will also transform both the marginals and the copula. Hence, the approach proposed by Mair et al. (2012) does not yield control neither over the marginals, nor over the underlying copula. For instance, it is not clear how to use their approach to simulate data whose univariate marginals are normally distributed and whose underlying copula is non-normal. This task is achievable with VITA, which gives complete control of the marginals and considerable control of the underlying copula.

Our simulation technique extends the NORTA methodology to the class of regular vines. While NORTA starts with a normal vector, our proposed method starts with a regular vine (R-vine). The original motivation for vines in Joe (1996) and Bedford & Cooke (2002) was partly to generalize certain properties of the Gaussian distribution and to nest the Gaussian distribution as a special case. It is therefore natural to connect covariance models with vines. However, this connection has not yet been fully studied for general non-elliptical copulas. Exceptions are Biller (2009) and Erhard & Czado (2009). Erhard & Czado (2009) use a vine-based data generating process with prescribed Pearson covariance matrix in the specific context of discrete Poisson margins and restrict attention to so-called C-vines. The copula-based Vector AutoRegressive To Anything (VARTA) method of Biller (2009) is much closer to the present work, but aims at simulating covariance-based time series models, and also restrict attention to C-vines and a small class of copulas. Since Biller (2009) is working in a time-series context, stationarity of the resulting vectorial time series must be fulfilled. This imposes difficult-to-check restrictions on the C-vine copula structure which limits the applicability of the technique in its current form. While the time-series focus of Biller (2009) presents some practical limitations for the VARTA method, it does logically nest the VITA method for simulating a multivariate time series with only one time-period, provided that VARTA is extended to R-vines. However, Biller (2009) does not provide theoretical results for the behaviour of the criterion function underlying the VARTA algorithm, in contrast to the analysis of VITA presented in this article.

Finally, we note that vine copulas are often employed in multivariate simulation studies where the focus is not on covariance modeling, but instead on properties such as tail dependency and asymmetry. In these studies a simulation approach for

vines is to fix values of a dependence measure, such as Spearman's ρ or Kendall's τ , on each edge of the vine, and then do sensitivity analyses of tail properties by choosing different bivariate copula families with the same value of the dependence measure. However, the present study has another focus, where covariance values are to be kept fixed. Our aim is to adapt the vine framework for this setting.

2. COVARIANCE MODEL SIMULATION USING REGULAR VINES

In Sections 2.1 and 2.2 preliminary theory is presented. Our algorithm is then presented in Section 2.3.

2.1. Basic copula theory and briefly on the NORTA method. We refer to the monograph by Joe (2014) for a thorough review of copula theory, and limit ourselves here to some basic definitions. Sklar (1959) noticed that a full description of a random vector $X = (X_1, X_2, \dots, X_d)^t$ may be obtained by separating the marginals from the dependence structure,

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)),$$

where C is a copula of X , that is, a distribution on $[0, 1]^d$ with uniform marginals. In the case of continuous marginal distributions the copula is unique.

Suppose given univariate marginals F_1, F_2, \dots, F_d and a family C_θ of copulas indexed by some parameter $\theta \in \Theta \subseteq \mathbb{R}^p$. Let $F_j^{-1}(t) := \inf\{x : F_j(x) \geq t\}$ for $1 \leq j \leq d$. If F_j is continuous, F_j^{-1} coincides with the unique quantile function. If

$$U := (U_1, U_2, \dots, U_d)' \sim C_\theta,$$

that is, if U has the cdf C_θ , then

$$X := (F_1^{-1}(U_1), F_2^{-1}(U_2), \dots, F_d^{-1}(U_d))'$$

has marginal distributions F_1, F_2, \dots, F_d and copula C_θ . Let μ_i denote the expectation of X_i , for $i = 1, \dots, d$. Then the covariances are given by $\sigma_{i,j}(\theta) := \mathbb{E}_\theta[F_i^{-1}(U_i)F_j^{-1}(U_j)] - \mu_i\mu_j$. When the univariate marginals are fixed, the covariance matrix $\Sigma(\theta) := (\sigma_{i,j}(\theta))_{1 \leq i,j \leq d}$ of X depends only on the bivariate copulas.

In particular, the Gaussian copula is defined as follows. Let (Z_1, Z_2) denote a bivariate Gaussian vector with standard normal marginals and correlation ρ , and let $\Phi(\cdot)$ denote the standard univariate Gaussian cumulative distribution function. Then the bivariate Gaussian copula with parameter $\theta = \rho$ is the distribution of the vector $(\Phi(Z_1), \Phi(Z_2))$. This naturally extends to higher dimension $d > 2$, in which case the Gaussian copula is parametrised by a vector θ containing all pairwise correlations between the elements of (Z_1, \dots, Z_d) . The fundamental idea of the NORTA method is to let $\{C_\theta : \theta \in \Theta\}$ be the class of all Gaussian copulas, and then select θ so that the covariance matrix $\Sigma(\theta)$ is correct.

We note that while any population covariance matrix can be reproduced using Gaussian distributions, this is no longer the case when one changes the marginals

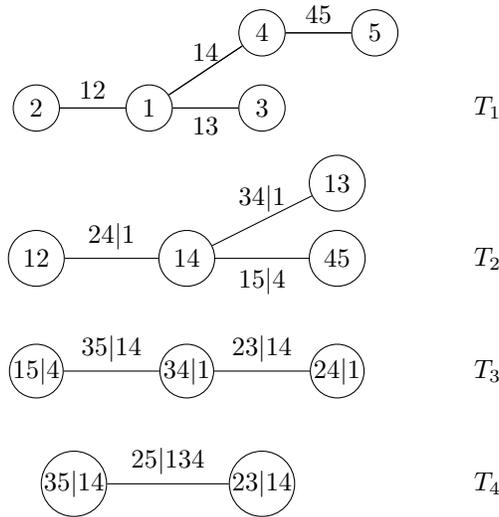


FIGURE 1. A five-dimensional regular vine.

even when still using a Gaussian copula. For a thorough discussion of the restrictions imposed on a covariance structure by fixing the univariate marginals, see Joe (2006).

2.2. On regular vines. R-vines form a flexible class of distributions, described and proved to exist in full generality in Bedford & Cooke (2002). A precursor to Bedford & Cooke (2002) is Joe (1996). Currently, R-vines is an active field of statistical research, with an increasing amount of applications in many fields. See Joe (2014) for an authoritative presentation.

In SEM, we are used to specify the pairwise dependencies among variables in terms of covariances. Instead of unconditional covariances, R-vines on d variables require the specification of $d - 1$ unconditional bivariate copulas and $d(d - 1)/2 - (d - 1)$ conditional bivariate copulas. Hence, R-vines involve specifying the same number, $d(d - 1)/2$, of bivariate connections, as there are pairwise covariances between the d variables (excluding variances). Hence we may determine an R-vine for a given covariance matrix by matching each bivariate copula in the R-vine to one and only one specific covariance.

We illustrate a specific R-vine on $d = 5$ variables in Figure 1. Note that there are ten edges in the R-vine, the same number as there are pairs of variables. That is, each pair of variables is represented by an edge in the R-vine. The edges in the first tree represent unconditional bivariate copulas. At the higher levels, the edges represent bivariate copulas of certain conditional distributions. Each edge corresponds to a copula linking two distributions on a lower level through conditioning. We can now fully specify an R-vine copula by supplementing the graphical representation

in Figure 1 with specifications of the five univariate marginals and the ten bivariate copulas, which are the copulas of conditional distributions. For instance, in tree T_4 , the edge 25|134 refers to the copula of (X_2, X_5) conditional on (X_1, X_3, X_4) . As an example, we can specify that all univariate marginals are standard normal, and that all bivariate copulas belong to the family of Clayton copulas. This family is parametrised by a single dependency parameter. The R-vine and the marginals are then combined to fully specify a 5-dimensional distribution after each of these dependency parameters are fixed.

We now proceed with a technical definition of R-vines. A *tree* on a set of nodes is an acyclic connected graph.

- Definition 1.** (1) $\mathcal{V} = (T_1, \dots, T_{d-1})$ is an R-vine on d elements if
- (A) T_1 is a tree with nodes $N_1 = \{1, \dots, d\}$ and a set of edges denoted E_1 .
 - (B) For $i = 2, \dots, d-1$, each T_i is a tree with nodes $N_i = E_{i-1}$ and edge set E_i .
 - (C) For $i = 2, \dots, d-1$ and $\{a, b\} \in E_i$ with $a = \{a_1, a_2\}$ and $b = \{b_1, b_2\}$, the proximity condition $\#(a \cap b) = 1$ must hold, where $\#$ denotes the cardinality of a set.
- (2) $(\mathbf{F}, \mathcal{V}, B)$ is an R-vine copula specification if $\mathbf{F} = (F_1, \dots, F_d)$ is a vector cumulative distribution functions, \mathcal{V} is a d -dimensional R-vine and $B = \{B_e : i = 1, \dots, d-1; e \in E_i\}$ is a set of copulas with B_e being a bivariate copula CDF.²
- (3) A *parametric R-vine copula specification* is a class of R-vine copula specifications $\{(\mathbf{F}, \mathcal{V}, B_\theta) : \theta \in \Theta\}$, where Θ is a subset of \mathbb{R}^q for some q .

Definition 2. The complete union of an edge $e_i \in E_i$ is the set

$$U_{e_i} := \{n_1 \in N_1 : \text{there exists a } e_j \in E_j, j = 1, \dots, i-1, \\ \text{with } n_1 \in e_1 \in e_2 \in \dots \in e_{i-1} \in e_i\} \subseteq N_1.$$

For an edge $e = \{a, b\}$, the conditioning set is $\mathfrak{D}_e := U_a \cap U_b$. The conditioned variables associated with e are $\mathbf{c}_{e,a} = U_a \setminus \mathfrak{D}_e$ and $\mathbf{c}_{e,b} = U_b \setminus \mathfrak{D}_e$.

We say that a joint distribution F of a random vector (X_1, \dots, X_d) realizes the R-vine copula specification $(\mathbf{F}, \mathcal{V}, B)$, if, firstly, the univariate marginals of F correspond to \mathbf{F} . Secondly, for each edge $e = \{a, b\}$, let $i = \mathbf{c}_{e_k, a}, j = \mathbf{c}_{e_k, b}$. Then B_e is the bivariate copula of the conditional distribution function $F_{e|\mathfrak{D}_e}$ of X_i and X_j , given $(X_k : k \in \mathfrak{D}_e)$. See Bedford & Cooke (2002) for the existence and uniqueness of F . A consequence of the R-vine construction is that vine distributions always fulfil the so-called “simplifying assumption”, namely that the copula of $F_{e|\mathfrak{D}_e}$ is

²The notation B is a mnemonic for bivariate.

functionally independent of the values in the conditioning variables in \mathfrak{D}_e , for all edges e . Some consequences of this is discussed in Hobæk Haff et al. (2010).

Finally, given an edge $e \in E_l$ in an R-vine \mathcal{V} , there corresponds a subvine $\mathcal{V}(e)$ whose edge set consists of all the edges in \mathcal{V} whose complete union is contained in the complete union of e . This result is given in the appendix.

2.3. The VITA algorithm. The VITA algorithm traverses the given R-vine edge by edge in a specified order that has to fulfil the following condition.

Definition 3. *A valid order \mathcal{O} is a sequence of unique elements $(e_k : k = 1, \dots, d(d-1)/2)$, where $e_k \in E_l$ for some l . Moreover, for $l \geq 2$ and $e_k = \{a, b\}$, there exists $1 \leq m, n \leq k-1$ such that $a = e_m$ and $b = e_n$.*

If this conditions holds, then each edge e in the sequence is connecting nodes that have already been processed by VITA. A canonical valid order is to first include all edges in T_1 and all edges in T_2 , and so forth, where the specific order within each tree is arbitrary.

The VITA algorithm runs through a loop of length $d(d-1)/2$, as dictated by a valid edge order \mathcal{O} . Each step corresponds to an edge e_k , and it is sufficient to consider the sub-vine associated with e_k . Let $i = \mathbf{c}_{e_k, a}, j = \mathbf{c}_{e_k, b}$. Suppose copula parameters $\theta_{e_1}, \dots, \theta_{e_k}$ are given for all edges up to k . From this sequence we can pick the copula parameters for each edge in the sub-vine $\mathcal{V}(e_k)$. Not all parameters in $\theta_{e_1}, \dots, \theta_{e_k}$ correspond to copulas in $\mathcal{V}(e_k)$, but for notational simplicity, we ignore this redundancy and write $C_{i,j}(u_i, u_j; \theta_{e_1}, \dots, \theta_{e_k})$ for the *unconditional* copula of (X_i, X_j) , i.e., the coordinates of a random vector that realizes the R-vine copula specification induced by $\mathcal{V}(e_k)$. Note that this unconditional copula is different from B_e , whenever e is not in T_1 . Let Θ_{e_k} be the parameter space of $B_{e_k, \theta_{e_k}}$. For fixed $\theta_{e_1}, \dots, \theta_{e_{k-1}}$, define the function $I_{e_k} : \Theta_{e_k} \mapsto \mathbb{R}$ via

$$(1) \quad I_{e_k}(\theta_{e_k}) = \int_0^1 \int_0^1 F_i^{-1}(u_i) F_j^{-1}(u_j) dC_{i,j}(u_i, u_j; \theta_{e_1}, \dots, \theta_{e_k}) - \mu_i \mu_j - \Sigma_{i,j}.$$

The crucial step in VITA is to solve the equation $I_{e_k}(\theta_{e_k}) = 0$, and hence ensure that the corresponding covariance is attained.

In Algorithm 1 the VITA algorithm is presented. It assumes that the user has specified univariate marginal CDFs F_i for each variable X_i , $i = 1, \dots, d$, with finite second order moments, and target covariance matrix Σ . Note that variances of each margin are given by the diagonal of Σ , and must be compatible with the marginals F_1, \dots, F_d . Also, the user has specified an R-vine \mathcal{V} with bivariate copula families B_{e, θ_e} for each edge e . These bivariate copulas have dependence parameters θ_e that needs to be determined by VITA, so that the resulting distribution has the user-specified target covariance matrix Σ . The order in which parameters are determined follows the user-specified valid order \mathcal{O} .

Algorithm 1 VITA

```

1: procedure VITA( $\mathbf{F}, \mathcal{V}, B, \mathcal{O}, \Sigma$ )
2:   for  $k \leftarrow 1, \dots, d(d-1)/2$  do
3:      $\theta_{e_k} \leftarrow \operatorname{argmin}_{\theta_{e_k}} |I_{e_k}(\theta_{e_k}; \theta_{e_{k-1}}, \dots, \theta_{e_1})|$  from eq. (1)
4:      $\epsilon_{e_k} \leftarrow |I_{e_k}(\theta_{e_k}; \theta_{e_{k-1}}, \dots, \theta_{e_1})|$ 
5:   end for
6:   return  $(\theta_{e_1}, \dots, \theta_{e_{d(d-1)/2}}), (\epsilon_{e_1}, \dots, \epsilon_{e_{d(d-1)/2}})$ 
7: end procedure

```

The following theorem, whose proof is presented in the appendix, establishes the continuity and monotonicity of the criterion function $\theta_{e,k} \mapsto I_{e_k}(\theta_{e_k}; \theta_{e_{k-1}}, \dots, \theta_{e_1})$, when the bivariate copulas all have just one parameter. Note that its increasing mapping requirement holds for all popular one-parameter bivariate copulas (Joe, 2014).

Theorem 1. *Suppose for any edge e in \mathcal{V} that $\Theta_e \subset \mathbb{R}$ and that $\theta_e \mapsto B_{e,\theta_e}(u_1, u_2)$ is an increasing mapping for each $(u_1, u_2) \in (0, 1)^2$. Suppose further the conditional quantile function of X_i for $i \in U_e$, conditioned on $X_j = x_j$ for $j \neq i$ and all $j \in U_e$, is changing continuously with respect to θ_e and $(x_j : j \neq i, j \in U_e)$. Finally, suppose that each marginal F_j has finite $2 + \varepsilon$ -moments for some $\varepsilon > 0$ and that we are given a valid order \mathcal{O} . Then the criterion function is continuous and monotonously increasing.*

A consequence of the continuity and monotonicity of the criterion functions is that any solution is unique, if it exists. Hence, if VITA can attain the target covariance matrix, then any valid search order \mathcal{O} will find the unique solution. That is, we may use a canonical search order.

If a solution exists that can be found by following the sequence in \mathcal{O} , then all errors $\epsilon_{e_1}, \dots, \epsilon_{e_{d(d-1)/2}}$ will be equal to zero. Exactly which covariance matrices that are supported by VITA is currently unknown. However VITA will always result in a distribution with a covariance matrix as close as possible to the given target. In the case where the VITA covariance matrix is not exactly equal to the target, i.e. not all the ϵ_{e_k} are equal to zero, a linear transformation can be applied to the VITA vector that guarantees correct target covariance (Mair et al., 2012). This will distort the marginals and change the copula. However, if the VITA covariance matrix is sufficiently close to the target covariance matrix, the distortion caused by the linear transformation is small, being close to the identity transformation.

Note that we cannot in practice exactly evaluate the integral in each VITA step. In our implementation, we evaluate this integral by Monte Carlo integration, using the VineCopula package (Schepsmeier et al., 2015). The technical details of using

Monte Carlo to approximate a criterion function in the optimisation step are non-trivial. Homem-de Mello & Bayraksan (2014) is one of many surveys on the topic, with many useful references. We consider the treatment of this issue outside the scope of the paper.

VITA determines the bivariate dependencies sequentially. As we move up in the tree structures, the numerical routine becomes more and more computationally demanding, even though the numerical optimisation is always with respect to a low-dimensional (often unidimensional) parameter. With current computer speed and for high dimensions (say, $d \geq 10$) VITA nevertheless needs considerable running time to determine the full vine distribution, in its current prototype implementation. Since the computationally demanding steps can easily be parallelised, VITA will tackle also higher-dimensional cases in future implementations. It is also important to notice that once VITA has terminated, fast simulation from the resulting vine distribution is available, provided the bivariate copulas in the R-vine have fast implementation of conditional quantile functions. This is the case for all commonly used copula families.

The generality of VITA will increase as new bivariate copula families are introduced in the literature. A long list of copulas is presented in Chapter 4 of the monograph Joe (2014). Classical copulas such as the Joe and Clayton copulas express only positive dependence. However, these copulas may be rotated to allow for negative dependence as well. Since all standard copula classes include the independence copula, it is possible to glue together unrotated and rotated copula families, such that the constructed copula family supports both positive and negative dependence, in a continuous manner.

Finally, we mention that starting values for the minimization step of VITA can easily be obtained for those one-parameter copula families that can be parametrized by a dependency measure such as Kendall's tau, through setting the starting value to that which is induced by the normal copula case: If the vine distribution were to consist solely of normal copulas and normal marginals, it is parametrized by the means and variances of the marginals and the partial correlations of $\mathbf{c}_{e,a}$, $\mathbf{c}_{e,b}$ conditioned on \mathfrak{D}_e for each edge e with algorithms to pass between the correlation matrix and the partial correlations given in Chapter 6.15 of Joe (2014). Once these partial correlations are fixed, one can compute the induced Kendall's tau τ_e of each of the partial correlations ρ_e parameterizing the bivariate normal copula for edge e via the formula $\tau_e = 2\pi^{-1} \arcsin(\rho_e)$, see Joe (2014, Section 4.3.1). The starting value for the search for θ_e of the actual copula B_e for edge e is then the parameter that induces a Kendall's tau equal to τ_e .

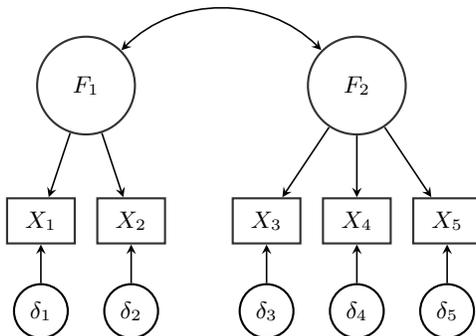


FIGURE 2. Factor model.

3. ILLUSTRATIONS

We provide two illustrations of the VITA method. In both illustrations, the obtained VITA distribution is then used to study the sensitivity of the NT-LR statistic to varying degrees of non-normality. The NT-LR statistic is used to evaluate model fit and is of central importance in SEM practice. It is reported by default in all SEM software packages. We use a new and simple interpolation technique that allows us to move from current simulation practices to VITA simulation, in a continuous manner. Model estimation and vine calibration were respectively done with the `lavaan` (Rosseel, 2012) and `VineCopula` (Schepsmeier et al., 2015) packages. Source code is provided as online supplementary material.

3.1. On the NT-LR test for a factor model with exactly normal margins.

Consider a simple factor model with two factors and five indicator variables, as depicted in Figure 2. A data-generating process is obtained from this model by fixing all five factor loadings to 0.95, the factor variances to one, and the correlation between F_1 and F_2 to 0.90. The choice of high factor loadings and factor correlation was motivated by the fact that the non-Gaussian copulas we will consider are farthest away from the Gaussian copula under strong dependence. Low factor loadings and correlation would lead to weak dependence, which would imply that the Gaussian and non-Gaussian cases are quite similar. The variances of the errors δ_i are fixed so that the X_i , $i = 1, \dots, 5$, have unit variance. This gives the following target covariance matrix for $X = (X_1, X_2, X_3, X_4, X_5)'$:

$$\Sigma = \begin{pmatrix} 1.000 & 0.902 & 0.812 & 0.812 & 0.812 \\ 0.902 & 1.000 & 0.812 & 0.812 & 0.812 \\ 0.812 & 0.812 & 1.000 & 0.902 & 0.902 \\ 0.812 & 0.812 & 0.902 & 1.000 & 0.902 \\ 0.812 & 0.812 & 0.902 & 0.902 & 1.000 \end{pmatrix}$$

Next we pre-specify the marginals of the X_i , $i = 1, \dots, 5$, to be standard normal. That is, $F_i = \Phi$ for $i = 1, \dots, 5$. There are many distributions on five variables with standard normal marginals and covariance matrix Σ . One such distribution is the Gaussian distribution with zero mean vector and covariance matrix Σ . This is the distribution that NORTA produces, in the special case of normal marginals. To the best of our knowledge, no existing simulation procedure, besides VITA, is able to produce distributions with normal marginals and a non-normal copula, targeting a given covariance matrix. The parametric R-vine copula specification we choose was discussed in Section 2.2. We constructed three R-vines based on three bivariate copula families. The first vine was defined with all ten bivariate copula families B_{e,θ_e} belonging to the Clayton family. The Clayton copula is parametrised by the dependence parameter $\theta_e \in (0, \infty)$. The VITA procedure runs through each edge in trees T_1, T_2, T_3 and T_4 and determines θ_e in each case. The R-vine is depicted in Figure 3, with calibrated Clayton parameters for each edge. The numerical integration routine is only approximate, so, for instance, the calibrated Clayton parameters for the unconditional pairs (1, 2) and (4, 5) in T_1 are not exactly equal. The true parameters for these edges must be equal, since the covariances in Σ corresponding to these pairs are identical. The Clayton copula is tail asymmetric, with lower, but not upper, tail dependence. In the second vine specification we reversed the situation with respect to the direction of tail dependence, setting all ten bivariate copulas to belong to the Gumbel family of copulas. This family has upper, but not lower, tail dependence. The third vine was specified with all bivariate copulas belonging to the Student t family of copulas with five degrees of freedom, which is tail symmetric, and has strong upper and lower tail dependence.

After calibration of the Clayton, Gumbel and Student t R-vines we assessed their accuracy in reproducing the target covariance matrix by evaluating the asymptotically distribution-free (ADF) statistic T_{ADF} (Browne, 1984) in large samples. The procedure, suggested by Mair et al. (2012), specifies a model with no free parameters, by fixing all parameters so that the model-implied covariance matrix is equal to Σ . The model is "fitted" in standard SEM software to large samples drawn from the calibrated vine, and T_{ADF} is calculated in each sample. If the vine reproduces Σ , then T_{ADF} is asymptotically distributed as a chi-square with 15 degrees of freedom. So we can test the calibrated VITA structure by simulating large samples from it, and check whether the rejection rate is close to the nominal level. We ran 1000 replications, each with sample size $n = 10^4$. In each sample, T_{ADF} was calculated using the `lavaan` package (Rosseel, 2012). The empirical rejection rate for the ADF test at the $\alpha = 0.05$ nominal level, was 0.065, 0.061 and 0.056 for the Clayton, Gumbel and Student t copulas, respectively, none of which are significantly different from 0.05. We conclude that all three calibrated R-vines reproduce Σ accurately.

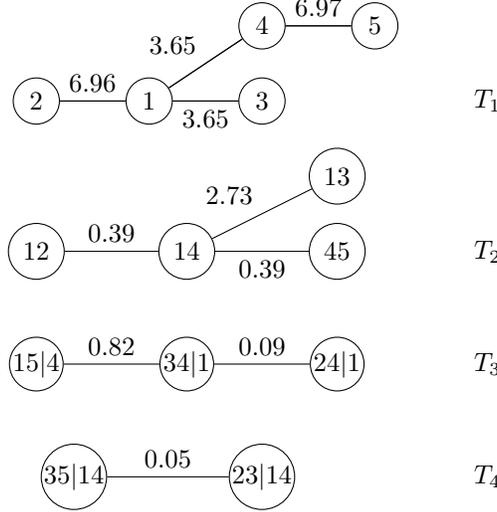


FIGURE 3. The calibrated VITA Clayton parameter θ_e , for each edge e .

Let us next evaluate the sensitivity of the NT-LR statistic to the underlying copula, as we move from a fully Gaussian copula to the non-Gaussian copula implied by each of the three R-vine specifications. We do this by the following interpolation scheme. Consider distributions whose stochastic representation are obtained as an interpolation between the Gaussian random vector X and the VITA random vector Y , both with standard normal marginals. That is, for some $0 \leq \alpha \leq 1$ and independently distributed X and Y , consider the following random vector:

$$(2) \quad Q(\alpha) = \sqrt{1-\alpha} \cdot X + \sqrt{\alpha} \cdot Y.$$

Due to the independence between X and Y , and the fact that both X and Y have the covariance matrix Σ , it follows that $Q(\alpha)$ also has covariance matrix Σ , with marginals that are standard normally distributed. At one end of the interpolation, for $\alpha = 0$, we obtain a normally distributed vector $Q(0)$. At the other end of the interpolation, with $\alpha = 1$, the resulting vector $Q(1)$ comes from VITA. Consequently, we have a smooth transition starting with the multivariate normal case for $\alpha = 0$, and moving increasingly toward the VITA vector at $\alpha = 1$. If the NT-LR statistic is sensitive to an increasingly non-normal copula, even while holding the marginals fixed to the univariate normal distribution, we would expect the rejection rates of the NT-LR statistic to deviate more and more from the nominal level as α increases from 0 to 1. We conducted a simulation study, where the factor model depicted in Figure 2 was estimated for each level of α and for each sample size. We let α move from 0 to 1 in 10 discrete steps, under two sample sizes: $n = 100$ and $n = 1000$. In each cell we replicated 2000 samples from each of the three distributions and collected the rejection rate of the NT-LR statistic at the nominal 5% level. Figure

4 depicts the rejection levels against α , together with 95% confidence intervals. At $\alpha = 0$, the multivariate normal case, rejection rates are close to the nominal 5% level in all conditions. For all three VITA distributions, the rejection rates increase as we approach $\alpha = 1$, where data are generated completely from Clayton, Gumbel or Student t R-vines. This shows that the NT-LR statistic is sensitive to the underlying copula. The Clayton and Gumbel R-vines result in rejection rates at $n = 100$ of 17% and 12%, respectively, and at $n = 1000$ of 13% and 12%, respectively. For the strongly tail dependent and tail symmetric Student t bivariate copula, however, the resulting R-vine yields much higher rejection rates of 24% and 28%, for $n = 100$ and $n = 1000$, respectively. This suggests that tail dependency may be more of a challenge to NT-LR robustness than tail asymmetry.

To accentuate this effect, we repeated the process with a Student t R-vine, specified in the same manner as above, but with a degree of freedom equal to 2.01. The resulting VITA was deemed to accurately reproduce the target covariance, with the ADF resulting in a rejection rate of 0.065% with 1000 replicated samples of sample size $2 \cdot 10^4$. The rejection rate of NT-LR at sample size 1000 was 73.7%. Note that while this copula is very far away from the Gaussian copula, and is a rather extreme example most likely never encountered in real data, the covariance model and the standard assumptions underlying its estimation via the (quasi) Gaussian MLE (e.g. described in Satorra, 1989) hold. Specifically, note that since the marginals have finite moments of all (positive integer) orders, the Hölder inequality $[\mathbb{E}|\prod_{j=1}^d X_j|^r]^{1/r} \leq \prod_{j=1}^d [\mathbb{E}|X_j|^{p_j}]^{1/p_j}$ where $r > 0$, $p_j > 0$ and $\sum_{j=1}^d p_j^{-1} = r^{-1}$, implies that also the full R-vine distribution has finite multivariate moments of all orders.

3.2. Revisiting the Holzinger-Swineford data. The data consists of mental ability test scores of 301 seventh- and eighth-grade children. Only 9 out of the original 26 tests are included. The CFA model consists of three latent variables, each with three indicators, as depicted in Figure 5. We wish to simulate data with a distribution that is close to the true, but unknown data-generating distribution, and with the same covariance matrix as the estimated model-implied covariance matrix when fitting the factor model to the Holzinger-Swineford data using the Gaussian maximum-likelihood estimator. We denote this target covariance matrix by Σ .

In order to identify a distribution which closely resemble those found in the original dataset and having the desired covariance matrix, the following procedure was followed. First, we determine univariate marginals that model the observed marginals fairly accurately. The main purpose here is to illustrate, so we proceed by simply visually inspecting the histograms of each X_i and fitting suitable distributions to the data using the marginal maximum likelihood estimator. The estimated distributions were then re-scaled to achieve the variance same as given by

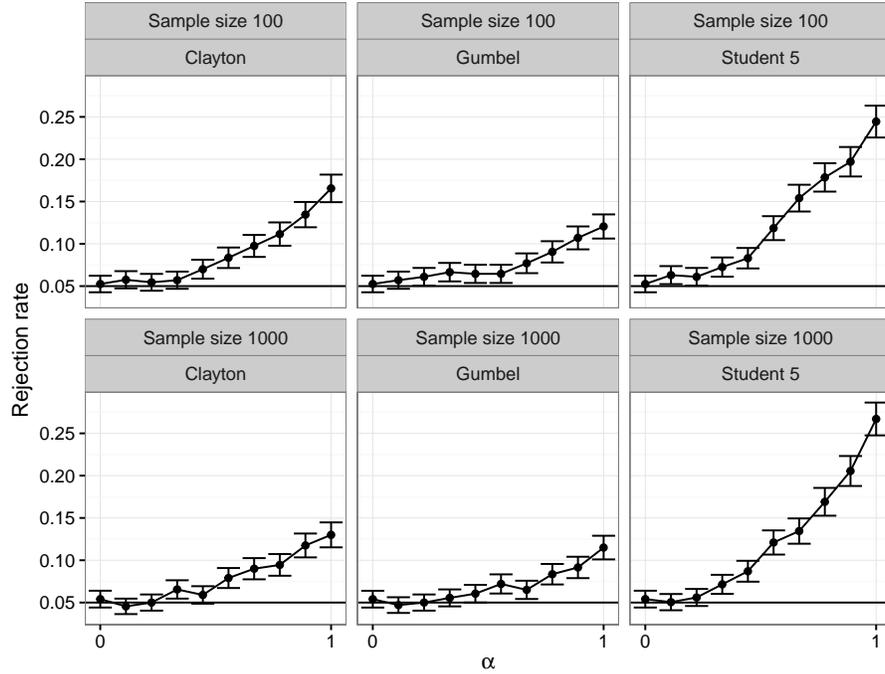


FIGURE 4. Rejection rates of the NT-LR statistic at sample size $n = 100$ and $n = 1000$. Error bars represent 95% confidence intervals. $\alpha = 0$ corresponds to a multivariate normal distribution, while $\alpha = 1$ corresponds to a VITA distribution with bivariate copulas belonging to families Clayton, Gumbel or Student t with 5 degrees of freedom.

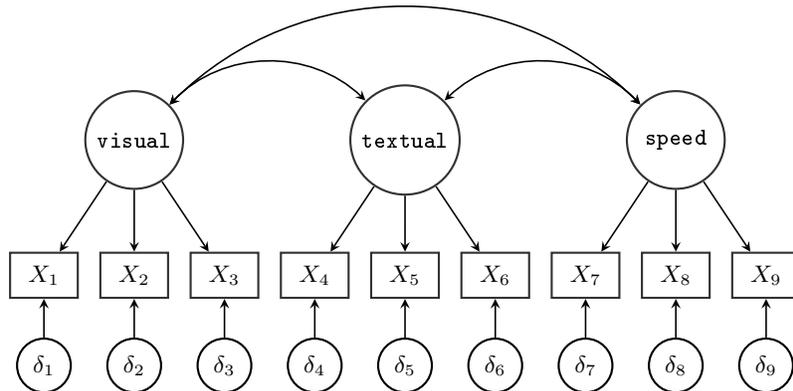


FIGURE 5. Holzinger-Swineford factor model.

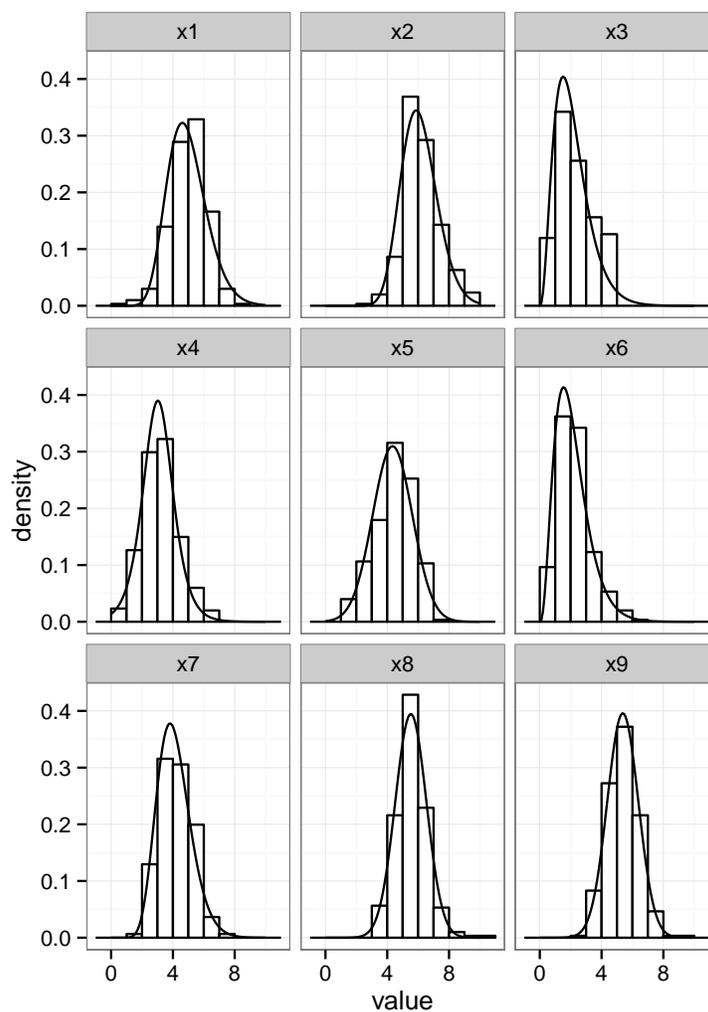


FIGURE 6. Histograms of nine variables from the Holzinger-Swineford sample, together with density curves of the fitted distributions.

Σ . These adjustments were all quite small. The result is shown in Figure 6, where the histograms of the observed data is shown together with density curves for the fitted distributions. For X_1, X_2, X_3, X_6 and X_7 we fitted a gamma distribution, while X_4 was fitted with a logistic distribution. The remaining variables, X_5, X_8 and X_9 , were fitted with a normal distribution.

Secondly, we fit an R-vine copula specification to the Holzinger-Swineford data. This involves to determine an appropriate R-vine, and then to choose a bivariate

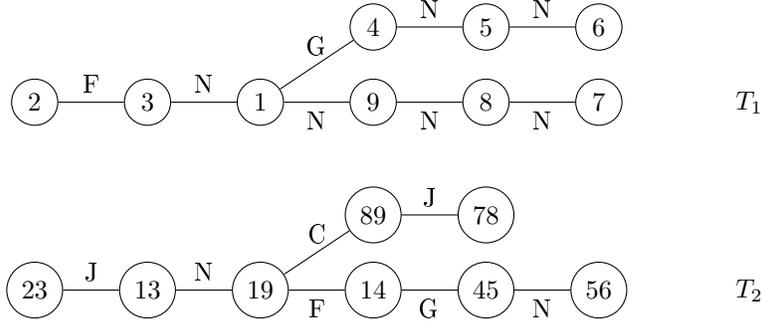


FIGURE 7. The first two trees in RV_1 . Copulas: N=Normal, J=Joe, C=Clayton, G=Gumbel, F=Frank.

copula family for each edge in the R-vine. The third step is to run VITA to calibrate the parameter for each of the bivariate copulas.

Dissmann et al. (2013) describe a way to carry out the second step, implemented in the `RVineStructureSelect` function from the `VineCopula` package (Schepsmeier et al., 2015). We used this routine with the following one-parameter copula families: Gaussian, Clayton, Gumbel, Frank and Joe. We refer to the resulting R-Vine as RV_1 , whose first two trees are depicted in Figure 7. The edges are annotated with the chosen family of bivariate copula. In total, RV_1 has 12 Gaussian, 6 Clayton, 3 Gumbel, 11 Frank and 4 Joe bivariate copulas.

The accuracy of RV_1 to reproduce the target covariance matrix Σ was assessed by the ADF test. We replicated 1000 samples from RV_1 , each with a sample size of $n = 10^4$. The empirical rejection rate was 0.062, fairly close to the nominal 0.05, so we deem RV_1 to accurately reproduce Σ . We conclude that VITA allows us to simulate from a distribution mimicking the empirical distribution in the Holzinger-Swineford sample, and whose covariance matrix is identical to the model-implied covariance matrix obtained from fitting the factor model in Figure 5 to this sample.

Finally, to compare RV_1 with NORTA, we construct a new R-vine RV_2 with the same vine structure and the same marginals as RV_1 , but with only normal bivariate copulas. This means that the aggregated nine-dimensional copula is normal, and hence that RV_2 is equivalent to NORTA. The accuracy of RV_2 was assessed with the ADF procedure, with an empirical rejection rate of 0.060 over one thousand replications.

RV_1 and RV_2 share the same covariance matrix and marginals, but differ with respect to underlying copula. However, this difference may not be substantial, given that many of the bivariate copulas in RV_1 are in fact normal. In particular, six of the eight unconditional pairs in T_1 are normal. To investigate whether the NT-LR statistic is affected by the transition from NORTA (RV_2) to VITA (RV_1), we again employed the interpolation scheme from eq. (2), where $\alpha = 0$ and $\alpha = 1$

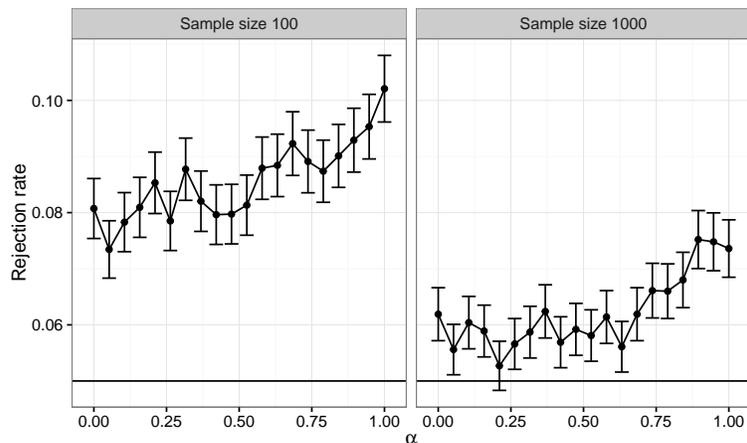


FIGURE 8. Rejection rates of the NT-LR statistic at sample sizes $n = 100$ and $n = 1000$. Error bars represent 95% confidence intervals. $\alpha = 0$ corresponds to NORTA, while $\alpha = 1$ corresponds to VITA.

corresponds to NORTA and VITA, respectively. We simulated 10^4 samples of sample size $n = 100$ and $n = 1000$, at each of the twenty discrete values of α , and estimated the rejection rate of the NT-LR test at the nominal 5% level. The empirical rejection rates as a function of α are depicted in Figure 8. As expected, the NT-LR test performs significantly worse under the VITA distribution than under the NORTA distribution, in both sample size conditions. However, the increase in rejection rates is not as substantial as in the previous illustration. Also, in contrast to the example in Section 3.1, NT-LR is here sensitive to sample size, with closer-to-nominal rejection rates for the larger sample size. One might speculate that this may be due to the asymptotic robustness of NT-LR in this particular situation. However, this seems not to be the case, since additional simulations with sample size $n = 2 \cdot 10^4$ and 10^4 replications, yielded rejection rates of 0.072 and 0.084 for NORTA and VITA, respectively.

4. DISCUSSION

The VITA method has applications in all fields where covariance simulation is of interest. In the present article we have restricted our attention to the continuous case, but it can be easily seen that the same methods apply to categorical ordinal data. Applications includes assessing the finite sample behaviour of asymptotic techniques, robustness studies, power studies and parametric bootstrap techniques.

In the present paper we have considered two illustrations on the robustness of the NT-LR statistic in SEM. Both illustrations demonstrate that specifying only

marginal skewness and kurtosis is insufficient in SEM robustness studies, contrary to common practice in many Monte Carlo studies. Also, even when fully specifying the marginal distributions, the NT-LR statistic is sensitive to the underlying copula. This illustrates that Monte Carlo studies in SEM would benefit from simulating from a more flexible class of distributions than the often-used Vale-Maurelli distribution class. With the Vale-Maurelli method, the univariate marginals are Fleishman polynomials, and the copula is basically normal (Foldnes & Grønneberg, 2014). The VITA method allows the specification of more varied study designs, with a high level of control of the simulated vector. This includes control of both the copulas and of the univariate marginal distributions. A user that wishes to control univariate skewness and kurtosis, may use VITA in conjunction with Fleishman polynomial marginals. However, VITA equally supports other distributions, e.g., the Pearson (Pearson, 1895) and Johnson families (Johnson, 1949), which also allows the specification of skewness and kurtosis.

It seems that some of the rules of thumb in current SEM practice, based on studies using the quite inflexible Vale-Maurelli simulation method, should be further tested and validated under more realistic conditions by using VITA. This is particularly important in conditions with highly correlated variables, where the copula of VITA will be far from the Gaussian copula. When the correlation is weak, results based on other simulation methods will likely be similar to results derived from using VITA. In our second example we show VITA can be used to emulate such a realistic condition.

The flexibility in VITA, that is, the choice of marginals, R-vine, copula families and a valid edge order, may appear daunting to a practitioner. Such a practitioner may either wish to simulate data with characteristics derived from a given real-world sample, or wish to define a simulation mechanism with pre-specified features described in a Monte Carlo design. The first case is illustrated in Section 3.2, and where the R-vine specification can be found through statistical methods, such as the model selection strategy of Dissmann et al. (2013). For the second case a starting point would be to use a canonical R-vine, known as a C-vine, or alternatively a so-called D-vine (Joe, 2014). Also a canonical edge order could be chosen. The choice of bivariate copulas and the marginals must then be specified in the Monte Carlo design.

Acknowledgements. *The simulation work was performed on the Abel Cluster, owned by the University of Oslo and the Norwegian metacenter for High Performance Computing (NOTUR). The idea for how to provide starting values for VITA is partially based on a question from an anonymous referee, who we thank for this contribution. We would also like to thank the associate editor and the referees for comments that improved the presentation of the paper.*

APPENDIX A. TECHNICAL RESULTS

Lemma 1. *Let the edge e belong to an R-vine \mathcal{V} . Then $\mathcal{V}(e)$ is an R-vine.*

Proof. The proximity condition is automatically satisfied, since each edge $e = \{a, b\}$ in $\mathcal{V}(e)$ also belongs to \mathcal{V} . Also, since the edges in $\mathcal{V}(e)$ also belongs to \mathcal{V} , there can be no cycles in $\mathcal{V}(e)$. Finally, let $e_1 = \{a, b\}$ be an edge in $T_j(e)$. Then, by the proximity condition, e_1 corresponds to a sub-tree T_1 , of $T_{j-1}(e)$ with two edges, a and b . Moreover, due to connectedness, e_1 shares a common node, say b , with some edge $e_2 = \{b, c\}$ in $T_j(e)$, unless $T_j(e)$ is exhausted. Corresponding to e_2 is a sub-tree T_2 of $T_{j-1}(e)$ that shares edge b with T_1 . Hence there is a path from a to c in $T_{j-1}(e)$. \square

Proof of Theorem 1. Let e_k be the k 'th edge in the order \mathcal{O} , and let $\mathbf{c}_{e_k, a} = i, \mathbf{c}_{e_k, b} = j$. We first show that $I_{e_k}(\theta_{e_k})$ is increasing. By Theorem 8.65 in Joe (2014), the CDF F_{ij} of (X_i, X_j) is increasing in θ_{e_k} since $\theta_{e_k} \mapsto B_{e_k, \theta_{e_k}}$ is increasing. Hence $I_{e_k}(\theta_{e_k})$ is increasing by the Hoeffding identity $\text{Cov}(X_i, X_j) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [F_{i,j}(x, y) - F_i(x)F_j(y)] dx dy$. Theorem 8.65 omits the lowest tree T_1 from its statement, but here the result is immediate by the Hoeffding identity since $C_{i,j} = B_e$.

We now show continuity by generalizing Lemma 3 in Cario & Nelson (1997) (see also Lemma A.2 in Cario & Nelson, 1996): they use a stochastic representation specifically for the normal case, while ours is general. Let $\vartheta_{e_k} = (\theta_{e_1}, \theta_{e_2}, \dots, \theta_{e_{k-1}})$, so that $(\theta_{e_1}, \theta_{e_2}, \dots, \theta_{e_k}) = (\vartheta_{e_k}, \theta_{e_k})$. Note that ϑ_{e_k} is algebraically independent of θ_{e_k} . Let $\delta = |U_{e_k}|$ be the cardinality of U_{e_k} . Suppose without loss of generality that $U_{e_k} = \{1, 2, \dots, \delta\}$ and that $i = \delta - 1$ and $j = \delta$. We now show continuity through showing that for any countable sequence $\{\theta_{e_k, n}\}_{n \geq 1} \subseteq \Theta_{e_k}$ with $\lim_{n \rightarrow \infty} \theta_{e_k, n} = \theta_{e_k}$ we have $\lim_{n \rightarrow \infty} \mathbb{E}_{(\vartheta_{e_k}, \theta_{e_k, n})} X_i X_j = \mathbb{E}_{(\vartheta_{e_k}, \theta_{e_k})} X_i X_j$. We use the multivariate quantile transform of O'Brien (1975) (see Section 3 of Rüschendorf (2009) for a helpful discussion of this transformation). Let $F_{1, \dots, \delta; \vartheta_{e_k}, \theta_{e_k, n}}^{-1}$ be the joint distribution induced by the parametric R-vine specification of $\mathcal{V}(e_k)$ with parameters $(\vartheta_{e_k}, \theta_{e_k, n})$, and let $F_{j|1, \dots, j-1; \vartheta_{e_k}, \theta_{e_k, n}}^{-1}$ be the conditional quantile function of $F_{1, \dots, \delta; \vartheta_{e_k}, \theta_{e_k, n}}^{-1}$ of the j 'th variable, conditioned on the 1st through $j - 1$ 'th variables.

Let $X(\vartheta_{e_k}, \theta_{e_k, n}) := (X_1(\vartheta_{e_k}, \theta_{e_k, n}), \dots, X_\delta(\vartheta_{e_k}, \theta_{e_k, n}))'$ where the coordinates are defined recursively through $X_1(\vartheta_{e_k}, \theta_{e_k, n}) := F_{1, \vartheta_{e_k}, \theta_{e_k, n}}^{-1}(U_1)$ and $X_j(\vartheta_{e_k, n}) := F_{j|1, \dots, j-1; \vartheta_{e_k}, \theta_{e_k, n}}^{-1}(U_j | X_1(\vartheta_{e_k}, \theta_{e_k, n}), \dots, X_{j-1}(\vartheta_{e_k}, \theta_{e_k, n}))$ for $2 \leq j \leq d$, where $U_1, U_2, \dots, U_\delta \sim U[0, 1]$ and independent. The multivariate quantile transform then gives $X(\vartheta_{e_k}, \theta_{e_k, n}) \sim F_{1, 2, \dots, \delta; \vartheta_{e_k}, \theta_{e_k, n}}$. Also let $X = (X_1, \dots, X_\delta) := X(\vartheta_{e_k}, \theta_{e_k})$ so that $X \sim F_{1, \dots, \delta; \vartheta_{e_k}, \theta_{e_k}}$. We now show that for each $1 \leq j \leq d$ we have $\lim_{n \rightarrow \infty} X_j(\vartheta_{e_k}, \theta_{e_k, n}) = X_j(\vartheta_{e_k}, \theta_{e_k})$ a.s. (almost surely, i.e. with probability one), which implies that $\lim_{n \rightarrow \infty} X(\vartheta_{e_k}, \theta_{e_k, n}) = X(\vartheta_{e_k}, \theta_{e_k}) = X$ a.s. We proceed by induction. Firstly, $\theta_e \mapsto F_{1; \vartheta_{e_k}, \theta_{e_k, n}}^{-1}(u_1)$ is assumed to be continuous. Hence,

$\lim_{n \rightarrow \infty} X_1(\vartheta_{e,n}) = \lim_{n \rightarrow \infty} F_{1;\vartheta_{e_k},\theta_{e_k,n}}^{-1}(U_1) = F_{1;\vartheta_{e_k},\theta_{e_k}}^{-1}(U_1) = X_1$ a.s. by continuity. Let $2 \leq j \leq \delta$ and assume $\lim_{n \rightarrow \infty} X_i(\vartheta_{e_k}, \theta_{e_k,n}) = X_i$ a.s. for $1 \leq i \leq j-1$. By assumption, $(\theta_e, x_1, \dots, x_{j-1}) \mapsto F_{j|1,\dots,j-1;\vartheta_{e_k},\theta_{e_k,n}}^{-1}(u_j|x_1, \dots, x_{j-1})$ is continuous. Hence, we complete the induction argument by concluding that $\lim_{n \rightarrow \infty} X_j(\vartheta_{e_k}, \theta_{e_k,n}) = \lim_{n \rightarrow \infty} F_{j|1,\dots,j-1;\vartheta_{e_k},\theta_{e_k,n}}^{-1}(U_j|X_1(\vartheta_{e,n}), \dots, X_{j-1}(\vartheta_{e,n})) = X_j$ a.s. Let $\xrightarrow[n \rightarrow \infty]{\mathscr{W}}$ denote convergence in distribution. The a.s. convergence implies $X(\vartheta_{e_k}, \theta_{e_k,n}) \xrightarrow[n \rightarrow \infty]{\mathscr{W}} X(\vartheta_{e_k}, \theta_{e_k})$, implying $(X_{\delta-1}(\vartheta_{e_k}, \theta_{e_k,n}), X_{\delta}(\vartheta_{e_k}, \theta_{e_k,n}))' \xrightarrow[n \rightarrow \infty]{\mathscr{W}} (X_{\delta-1}, X_{\delta})'$. The continuous mapping theorem (Billingsley, 1995, Theorem 29.2) implies that $X_{\delta-1}(\vartheta_{e_k}, \theta_{e_k,n})X_{\delta}(\vartheta_{e_k}, \theta_{e_k,n}) \xrightarrow[n \rightarrow \infty]{\mathscr{W}} X_{\delta-1}X_{\delta}$. Theorem 25.12 of Billingsley (1995) (i.e. uniform integrability and weak convergence implies moment convergence) and his eq. (25.13) (a moment condition which implies uniform integrability) shows that if $\sup_n \mathbb{E}|X_{\delta-1}(\vartheta_{e_k}, \theta_{e_k,n})X_{\delta}(\vartheta_{e_k}, \theta_{e_k,n})|^{1+\varepsilon} < \infty$, then we can transfer convergence in distribution to moment-convergence, and we get $\lim_{n \rightarrow \infty} \mathbb{E}_{(\vartheta_{e_k}, \theta_{e_k,n})} X_i X_j = \mathbb{E}_{(\vartheta_{e_k}, \theta_{e_k})} X_i X_j$. And by the Cauchy-Schwarz inequality and the moment assumption, we have

$$\begin{aligned}
& \sup_n \mathbb{E}|X_{\delta-1}(\vartheta_{e_k}, \theta_{e_k,n})X_{\delta}(\vartheta_{e_k}, \theta_{e_k,n})|^{1+\varepsilon/2} \\
& \leq \sup_{\theta_{e_k} \in \Theta_{e_k}} \int_{[0,1]^d} |F_i^{-1}(u_i)F_j^{-1}(u_j)|^{1+\varepsilon/2} dC_{i,j}(u; \vartheta_{e_k}, \theta_{e_k}) \leq \sup_{\theta \in \Theta} \mathbb{E}_{\theta}|X_i X_j|^{1+\varepsilon/2} \\
& \leq \sup_{\theta \in \Theta} \sqrt{\mathbb{E}_{\theta}|X_i|^{2+\varepsilon}} \sqrt{\mathbb{E}_{\theta}|X_j|^{2+\varepsilon}} = \sqrt{\int_{\mathbb{R}} |x|^{2+\varepsilon} dF_i(x)} \sqrt{\int_{\mathbb{R}} |x|^{2+\varepsilon} dF_j(x)} < \infty,
\end{aligned}$$

where the equality follows since the copula-parametrisation does not change the marginals. This finishes the proof. \square

REFERENCES

- BEDFORD, T. & COOKE, R. M. (2002). Vines: A new graphical model for dependent random variables. *Annals of Statistics*, 1031–1068.
- BENTLER, P. (2006). Eqs 6 structural equations program manual.
- BENTLER, P. M. & TONG, X. (2013). Evaluation of a New Mean Scaled and Moment Adjusted Test Statistic for SEM. *Structural Equation Modeling: A Multidisciplinary Journal* **20**, 148–156.
- BILLER, B. (2009). Copula-based multivariate input models for stochastic simulation. *Operations research* **57**, 878–892.
- BILLINGSLEY, P. (1995). *Probability and measure*. John Wiley & Sons.
- BOOMSMA, A. (1983). *On the robustness of LISREL (maximum likelihood estimation) against small sample size and non-normality*. Ph.D. thesis, University of Groningen.

- BROWNE, M. W. (1984). Asymptotically distribution-free methods for the analysis of covariance structures. *British Journal of Mathematical and Statistical Psychology* **37**, 62–83.
- CARIO, M. C. & NELSON, B. L. (1996). Autoregressive to anything: Time-series input processes for simulation. *Operations Research Letters* **19**, 51–58.
- CARIO, M. C. & NELSON, B. L. (1997). Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Tech. rep., Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois.
- CURRAN, P. J., WEST, S. G. & FINCH, J. F. (1996). The robustness of test statistics to nonnormality and specification error in confirmatory factor analysis. *Psychological Methods* **1**, 16–29.
- DISSMANN, J., BRECHMANN, E. C., CZADO, C. & KUROWICKA, D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis* **59**, 52–69.
- ERHARD, V. & CZADO, C. (2009). Sampling count variables with specified pearson correlation: A comparison between a naive and a c-vine sampling approach. *Dependence Modeling: Vine Copula Handbook* , 73–87.
- FOLDNES, N. & GRØNNEBERG, S. (2014). How general is the vale–maurelli simulation approach? *Psychometrika* , 1–18.
- FOLDNES, N. & OLSSON, U. H. (2016). A simple simulation technique for non-normal data with prespecified skewness, kurtosis, and covariance matrix. *Multivariate Behavioral Research* **51**, 207–219. PMID: 27014851.
- FOULADI, R. (2000). Performance of modified test statistics in covariance and correlation structure analysis under conditions of multivariate nonnormality. *Structural Equation Modeling: A Multidisciplinary Journal* **7**, 356–410.
- HOBÆK HAFF, I., AAS, K. & FRIGESSI, A. (2010). On the simplified pair-copula construction - simply useful or too simplistic? *Journal of Multivariate Analysis* **101**, 1296–1310.
- HOMEM-DE MELLO, T. & BAYRAKSAN, G. (2014). Monte carlo sampling-based methods for stochastic optimization. *Surveys in Operations Research and Management Science* **19**, 56–85.
- JOE, H. (1996). Families of m-variate distributions with given margins and m(m-1)/2 bivariate dependence parameters. *Lecture Notes-Monograph Series* , 120–141.
- JOE, H. (2006). Range of correlation matrices for dependent random variables with given marginal distributions. In *Advances in Distribution Theory, Order Statistics, and Inference*. Springer, pp. 125–142.
- JOE, H. (2014). *Dependence modeling with copulas*. CRC Press.

- JOHNSON, N. (1949). Systems of frequency curves generated by methods of translation. *Biometrika* **36**, 149–176.
- LI, S. T. & HAMMOND, J. L. (1975). Generation of pseudorandom numbers with specified univariate distributions and correlation coefficients. *Systems, Man and Cybernetics, IEEE Transactions on* , 557–561.
- MAIR, P., SATORRA, A. & BENTLER, P. M. (2012). Generating Nonnormal Multivariate Data Using Copulas: Applications to SEM. *Multivariate Behavioral Research* **47**, 547–565.
- MUTHÉN, B. & MUTHÉN, L. (2012). Mplus version 7: User’s guide. *Los Angeles, CA: Muthén & Muthén* .
- O’BRIEN, G. (1975). The comparison method for stochastic processes. *The Annals of Probability* , 80–88.
- PEARSON, K. (1895). Contributions to the mathematical theory of evolution. ii. skew variation in homogeneous material. *Philosophical Transactions of the Royal Society of London. A* , 343–414.
- R ÜSCHENDORF, L. (2009). On the distributional transform, Sklar’s theorem, and the empirical copula process. *Journal of Statistical Planning and Inference* **139**, 3921–3927.
- R CORE TEAM (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- ROSSEEL, Y. (2012). lavaan: An r package for structural equation modeling. *Journal of Statistical Software* **48**, 1–36.
- SATORRA, A. (1989). Alternative test criteria in covariance structure analysis: A unified approach. *Psychometrika* **54**, 131–151.
- SCHEPSMEIER, U., STOEBER, J., BRECHMANN, E. C. & GRAELER, B. (2015). *VineCopula: Statistical Inference of Vine Copulas*. R package version 1.4.
- SKLAR, M. (1959). *Fonctions de répartition à n dimensions et leurs marges*. Université Paris 8.
- VALE, C. & MAURELLI, V. (1983). Simulating nonnormal distributions. *Psychometrika* **48**.

DEPARTMENT OF ECONOMICS, BI NORWEGIAN BUSINESS SCHOOL, OSLO, NORWAY 0484
E-mail address: `steffeng@gmail.com`

DEPARTMENT OF ECONOMICS, BI NORWEGIAN BUSINESS SCHOOL, STAVANGER, NORWAY 4014
E-mail address: `njal.foldnes@bi.no`