Selection of Vine Copulas

Claudia Czado and Eike Christian Brechmann and Lutz Gruber

Abstract Vine copula models have proven themselves as a very flexible class of multivariate copula models with regard to symmetry and tail dependence for pairs of variables. The full specification of a vine model requires the choice of vine tree structure, copula families for each pair copula term and their corresponding parameters. In this survey we discuss the different approaches, both frequentist as well as Bayesian, for these model choices so far and point to open problems.

1 Introduction

The analysis of high dimensional data sets requires flexible multivariate stochastic models which can capture the inherent dependency patterns. The copula approach, which separates the modeling of the marginal distributions from modeling the dependence characteristics, is a natural one to follow in this context. This development has spawned a tremendous increase of copula-based applications in the last 10 years, especially in the area of finance, economics and hydrology.

Considerable efforts have been undertaken to increase the flexibility of multivariate copula models beyond the scope of elliptical and Archimedean copulas. Vine copulas are among the best-received of such efforts. Vine copulas use (conditional) bivariate copulas as the so-called "pair copula" building blocks to describe a mul-

Claudia Czado
Zentrum Mathematik, Technische Universität München, München, Germany e-mail: cczado@ma.tum.de

Eike Christian Brechmann
Zentrum Mathematik, Technische Universität München, München, Germany e-mail: brechmann@ma.tum.de

Lutz Gruber
Zentrum Mathematik, Technische Universität München, München, Germany e-mail: lutz.gruber@ma.tum.de

tivariate distribution (see [36]). A set of linked trees—the "vine"—describes a vine copula's factorization of the multivariate copula density function into the density functions of its pair copulas (see [8, 9]). The article by [1] illustrates a first application of the vine copula concept using non-Gaussian pair copulas to financial data. The first comprehensive account of vine copulas is found in [44], a recent survey in [20], and the current developments of this active research area in [45] .

Elliptical copulas as well as Archimedean copulas have been shown to be inadequate models to describe the dependence characteristics of real data applications (see for example [1, 24, 23]). As a pair copula construction, vine copulas allow different structural behaviors of pairs of variables to be modeled suitably, in particular so with regard to their symmetry, or lack thereof, strength of dependence, and tail dependencies. Such flexibility requires well-designed model selection procedures to realize the full potential of vine copulas as dependence models. Successful applications of vines can be found, amongst others, in [10, 12, 17, 22, 24, 31, 47, 50, 53, 60].

A parametric vine copula consists of three components: a set of linked trees identifying the pairs of variables and their conditioning variables, the copula families for each pair copula term given by the tree structure, and the corresponding copula parameters. The three-layered definition leads to three fundamental estimation and selection tasks: 1. Estimation of copula parameters for a chosen vine tree structure and pair copula families, 2. Selection of the parametric copula family for each pair copula term and estimation of the corresponding parameters for a chosen vine tree structure, and 3. Selection and estimation of all three model components. In this survey we address these tasks and give an overview of the statistical approaches taken so far. These range from frequentist to Bayesian methods.

The remainder of this paper is structured as follows. In Section 2 we provide the necessary methodical background on regular vines and regular vine copulas. We then discuss estimation of parameters of regular vine copulas in Section 3 and the selection of appropriate pair copulas in Section 4. Section 5 treats the joint selection of the regular vine tree structure, the copula families, and their parameters. Section 6 concludes with a discussion of available software and open problems.

2 Regular Vine Copulas

Copulas describe a statistical model's dependence behavior separately from its marginal distributions [59]. As such, a copula is a multivariate distribution function with all marginal distributions being uniform: i.e. the copula associated with an n-variate cumulative distribution function $F_{1:n}$ with univariate marginal distribution functions $F_1, \ldots, F_n$ is a distribution function $C : [0,1]^n \to [0,1]$ satisfying

$$F_{1:n}(\mathbf{x}) = C(F_1(x_1), \ldots, F_n(x_n)), \quad \mathbf{x} = (x_1, ..., x_n)' \in \mathbb{R}^n.$$

If C is absolutely continuous, its density is denoted by c.

The factorization of multivariate copula densities into (conditional) bivariate copula densities is due to [36] and [8]. The details of these factorizations are represented by the graph theoretical construction called regular vine to organize different decompositions. Graphs are defined in terms of a set of nodes N and a set of edges E connecting these nodes, i.e. $E \subset N \times N$. Vines are based on trees which are particular graphs where there is a unique sequence of edges between each two nodes.

Definition 2.1 (Regular Vine Tree Sequence). A set of linked trees $\mathscr{V} = (T_1, T_2, ..., T_{n-1})$ is a regular vine (R-vine) on n elements if

1. T_1 is a tree with nodes $N_1 = \{1, ..., n\}$ and a set of edges denoted by E_1.
2. For $i = 2, ..., n-1$, T_i is a tree with nodes $N_i = E_{i-1}$ and edge set E_i.
3. For $i = 2, ..., n-1$, if $a = \{a_1, a_2\}$ and $b = \{b_1, b_2\}$ are two nodes in N_i connected by an edge, then exactly one of the a_i equals one of the b_i (proximity condition).

In other words, the proximity condition requires that the edges corresponding to two connected nodes in tree T_i share a common node in tree T_{i-1}. This ensures that the decomposition into bivariate copulas which is given below is well-defined.

Two sub-classes of regular vines have been studied extensively in the literature: canonical vines (C-vines) and drawable vines (D-vines) (see [44, 1]). C-vines are characterized by a root node in each tree T_i, $i \in \{1, ..., n-1\}$, which has degree $n-i$, in other words the root node is connected to all other nodes of the tree. D-vines, on the other hand, are uniquely characterized through their first tree which is, in graph theoretical terms, a path; this means that each node has degree of at most 2. Therefore the order of variables in the first tree defines the complete D-vine tree sequence.

Some more definitions are needed to introduce regular vine copulas: the complete union A_e of an edge $e = \{a, b\} \in E_i$ in tree T_i of a regular vine $\mathscr{V}$ is defined by

$$A_e = \{v \in N_1 : \exists e_m \in E_m,\ m = 1, \ldots, i-1, \text{ such that } v \in e_1 \in \cdots \in e_{i-1} \in e\}. \tag{1}$$

The conditioning set associated with $e = \{a, b\}$ is defined as $D_e := A_a \cap A_b$ and the conditioned sets associated with $e = \{a, b\}$ are defined as $\mathscr{C}_{e,a} := A_a \setminus D_e$ and $\mathscr{C}_{e,b} := A_b \setminus D_e$. [8] showed that the conditioned sets are singletons, and we will therefore refer to edges by their labels $\{\mathscr{C}_{e,a}, \mathscr{C}_{e,b} | D_e\} \hat{=} \{i(e), j(e) | D(e)\}$. An exemplary regular vine on five elements is shown in Figure 1.

Given these sets, we can specify a regular vine copula by associating a (conditional) bivariate copula to each edge of the regular vine, a so-called pair copula.

Definition 2.2 (Regular Vine Copula). A regular vine copula $C = (\mathscr{V}, \mathscr{B}(\mathscr{V}), \theta(\mathscr{B}(\mathscr{V})))$ in n dimensions is a multivariate distribution function such that for a random vector $\mathbf{U} = (U_1, \ldots, U_n)' \sim C$ with uniform margins

1. $\mathscr{V}$ is a regular vine on n elements,
2. $\mathscr{B}(\mathscr{V}) = \left\{C_{i(e),j(e)|D(e)} | e \in E_m,\ m = 1, \ldots n-1\right\}$ is a set of $n(n-1)/2$ copula families identifying the conditional distributions of $U_{i(e)}, U_{j(e)} | \mathbf{U}_{D(e)}$,
3. $\theta(\mathscr{B}(\mathscr{V})) = \left\{\theta_{i(e),j(e)|D(e)} | e \in E_m,\ m = 1, \ldots n-1\right\}$ is the set of parameter vectors, corresponding to the copulas in $\mathscr{B}(\mathscr{V})$.

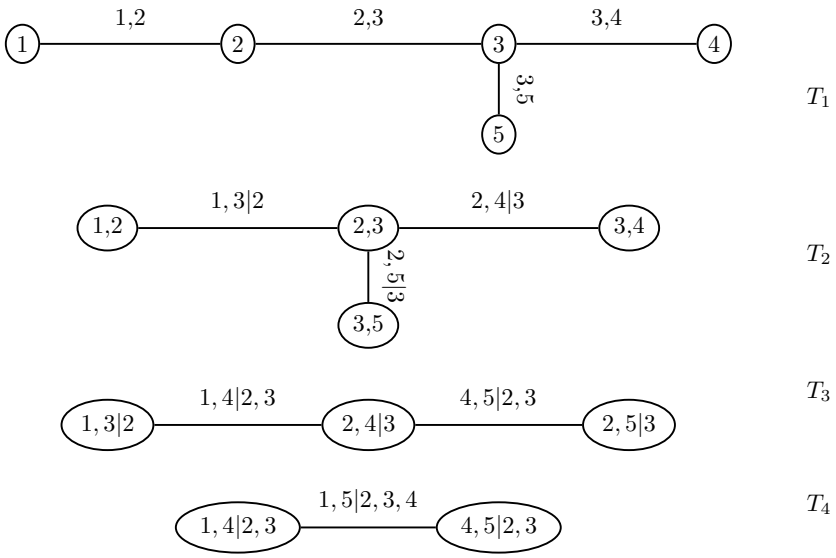

Fig. 1 Regular vine on five elements.

Therefore the full specification of a regular vine copula consists of three layers: the regular vine tree structure $\mathscr{V}$, the pair copula families $\mathscr{B} = \mathscr{B}(\mathscr{V})$ and the pair copula parameters $\theta = \theta(\mathscr{B}(\mathscr{V}))$. Regular vine copulas which differ in the tree structure or in at least one pair copula family represent in general different statistical models. Notable exceptions from this are the multivariate Gaussian, Student's t or Clayton copulas, which can be decomposed into bivariate Gaussian, Student's t or Clayton copulas, respectively, in multiple ways (see [63]). The probability density function $f_{1:n}$ at point $\mathbf{x} = (x_1, \ldots, x_n)' \in \mathbb{R}^n$ of an n-dimensional regular vine dependent distribution $F_{1:n}$ is easily calculated as

$$f_{1:n}(\mathbf{x}|\mathscr{V}, \mathscr{B}, \theta) = \left(\prod_{m=1}^{n-1} \prod_{e=\{a,b\} \in E_m} c_{i(e),j(e)|D(e)} \left(F_{i(e)|D(e)}, F_{j(e)|D(e)} | \theta_{i(e),j(e)|D(e)} \right) \right) \times f_1(x_1) \cdots f_n(x_n), \tag{2}$$

where $F_{i(e)|D(e)} := F_{i(e)|D(e)}\left(x_{i(e)}|\mathbf{x}_{D(e)}\right)$ and $F_{j(e)|D(e)} := F_{j(e)|D(e)}\left(x_{j(e)}|\mathbf{x}_{D(e)}\right)$ (see [8]). These conditional distribution functions can be determined recursively tree-by-tree using the following relationship

$$\begin{aligned} F_{i(e)|D(e)}\left(x_{i(e)}|\mathbf{x}_{D(e)}\right) &= F_{\mathscr{C}_{e,a}|D_e}\left(x_{\mathscr{C}_{e,a}}|\mathbf{x}_{D_e}\right) \\ &= \frac{\partial C_{\mathscr{C}_{a,a_1},\mathscr{C}_{a,a_2}|D_a}\left(F_{\mathscr{C}_{a,a_1}|D_a}\left(x_{\mathscr{C}_{a,a_1}}|\mathbf{x}_{D_a}\right), F_{\mathscr{C}_{a,a_2}|D_a}\left(x_{\mathscr{C}_{a,a_2}}|\mathbf{x}_{D_a}\right)\right)}{\partial F_{\mathscr{C}_{a,a_2}|D_a}\left(x_{\mathscr{C}_{a,a_2}}|\mathbf{x}_{D_a}\right)}, \end{aligned} \tag{3}$$

where $e = \{a,b\}$ with $a = \{a_1,a_2\}$ as before (see [23] for details).

To facilitate inference, it is assumed that the pair copulas $C_{i(e),j(e)|D(e)}$ only depend on the variables with indices in $D(e)$ through the arguments $F_{i(e)|D(e)}$ and $F_{j(e)|D(e)}$. This so-called simplifying assumption has been investigated by [34] and [63]. In particular, [34] show examples where the assumption is not severe and [63] show that the multivariate Clayton copula is the only Archimedean copula which can be represented as a simplified vine. A critical look on the subject can be found in [2] who take a first step in building vine copulas of non-simplified structure.

Following Expression (2), the likelihood L of a regular vine copula $C = (\mathscr{V},\mathscr{B},\theta)$ given the observed data $\mathbf{x} = (\mathbf{x}_1,\ldots,\mathbf{x}_N)' \in \mathbb{R}^{N\times n}$ can be calculated as

$$L(\mathscr{V},\mathscr{B},\theta|\mathbf{x}) = \prod_{k=1}^{N} f_{1:n}(\mathbf{x}_k|\mathscr{V},\mathscr{B},\theta). \tag{4}$$

The corresponding log-likelihood is denoted by ℓ.

A regular vine copula is said to be *truncated at level M* if all pair copulas conditioning on M or more variables are set to bivariate independence copulas [14]. As a result, the iteration index m of the first product in Equation (2) runs only up to $m = M$ and the distribution is fully specified by $\big(\mathscr{V} = (T_1,\ldots,T_M),\mathscr{B}(\mathscr{V}),\theta(\mathscr{B}(\mathscr{V}))\big)$.

In the following, we now discuss in reverse order how the components of a regular vine copula can be selected and estimated. That is, we begin with estimation of the parameters θ, then treat the selection of appropriate copula families $\mathscr{B}$, and finally discuss the selection of vine trees $\mathscr{V}$.

3 Parameter Estimation for Given Vine Tree Structure and Pair Copula Families

Given a vine tree structure $\mathscr{V}$ and pair copula families $\mathscr{B} = \mathscr{B}(\mathscr{V})$, the challenge is to estimate the parameters $\theta = \theta(\mathscr{B}(\mathscr{V}))$ of a regular vine copula for observed data $\mathbf{x} \in \mathbb{R}^{N\times n}$. The crucial point of this task are the conditional distribution functions $F_{i(e)|D(e)}$, which depend on copulas of previous trees (see (3)).

3.1 Maximum Likelihood Estimation

Classically, parameters of a statistical model are often estimated using maximum likelihood techniques. Here, this means that regular vine copula parameters θ and parameters of the marginal distributions are estimated by maximizing the likelihood (4) in terms of these parameters. For copulas, in particular if $n > 2$, the number of parameters to be estimated may however be too large, so that one typically either uses empirical distribution functions for the margins as proposed by [25, 26] or estimates parameters of the marginal distributions in a first step and then fixes

these marginal parameters to their estimated values in the estimation of the copula parameters. The latter method is called inference functions for margins (IFM) by [38, 37].

But even when estimation of marginal and dependence parameters is separated, joint maximum likelihood estimation of regular vine copula parameters can be very challenging, since the vine decomposition involves $n(n-1)/2$ bivariate copulas with corresponding parameters. [1] therefore proposed a sequential method: starting with the copulas of the first tree, this method proceeds tree-wise and estimates the parameters of the copulas in a tree by fixing the parameters of copulas in all previous trees.

Example 3.1 (Sequential Estimation). Let the 5-dimensional regular vine $\mathscr{V}$ of Figure 1 be given with copulas $\mathscr{B}(\mathscr{V})$. In the first step, we estimate the parameters of the copulas $C_{1,2}$, $C_{2,3}$, $C_{3,4}$ and $C_{3,5}$ using maximum likelihood based on the transformed observations $F_j(x_{kj})$ for $x_{kj},\ k = 1,...,N,\ j = 1,...,5$. In the second tree, we then have to estimate for instance the parameter(s) of the copula $C_{1,3|2}$. For this we form pseudo-observations $F_{1|2}(x_{k1}|x_{k2},\hat{\theta}_{1,2})$ and $F_{3|2}(x_{k3}|x_{k2},\hat{\theta}_{2,3}),\ k = 1,...,N$, according to Expression (3) and using the estimated parameters of copulas $C_{1,2}$ and $C_{2,3}$, respectively. Based on these pseudo-observations estimation of $\theta_{1,3|2}$ is again straightforward. Similarly for all other copulas.

Note that this strategy only involves estimation of bivariate copulas and therefore is computationally much simpler than joint maximum likelihood estimation of all parameters at once. For more details see [1] as well as [32] who investigates the asymptotic properties of this sequential approach. A comparison study of estimators for regular vine copulas can be found in [33].

If joint maximum likelihood estimates are desired, the sequential method can be used to obtain starting values for the numerical optimization. A detailed discussion how to compute score functions and the observed information matrix for regular vine copulas is provided by [64].

3.2 Bayesian Posterior Estimation

Bayesian statistics considers parameters θ as random variables. As such, inference focuses on estimating the entire distribution of the parameters instead of only finding a point estimate. In particular, the so-called posterior distribution $p(\theta) := p(\theta|\mathbf{x})$, the distribution of the parameters θ given the observed data $\mathbf{x}$, is the main object of interest. The unnormalized posterior density factorizes into the product of the likelihood function $L(\theta|\mathbf{x})$ (4) and the prior density function $\pi(\theta)$:

$$p(\theta) := p(\theta|\mathbf{x}) = \frac{L(\theta|\mathbf{x}) \cdot \pi(\theta)}{f(\mathbf{x})} \propto L(\theta|\mathbf{x}) \cdot \pi(\theta).$$

The prior distribution incorporates a priori beliefs about the distribution of the parameters. It must not depend on, or be conditional upon, the observed data.

Markov chain Monte Carlo (MCMC) procedures clear the remaining obstacle of sampling from a distribution which is known only up to a constant. The trick is to simulate the run of a Markov chain whose equilibrium distribution is the targeted posterior distribution $p(\theta)$ of the parameters θ. Upon convergence of the Markov chain, its states represent draws from the desired distribution.

The Metropolis–Hastings algorithm [48, 30] implements the simulation of such a Markov chain through an acceptance/rejection mechanism for the updates of the chain: in a first step, an update of the chain from its current state θ to $\theta^* \sim q(\cdot|\theta)$ is proposed. The proposal distribution q can be chosen almost arbitrarily. In the second step, the proposal is accepted with probability $\alpha := \alpha(\theta, \theta^*)$. The acceptance probability α is chosen such that convergence of the Markov chain to the targeted distribution is guaranteed by Markov chain theory [49]. The Metropolis–Hastings acceptance probability for convergence to the target distribution $p(\cdot)$ is

$$\alpha(\theta, \theta^*) = \frac{p(\theta^*)}{p(\theta)} \cdot \frac{q(\theta|\theta^*)}{q(\theta^*|\theta)}.$$

The target distribution goes into the Metropolis–Hastings algorithm only at the evaluation of the acceptance probability α. It can be easily seen that it is sufficient to know the density function of the target distribution only up to a constant for this algorithm to work, given that α only depends on the ratio of two densities.

The following Metropolis–Hastings scheme to sample from the posterior distribution of the parameters θ of a regular vine copula is proposed in [61] and [28]. The authors suggest normally distributed random walk proposals be used in the update step.

Algorithm 3.1 (Metropolis–Hastings Sampler for Parameter Estimation)

Choose arbitrary starting values $\theta^0 = (\theta_1^0, \ldots, \theta_D^0)'$, *where* $D := n(n-1)/2$.
for *each iteration* $r = 1, \ldots, R$ ***do***
  Set $\theta^r = \theta^{r-1}$.
  for $i = 1, ..., D$ ***do***
    Draw θ_i^r *from a* $N(\theta_i^{r-1}, \sigma_i^2)$ *distribution with probability density function* $\phi_{(\theta_i^{r-1}, \sigma_i^2)}(\cdot)$.
    Evaluate the Metropolis–Hastings acceptance probability of the proposal

$$\alpha := \alpha(\theta_i^{r-1}, \theta_i^r) = \frac{p(\theta^*)}{p(\theta)} \cdot \frac{\phi_{(\theta_i^r, \sigma_i^2)}(\theta_i^{r-1})}{\phi_{(\theta_i^{r-1}, \sigma_i^2)}(\theta_i^r)} = \frac{L(\theta^*|\mathbf{x})}{L(\theta|\mathbf{x})} \cdot \frac{\pi(\theta^*)}{\pi(\theta)},$$

    where $\theta^* = (\theta_1^r, \ldots, \theta_i^r, \theta_{i+1}^{r-1}, \ldots, \theta_D^{r-1})'$ *and* $\theta = (\theta_1^r, \ldots, \theta_{i-1}^r, \theta_i^{r-1}, \ldots, \theta_D^{r-1})'$.
    Accept the proposal θ_i^r *with probability* α*; if rejected, set* $\theta_i^r = \theta_i^{r-1}$.
  end for
end for
return $(\theta^1, \ldots, \theta^R)$

4 Selection of the Pair Copula Families and their Parameters for a Known Vine Tree Structure

An n-dimensional regular vine copula with tree structure $\mathscr{V}$ is based on a set $\mathscr{B}(\mathscr{V})$ of $n(n-1)/2$ bivariate copulas and their corresponding parameters. The copula families can be chosen arbitrarily, e.g., from the popular classes of Archimedean, elliptical or extreme-value copulas. Assuming that an appropriate vine structure $\mathscr{V}$ is chosen, the question therefore is how to select adequate (conditional) pair copulas $C_{i(e),j(e)|D(e)}$ and their parameters for given data $\mathbf{x}$.

4.1 Sequential Selection

If a regular vine is truncated at level 1—that is the copulas in trees T_2 to T_{n-1} are set to independence—, the corresponding regular vine copula reduces to the product of unconditional bivariate copulas $C_{i(e),j(e)}$ which can be selected based on data $F_{i(e)}(x_{k,i(e)})$ and $F_{j(e)}(x_{k,j(e)})$, $k = 1,...,N$. Typical criteria for copula selection from a given set of families are information criteria such as the AIC as proposed by [46] and [11]. The latter compares AIC-based selection to three alternative selection strategies: selection of family with highest p-value of a copula goodness of fit test based on the Cramér-von Mises statistic, with smallest distance between empirical and modeled dependence characteristics (Kendall's τ, tail dependence), or with highest number of wins in pairwise comparisons of families using the test by [65]. In a large-scale Monte Carlo study the AIC turned out to be the most reliable selection criterion.

In a general regular vine, the selection of a pair copula $C_{i(e),j(e)|D(e)}$ however depends on the choices for copulas in previous trees due to its arguments (3). Since a joint selection seems infeasible because of the many possibilities, one typically proceeds tree by tree as in the sequential estimation method. That is, instead of estimating the parameters $\theta_{i(e),j(e)|D(e)}$ of $C_{i(e),j(e)|D(e)}$, the copula is selected first and then estimated, which usually coincides for most selection strategies. Given the selected and estimated copulas of previous trees, one then selects the copulas of the next tree.

Example 4.1 (Sequential Selection). Again consider the five-dimensional regular vine $\mathscr{V}$ of Figure 1 but now with unknown copulas $\mathscr{B}(\mathscr{V})$. In the first tree, we select (and then estimate) the copulas $C_{1,2}$, $C_{2,3}$, $C_{3,4}$ and $C_{3,5}$ using our method of choice based on $F_j(x_{kj})$, $k = 1,...,N$, $j = 1,...,5$. Given these copulas, we then have to select conditional copulas in the second tree. In case of the copula $C_{1,3|2}$, we therefore again form the pseudo-observations $F_{1|2}(x_{k1}|x_{k2},\hat{\theta}_{1,2})$ and $F_{3|2}(x_{k3}|x_{k2},\hat{\theta}_{2,3})$, $k = 1,...,N$, according to Expression (3) and then select (and estimate) $C_{1,3|2}$ based on them. This can be iterated for all trees.

Clearly this sequential selection strategy accumulates uncertainty in the selection and therefore the final model has to be carefully checked and compared to alternative

models. For the latter, the tests for non-nested model comparison by [65] and [18] may be used.

4.2 Reversible Jump MCMC-Based Bayesian Selection

Bayesian copula family and parameter selection aims at estimating the joint posterior distribution of the pair copula families $\mathscr{B} = \mathscr{B}(\mathscr{V})$ and parameters $\theta = \theta(\mathscr{B}(\mathscr{V}))$. The posterior density function factorizes into the product of the likelihood function L (4) and the prior density function π: $p(\mathscr{B}, \theta) \propto L(\mathscr{B}, \theta|\mathbf{x}) \cdot \pi(\mathscr{B}, \theta)$.

As in Section 4.1, family selection is understood in the context of choosing from a pre-specified set $\mathbf{B}$ of parametric pair copula families. Model sparsity can be induced through the choice of prior distributions which favor regular vine copulas with fewer parameters over those with more parameters.

Bayesian techniques to select the pair copula families of D-vine copulas are covered in [50], [51] and [60]. The reversible jump MCMC algorithm presented in this section follows the latest developments of family selection methods for regular vine copulas laid out in [28].

Reversible jump Markov chain Monte Carlo is an extension of ordinary MCMC to sample from discrete-continuous posterior distributions. The sampling algorithm and the mathematics underpinning the convergence statements are an immediate generalization of the Metropolis–Hastings algorithm [27]. A reversible jump MCMC algorithm functions like an ordinary MCMC algorithm with one extra step built in: before, or after, updating the parameters, a "jump" to another model is attempted. That is called the "between-models move," while the updating of the parameters only is called the "within-model move" [16].

Algorithm 4.1 implements a reversible jump MCMC sampler which samples from the joint posterior distribution of the pair copula families $\mathscr{B}$ and the parameters θ, given a regular vine tree structure $\mathscr{V}$. In each iteration, it first updates all parameters as in Algorithm 3.1. Then the pair copula families are updated along with their parameters one-by-one. The pair copula family proposals are drawn from a uniform distribution over $\mathbf{B}$, while the parameter proposals are sampled from a (multivariate) normal distribution centered at the maximum likelihood estimates of the parameters.

Algorithm 4.1 (Reversible Jump MCMC Sampler for Family Selection)

1: *Choose arbitrary copula family and parameter starting values* $(\mathscr{B}^0, \theta^0) = ((\mathscr{B}_1^0, \theta_1^0), \ldots, (\mathscr{B}_D^0, \theta_D^0))$*, where* $D := n(n-1)/2$.
2: ***for*** *each iteration* $r = 1, \ldots, R$ ***do***
3: *Set* $(\mathscr{B}^r, \theta^r) = (\mathscr{B}^{r-1}, \theta^{r-1})$.
4: *Perform one update step of Algorithm 3.1 for the parameters* θ^r*; denote the updated parameter entries by* θ^+.
5: ***for*** $i = 1, ..., D$ ***do***
6: *Draw* $\mathscr{B}_i^r$ *from a* $Unif(\mathbf{B})$ *distribution.*

7: *Draw θ_i^r from a multivariate $N(\hat{\theta}_i^r, \hat{\Sigma}_i^r)$ distribution with probability density function $\phi_{(\hat{\theta}_i^r, \hat{\Sigma}_i^r)}(\cdot)$. Here $\hat{\theta}_i^r$ denotes the MLE of the copula parameters of the pair copula $\mathscr{B}_i^r$ and $\hat{\Sigma}_i^r$ denotes the estimated approximative covariance matrix of the parameter estimates $\hat{\theta}_i^r$.*

8: *Evaluate the generalized Metropolis–Hastings acceptance probability of the proposal*

$$\begin{aligned}\alpha &:= \alpha((\mathscr{B}_i^{r-1}, \theta_i^+), (\mathscr{B}_i^r, \theta_i^r)) \\ &= \frac{p(\mathscr{B}^*, \theta^*)}{p(\mathscr{B}, \theta)} \cdot \frac{\phi_{(\hat{\theta}_i^{r-1}, \hat{\Sigma}_i^{r-1})}(\theta_i^+)}{\phi_{(\hat{\theta}_i^r, \hat{\Sigma}_i^r)}(\theta_i^r)} \\ &= \frac{\pi(\mathscr{B}^*, \theta^*)}{\pi(\mathscr{B}, \theta)} \cdot \frac{L(\mathscr{B}^*, \theta^* | \mathbf{x})}{L(\mathscr{B}, \theta | \mathbf{x})} \cdot \frac{\phi_{(\hat{\theta}_i^{r-1}, \hat{\Sigma}_i^{r-1})}(\theta_i^+)}{\phi_{(\hat{\theta}_i^r, \hat{\Sigma}_i^r)}(\theta_i^r)},\end{aligned}$$

where $\mathscr{B}^* = (\mathscr{B}_1^r, \ldots, \mathscr{B}_i^r, \mathscr{B}_{i+1}^{r-1}, \ldots, \mathscr{B}_D^{r-1})$, $\mathscr{B} = (\mathscr{B}_1^r, \ldots, \mathscr{B}_{i-1}^r, \mathscr{B}_i^{r-1}, \ldots, \mathscr{B}_D^{r-1})$, $\theta^* = (\theta_1^r, \ldots, \theta_i^r, \theta_{i+1}^+, \ldots, \theta_D^+)$, and $\theta = (\theta_1^r, \ldots, \theta_{i-1}^r, \theta_i^+, \ldots, \theta_D^+)$.

9: *Accept the proposal $(\mathscr{B}_i^r, \theta_i^r)$ with probability α; if rejected, set $(\mathscr{B}_i^r, \theta_i^r) = (\mathscr{B}_i^{r-1}, \theta_i^+)$.*

10: ***end for***

11: ***end for***

12: ***return*** $((\mathscr{B}^1, \theta^1), \ldots, (\mathscr{B}^R, \theta^R))$

5 Selection of Vine Tree Structure, Pair Copula Families and Parameters

As pair copula families $\mathscr{B} = \mathscr{B}(\mathscr{V})$ and parameters $\theta = \theta(\mathscr{B}(\mathscr{V}))$ both depend on the vine tree structure $\mathscr{V}$, the identification of adequate trees is crucial to the model building process using vine copulas. As it was already the case for pair copula selection in Section 4.1, it is again not feasible to simply try and fit all possible regular vine copulas $C = (\mathscr{V}, \mathscr{B}, \theta)$ and then choose the 'best' one. In particular, the number of possible regular vines on n variables is $\frac{n!}{2} \times 2^{\binom{n-2}{2}}$ as shown by [52]. This means that even if pair copulas and parameters were known, the number of different models would still be enormous.

This remains true even when the selection is restricted to the sub-classes of C- and D-vines, since there are still $n!/2$ different C- and D-vines in n dimensions, respectively (see [1]). It should however be noted that C- and D-vine copulas are most appropriate if their structure is explicitly motivated by the data. In particular, C-vine copulas may be used if there is a set of pivotal variables such as stock indices (see [31, 12]) and D-vine copulas are particularly attractive to model variables with temporal order (see [60, 13]). Nevertheless we describe how C- and D-vines can be selected for arbitrary data sets.

5.1 Top-Down and Bottom-Up Selection

Due to the proximity condition (see Definition 2.2) regular vine trees are closely linked to each other and have to be constructed carefully. Two construction strategies have been proposed in the literature: a top-down approach by [23] and a bottom-up method by [42]. Both strategies proceed sequentially tree by tree and respect the proximity condition in each step. We first describe the top-down, then the bottom-up method.

5.1.1 Top-Down Selection

Selecting regular vine trees top-down means that one starts with the selection of the first tree T_1 and continues tree by tree up to the last tree T_{n-1}. The first tree T_1 can be selected as an arbitrary spanning tree. Given that a tree $T_m,\ m \in \{1,...,n-2\}$, has been selected, the next tree T_{m+1} is chosen respecting the proximity condition (see Definition 2.2). In other words, T_{m+1} can only be formed by (conditional) pairs $\{i(e), j(e)|D(e)\}$ which fulfill the proximity condition.

Example 5.1 (Top-Down Tree Selection). Assume that we have selected the first tree T_1 as shown in Figure 1. Then the question is which pairs $\{i(e), j(e)|D(e)\}$ are eligible for tree construction in the second tree T_2. According to the proximity condition these are $\{1,3|2\}$, $\{2,4|3\}$, $\{2,5|3\}$ and $\{4,5|3\}$. Obviously, the last three pairs form a cycle and therefore only two of them can be selected for T_2. One of the three possibilities is shown in Figure 1.

To perform this iterative selection strategy a criterion is needed to select a spanning tree among the set of eligible edges, where a spanning tree simply denotes a tree on all nodes. Clearly, the log-likelihood ℓ_m of the pair copulas in tree T_m of a regular vine copula (see Expression (4)) is given by

$$\ell_m\left(T_m, \mathscr{B}_{T_m}, \boldsymbol{\theta}_{T_m}|\mathbf{x}\right) = \sum_{k=1}^{N} \sum_{e \in E_m} \log\left(c_{i(e),j(e)|D(e)}\left(F_{i(e)|D(e)}, F_{j(e)|D(e)}|\boldsymbol{\theta}_{i(e),j(e)|D(e)}\right)\right), \tag{5}$$

where we write $\mathscr{B}_{T_m} := \mathscr{B}(T_m)$ and $\boldsymbol{\theta}_{T_m} := \boldsymbol{\theta}(\mathscr{B}(T_m))$.

A straightforward solution therefore would be to choose the tree such that (5) is maximized after having selected pair copulas with high log-likelihood for each (conditional) pair $\{i(e), j(e)|D(e)\}$ which fulfills the proximity condition. This solution however leads to highly over-parameterized models, since models with more parameters in which simpler models are nested will always give a higher likelihood. For instance, the Student's t copula always has a higher likelihood value than the Gaussian copula which is a special case of the Student's t as the degrees of freedom go to infinity.

Therefore we formulate the following algorithm in terms of a general weight ω assuming that we want to maximize it for each tree. The previously discussed strategy corresponds to choosing the pair copula log-likelihoods as weights.

Algorithm 5.1 (Sequential Top-Down Selection Based on Weights)

1: *Calculate the weight* $\omega_{i,j}$ *for all possible variable pairs* $\{i, j\}, 1 \leq i < j \leq n$.
2: *Select the maximum spanning tree, i.e.*

$$T_1 = \underset{T=(N,E) \text{ spanning tree}}{\operatorname{argmax}} \sum_{e \in E} \omega_{i(e),j(e)}.$$

3: ***for*** *each edge* $e \in E_1$ ***do***
4: *Select a copula* $C_{i(e),j(e)}$.
5: *Estimate the corresponding parameter(s)* $\theta_{i(e),j(e)}$.
6: *For* $k = 1, ..., N$ *transform* $F_{i(e)|j(e)}(x_{k,i(e)}|x_{k,j(e)}, \hat{\theta}_{i(e),j(e)})$ *and* $F_{j(e)|i(e)}(x_{k,j(e)}|x_{k,i(e)}, \hat{\theta}_{i(e),j(e)})$ *using* (3).
7: ***end for***
8: ***for*** $m = 2, \ldots, n-1$ ***do***
9: *Calculate the weight* $\omega_{i(e),j(e)|D(e)}$ *for all conditional variable pairs* $\{i(e), j(e)|D(e)\}$ *that can be part of tree* T_m*, i.e., all edges* E_P *fulfilling the proximity condition.*
10: *Among these edges, select the maximum spanning tree, i.e.,*

$$T_m = \underset{T=(N,E) \text{ spanning tree with } E \subset E_P}{\operatorname{argmax}} \sum_{e \in E} \omega_{i(e),j(e)|D(e)}.$$

11: ***for*** *each edge* $e \in E_m$ ***do***
12: *Select a conditional copula* $C_{i(e),j(e)|D(e)}$.
13: *Estimate the corresponding parameter(s)* $\theta_{i(e),j(e)|D(e)}$.
14: *For* $k = 1, ..., N$ *transform* $F_{i(e)|j(e)\cup D(e)}(x_{k,i(e)}|x_{k,j(e)}, \mathbf{x}_{k,D(e)}, \hat{\theta}_{i(e),j(e)|D(e)})$ *and* $F_{j(e)|i(e)\cup D(e)}(x_{k,j(e)}|x_{k,i(e)}, \mathbf{x}_{k,D(e)}, \hat{\theta}_{i(e),j(e)|D(e)})$ *using* (3).
15: ***end for***
16: ***end for***
17: ***return*** *the sequential model estimate* $(\hat{\mathscr{V}}, \hat{\mathscr{B}}, \hat{\theta})$.

Clearly, this algorithm only makes a locally optimal selection in each step, since the impact on previous and subsequent trees is ignored. The strategy is however reasonable in light of the definition of regular vines (see Definition 2.2). The maximum spanning tree in lines 2 and 10 can be found, e.g., using the classical algorithms by Prim or Kruskal (see [19, Section 23.2]). Possible choices for the weight ω are, for example,

- the absolute empirical Kendall's τ as proposed by [23] and [22];
- the AIC of each pair copula corresponding to the discussion in Section 4.1;
- the (negative) estimated degrees of freedom of Student's t pair copulas (see [47]);
- the p-value of a copula goodness of fit test and variants as proposed by [21].

Some remarks: First, taking the empirical Kendall's τ as weight does not require to select and estimate pair copulas prior to the tree selection step. The other three weights in fact require this, so that lines 4-5 and 12-13 in Algorithm 5.1 may be redundant in this case. Second, the AIC as weight also maximizes the AIC of the complete tree, since the individual AICs can be summed up like the log-likelihood.

Third, the strategy proposed by [47] concentrates on tail dependence, since the algorithm will select pairs with estimated small degrees of freedom corresponding to a stronger deviation from Gaussianity with no tail dependence.

Copula goodness of fit tests based on the Cramér-von Mises statistic are considered in [21] to take into account the uncertainty of the pair copula fit. Corresponding p-values are calculated using fast bootstrap methods based on the multiplier approach developed and implemented in [39, 40, 41].

Finally, we like to note that the empirical Kendall's τ as weight also approximately maximizes (5), since it can be checked that typically log-likelihoods of bivariate copulas increase with increasing absolute value of Kendall's τ. It however does not lead to over-parameterization, since copula selection and tree selection in a particular tree are independent. The strategy by [23] has therefore already been used successfully in applications with up to 52 variables (see [12]).

Recently, [35] proposed the construction of vines with non-parametric pair copulas based on empirical copula estimators. For vine tree selection [35] also advocate the use of empirical dependence measures such as Spearman's ρ, since it does not require any parametric assumption on the pair copulas. Other dependence measures such as Blomqvist's β could of course be used instead in the approaches by [23] and [35].

Especially the strategy based on Kendall's τ leads to regular vine copulas with decreasing dependence in higher order trees. This is advantageous to the vine structure, since numerical errors may add up and estimates become less precise as the number of conditioning variables increases. It also has been exploited by [14] to identify appropriate truncation levels of regular vine copulas.

Algorithm 5.1 can easily be modified to select C- or D-vines instead of general regular vines. For C-vines the root node in each tree can simply be identified as the node with maximal sum of weights to all other nodes (see [22]). In the case of D-vines only the order of variables in the first tree has to be chosen. Since D-vine trees are paths on all nodes, so-called Hamiltonian paths, a maximum Hamiltonian path has to be identified. This problem is equivalent to a traveling salesman problem as discussed by [11]. As an NP-hard problem, there is no known efficient algorithm to find a solution.

5.1.2 Bottom-up selection

Rather than beginning with the first tree T_1, [42] proposed a bottom-up selection strategy which starts with tree T_{n-1} and then sequentially selects trees T_m for $m = n-2, \ldots, 1$. Similar to top-down selection using Kendall's τ by [23], this approach is also motivated by choosing a regular vine with weaker dependence in later trees. Here, dependence is however measured in terms of partial correlations. More precisely, each tree is selected such that the sum of absolute partial correlations as edge weights is minimized. This is feasible, since partial correlation estimates for each combination of variables can be obtained from the data without any parametric assumptions (see, e.g., [66]).

The bottom-up strategy by [42] thus selects tree T_{n-1} as the edge corresponding to the pair of variable with lowest absolute partial correlation given all other variables. [42] then provides conditions such that the proximity condition is satisfied when a tree $T_m,\ m \in \{1,...,n-2\}$, is selected given trees $T_{m+1},...,T_{n-1}$. As before there may be several choices possible in the set of eligible edges.

Example 5.2 (Bottom-Up Tree Selection). Assume in Figure 1 that we have selected T_4 with edge $e = \{a,b\} = \{1,5|2,3,4\}$ as shown. Then $A_a = \{1,2,3,4\}$ and $A_b = \{2,3,4,5\}$ (see (1)). Any choice of edges from A_a and A_b in tree T_3 leads to a valid tree T_4. For instance, the edges $\{1,2|3,4\}$ and $\{4,5|2,3\}$ are compatible with A_a and A_b, respectively, and would lead to another vine than the one shown in Figure 1. In fact, it could be constructed as a D-vine with order $1,4,3,2,5$.

Since partial correlations are equal to conditional correlations for elliptical distributions (see [5]), this strategy is particularly appropriate if all pair copula families are elliptical. Contrariwise to the top-down approach discussed above which requires the selection and estimation of pair copulas in each tree, this strategy chooses the vine tree structure without any assumption on the pair copula families. Having selected the vine tree structure, the selection of pair copulas can therefore proceed as discussed in Section 4.1.

5.2 Sequential Bayesian Tree Selection

Bayesian approaches to estimating the posterior distribution of the tree structure of a regular vine copula have only recently been developed. The sheer number of possible regular vine tree structures poses a relevant challenge to computationally efficient posterior evaluation.

This section presents a reversible jump MCMC-based approach proposed by [28] to obtain a sequential estimate of the posterior distribution of the regular vine tree structure $\mathscr{V}$, the pair copula families $\mathscr{B}$, and the copula parameters θ. The term "sequential estimation" is understood analogously to the notion discussed in Section 5.1.1 as a tree-by-tree procedure.

As in Section 4.2, model priors $\pi(\mathscr{V},\mathscr{B},\theta)$ which favor sparse models can serve to guard against selecting models with runaway complexity. On the other hand, the use of non-informative flat priors allows for tree-by-tree maximum likelihood estimation of the regular vine tree structure $\mathscr{V}$, the pair copula families $\mathscr{B}$, and the copula parameters θ. As a result, the posterior mode estimate of this procedure will agree with the model estimate of Algorithm 5.1, if flat priors are used. Again, the pair copula families are chosen from a set $\mathbf{B}$ of parametric pair copula families.

For the sake of notational convenience and enhanced readability, the procedure is presented in two algorithms. Algorithm 5.2 gives the general outline of the sampling procedure. Algorithm 5.3 details the reversible jump MCMC algorithm to sample from the posterior distribution of tree T_m of the regular vine tree structure $\mathscr{V} = (T_1,\ldots,T_{n-1})$.

Algorithm 5.2 implements the tree-by-tree selection procedure. As such, it calls Algorithm 5.3 for the actual posterior estimation. However, it condenses the posterior sample into the posterior mode estimate and organizes the move from selecting tree T_m to selecting tree T_{m+1}. The posterior mode estimate $(\hat{T}_m, \hat{\mathscr{B}}_{T_m}, \hat{\theta}_{T_m})$ of tree T_m is the most frequently sampled combination of tree structure T_m and pair copula families $\mathscr{B}_{T_m}$ parameterized at the mode of the posterior sample of the parameters θ_{T_m} of this model.

Algorithm 5.2 (Outline of the Tree-by-Tree Sampling Algorithm)

1: *Sample from the posterior distribution of the first tree,* $(T_1, \mathscr{B}_{T_1}, \theta_{T_1})$*: see Algorithm 5.3.*
2: *Set the tree estimate* $(\hat{T}_1, \hat{\mathscr{B}}_{T_1}, \hat{\theta}_{T_1})$ *to the posterior mode.*
3: ***for*** $m = 2, ..., n-1$ ***do***
4: *Sample from the posterior distribution of the m-th tree,* $(T_m, \mathscr{B}_{T_m}, \theta_{T_m})$*, given the previous trees' estimates* $(\hat{T}_1, \hat{\mathscr{B}}_{T_1}, \hat{\theta}_{T_1}), \ldots, (\hat{T}_{m-1}, \hat{\mathscr{B}}_{T_{m-1}}, \hat{\theta}_{T_{m-1}})$*: see Algorithm 5.3.*
5: *Set the tree estimate* $(\hat{T}_m, \hat{\mathscr{B}}_{T_m}, \hat{\theta}_{T_m})$ *to the posterior mode.*
6: ***end for***
7: ***return*** *the sequential Bayesian model estimate*

$$(\hat{\mathscr{V}}, \hat{\mathscr{B}}, \hat{\theta}) = ((\hat{T}_1, \ldots, \hat{T}_{n-1}), (\hat{\mathscr{B}}_{T_1}, \ldots, \hat{\mathscr{B}}_{T_{n-1}}), (\hat{\theta}_{T_1}, \ldots, \hat{\theta}_{T_{n-1}})).$$

Algorithm 5.3 samples from the posterior distribution of $(T_m, \mathscr{B}_{T_m}, \theta_{T_m})$ given the previously selected trees $((\hat{T}_l, \hat{\mathscr{B}}_{T_l}, \hat{\theta}_{T_l}), l = 1, \ldots, m-1)$. The within-model move of this algorithm to update the pair copula parameters θ follows Algorithm 3.1. In the between-models move, the proposal trees are sampled as spanning trees which satisfy the proximity condition. Practitioners may note that this functionality is implemented, e.g., in the C++ `boost` library [58]. Edge weights can be used to fine-tune the variance of the proposal distribution: in our implementation, higher values for $p \in (0,1)$ increase the probability that the proposals for the tree structure, T_m^r, are similar to the current state T_m^{r-1} of the sampling chain. As in Algorithm 4.1, the pair copula family proposals are sampled from a uniform distribution over $\mathbf{B}$, and the parameter proposals are drawn from a (multivariate) normal distribution centered at the maximum likelihood estimates of the parameters.

Algorithm 5.3 (Reversible Jump MCMC Algorithm to Sample $(T_m, \mathscr{B}_{T_m}, \theta_{T_m})$**)**

Denote the pair copula families of tree $T_m = (N_m, E_m)$ *by* $\mathscr{B}_{T_m} = (\mathscr{B}_e, e \in E_m)$*, and the corresponding pair copula parameters by* $\theta_{T_m} = (\theta_e, e \in E_m)$*.*

1: *Choose arbitrary, but valid, tree structure, copula family and parameter starting values* $(T_m^0, \mathscr{B}_{T_m}^0, \theta_{T_m}^0)$*. If* $m \geq 2$*, observe the proximity condition imposed on* T_m^0*.*
2: ***for*** *each iteration* $r = 1, \ldots, R$ ***do***
3: *Set* $(T_m^r, \mathscr{B}_{T_m}^r, \theta_{T_m}^r) = (T_m^{r-1}, \mathscr{B}_{T_m}^{r-1}, \theta_{T_m}^{r-1})$*.*
4: *Perform one update step of Algorithm 3.1 for the pair copula parameters of tree* T_m*,* $\theta_{T_m}^r$*; denote the updated parameter entries by* $\theta_{T_m}^+$*.*
5: *Draw a spanning tree* T_m^r *which satisfies the proximity condition from the proposal distribution*

$$q(T_m^{r-1} = (N_m, E_m^{r-1}) \to T_m^r = (N_m, E_m^r)) \propto p^{|E_m^r \cap E_m^{r-1}|}(1-p)^{|E_m^r \setminus E_m^{r-1}|}.$$

6: ***for*** *each edge* $e \in E_m^r$ ***do***

7: *Draw* $\mathscr{B}_e^r$ *from a* $Unif(\mathbf{B})$ *distribution.*

8: *Draw* θ_e^r *from a multivariate* $N(\hat{\theta}_e^r, \hat{\Sigma}_e^r)$ *distribution with probability density function* $\phi_{(\hat{\theta}_e^r, \hat{\Sigma}_e^r)}(\cdot)$. *Here* $\hat{\theta}_e^r$ *denotes the MLE of the copula parameters of the pair copula* $\mathscr{B}_e^r$ *and* $\hat{\Sigma}_e^r$ *denotes the estimated approximative covariance matrix of the parameter estimates* $\hat{\theta}_e^r$.

9: ***end for***

10: *Set* $\mathscr{B}_{T_m}^r = (\mathscr{B}_e^r, e \in E_m^r)$, $\theta_{T_m}^r = (\theta_e^r, e \in E_m^r)$.

11: *Evaluate the generalized Metropolis–Hastings acceptance probability of the proposal*

$$\begin{aligned}
\alpha &:= \alpha((T_m^{r-1}, \mathscr{B}_{T_m}^{r-1}, \theta_{T_m}^+), (T_m^r, \mathscr{B}_{T_m}^r, \theta_{T_m}^r)) \\
&= \frac{p(\mathscr{V}^*, \mathscr{B}^*, \theta^*)}{p(\mathscr{V}, \mathscr{B}, \theta)} \cdot \frac{\prod_{e \in E_m^{r-1}} \phi_{(\hat{\theta}_e^{r-1}, \hat{\Sigma}_e^{r-1})}(\theta_e^+)}{\prod_{e \in E_m^r} \phi_{(\hat{\theta}_e^r, \hat{\Sigma}_e^r)}(\theta_e^r)} \\
&= \frac{\pi(\mathscr{V}^*, \mathscr{B}^*, \theta^*)}{\pi(\mathscr{V}, \mathscr{B}, \theta)} \cdot \frac{L(\mathscr{V}^*, \mathscr{B}^*, \theta^* | \mathbf{x})}{L(\mathscr{V}, \mathscr{B}, \theta | \mathbf{x})} \cdot \frac{\prod_{e \in E_m^{r-1}} \phi_{(\hat{\theta}_e^{r-1}, \hat{\Sigma}_e^{r-1})}(\theta_e^+)}{\prod_{e \in E_m^r} \phi_{(\hat{\theta}_e^r, \hat{\Sigma}_e^r)}(\theta_e^r)},
\end{aligned}$$

where $\mathscr{V}^* = (\hat{T}_1, \ldots, \hat{T}_{m-1}, T_m^r)$, $\mathscr{V} = (\hat{T}_1, \ldots, \hat{T}_{m-1}, T_m^{r-1})$, $\mathscr{B}^* = (\hat{\mathscr{B}}_{T_1}, \ldots, \hat{\mathscr{B}}_{T_{m-1}}, \mathscr{B}_{T_m}^r)$, $\mathscr{B} = (\hat{\mathscr{B}}_{T_1}, \ldots, \hat{\mathscr{B}}_{T_{m-1}}, \mathscr{B}_{T_m}^{r-1})$, $\theta^* = (\hat{\theta}_{T_1}, \ldots, \hat{\theta}_{T_{m-1}}, \theta_{T_m}^r)$, *and* $\theta = (\hat{\theta}_{T_1}, \ldots, \hat{\theta}_{T_{m-1}}, \theta_{T_m}^+)$.

12: *Accept the proposal* $(T_m^r, \mathscr{B}_{T_m}^r, \theta_{T_m}^r)$ *with probability* α*; if rejected, set* $(T_m^r, \mathscr{B}_{T_m}^r, \theta_{T_m}^r) = (T_m^{r-1}, \mathscr{B}_{T_m}^{r-1}, \theta_{T_m}^+)$.

13: ***end for***

14: ***return*** $((T_m^1, \mathscr{B}_{T_m}^1, \theta_{T_m}^1), \ldots, (T_m^R, \mathscr{B}_{T_m}^R, \theta_{T_m}^R))$

6 Conclusions and Outlook

We have focused on the various model selection and estimation methods for regular vine copulas. Since a regular vine model is specified by three linked components, this results in three fundamental tasks with increasing complexity. We discussed frequentist and Bayesian approaches for each of these tasks. In particular this involved sequential approaches starting from the top tree until the last tree. The frequentist approaches are implemented in the R-packages `CDVine` [15, 56] for D- and C-vines and `VineCopula` [57] for regular vines, respectively.

In view of the linked nature of the vine tree structure the sequential approach is a natural approach. However as in the case of covariate selection in linear models, this might not yield the best fit to the data. In addition the approaches so far primarily considered in-sample fit and model comparisons, more empirical work is needed to validate the models in an out-of-sample performance study. However this is now feasible and tractable and is the subject of current investigations.

It has been recently recognized by [2] that even the flexible class of simplified regular vines might be insufficient for some data sets. This might be the result that the underlying joint density is not well approximated by a simplified regular vine density, where the conditional copula family terms are chosen to be independent of the conditioning value. Currently the proposed solution of [2] using non-parametric two-dimensional smoothing methods is limited to three dimensions and it will be a major challenge to extend the model and the selection methods to higher dimensions.

Other non-standard vine models occur when the pair copula terms depend on the conditioning time point, thus yielding models with time varying copula parameters. These can also be seen as non-simplified vines in the special case that the conditioning variables follow a functional relationship to time. First parameter driven time dependence was considered using an AR(1) dynamics in the copula parameters in the papers by [3] and [4], while [61] follow a regime switching approach. Here only parameter estimation and assessment of the uncertainty of the parameters are considered so far. It has to be investigated if the additional flexibility of the copula families and the different vine tree structures is needed here.

This survey primarily focused on the selection and estimation problem of regular vines with parametric pair copula families. The approach of [35] can be compared to ones which are based on kernel methods as proposed in [7] and [55].

Other data structures than multivariate continuous data have also gained by models based on pair copula constructions. In particular network structures were considered by [29] and [43] in a primarily Gaussian set-up and by [7] in a non-Gaussian setting. Here the network is modeled by a directed acyclic graph (DAG) model. While learning the network structure from data is a very complex task even in Gaussian DAG models, non-Gaussian learning algorithms are currently developed and investigated in [6].

Another very interesting multivariate data structure are discrete and mixed discrete-continuous outcomes which occur most often in the life sciences. Pair copula constructions using D-vines for discrete outcomes were recently developed in [54] together with highly efficient parameter estimation techniques. Current research is conducted to allow for mixed outcomes [62].

While these extensions are still concerned primarily with parameter estimation, it remains an important open challenge to find non-sequential solutions to the selection of regular vines when all three components have to be selected.

Acknowledgements The second and third author gratefully acknowledge the support of the TUM Graduate School's International School of Applied Mathematics. The second author is also supported by a research stipend from Allianz Deutschland AG.

References

1. Aas, K., Czado, C., Frigessi, A., Bakken, H.: Pair-copula constructions of multiple dependence. Insurance: Mathematics and Economics **44**, 182–198 (2009)
2. Acar, E.F., Genest, C., Nešlehová, J.: Beyond simplified pair-copula constructions. Journal of Multivariate Analysis **110**, 74–90 (2012)
3. Almeida, C., Czado, C.: Efficient Bayesian inference for stochastic time-varying copula models. Computational Statistics & Data Analysis **56**(6), 1511–1527 (2012)
4. Almeida, C., Czado, C., Manner, H.: Modeling high dimensional time-varying dependence using D-vine SCAR models (2012). Preprint, http://arxiv.org/abs/1202.2008
5. Baba, K., Shibata, R., Sibuya, M.: Partial correlation and conditional correlation as measures of conditional independence. Australian and New Zealand Journal of Statistics **46(4)**, 657–664 (2004)
6. Bauer, A., Czado, C.: Pair copula Bayesian networks (2012). Preprint, http://arxiv.org/abs/1211.5620
7. Bauer, A., Czado, C., Klein, T.: Pair-copula constructions for non-Gaussian DAG models. The Canadian Journal of Statistics **40**(1), 86–109 (2012)
8. Bedford, T., Cooke, R.M.: Probability Density Decomposition for Conditionally Dependent Random Variables Modeled by Vines. Annals of Mathematics and Artificial Intelligence **32**, 245–268 (2001)
9. Bedford, T., Cooke, R.M.: Vines: A New Graphical Model for Dependent Random Variables. Annals of Statistics **30**(4), 1031–1068 (2002)
10. Berg, D., Aas, K.: Models for construction of higher-dimensional dependence: A comparison study. European Journal of Finance **15**, 639–659 (2009)
11. Brechmann, E.C.: Truncated and simplified regular vines and their applications. Master's thesis, Technische Universität München (2010)
12. Brechmann, E.C., Czado, C.: Risk management with high-dimensional vine copulas: An analysis of the Euro Stoxx 50. (2011). Preprint, http://mediatum.ub.tum.de/doc/1079276/1079276.pdf
13. Brechmann, E.C., Czado, C.: COPAR - Multivariate time series modeling using the COPula AutoRegressive model. (2012). Preprint, http://arxiv.org/abs/1203.3328
14. Brechmann, E.C., Czado, C., Aas, K.: Truncated regular vines in high dimensions with application to financial data. Canadian Journal of Statistics **40**(1)
15. Brechmann, E.C., Schepsmeier, U.: Modeling dependence with C- and D-vine copulas: The R-package CDVine. To appear in the Journal of Statistical Software (2012)
16. Brooks, S., Gelman, A., Jones, G.L., Meng, X.L.: Handbook of Markov Chain Monte Carlo. Chapman & Hall/CRC (2011)
17. Chollete, L., Heinen, A., Valdesogo, A.: Modeling international financial returns with a multivariate regime switching copula. Journal of Financial Econometrics **7**, 437–480 (2009)
18. Clarke, K.A.: A simple distribution-free test for nonnested model selection. Political Analysis **15**(3), 347–363 (2007)
19. Cormen, T., Leiserson, E., Rivest, R., Stein, C.: Introduction to Algorithms. The MIT Press, Camebridge (2009)
20. Czado, C.: Pair copula constructions of multivariate copulas. In: P. Jaworki, F. Durante, W. Härdle, W. Rychlik (eds.) Workshop on Copula Theory and its Applications. Springer (2010)
21. Czado, C., Jeske, S., Hofmann, M.: Selection strategies for regular vine copulae (2012). Preprint
22. Czado, C., Schepsmeier, U., Min, A.: Maximum likelihood estimation of mixed C-vines with application to exchange rates. Statistical Modelling **12**(3), 229–255 (2012)
23. Dißmann, J., Brechmann, E.C., Czado, C., Kurowicka, D.: Selecting and estimating regular vine copulae and application to financial returns. Computational Statistics & Data Analysis **59**, 52–69 (2013)

24. Fischer, M., Köck, C., Schlüter, S., Weigert, F.: An empirical analysis of multivariate copula models. Quantitative Finance **9**(7), 839–854 (2009)
25. Genest, C., Ghoudi, K., Rivest, L.P.: A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. Biometrika **82**, 543–552 (1995)
26. Genest, C., Werker, B.J.M.: Conditions for the asymptotic semiparametric efficiency of an omnibus estimator of dependence parameters in copula models. In: C.M. Cuadras, J. Fortiana, J.A. Rodríguez Lallena (eds.) Distributions with Given Marginals and Statistical Modelling, pp. 103–112. Kluwer, Dordrecht (2002)
27. Green, P.J.: Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. Biometrika **82**, 711–732 (1995)
28. Gruber, L.F., Stöber, J., Czado, C.: Bayesian model selection for regular vine copulas using reversible jump MCMC (2012). Preprint
29. Hanea, A., Kurowicka, D., Cooke, R.: Hybrid method for quantifying and analyzing Bayesian belief nets. Quality and Reliability Engineering International **22**(6), 613–729 (2006)
30. Hastings, W.: Monte Carlo sampling methods using Markov chains and their applications. Biometrika **57**, 97–109 (1970)
31. Heinen, A., Valdesogo, A.: Asymmetric CAPM dependence for large dimensions: The Canonical Vine Autoregressive Model. CORE discussion papers 2009069, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE) (2009)
32. Hobæk Haff, I.: Parameter estimation for pair-copula constructions. To appear in Bernoulli (2011)
33. Hobæk Haff, I.: Comparison of estimators for pair-copula constructions. Journal of Multivariate Analysis **110**, 91–105 (2012)
34. Hobæk Haff, I., Aas, K., Frigessi, A.: On the simplified pair-copula construction - simply useful or too simplistic? Journal of Multivariate Analysis **101**(5), 1296–1310 (2010)
35. Hobæk Haff, I., Segers, J.: Nonparametric estimation of pair-copula constructions with the empirical pair-copula (2012). Preprint, `http://arxiv.org/abs/1201.5133`
36. Joe, H.: Families of m-variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. In: L. Rüschendorf, B. Schweizer, M.D. Taylor (eds.) Distributions with fixed marginals and related topics, pp. 120–141. Institute of Mathematical Statistics, Hayward (1996)
37. Joe, H.: Asymptotic efficiency of the two-stage estimation method for copula-based models. Journal of Multivariate Analysis **94**, 401–419 (2005)
38. Joe, H., Xu, J.: The estimation method of inference functions for margins for multivariate models. Tech. Rep. 166, Department of Statistics, University of British Columbia (1996)
39. Kojadinovic, I., Yan, J.: Modeling multivariate distributions with continuous margins using the copula R package. Journal of Statistical Software **34(9)**, 1–20 (2010)
40. Kojadinovic, I., Yan, J.: A goodness-of-fit test for multivariate multiparameter copulas based on multiplier central limit theorems. Statistics and Computing **21**, 17–30 (2011)
41. Kojadinovic, I., Yan, J., Holmes, M.: Fast large-sample goodness-of-fit test for copulas. Statistica Sinica **21**, 841–871 (2011)
42. Kurowicka, D.: Optimal truncation of vines. In: D. Kurowicka, H. Joe (eds.) Dependence Modeling: Vine Copula Handbook. World Scientific Publishing Co. (2011)
43. Kurowicka, D., Cooke, R.: Distribution-Free Continuous Bayesian Belief Nets. In: A. Wilson, N. Limnios, S. Keller-McNulty, Y. Armojo (eds.) Modern Statistical and Mathematical Methods in Reliability, pp. 309–323 (2005)
44. Kurowicka, D., Cooke, R.M.: Uncertainty Analysis with High Dimensional Dependence Modelling. John Wiley & Sons Ltd. (2006)
45. Kurowicka, D., Joe, H.: Dependence Modeling: Vine Copula Handbook. World Scientific Publishing Co., Singapore (2011)
46. Manner, H.: Estimation and model selection of copulas with an application to exchange rates. METEOR research memorandum 07/056, Maastricht University (2007)
47. Mendes, B.V.d.M., Semeraro, M.M., Leal, R.P.C.: Pair-copulas modeling in finance. Financial Markets and Portfolio Management **24**(2), 193–213 (2010)

48. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of State Calculations by Fast Computing Machines. Journal of Chemical Physics **21**, 1087–1092 (1953)
49. Meyn, S., Tweedie, R.L.: Markov Chain and Stochastic Stability, second edn. Cambridge University Press (2009)
50. Min, A., Czado, C.: Bayesian inference for multivariate copulas using pair-copula constructions. Journal of Financial Econometrics **8(4)**, 511–546 (2010)
51. Min, A., Czado, C.: Bayesian model selection for D-vine pair-copula constructions. Canadian Journal of Statistics **39**, 239–258 (2011)
52. Morales-Nápoles, O.: Counting vines. In: D. Kurowicka, H. Joe (eds.) Dependence Modeling: Vine Copula Handbook. World Scientific Publishing Co. (2011)
53. Nikoloulopoulos, A.K., Joe, H., Li, H.: Vine copulas with asymmetric tail dependence and applications to financial return data. Computational Statistics & Data Analysis **56**(11), 3659–3673 (2012)
54. Panagiotelis, A., Czado, C., Joe, H.: Pair copula constructions for multivariate discrete data. Journal of the American Statistical Association **107**(499), 1063–1072 (2012)
55. Schellhase, C., Kauermann, G.: Flexible pair-copula estimation in D-vines with penalized splines (2012). Preprint
56. Schepsmeier, U., Brechmann, E.C.: CDVine: Statistical inference of C- and D-vine copulas (2012). R package version 1.1-9, `http://CRAN.R-project.org/package=CDVine`
57. Schepsmeier, U., Stoeber, J., Brechmann, E.C.: VineCopula: Statistical inference of vine copulas (2012). R package version 1.0, `http://CRAN.R-project.org/package=VineCopula`
58. Siek, J., Lee, L.Q., Lumsdaine, A.: Boost Graph Library: Random Spanning Tree - 1.49.0 (2012). `http://www.boost.org/doc/libs/1_49_0/libs/graph/doc/random_spanning_tree.html`
59. Sklar, A.: Fonctions de répartition à n dimensions et leurs marges. Publications de l'Institut de Statistique de l'Université de Paris **8**, 229–231 (1959)
60. Smith, M., Min, A., Almeida, C., Czado, C.: Modeling Longitudinal Data Using a Pair-Copula Decomposition of Serial Dependence. Journal of the American Statistical Association **105**(492), 1467–1479 (2010)
61. Stöber, J., Czado, C.: Detecting regime switches in the dependence structure of high dimensional financial data (2011). Preprint, `http://arxiv.org/abs/1202.2009`
62. Stöber, J., Czado, C., Hong, H., Ghosh, P.: Interrelation among chronic diseases in the elderly (2012). Working Paper
63. Stöber, J., Czado, C., Joe, H.: Simplified Pair Copula Constructions - Limits and Extensions (2012). Preprint, `http://arxiv.org/abs/1205.4844`
64. Stöber, J., Schepsmeier, U.: Is there significant time-variation in multivariate copulas? (2012). Preprint, `http://arxiv.org/abs/1205.4841`
65. Vuong, Q.H.: Ratio tests for model selection and non-nested hypotheses. Econometrica **57**(2), 307–333 (1989)
66. Whittaker, J.: Graphical Models in Applied Multivariate Statistics. Wiley (1990)