

Bayesian inference for dynamic vine copulas in higher dimensions

Alexander Kreuzer ^{*} Claudia Czado

Zentrum Mathematik, Technische Universität München

November 3, 2019

Abstract

We propose a class of dynamic vine copula models. This is an extension of static vine copulas and a generalization of dynamic C-vine and D-vine copulas studied by Almeida et al (2016) and Goel and Mehra (2019). Within this class, we allow for time-varying dependence by driving the vine copula parameters with latent AR(1) processes. This modeling approach is very flexible but estimation is not straightforward due to the high-dimensional parameter space. We propose a Bayesian estimation approach, which relies on a novel approximation of the posterior distribution. This approximation allows to use Markov Chain Monte Carlo methods, such as elliptical slice sampling, in a sequential way. In contrast to other Bayesian sequential estimation procedures for vine copula models as proposed by Gruber and Czado (2015), there is no need to collapse copula parameters to point estimates before proceeding to the next tree. Thus more information and uncertainty is propagated from lower to higher trees. A simulation study shows satisfactory performance of the Bayesian procedure. This dynamic modeling and inference approach can be applied in various fields, where static vine copulas have already proven to be successful, including environmental sciences, medicine and finance. Here we study the dependence among 21 exchange rates. For comparison we also estimate a static vine copula model and dynamic C-vine and D-vine copula models. This comparison shows superior performance of the proposed dynamic vine copula model with respect to one day ahead forecasting accuracy.

Keywords: vine copula, financial time series, time-varying parameters, Bayesian inference

1 Introduction

Appropriate models for dependence are crucial in many areas. For example, the risk of a portfolio, consisting of several financial assets, is influenced by the dependence among those assets (Embrechts et al (2002)). The multivariate normal distribution is often not sufficient to describe complex dependence structures. It does not allow for tail dependence or asymmetry in the tails. A more flexible framework is provided by copulas. Sklar's theorem (Sklar (1959)) states that any multivariate distribution can be decomposed into its marginal distributions and the corresponding copula. So the modeling process can be divided into two parts: the marginals and the dependence model. After modeling the margins we can focus entirely on the dependence structure, i.e. on finding an appropriate copula model. In higher dimensions, the class of vine copulas (Bedford and Cooke (2001), Aas et al (2009), Czado (2019)) is especially useful. Vine copulas are constructed from bivariate copula models. These bivariate copula models can be chosen individually from different classes of copulas, such as elliptical and Archimedean copulas. Due to their flexibility, vine copula models have gained huge popularity. They have been successfully applied in many fields, including environmental sciences (Erhardt et al (2015), Möller et al (2018)), medicine (Barthel et al (2018)) and finance (Brechmann and Czado (2013), Aas (2016), Stübinger et al (2018)).

^{*}Corresponding author: a.kreuzer@tum.de, Boltzmannstr. 3, 85748 Garching b. München, Germany

Although in many of these applications it is assumed that the dependence does not change over time, this assumption is often not appropriate. For example, there is evidence that the correlations between the returns of stocks and bonds change over time (Baele et al (2010)). Popular models for financial data that account for dynamic dependence are multivariate GARCH models with time-varying correlations, such as the DCC-GARCH (Engle (2002)) and multivariate factor stochastic volatility models (Harvey et al (1994), Pitt and Shephard (1999), Kastner et al (2017)). But these models again rely (conditionally) on multivariate normal distributions. New models have been proposed to overcome the shortcomings of the multivariate normal distribution and to allow for more flexible time-varying dependence structures. One example for such a model is the dynamic copula model of Oh and Patton (2018). Another approach to construct dynamic dependence models in higher dimensions is to extend the flexible class of vine copulas. As already mentioned above, vine copulas are constructed from bivariate copulas. So the vine copula framework allows us to scale dynamic bivariate copula models to arbitrary dimensions. Acar et al (2019) use nonparametric smoothing techniques to allow for time-variation in bivariate copula models and extend this bivariate approach to higher dimensions using vine copulas. Vatter and Chavez-Demoulin (2015) propose a bivariate copula model, where the copula parameters depend on covariates through generalized additive models. Using the vine copula framework, this bivariate model is extended to higher dimensions by Vatter and Nagler (2018). Similarly the bivariate dynamic copula model as proposed by Almeida and Czado (2012) and Hafner and Manner (2012) is extended to dynamic D-vine copulas by Almeida et al (2016) and later to dynamic C-vine copulas by Goel and Mehra (2019).

The bivariate dynamic copula model of Almeida and Czado (2012) provides a flexible building block by modeling time-varying dependencies with latent AR(1) processes. However, estimation is no longer straightforward since the likelihood involves high-dimensional integration. Goel and Mehra (2019) follow Almeida et al (2016), who use a frequentist estimation approach with approximation of the likelihood utilizing efficient importance sampling (Richard and Zhang (2007)). In this approach, parameters are estimated sequentially tree by tree. For estimating parameters of higher trees, the parameters of lower trees are fixed at point estimates. Thus uncertainty of parameters in lower trees is ignored and therefore uncertainty quantification cannot be provided.

The paper contains two major contributions: So far dynamic vine copula models, as generalization of the dynamic bivariate copula model of Almeida and Czado (2012), were restricted to D-vine (Almeida et al (2016)) and C-vine (Goel and Mehra (2019)) structures. First we develop an approach to allow for general regular vine tree structures. D-vine structures are especially suited to describe temporal dependence. But when it comes to cross-sectional dependence structures, such as the dependence among several stocks, the D-vine structure might be too restrictive. General regular vine tree structures are more flexible and include C-vine and D-vine structures as special cases.

Second, we present a novel Bayesian estimation approach. To our knowledge Bayesian estimation of vine copula models, including structure selection, was only tackled by Gruber and Czado (2015) and Gruber and Czado (2018). These approaches only allow for static pair copulas and have not been applied in more than 10 dimensions, while our approach allows for static as well as dynamic pair copulas and can handle higher dimensions.

Our methodology is based on an approximation of the posterior distribution. Approximations to the posterior are also used in variational Bayesian inference (Wainwright et al (2008)) and have become popular since they make estimation feasible in high-dimensional settings. Variational Bayesian approaches assume that the posterior distribution belongs to some family of distributions, such as the multivariate normal distribution. Our approach does not rely on such an assumption. We propose an approximation, which uses ideas of the frequentist sequential procedure of Dissmann et al (2013) and which allows to estimate pair copulas of one tree independently of each other using Markov Chain Monte Carlo (MCMC) schemes developed in Kreuzer and Czado (2019). The MCMC schemes rely on elliptical slice sampling (Murray et al (2010)) to exploit the underlying autoregressive dependence structure and on an ancillarity-sufficiency interweaving strategy (Yu and Meng (2011)). The posterior approximation also enables the user to run several MCMC chains in parallel leading to faster computation and making the approach applicable to higher dimensions than the approach of Gruber and Czado (2018). Further our Bayesian approach includes pair copula family selection based on

a set of candidate families. Here we exploit the fact that for several copula families, there is a one-to-one correspondence between the copula parameter and Kendall's τ . This allows to share the Kendall's τ parameter among different copula families, which reduces the parameter space and simplifies estimation. Additionally, our approach also contributes to the selection of sparse models by assessing, whether a pair copula term needs to be modeled dynamically or not. For this the information criteria of Watanabe (2010) is adapted. Another advantage of our Bayesian approach is that it is not necessary to fix copula parameters at point estimates although our procedure is sequential. Uncertainty of parameter estimates and information in lower trees is no longer ignored, but propagated as we move up to higher trees in the estimation procedure. All ideas are investigated through simulation and illustrated with real data.

The outline of the paper is as follows: Section 2 discusses the bivariate building blocks that are needed to construct the dynamic vine copula model. We introduce dynamic and static building blocks and show how to select among them. The selection procedure is illustrated with a small simulation study. Section 3 introduces the dynamic vine copula model and a novel algorithm for parameter estimation. The performance of the estimation procedure is evaluated with simulated data. In Section 4 we model the dependence among 21 exchange rates. Within this application, we compare the predictive accuracy of the proposed dynamic vine copula model to competitor models: a static vine copula, a dynamic C-vine copula and a dynamic D-vine copula. We conclude with providing ideas for future research in Section 5.

2 Bivariate building blocks for the dynamic vine copula model

The dynamic vine copula model, we introduce in Section 3, relies on dynamic and static bivariate copula models. These bivariate models are introduced in Sections 2.1 and 2.2. Section 2.3 discusses selection among different bivariate copula models.

2.1 Dynamic bivariate copulas

Model specification

We extend the (time-) dynamic bivariate copula model introduced by Almeida and Czado (2012) and Hafner and Manner (2012) by a model indicator m to allow for Bayesian copula family selection. We consider a set \mathcal{M} of single-parameter copula families for which there is a one-to-one correspondence between the copula parameter and Kendall's τ . The mapping from the copula parameter to the corresponding Kendall's τ is denoted by $g_m(\cdot)$ depending on the copula family $m \in \mathcal{M}$, e.g. for the Gumbel copula it holds that $g_{Gumbel}(\theta) = 1 - \frac{1}{\theta}$. It is more convenient to work with the corresponding values of Kendall's τ , since parameters of different families may live on different domains and may possess different interpretations. Further, we transform Kendall's τ , which is restricted to the interval $(-1, 1)$, to the unconstrained scale using the Fisher's Z transformation $F_Z(x) = \frac{1}{2} \log\left(\frac{1+x}{1-x}\right)$. The transformed Kendall's τ at time t is denoted by s_t and its dynamics is modeled by an AR(1) process. More precisely, the AR(1) process with states s_0, \dots, s_T , mean $\mu \in \mathbb{R}$, persistence parameter $\phi \in (-1, 1)$ and standard deviation parameter $\sigma \in (0, \infty)$ satisfies

$$s_t = \mu + \phi(s_{t-1} - \mu) + \sigma\eta_t \quad (1)$$

for $t = 1, \dots, T$. Here we assume that the errors are $\eta_t \sim N(0, 1)$ iid and the initial distribution is $s_0 | \mu, \phi, \sigma \sim N\left(\mu, \frac{\sigma^2}{1-\phi^2}\right)$. The state s_t is mapped to the parameter θ_t^m of the copula family m at time t as follows

$$\theta_t^m = g_m^{-1}(F_Z^{-1}(s_t)). \quad (2)$$

We now study the following model for T bivariate random vectors $(U_{t1}, U_{t2})_{t=1, \dots, T} \in [0, 1]^{T \times 2}$

$$(U_{t1}, U_{t2}) | m, s_t \sim c^m(u_{t1}, u_{t2}; \theta_t^m) \text{ independently,} \quad (3)$$

for $t = 1, \dots, T$. Here $c^m(\cdot, \cdot; \theta_t^m)$ is the bivariate density of copula family m with parameter θ_t^m as specified in (1) and (2). The copula parameter θ_t^m is a function of the model indicator m and

the state s_t . The state s_t has the same interpretation for different copula families as the Fisher's Z transform of the corresponding Kendall's τ value. This allows us to share the parameters $\mathbf{s}_{0:T} = (s_0, \dots, s_T), \mu, \phi, \sigma$ among different copula families, which keeps the parameter space smaller and simplifies estimation. More details about parameter sharing are given in Appendix A.

A Bayesian model specification is complete by introducing priors for the model parameters. Let $\varphi(\cdot|\mu_{Normal}, \sigma_{Normal}^2)$ denote the density of the univariate normal distribution with mean μ_{Normal} and variance σ_{Normal}^2 . This allows us to express the prior for $\mathbf{s}_{0:T}$ conditional on μ, ϕ and σ , implied by the AR(1) process in (1), as

$$\pi(\mathbf{s}_{0:T}|\mu, \phi, \sigma) = \varphi(s_0|\mu, \sigma^2(1 - \phi^2)^{-1}) \prod_{t=1}^T \varphi(s_t|\mu + \phi(s_{t-1} - \mu), \sigma^2). \quad (4)$$

Further, we assume that

$$\mu \sim N(0, 100), \quad \frac{\phi + 1}{2} \sim Beta(5, 1.5), \quad \sigma^2 \sim Gamma\left(\frac{1}{2}, \frac{1}{2}\right) \quad (5)$$

as in Kreuzer and Czado (2019). These are the same priors that Kastner (2016) recommend for the latent AR(1) process of the stochastic volatility model. For μ we utilize a vague prior. But for ϕ we use a rather informative prior, which gives more prior probability to larger values within $(-1, 1)$. Thus we favor trajectories of Kendall's τ , where there is positive dependence among two subsequent time points ($\text{cor}(s_t, s_{t-1}) = \phi$). The Gamma prior for the variance parameter σ^2 is different to the frequently used inverse Gamma prior. The Gamma prior has more mass close to zero compared to the inverse Gamma prior. So we give more probability to smoother trajectories that do not oscillate a lot. For $m \in \mathcal{M}$ we assume a discrete uniform prior, i.e.

$$\pi(m) = \frac{1}{|\mathcal{M}|}. \quad (6)$$

We further assume prior independence among the parameters μ, ϕ, σ and m .

Bayesian inference

Model selection procedures often have to deal with model specific parameters. They might have different dimensions. A popular approach in this context is reversible jump MCMC (see Green (1995)), which requires dimension matching. Min and Czado (2011) and Gruber and Czado (2015) use reversible jump MCMC for selection among vine copula models, while Tan et al (2019) apply it for model selection among static single factor copula models. The way we constructed our model, dimension matching is not needed, since we share the parameters $s_0, \dots, s_T, \mu, \phi$ and σ among different models. This allows us to use an efficient Gibbs approach as outlined in the following.

Here the likelihood given the data $U = (u_{t1}, u_{t2})_{t=1, \dots, T}$ can be expressed as

$$\ell(\mu, \phi, \sigma, \mathbf{s}_{0:T}, m|U) = \prod_{t=1}^T \ell_t(s_t, m|u_{t1}, u_{t2}) = \prod_{t=1}^T c^m(u_{t1}, u_{t2}; \theta_t^m). \quad (7)$$

The quantity $\ell_t(s_t, m|u_{t1}, u_{t2}) = c^m(u_{t1}, u_{t2}; \theta_t^m)$ is the contribution to the likelihood at time t and $\theta_t^m = g_m^{-1}(F_Z^{-1}(s_t))$.

We now employ a Gibbs sampler for parameter estimation. The indicator m is sampled from its full conditional, given by

$$\begin{aligned} P(m|U, \mu, \phi, \sigma, \mathbf{s}_{0:T}) &= \frac{f(U|\mu, \phi, \sigma, \mathbf{s}_{0:T}, m)\pi(\mu, \phi, \sigma, \mathbf{s}_{0:T})\pi(m)}{\sum_{m' \in \mathcal{M}} f(U|\mu, \phi, \sigma, \mathbf{s}_{0:T}, m')\pi(\mu, \phi, \sigma, \mathbf{s}_{0:T})\pi(m')} \\ &= \frac{\ell(\mu, \phi, \sigma, \mathbf{s}_{0:T}, m|U)}{\sum_{m' \in \mathcal{M}} \ell(\mu, \phi, \sigma, \mathbf{s}_{0:T}, m'|U)} \\ &= \frac{\prod_{t=1}^T c^m(u_{t1}, u_{t2}; \theta_t^m)}{\sum_{m' \in \mathcal{M}} \prod_{t=1}^T c^{m'}(u_{t1}, u_{t2}; \theta_t^{m'})}, \end{aligned} \quad (8)$$

where $\pi(\mu, \phi, \sigma, \mathbf{s}_{0:T}) = \pi(\mathbf{s}_{0:T}|\mu, \phi, \sigma)\pi(\mu)\pi(\phi)\pi(\sigma)$ and $\pi(\cdot)$ as specified in (4), (5) and (6). Here the presence of the shared parameters simplifies the updates of the copula family indicator m considerably. To sample $\mu, \phi, \sigma, \mathbf{s}_{0:T}$ conditioned on the indicator m , we use the same approach as described in Kreuzer and Czado (2019) for multivariate state space models with a univariate autoregressive state equation. This sampler relies on an interweaving strategy (Yu and Meng (2011)), elliptical slice sampling (Murray et al (2010)) and adaptive Metropolis-Hastings updates (Garthwaite et al (2016)).

2.2 Static bivariate copulas

Model specification

To allow for static (time-constant) copulas we consider a static state $s \in \mathbb{R}$ which is mapped to the copula parameter similar to (2), i.e.

$$\theta^m = g_m^{-1}(F_Z^{-1}(s)) \quad (9)$$

for a copula family $m \in \mathcal{M}$. We assume that T bivariate random vectors $(U_{t1}, U_{t2})_{t=1, \dots, T}$ are generated as follows

$$(U_{t1}, U_{t2})|m, s \sim c^m(\cdot, \cdot; \theta^m), \text{ independently.} \quad (10)$$

The prior for the state parameter s is chosen such that the corresponding Kendall's τ is uniformly distributed on $(-1, 1)$. The prior for m is chosen as above in (6). The priors reflect the fact that we do not have any prior information about the parameters.

Bayesian inference

The parameters of this reduced model are also estimated utilizing a Gibbs sampling approach. Here we sample $m|U, s$ directly from its full conditional, which can be derived similar to (8). The parameter s is updated conditional on (U, m) with random walk Metropolis-Hasting with Gaussian proposal and adaptive proposal variance as in Garthwaite et al (2016).

2.3 Model selection among dynamic and static pair copulas using the widely applicable information criteria (WAIC)

One might be interested in how the dynamic copula model compares to the static copula model and to the independence model. We refer to these model classes as dynamic, static and zero dependence, respectively. To also allow for bivariate static and independence copulas is especially important, if the bivariate dynamic copula is used as a building block for vine copula models in higher dimensions. For vine copula models, independence copulas might be useful in higher trees to avoid overfitting.

Bayes factors or the commonly used information criteria AIC and BIC are here not tractable choices for selecting the type of dependence (dynamic, static or zero), since their evaluation would require high-dimensional integration. Instead, we rely on the widely applicable information criteria (WAIC) introduced by Watanabe (2010). For the proposed dynamic copula model it is given by

$$\text{WAIC} = -2 \left(\sum_{t=1}^T \ln(\mathbb{E}(\ell_t(s_t, m|u_{t1}, u_{t2}))) - \sum_{t=1}^T \text{Var}(\ln(\ell_t(s_t, m|u_{t1}, u_{t2}))) \right) \quad (11)$$

with ℓ_t as in (7). The expectation and variance are taken with respect to $P(s_t, m)$, the probability measure of the posterior distribution of s_t and m , i.e.

$$\mathbb{E}(\ell_t(s_t, m|u_{t1}, u_{t2})) = \int_{\mathbb{R} \times \mathcal{M}} \ell_t(s_t, m|u_{t1}, u_{t2}) dP(s_t, m)$$

and

$$\text{Var}(\ln(\ell_t(s_t, m|u_{t1}, u_{t2}))) = \int_{\mathbb{R} \times \mathcal{M}} (\ln(\ell_t(s_t, m|u_{t1}, u_{t2})))^2 dP(s_t, m) - (\mathbb{E}(\ln(\ell_t(s_t, m|u_{t1}, u_{t2}))))^2.$$

The WAIC can be seen as a Bayesian version of the AIC, where $\sum_{t=1}^T \text{Var}(\ln(\ell_t(s_t, m|u_{t1}, u_{t2})))$ is used as a penalty instead of the number of parameters.

For R observed quantities $(x^r)_{r=1, \dots, R}$ we denote by $\hat{E}((x^r)_{r=1, \dots, R}) = \frac{1}{R} \sum_{r=1}^R x^r$ the sample mean and by $\widehat{\text{Var}}((x^r)_{r=1, \dots, R}) = \frac{1}{R-1} \sum_{r=1}^R (x^r - \hat{E}((x^r)_{r=1, \dots, R}))^2$ the sample variance. Following Vehtari et al (2017), the WAIC can be estimated from R samples of the posterior distribution $(s_t^1, m^1), \dots, (s_t^R, m^R)$ with $t = 1, \dots, T$, by

$$\widehat{\text{WAIC}} = -2 \left(\sum_{t=1}^T \ln(\hat{E}((\ell_t^r)_{r=1, \dots, R})) - \sum_{t=1}^T \widehat{\text{Var}}((\ln(\ell_t^r))_{r=1, \dots, R}) \right), \quad (12)$$

where $\ell_t^r = \ell_t(s_t^r, m^r|u_{t1}, u_{t2})$. By setting

$$\widehat{\text{WAIC}}_t = -2 \left(\ln(\hat{E}((\ell_t^r)_{r=1, \dots, R})) - \widehat{\text{Var}}((\ln(\ell_t^r))_{r=1, \dots, R}) \right),$$

we can express $\widehat{\text{WAIC}}$ as $\widehat{\text{WAIC}} = \sum_{t=1}^T \widehat{\text{WAIC}}_t$.

To compare between two models with estimated WAIC values $\widehat{\text{WAIC}}^A$ and $\widehat{\text{WAIC}}^B$, Vehtari et al (2017) suggest to consider the difference in the estimated WAIC given by

$$\widehat{\text{WAIC}}^A - \widehat{\text{WAIC}}^B = \sum_{t=1}^T (\widehat{\text{WAIC}}_t^A - \widehat{\text{WAIC}}_t^B) \quad (13)$$

with corresponding standard error estimate

$$\widehat{se}(\widehat{\text{WAIC}}^A - \widehat{\text{WAIC}}^B) = \sqrt{T \cdot \widehat{\text{VAR}}((\widehat{\text{WAIC}}_t^A - \widehat{\text{WAIC}}_t^B)_{t=1, \dots, T})}. \quad (14)$$

To estimate the standard error, Vehtari et al (2017) assume independence among the components $\widehat{\text{WAIC}}_t, t = 1, \dots, T$. For the static copula model we use $\ell_t(s, m|u_{t1}, u_{t2}) = c^m(u_{t1}, u_{t2}; \theta^m)$. Further, WAIC is zero for the independence copula. In our framework the dynamic model is considered to be more complex than the static model. Similarly, the static and the dynamic model are considered to be more complex than the independence model. Here, we select the more complex model if its estimated WAIC is at least 2 standard errors smaller than the WAIC of the other model.




Figure 1: Model selection procedure for bivariate copula models.

Alternatively we could have incorporated the independence, the static and the dynamic copula within one sampler. In this case we would need to move between models with different dimensions by employing e.g. reversible jump MCMC. But proposing moves efficiently from the

parameter free independence copula or from the static copula to dynamic copulas with more than T parameters is difficult and chains might take very long to converge. Another alternative is to select among all models, including family choices, with WAIC. But this would require to estimate a dynamic bivariate copula model for each copula family, which is computationally expensive. So we propose to use the Gibbs sampler to move between models with the same dimension, where the parameters can be shared among the different models. The WAIC is used to select between models with different parameter dimensions where parameters can not be shared. The whole selection procedure for the bivariate copula models is visualized in Figure 1.

2.4 Simulation study

We conduct a small simulation study to investigate the ability of WAIC to select the type of dependence. We consider five scenarios specified in Table 1. In Scenarios 1 and 2 we simulate from the dynamic model specified in (3), in Scenarios 3 and 4 from the static one specified in (10) and in Scenario 5 from the bivariate independence copula. For each scenario we simulate $T = 1000$ observations. Based on these observations we fit the dynamic and the static model and select among different models with WAIC as explained in Section 2.3. We consider the following set for copula family selection $\mathcal{M} = \{\text{Independence, Gaussian, Student } t(\text{df}=4), \text{eClayton, eGumbel}\}$. Here eClayton and eGumbel are extended Clayton and Gumbel copulas that also allow for negative Kendall's τ values as in Kreuzer and Czado (2019). We repeat the simulations for each scenario 100 times. From Table 2 we see that in each scenario, the correct type of dependence was selected at least in 84 out of 100 cases. The correct family was selected in at least 98 out of 100 cases according to Table 2. We conclude that our Bayesian family selection procedure performs well and that WAIC can be utilized to select the appropriate type of dependence.

	Dynamic		Static		Independence
	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5
m	Gaussian	eClayton	Student t(df=4)	eGumbel	Independence
μ	0.4	0.4			
ϕ	0.95	0.8			
σ	0.1	0.2			
s			1	0.4	

Table 1: Parameter specification for the simulation study for bivariate copula models.

	Scenario	Dynamic		Static		Independence
		1	2	3	4	5
Copula family	Independence	0	0	0	0	100
	Gaussian	100	2	1	0	0
	Student t(df=4)	0	0	99	0	0
	eClayton	0	98	0	0	0
	eGumbel	0	0	0	100	0
Type of dependence	Dynamic	100	100	16	0	1
	Static	0	0	84	100	0
	Zero	0	0	0	0	100

Table 2: For each scenario of the simulation study, we show how often the different copula families were selected and how often each type of dependence (dynamic, static and zero) was selected. The selected copula family is the marginal posterior mode estimate of m , i.e. the family that occurs most frequently among the posterior samples for m . The true copula family and the true type of dependence, which we used for simulation, is marked in bold.

3 Dynamic vine copulas

3.1 Model specification

Vine copulas (Bedford and Cooke (2001), Aas et al (2009), Joe (2014), Czado (2019)) are a popular class in dependence modeling. They allow for great flexibility by constructing a density of arbitrary dimension from two-dimensional densities.

For this construction, vine copulas are represented as graphical models. A regular vine (R-vine) tree sequence is a sequence of trees $\mathcal{V} = T_1, \dots, T_{d-1}$ satisfying the following conditions

- each tree is connected,
- T_1 is a tree with nodes $N_1 = \{1, \dots, d\}$ and set of edges E_1 ,
- T_j with $j \geq 2$ is a tree with nodes $N_j = E_{j-1}$ and edges E_j and
- for $j = 2, \dots, d-1$ and $\{a, b\} \in E_j$, it holds that a and b as edges in tree T_{j-1} share a common node (proximity condition).

Furthermore, the complete union of an edge $e \in E_i$ is defined as

$$\mathbf{A}_e := \{j \in N_1 | \exists e_1 \in E_1, \dots, e_{i-1} \in E_{i-1}; j \in e_1 \in \dots \in e_{i-1} \in e\}.$$

The conditioning set of edge $e = \{a, b\}$ is obtained as

$$\mathbf{D}_e := \mathbf{A}_a \cap \mathbf{A}_b,$$

and the conditioned sets are given by

$$a_e := \mathbf{A}_a \setminus \mathbf{D}_e, \quad b_e := \mathbf{A}_b \setminus \mathbf{D}_e.$$

Bivariate copulas of conditional distributions of the marginally uniformly distributed random vector (U_1, \dots, U_d) can be identified with the conditioning and the conditioned sets. For an edge e we denote by $c_{a_e, b_e; \mathbf{D}_e}$ the density of the bivariate copula of $(U_{a_e} U_{b_e}) | \mathbf{U}_{\mathbf{D}_e} = \mathbf{u}_{\mathbf{D}_e}$, where $\mathbf{u}_{\mathbf{F}} = (u_i)_{i \in \mathbf{F}}$ for a set \mathbf{F} . Many researchers assume that $c_{a_e, b_e; \mathbf{D}_e}$ does not depend on $\mathbf{u}_{\mathbf{D}_e}$, which is called the simplifying assumption (Haff et al (2010), Stoeber et al (2013)). This assumption allows for sequential estimation and selection of vine copula models (Brechmann and Czado (2013), Dissmann et al (2013)).

Based on these graphical definitions, Bedford and Cooke (2001) build a d -dimensional vine copula model with joint density

$$c(u_1, \dots, u_d) = \prod_{i=1}^{d-1} \prod_{e \in E_i} c_{a_e, b_e; \mathbf{D}_e}(u_{a_e | \mathbf{D}_e}, u_{b_e | \mathbf{D}_e}). \quad (15)$$

It is a simplified vine copula model, since $c_{a_e, b_e; \mathbf{D}_e}$ does not depend on the conditioning value $\mathbf{u}_{\mathbf{D}_e}$. Here, $u_{a_e | \mathbf{D}_e}$ and $u_{b_e | \mathbf{D}_e}$ are called pseudo data. They are obtained as $u_{a_e | \mathbf{D}_e} = C_{a_e | \mathbf{D}_e}(u_{a_e} | \mathbf{u}_{\mathbf{D}_e})$ and $u_{b_e | \mathbf{D}_e} = C_{b_e | \mathbf{D}_e}(u_{b_e} | \mathbf{u}_{\mathbf{D}_e})$, where $C_{a_e | \mathbf{D}_e}$ and $C_{b_e | \mathbf{D}_e}$ are the conditional distribution functions of $U_{a_e} | \mathbf{U}_{\mathbf{D}_e} = \mathbf{u}_{\mathbf{D}_e}$ and $U_{b_e} | \mathbf{U}_{\mathbf{D}_e} = \mathbf{u}_{\mathbf{D}_e}$, respectively. In the first tree \mathbf{D}_e is the empty set and the pseudo data of the first tree is just u_1, \dots, u_d . Note that the class of simplified vine copulas is broad, including multivariate Gaussian and Student t copulas (Joe (2014), Chapter 3).

To evaluate the conditional distribution functions and to obtain the corresponding pseudo data for all trees, the h functions (Aas et al (2009)) for an edge e are defined as

$$\begin{aligned} h_{a_e | b_e; \mathbf{D}_e}(u_1 | u_2) &= \frac{d}{du_2} C_{a_e, b_e; \mathbf{D}_e}(u_1, u_2) \\ h_{b_e | a_e; \mathbf{D}_e}(u_2 | u_1) &= \frac{d}{du_1} C_{a_e, b_e; \mathbf{D}_e}(u_1, u_2). \end{aligned} \quad (16)$$

If the copula $C_{a_e, b_e; \mathbf{D}_e}$ depends on a set of parameters $\boldsymbol{\delta}$ we write $h_{a_e | b_e; \mathbf{D}_e}(u_1 | u_2; \boldsymbol{\delta})$ and $h_{b_e | a_e; \mathbf{D}_e}(u_2 | u_1; \boldsymbol{\delta})$. Based on pseudo data $u_{a_e | \mathbf{D}_e}$ and $u_{b_e | \mathbf{D}_e}$, we obtain pseudo data for the next tree as

$$\begin{aligned} u_{a_e | b_e \cup \mathbf{D}_e} &= h_{a_e | b_e; \mathbf{D}_e}(u_{a_e | \mathbf{D}_e} | u_{b_e | \mathbf{D}_e}), \\ u_{b_e | a_e \cup \mathbf{D}_e} &= h_{b_e | a_e; \mathbf{D}_e}(u_{b_e | \mathbf{D}_e} | u_{a_e | \mathbf{D}_e}). \end{aligned} \quad (17)$$

So for the calculation of pseudo data only bivariate copulas from lower trees are involved.

Figure 2 visualizes the first three trees of a six-dimensional vine copula. Assuming that there are only independence copulas in trees higher than Tree 3, the associated density is given by

$$\begin{aligned} c(u_1, \dots, u_6) = & c_{12}(u_1, u_2) \cdot c_{26}(u_2, u_6) \cdot c_{36}(u_3, u_6) \cdot c_{46}(u_4, u_6) \cdot c_{56}(u_5, u_6) \\ & \cdot c_{45;6}(u_{4|6}, u_{5|6}) \cdot c_{35;6}(u_{3|6}, u_{5|6}) \cdot c_{25;6}(u_{2|6}, u_{5|6}) \cdot c_{16;2}(u_{1|2}, u_{6|2}) \\ & \cdot c_{3,4;5,6}(u_{3|5,6}, u_{4|5,6}) \cdot c_{2,4;5,6}(u_{2|5,6}, u_{4|5,6}) \cdot c_{1,5;2,6}(u_{1|2,6}, u_{5|2,6}). \end{aligned} \quad (18)$$

The pseudo data of the vine copula model can be determined as outlined in (17). E.g. $u_{4|6} = h_{4|6}(u_4|u_6)$, $u_{5|6} = h_{5|6}(u_5|u_6)$ and $u_{4|56} = h_{4|5;6}(u_{4|6}|u_{5|6}) = h_{4|5;6}(h_{4|6}(u_4|u_6)|h_{5|6}(u_5|u_6))$.

Such vine copulas, where pair copulas above a certain tree level are set to the independence copula, are called truncated (Brechmann et al (2012)). With truncation we can achieve different levels of sparsity.




Figure 2: Tree structure of a vine copula model

Within the vine copula framework, we only need to specify bivariate copulas and can scale to arbitrary dimensions. Here, we replace each bivariate copula in (15) by the dynamic bivariate copula model specified in (3), the static bivariate copula model specified in (10) or the independence copula. We denote by E_i^{dyn} , E_i^{static} and E_i^{ind} the set of edges in the i -th tree where the corresponding pair copulas are dynamic, static or independence copulas, respectively. For each edge $e \in E_i^{dyn}$, $i = 1, \dots, d-1$ we have a corresponding family indicator m_e and a corresponding latent AR(1) process given by

$$s_{t,e} = \mu_e + \phi_e(s_{t-1,e} - \mu_e) + \sigma_e \eta_{t,e}, \eta_{t,e} \sim N(0, 1) \text{ iid}, \quad (19)$$

with μ_e , ϕ_e , σ_e and $s_{0,e}$ as in Section 2.1. For each edge e in E_i^{static} we have a corresponding family indicator m_e and a state s_e as in Section 2.2. The states $s_{t,e}$ and s_e are mapped to the copula parameter as in (2) and (9), i.e.

$$\theta_{t,e}^{m_e} = g_{m_e}^{-1}(F_Z^{-1}(s_{t,e})) \text{ and } \theta_e^{m_e} = g_{m_e}^{-1}(F_Z^{-1}(s_e)), \text{ respectively.} \quad (20)$$

This yields the following parameter set for a d -dimensional dynamic vine copula model

$$\theta^V = \{\mu_e, \phi_e, \sigma_e, s_{0,e}, \dots, s_{T,e}, m_e | e \in E_i^{dyn}, i = 1, \dots, d-1\} \cup \{s_e, m_e | e \in E_i^{static}, i = 1, \dots, d-1\}$$

Within the dynamic vine copula model, we assume that T random vectors of dimension d , $(U_{t1}, \dots, U_{td})_{t=1, \dots, T}$ are generated as follows

$$\begin{aligned}
(U_{t1}, \dots, U_{td}) | \boldsymbol{\theta}^V &\sim \prod_{i=1}^{d-1} \prod_{e \in E_i^{dyn}} c_{a_e, b_e; D_e}^{m_e}(u_{t, a_e} | D_e, u_{t, b_e} | D_e; \boldsymbol{\theta}_{t, e}^{m_e}). \\
&\cdot \prod_{i=1}^{d-1} \prod_{e \in E_i^{static}} c_{a_e, b_e; D_e}^{m_e}(u_{t, a_e} | D_e, u_{t, b_e} | D_e; \boldsymbol{\theta}_e^{m_e}), \text{ independently}
\end{aligned} \tag{21}$$

for $t = 1, \dots, T$. Further, we assume that parameters for different edges are a priori independent and we use the same priors as specified in Sections 2.1 and 2.2 for the parameters of dynamic and static pair copulas, respectively. Here, the conditioned and conditioning sets a_e, b_e and D_e do not depend on the time, i.e. the tree structure does not change over time.

3.2 Sequential estimation

Since there exist $\frac{d!}{2} \cdot 2^{\binom{d-2}{2}}$ different regular vine tree structures in d dimensions (Morales-Nápoles (2010)), model selection is complex and it is not possible to take all possible structures into account as the dimension grows. Gruber and Czado (2018) estimate all trees and parameters jointly using a Bayesian approach, but their procedure is only suitable in lower dimensions and requires substantial computations. Earlier, in a frequentist setup, Dissmann et al (2013) proposed a sequential selection and estimation approach for static copula parameters. This sequential approach makes model selection feasible in higher dimensions and therefore this idea is often used for vine copula based models. Gruber and Czado (2015) employ a Bayesian approach, where they estimate parameters of static vine copula models tree by tree and fix parameters at point estimates before proceeding to the next tree. Vatter and Nagler (2018) sequentially estimate the parameters of a vine copula model, where the copula parameters are modeled by generalized additive models. We propose a Bayesian procedure, where the copula parameters do not need to be collapsed to point estimates before proceeding to the next tree. The procedure is based on an approximation of the posterior density inspired by the frequentist sequential approach of Dissmann et al (2013).

To simplify notation we denote by $\boldsymbol{\theta}_{T_i}$ the parameters of a dynamic vine copula corresponding to tree T_i , i.e.

$$\boldsymbol{\theta}_{T_i} = \{\mu_e, \phi_e, \sigma_e, s_{0,e}, \dots, s_{T,e}, m_e, e \in E_i^{dyn}\} \cup \{s_e, m_e, e \in E_i^{static}\}. \tag{22}$$

Further, we define by $U = (u_{tj})_{t=1, \dots, T, j=1, \dots, d}$ the data matrix and the likelihood contribution of the i -th tree is given by

$$\begin{aligned}
\ell_i(\boldsymbol{\theta}_{T_1}, \dots, \boldsymbol{\theta}_{T_i} | U) &= \prod_{e \in E_i^{dyn}} c_{a_e, b_e; D_e}^{m_e}(u_{t, a_e} | D_e, u_{t, b_e} | D_e; \boldsymbol{\theta}_{t, e}^{m_e}) \\
&\cdot \prod_{e \in E_i^{static}} c_{a_e, b_e; D_e}^{m_e}(u_{t, a_e} | D_e, u_{t, b_e} | D_e; \boldsymbol{\theta}_e^{m_e}).
\end{aligned} \tag{23}$$

The complete likelihood is obtained as

$$\ell(\boldsymbol{\theta}_{T_1}, \dots, \boldsymbol{\theta}_{T_{d-1}} | U) = \prod_{i=1}^{d-1} \ell_i(\boldsymbol{\theta}_{T_1}, \dots, \boldsymbol{\theta}_{T_i} | U), \tag{24}$$

and the posterior density is proportional to

$$f(\boldsymbol{\theta}_{T_1}, \dots, \boldsymbol{\theta}_{T_{d-1}} | U) \propto \prod_{i=1}^{d-1} \ell_i(\boldsymbol{\theta}_{T_1}, \dots, \boldsymbol{\theta}_{T_i} | U) \pi(\boldsymbol{\theta}_{T_i}). \tag{25}$$

Note that the posterior density is a joint density of continuous and discrete parameters. For discrete parameters $\boldsymbol{\delta}^{disc}$ and continuous parameters $\boldsymbol{\delta}^{cont}$ the joint density is obtained as

$$f(\boldsymbol{\delta}^{cont}, \boldsymbol{\delta}^{disc}) = f(\boldsymbol{\delta}^{cont} | \boldsymbol{\delta}^{disc}) f(\boldsymbol{\delta}^{disc}),$$

where $f(\boldsymbol{\delta}^{cont} | \boldsymbol{\delta}^{disc})$ is a joint probability density function and $f(\boldsymbol{\delta}^{disc})$ is a joint probability mass function.

Four-dimensional illustration for a model with static pair copulas and known tree structure

To illustrate our idea in four dimensions, we consider only static pair copulas and the tree structure to be known. The tree structure contains the pair copulas as specified in (26), (27) and (28). We assume, that we observe data $U = (u_{tj})_{t=1, \dots, T, j=1, \dots, 4}$. The contributions to the likelihood corresponding to trees 1, 2 and 3 are given by

$$\ell_1(\boldsymbol{\theta}_{\mathcal{T}_1}|U) = \prod_{t=1}^T c_{13}^{m_{13}}(u_{t1}, u_{t3}; \theta_{13}^{m_{13}}) c_{23}^{m_{23}}(u_{t2}, u_{t3}; \theta_{23}^{m_{23}}) c_{34}^{m_{34}}(u_{t3}, u_{t4}; \theta_{34}^{m_{34}}), \quad (26)$$

$$\ell_2(\boldsymbol{\theta}_{\mathcal{T}_1}, \boldsymbol{\theta}_{\mathcal{T}_2}|U) = \prod_{t=1}^T c_{12;3}^{m_{12;3}}(u_{t1|3}, u_{t2|3}; \theta_{12;3}^{m_{12;3}}) c_{24;3}^{m_{24;3}}(u_{t2|3}, u_{t4|3}; \theta_{24;3}^{m_{24;3}}), \quad (27)$$

$$\ell_3(\boldsymbol{\theta}_{\mathcal{T}_1}, \boldsymbol{\theta}_{\mathcal{T}_2}, \boldsymbol{\theta}_{\mathcal{T}_3}|U) = \prod_{t=1}^T c_{14;23}^{m_{14;23}}(u_{t1|23}, u_{t4|23}; \theta_{14;23}^{m_{14;23}}). \quad (28)$$

In a frequentist sequential procedure the parameters of the first tree are estimated by considering the part of the likelihood corresponding to the first tree as given in (26), ignoring the likelihood contributions of higher trees (27) and (28) to the parameters of the first tree. This allows to maximize $\prod_{t=1}^T c_{13}^{m_{13}}(u_{t1}, u_{t3}; \theta_{13}^{m_{13}})$, $\prod_{t=1}^T c_{23}^{m_{23}}(u_{t2}, u_{t3}; \theta_{23}^{m_{23}})$ and $\prod_{t=1}^T c_{34}^{m_{34}}(u_{t3}, u_{t4}; \theta_{34}^{m_{34}})$ independently.

In the Bayesian setup the marginal posterior density of the parameters corresponding to the first tree is obtained by integrating out parameters of higher trees, i.e.

$$f(\boldsymbol{\theta}_{\mathcal{T}_1}|U) \propto \ell_1(\boldsymbol{\theta}_{\mathcal{T}_1}|U) \pi(\boldsymbol{\theta}_{\mathcal{T}_1}) \int_{\text{domain}(\boldsymbol{\theta}_{\mathcal{T}_2}, \boldsymbol{\theta}_{\mathcal{T}_3})} \prod_{i=2}^3 \ell_i(\boldsymbol{\theta}_{\mathcal{T}_1}, \dots, \boldsymbol{\theta}_{\mathcal{T}_i}|U) \pi(\boldsymbol{\theta}_{\mathcal{T}_i}) d\boldsymbol{\theta}_{\mathcal{T}_2} d\boldsymbol{\theta}_{\mathcal{T}_3}, \quad (29)$$

where $\text{domain}(\boldsymbol{\theta}_{\mathcal{T}_2}, \boldsymbol{\theta}_{\mathcal{T}_3})$ is the domain of the parameters $\boldsymbol{\theta}_{\mathcal{T}_2}, \boldsymbol{\theta}_{\mathcal{T}_3}$ (For the family indicator m_e the integral is replaced by a sum). While in this illustrative example it might be possible to work with the marginal posterior (29), its complexity grows fast if we consider more dimensions or allow for dynamic copulas. For example, if the second and third tree were modeled with dynamic bivariate copulas, we would need to integrate out several thousand parameters for $T = 1000$. To reduce complexity, we approximate $f(\boldsymbol{\theta}_{\mathcal{T}_1}|U)$. We make use of the following notation

$$g(\boldsymbol{\delta}) \approx h(\boldsymbol{\delta})$$

to denote that a density g is approximately proportional to a non negative and integrable function h . This means that the density g is approximated by the density $h^{\text{normalized}}$ given by $h^{\text{normalized}}(\boldsymbol{\delta}) = h(\boldsymbol{\delta}) \left(\int_{\text{domain}(\boldsymbol{\delta})} h(\boldsymbol{\delta}) d\boldsymbol{\delta} \right)^{-1}$.

Following the idea of a frequentist sequential estimation we approximate the marginal posterior $\boldsymbol{\theta}_{\mathcal{T}_1}|U$ by considering only the part of the likelihood corresponding to the first tree, i.e.

$$\begin{aligned} f(\boldsymbol{\theta}_{\mathcal{T}_1}|U) &\approx \ell_1(\boldsymbol{\theta}_{\mathcal{T}_1}|U) \pi(\boldsymbol{\theta}_{\mathcal{T}_1}) = \left(\prod_{t=1}^T c_{13}^{m_{13}}(u_{t1}, u_{t3}; \theta_{13}^{m_{13}}) \right) \pi(s_{13}) \pi(m_{13}) \\ &\quad \left(\prod_{t=1}^T c_{23}^{m_{23}}(u_{t2}, u_{t3}; \theta_{23}^{m_{23}}) \right) \pi(s_{23}) \pi(m_{23}) \left(\prod_{t=1}^T c_{34}^{m_{34}}(u_{t3}, u_{t4}; \theta_{34}^{m_{34}}) \right) \pi(s_{34}) \pi(m_{34}). \end{aligned} \quad (30)$$

This approximation simplifies sampling enormously. We do not only get rid of the integral, but in addition, the parameters corresponding to different edges are independent. To obtain samples from the posterior approximation in (30) we can sample parameters of different static bivariate copula models independently by utilizing the algorithm of Section 2.2. In particular the parameters m_{13}, s_{13} are sampled from a static bivariate copula model with corresponding posterior density proportional to $\left(\prod_{t=1}^T c_{13}^{m_{13}}(u_{t1}, u_{t3}; \theta_{13}^{m_{13}}) \right) \pi(s_{13}) \pi(m_{13})$. Approximations of the posterior distribution that induce independence among parameters are commonly used in

variational Bayesian approaches (Wainwright et al (2008)). For example in mean field variational inference it is assumed that all parameters are independent in the posterior distribution. Our assumptions are less restrictive, since we do not assume that parameters corresponding to one pair copula are independent. Further, these parameters are updated jointly.

When estimating parameters of higher trees in a sequential frequentist procedure, we condition on estimates from lower trees. Therefore we consider now the following density

$$f(\boldsymbol{\theta}_{T_2}|\boldsymbol{\theta}_{T_1}, U) \propto \ell_2(\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}|U)\pi(\boldsymbol{\theta}_{T_2}) \int_{\text{domain}(\boldsymbol{\theta}_{T_3})} \ell_3(\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}, \boldsymbol{\theta}_{T_3}|U)\pi(\boldsymbol{\theta}_{T_3})d\boldsymbol{\theta}_{T_3}. \quad (31)$$

We utilize a similar approximation as in (30) and obtain

$$f(\boldsymbol{\theta}_{T_2}|\boldsymbol{\theta}_{T_1}, U) \approx \ell_2(\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}|U)\pi(\boldsymbol{\theta}_{T_2}) = \left(\prod_{t=1}^T c_{12;3}^{m_{12;3}}(u_{t1|3}, u_{t2|3}; \theta_{12;3}^{m_{12;3}}) \right) \pi(s_{12;3})\pi(m_{12;3}) \\ \left(\prod_{t=1}^T c_{24;3}^{m_{24;3}}(u_{t2|3}, u_{t4|3}; \theta_{24;3}^{m_{24;3}}) \right) \pi(s_{24;3})\pi(m_{24;3}). \quad (32)$$

The pseudo data $u_{t1|3} = h_{1|3}(u_{t1}|u_{t3}; \theta_{13}^{m_{13}}, m_{13}) = \frac{d}{du_3} C_{13}^{m_{13}}(u_{t1}, u_3; \theta_{13}^{m_{13}}) \Big|_{u_3=u_{t3}}$, $u_{t2|3} = h_{2|3}(u_{t2}|u_{t3}; \theta_{23}^{m_{23}}, m_{23})$ and $u_{t4|3} = h_{4|3}(u_{t4}|u_{t3}, \theta_{34}^{m_{34}}, m_{34})$, $t = 1, \dots, T$ only depend on parameters of the first tree, on which we condition on. Further, the posterior density factorizes as in (30) and we can sample parameters corresponding to different edges independently. In particular $s_{12;3}, m_{12;3}$ are sampled from a static bivariate copula model with posterior density proportional to $\left(\prod_{t=1}^T c_{12;3}^{m_{12;3}}(u_{t1|3}, u_{t2|3}; \theta_{12;3}^{m_{12;3}}) \right) \pi(s_{12;3})\pi(m_{12;3})$, where $\{u_{t1|3}, u_{t2|3}, t = 1, \dots, T\}$ is interpreted as observed data. For the third tree we obtain

$$f(\boldsymbol{\theta}_{T_3}|\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}, U) \propto \ell_3(\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}, \boldsymbol{\theta}_{T_3}|U)\pi(\boldsymbol{\theta}_{T_3}) \\ = \left(\prod_{t=1}^T c_{14;23}^{m_{14;23}}(u_{t1|23}, u_{t4|23}; \theta_{14;23}^{m_{14;23}}) \right) \pi(s_{14;23})\pi(m_{14;23}). \quad (33)$$

As before, $u_{t1|23}$ and $u_{t4|23}$ only depend on parameters from lower trees, on which we condition on. Interpreting $\{u_{t1|23}, u_{t4|23}, t = 1, \dots, T\}$ as observed data, (33) is the posterior density of a static bivariate copula model as introduced in Section 2.2.

To obtain samples from the posterior density $f(\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}, \boldsymbol{\theta}_{T_3}|U)$, we utilize the approximations above to first sample $\boldsymbol{\theta}_{T_1}$ from $f(\boldsymbol{\theta}_{T_1}|U)$, then $\boldsymbol{\theta}_{T_2}$ from $f(\boldsymbol{\theta}_{T_2}|\boldsymbol{\theta}_{T_1}, U)$ and then $\boldsymbol{\theta}_{T_3}$ from $f(\boldsymbol{\theta}_{T_3}|\boldsymbol{\theta}_{T_1}, \boldsymbol{\theta}_{T_2}, U)$, i.e. we employ a collapsed Gibbs sampler (Liu (1994)). To update the parameters we use the sampling procedure of Section 2.2. This sampler simulates a Markov chain, where subsequent draws are autocorrelated. So, by applying this sampler we obtain a sample of $\boldsymbol{\theta}_{T_i}$ in the r -th iteration, denoted by $\boldsymbol{\theta}_{T_i}^r$, which depends on the previous value $\boldsymbol{\theta}_{T_i}^{r-1}$. While this is not a problem for conventional Gibbs samplers, as in Metropolis-Hastings within Gibbs, this can lead to undesired samples for collapsed Gibbs schemes as shown by Van Dyk and Jiao (2015). Following Van Dyk and Jiao (2015), we can circumvent this problem by running the updates for $\boldsymbol{\theta}_{T_i}$ with starting value $\boldsymbol{\theta}_{T_i}^{r-1}$ for k iterations. We set $\boldsymbol{\theta}_{T_i}^r$ equal to the update obtained in the k -th step. If we choose k large enough, $\boldsymbol{\theta}_{T_i}^r$ will be almost independent of $\boldsymbol{\theta}_{T_i}^{r-1}$. Thus, in total, $R \cdot k$ draws are obtained and R draws are stored for each parameter. We obtain the following procedure

- Set starting values $\boldsymbol{\theta}_{T_1}^0, \boldsymbol{\theta}_{T_2}^0, \boldsymbol{\theta}_{T_3}^0$
- For $r = 1, \dots, R$ do
 - For $i = 1, \dots, 3$ do
 - Use the sampler of Section 2.2 to sample from $f(\boldsymbol{\theta}_{T_1}|U)$ if $i = 1$, from $f(\boldsymbol{\theta}_{T_2}|\boldsymbol{\theta}_{T_1}^r, U)$ if $i = 2$ or from $f(\boldsymbol{\theta}_{T_3}|\boldsymbol{\theta}_{T_1}^r, \boldsymbol{\theta}_{T_2}^r, U)$ if $i = 3$. The sampler is run for k iterations using $\boldsymbol{\theta}_{T_i}^{r-1}$ as starting value. We set $\boldsymbol{\theta}_{T_i}^r$ equal to the sample obtained in the k -th iteration.

- The pseudo data for the next tree is constructed utilizing the h functions defined in (16). For $i = 1$ the pseudo data is determined as $u_{t1|3}^r = h_{1|3}(u_{t1}|u_{t3}; (\theta_{13}^{m_{13}})^r, m_{13}^r)$, $u_{t2|3}^r = h_{2|3}(u_{t2}|u_{t3}; (\theta_{23}^{m_{23}})^r, m_{23}^r)$ and $u_{t4|3}^r = h_{4|3}(u_{t4}|u_{t3}; (\theta_{34}^{m_{34}})^r, m_{34}^r)$, $t = 1, \dots, T$. For $i = 2$ the pseudo data is determined as $u_{t1|23}^r = h_{1|2;3}(u_{t1|3}^r, u_{t2|3}^r; (\theta_{12;3}^{m_{12;3}})^r, m_{12;3}^r)$ and $u_{t4|23}^r = h_{4|2;3}(u_{t4|3}^r, u_{t2|3}^r; (\theta_{24;3}^{m_{24;3}})^r, m_{24;3}^r)$, $t = 1, \dots, T$.

In this procedure, no point estimates of the parameters $\theta_{T_1}, \theta_{T_2}, \theta_{T_3}$ are required. We can further extend the procedure in the following ways.

- a) The loops over i and r can be exchanged. We can first obtain R samples from the first tree, then obtain R samples from the second tree and then all R samples from the third tree. This is visualized in Figure 3. If the tree structure was not known we could select the tree structure of the first tree, then obtain R samples from the parameters of the first tree. Based on these samples, we can construct the pseudo data, which can be used to select the tree structure of the second tree and so on. Based on the (pseudo) data of a certain tree level, the corresponding structure can be selected as a maximum spanning tree. This is similar to the algorithm of Dissmann et al (2013).
- b) The parameters of different edges of a tree are sampled independently by utilizing the sampler of Section 2.2 for the static copula model. If we did not know that Kendall's τ was static we could in addition run the sampler of Section 2.1 for the dynamic bivariate copula model. We can decide between the dynamic, the static and the independence model as outlined in Section 2.3. Here it is important that these decisions for the type of dependence can be made independently for each edge of the tree.




Figure 3: Graphical representation of the proposed sampler without selection of the type of dependence and without structure selection. Here $U \in [0, 1]^{T \times d}$ denotes the data matrix used for fitting the model and U_l^r denotes the pseudo data for tree l obtained from parameter draws of the previous tree (tree $l - 1$) in iteration r .

The general procedure in d dimensions with vine structure selection

Based on the four-dimensional illustration, we now formulate our procedure for a d -dimensional dynamic vine copula as introduced in Section 3.1, incorporating extensions a) and b). The tree

structure and the sets $E_i^{dyn}, E_i^{static}, E_i^{ind}$ are selected sequentially as we move up the trees and are fixed at point estimates. In Gruber and Czado (2018) searching among different structures within a full Bayesian procedure resulted in very long computation times for static copula models. Here it would be even worse, since we deal with more complex dynamic pair copulas. Note that both, the sets $E_i^{dyn}, E_i^{static}, E_i^{ind}$ and the tree structure do not change over time.

We propose the following approach with iterations parameter R , burn-in parameter *burnin* and thinning parameter k for structure selection and parameter estimation. Note that, as mentioned above, $R \cdot k$ draws are obtained in total and R iterations are stored for each parameter.

1. Select the tree structure of tree T_1 : For all edges e that are allowed in the first tree T_1 , i.e. for all pairs (a_e, b_e) with $1 \leq a_e < b_e \leq d$, estimate τ_{a_e, b_e} by the empirical Kendall's τ using $\{u_{t, a_e}, u_{t, b_e}, t = 1, \dots, T\}$. The structure of tree T_1 is selected as the maximum spanning tree among those edges, where the absolute value of empirical Kendall's τ serves as the corresponding weight.
2. (a) For each edge $e \in E_1$ in tree T_1 , with corresponding observations $\{u_{t, a_e}, u_{t, b_e}, t = 1, \dots, T\}$, run the samplers of Sections 2.1 and 2.2 for the bivariate dynamic and static copula models. The samplers are run for $R \cdot k$ iterations and we thin the samples with factor k .
- (b) For each edge $e \in E_1$, we select among the bivariate dynamic, static and the bivariate independence copula model as discussed in Section 2.3.
- (c) For each edge $e \in E_1$, the pseudo data for the next tree is obtained as

$$\begin{aligned} u_{t, a_e | b_e}^r &= h_{a_e | b_e}(u_{t, a_e} | u_{t, b_e}; (\theta_{t, e}^{m_e})^r, m_e^r), \\ u_{t, b_e | a_e}^r &= h_{b_e | a_e}(u_{t, b_e} | u_{t, a_e}; (\theta_{t, e}^{m_e})^r, m_e^r), \end{aligned} \quad (34)$$

for $r = 1, \dots, R, t = 1, \dots, T$, if the dynamic copula was selected for edge e . If the static copula was selected we replace $\theta_{t, e}^{m_e}$ by $\theta_e^{m_e}$ in (34). For the independence copula model we use $u_{t, a_e | b_e}^r = u_{t, a_e}$ and $u_{t, b_e | a_e}^r = u_{t, b_e}$.

3. Set $l = 2$.
4. Select the tree structure of tree T_l : For all edges that are allowed in tree T_l according to the proximity condition, estimate Kendall's τ of edge $e = (a_e, b_e; \mathbf{D}_e)$ denoted by $\tau_{a_e, b_e; \mathbf{D}_e}$ by the empirical Kendall's τ . Therefore we use posterior mode estimates of the pseudo data $\{\hat{u}_{t, a_e | \mathbf{D}_e}, \hat{u}_{t, b_e | \mathbf{D}_e}, t = 1, \dots, T\}$, where $\hat{u}_{t, a_e | \mathbf{D}_e}$ is the mode of the univariate kernel density estimate of $\{u_{t, a_e | \mathbf{D}_e}^r, r = \text{burnin} + 1, \dots, R\}$ and $\hat{u}_{t, b_e | \mathbf{D}_e}$ is obtained similarly. Here the posterior mode pseudo data $\{\hat{u}_{t, a_e | \mathbf{D}_e}, \hat{u}_{t, b_e | \mathbf{D}_e}, t = 1, \dots, T\}$ are treated as an i.i.d. sample for the estimation of $\tau_{a_e, b_e; \mathbf{D}_e}$. The structure of tree T_l is selected as the maximum spanning tree among those edges, where the absolute value of empirical Kendall's τ serves as the corresponding weight.
5. (a) For each edge $e \in E_l$ in tree T_l , with corresponding pseudo data $\{u_{t, a_e | \mathbf{D}_e}^r, u_{t, b_e | \mathbf{D}_e}^r, t = 1, \dots, T, r = 1, \dots, R\}$, obtain R samples (based on a total of $R \cdot k$ MCMC draws) from the bivariate dynamic and static copula models utilizing the approaches of Sections 2.1 and 2.2. For the static bivariate copula we proceed as follows for an edge e .
 - Set starting values s_e^0, m_e^0 .
 - For $r = 1, \dots, R$: obtain k samples of s_e, m_e from a static bivariate copula model based on data $\{u_{t, a_e | \mathbf{D}_e}^r, u_{t, b_e | \mathbf{D}_e}^r, t = 1, \dots, T\}$. We use s_e^{r-1}, m_e^{r-1} as starting value and set s_e^r, m_e^r to the sample obtained in the k -th iteration.

For the dynamic copula model we proceed similarly.

- (b) We select for each edge $e \in E_l$ among the bivariate dynamic, static and the bivariate independence copula model as explained in Section 2.3.
- (c) For each edge $e \in E_l$, the pseudo data for the next tree is obtained as

$$\begin{aligned} u_{t, a_e | b_e \cup \mathbf{D}_e}^r &= h_{a_e | b_e; \mathbf{D}_e}(u_{t, a_e | \mathbf{D}_e}^r | u_{t, b_e | \mathbf{D}_e}^r; (\theta_{t, e}^{m_e})^r, m_e^r), \\ u_{t, b_e | a_e \cup \mathbf{D}_e}^r &= h_{b_e | a_e; \mathbf{D}_e}(u_{t, b_e | \mathbf{D}_e}^r | u_{t, a_e | \mathbf{D}_e}^r; (\theta_{t, e}^{m_e})^r, m_e^r), \end{aligned} \quad (35)$$

for $r = 1, \dots, R, t = 1, \dots, T$, if the dynamic copula was selected for edge e . If the static copula was selected we replace $\theta_{t,e}^{m_e}$ by $\theta_e^{m_e}$ in (35). For the independence copula model we set $u_{t,a_e|b_e \cup \mathcal{D}_e}^r = u_{t,a_e|\mathcal{D}_e}^r$ and $u_{t,b_e|a_e \cup \mathcal{D}_e}^r = u_{t,b_e|\mathcal{D}_e}^r$.

6. If $l < d - 1$, set $l = l + 1$ and go to 4.

Runtime and scalability

The MCMC samplers in step 2.(a) and the ones in step 5.(a) can be run in parallel, respectively. The MCMC samplers are the main drivers for the runtime and therefore parallelization speeds up computation a lot. When enough cores, i.e. at least $d - 1$ cores for d -dimensional data, are available, we observed that the computation time for one tree is no more than 40 minutes for time series data of length $T = 1000$ with $R = 1100, burnin = 100, k = 25$, independently of d . For estimating a full vine (i.e. a vine without truncation), we expect that the computation time grows roughly linearly with the dimension d and a full vine in 11 dimension (containing 10 trees) should take no more than $10 \cdot 40$ minutes. But in higher dimensions it is often not necessary to estimate all trees, e.g. we expect the runtime for a 100-dimensional vine, truncated after the 10-th tree, to be not much more than $10 \cdot 40$ minutes. Thus, in combination with truncation, we expect our method to scale very well to higher dimensions.

3.3 Simulation study

With this simulation study, we aim to obtain a first impression of the ability of the procedure proposed in Section 3.2 to recover trajectories of Kendall's τ and of the ability to select copula families and the type of dependence (dynamic, static, zero). A more extensive simulation study to investigate the potential of the novel approach in more detail is planned for the future.

First, we assume the tree structure to be known and the steps for the vine structure selection in our procedure are left out. This allows to compare the true and estimated Kendall's τ values for each pair copula. Afterwards we allow for vine structure selection. In this case our procedure might select different tree structures. Thus the true trajectories of Kendall's τ and the copula family for some pair copulas included in the selected vine structure may not be directly known. We deal with this case by comparing simulations from the true and the estimated model. In addition we compare average log likelihoods of true and estimated models.

Known tree structure

We consider the tree structure presented in Figure 2. The corresponding families are chosen from the set: {Independence, Gaussian, Student t(df=2), Student t(df=4), Student t(df=8), eGumbel, eClayton}. For each pair copula we simulate one trajectory of length $T = 1000$ for Kendall's τ from an AR(1) process. The chosen copula families and the parameters of the AR(1) processes are specified in Appendix B. We keep the tree structure, the choice of the families and the trajectories for Kendall's τ fixed and simulate 100 times from this model.

For each of the 100 simulated data sets we run the algorithm proposed in Section 3.2. Within our sequential procedure, we set $R = 1100, k = 25$ and $burnin = 100$. This means that within the procedure $1100 \cdot 25$ draws have been obtained for each parameter, whereas 1100 iterations are stored. Of these 1100 stored iterations the first 100 are discarded for burn-in. From Table 3 we see that for each pair copula in the first two trees the correct family was selected in at least 94 out of 100 cases. In Tree 3, the two independence copulas were detected in 69 and 100 out of 100 cases. The static copula in Tree 3 was detected in 88 out of 100 cases. The independence copulas in Trees 4 and 5 were detected every time. Here we count a Student t copula as correctly detected if it was selected as a Student t copula, independently of the degrees of freedom parameter. Table 3 also shows how often the correct type of dependence (dynamic, static, zero) was selected. For one pair copula in the first tree the correct type was only detected in 75 out of 100 cases. The corresponding Kendall's τ is shown in the fifth row, third column in Figure 4. We see that this Kendall's τ does not change a lot over time. So it is difficult to distinguish between the dynamic and the static model for this pair copula. Except for this pair copula, the type of dependence of pair copulas in the first two trees was detected in at least 93 out of 100 cases. In trees, higher than tree 2, the correct type was selected in at least

69 out of 100 cases. We think that these are reasonable results for the selection of the family and of the type of dependence for the pair copulas, that make up the dynamic vine copula. In addition, Figures 4 and 5 illustrate that our procedure can recover the simulated trajectories of Kendall's τ . In these figures, we show marginal (univariate) posterior mode estimates of Kendall's τ parameters, which will be utilized later (Section 4) as point estimates.

Tree	Copula family					Type of dependence				
5	100					100				
4	100	100				100	100			
3	69	88	100			69	78	100		
2	100	100	94	100		100	96	97	100	
1	94	100	100	99	100	100	93	75	98	97

Table 3: This table shows how often the correct copula family and how often the correct type of dependence (dynamic, static, zero) was selected out of the 100 simulations for each pair copula of the dynamic vine copula. There are $6 - i$ pair copulas in the i -th tree. The selected copula family for an edge e is the marginal posterior mode estimate of m_e , i.e. the family that occurs most frequently among the posterior samples for m_e .




Figure 4: This plot corresponds to a fitted model for one simulated data set. Posterior mode estimates of Kendall's τ at time t are plotted against t for each pair copula (black lines). The posterior mode estimates are obtained from marginal (univariate) kernel density estimates of the corresponding Kendall's τ parameter. A 90% credible region constructed from the estimated 5% and 95% posterior quantiles is added in grey. True values of Kendall's τ are added in red.




Figure 5: In this plot we consider 100 estimated models. The mean of 100 posterior mode estimates of Kendall's τ at time t is plotted against t for each pair copula (black line). The posterior mode estimates are obtained from marginal (univariate) kernel density estimates of the corresponding Kendall's τ parameter. The blue region is constructed from the empirical 5% and 95% quantiles of the 100 posterior mode estimates. True values of Kendall's τ are added in red.

Unknown tree structure

We use the same 100 simulated data sets as in the case with known tree structure but allow here for structure selection within our procedure. As already mentioned, evaluating our results is not straightforward in this case. Our estimated model may contain pair copulas for which we do not know the true copula families and Kendall's τ values directly. In this case we simulate 500 times from the true model and from the estimated model. Then we can calculate empirical Kendall's τ values for each of the $\frac{6 \cdot 5}{2} = 15$ pairs $(U_1, U_2), (U_1, U_3), \dots$. We compare trajectories of the empirical Kendall's τ values in Figures 6 and 7. These trajectories look similar for the true and the estimated models. We see that also our procedure with structure selection is able to recover the simulated trajectories of Kendall's τ .

For further evaluation of the proposed procedure we compare log-likelihoods of estimated and true models, as in Gruber and Czado (2015). To save computation time, we evaluate the likelihoods of estimated models based on point estimates (marginal posterior mode estimates) of the parameters, instead of evaluating the likelihoods for all posterior draws. The average log-likelihood of models without structure selection was 94% of the log-likelihood of the true model, whereas the log-likelihood of the models estimated with structure selection was on average 89% of the log-likelihood of the true model. It is not surprising that we perform a bit better if we assume the vine structure to be known. But the difference is not very big and in both cases, with and without vine structure selection, we obtain reasonable results. For further comparison we also estimated dynamic C-vine and D-vine copulas. Therefore we just restrict our structure selection procedure in Section 3.2 to C-vine and D-vine structures, respectively. The dynamic C-vine and D-vine copulas achieved 86% and 88% of the log-likelihood of the true model, respectively. Thus, in this scenario, allowing for general vine structures improves the fit compared to restricting the structure to C-vines or D-vines.




Figure 6: This plot corresponds to a fitted model for one simulated data set. The empirical unconditional Kendall's τ estimate at time t , obtained from simulations from the fitted model, is plotted against t for each pair $(U_i, U_j), i, j \in \{1, \dots, 6\}, i < j$ (black line). True Kendall's τ values determined by simulating from the true model 500 times are added in red.

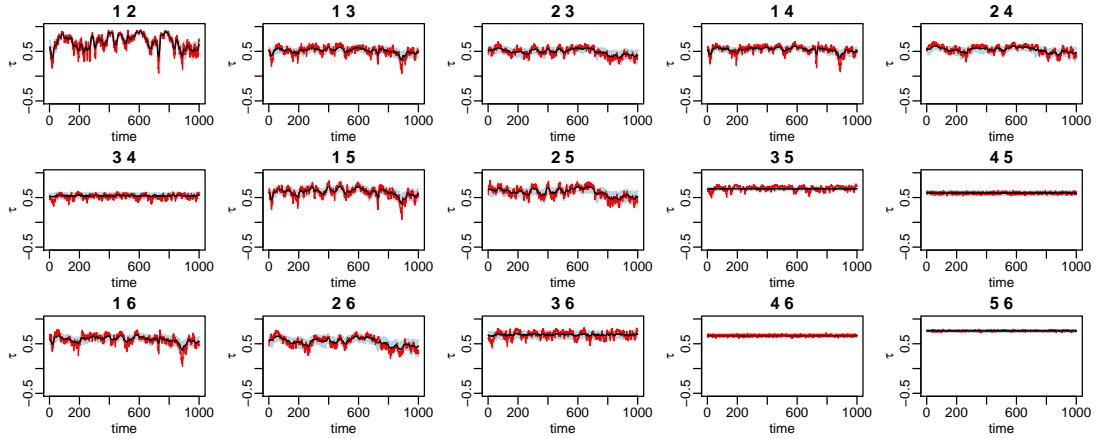


Figure 7: In this plot we consider 100 fitted models. The mean of 100 empirical estimates of the unconditional Kendall's τ at time t , obtained from simulations of the fitted model, is plotted against t for each pair $(U_i, U_j), i, j \in \{1, \dots, 6\}, i < j$ (black line). The blue region is constructed from the 5% and 95% empirical quantiles of the 100 empirical estimates of Kendall's τ . True Kendall's τ values determined by simulating from the true model 500 times are added in red.

4 Application: Dynamic exchange rates dependence

We employ the proposed dynamic vine copula model to model the dependence among 21 exchange rates with respect to the US Dollar (USD). For this we use data obtained from the FRED database of the Federal Reserve Bank of St. Louis (<https://fred.stlouisfed.org/categories/94>) which comprises daily log returns of 21 exchange rates with respect to the USD from 2007 to 2018, resulting in 3130 observations. The 21 currencies and their ticker symbols are summarized in Appendix C. We estimate our model based on the first 1500 observations and evaluate its predictive performance based on the remaining 1630 observations. First the data is demeaned based on the first 1500 observations and we collect the demeaned log returns in the data matrix $Y = (y_{tj})_{t=1, \dots, 3130, j=1, \dots, 21} \in \mathbb{R}^{3130 \times 21}$.

For the marginals we use skew Student t stochastic volatility models, i.e. we assume that

$$\begin{aligned}
Y_{tj} &= \exp\left(\frac{s_{tj}^{st}}{2}\right)\epsilon_{tj}^{st} \\
s_{tj}^{st} &= \mu_j^{st} + \phi_j^{st}(s_{t-1j}^{st} - \mu_j^{st}) + \sigma_j^{st}\eta_{tj}^{st}
\end{aligned} \tag{36}$$

with $\eta_{tj} \sim N(0, 1)$ independently, $\mu_j^{st} \in \mathbb{R}, \phi_j^{st} \in (-1, 1), \sigma_j^{st} \in (0, \infty)$ for $t = 1, \dots, 3130$. The error ϵ_{tj} follows marginally a standardized skew Student t distribution with skewness parameter $\alpha_j \in \mathbb{R}$ and degrees of freedom parameter $df_j \in (2, \infty)$ (see Kreuzer and Czado (2019)). The corresponding density and distribution functions are denoted by $sst(\cdot|\alpha_j, df_j)$ and $SST(\cdot|\alpha_j, df_j)$, respectively. The joint distribution among the errors is modeled by the proposed dynamic vine copula model.

We follow ideas of the two step approach, commonly used in copula modeling, and assume independence among the errors ϵ_{tj}^{st} for estimating the margins. But instead of collapsing parameters of the marginal skew Student t stochastic volatility models to point estimates and obtain the copula data based on these point estimates we follow ideas from Section 3.2. For each of the 21 marginal time series y_{1j}, \dots, y_{1500j} we estimate a skew Student t stochastic volatility model as explained in Kreuzer and Czado (2019). The sampler is run for $1100 \cdot 25$ iterations and then we thin the samples with factor 25. The parameter draws of the skewness parameters, of the degrees of freedom parameters and of the latent log variances are denoted by $(\alpha_j^{st})^r, (df_j^{st})^r, (s_{0j}^{st})^r, \dots, (s_{1500j}^{st})^r, r = 1, \dots, 1100$. For each parameter draw, we obtain pseudo copula data as follows

$$u_{tj}^r = SST\left(y_{tj} \exp\left(-\frac{(s_{tj}^{st})^r}{2}\right) \middle| (\alpha_j^{st})^r, (df_j^{st})^r\right) \tag{37}$$

for $t = 1, \dots, 1500, j = 1, \dots, 21, r = 1, \dots, 1100$.

Based on these pseudo copula data sets we fit a dynamic vine copula model. The algorithm of Section 3.2 is slightly modified. We start with Step 3. and set $l = 1$ since we fit our model with a collection of copula data sets instead of only one copula data set. Further we set $R = 1100, k = 25$ and $burnin = 100$. The copula families are selected from the following set $\mathcal{M} = \{\text{Independence, Gaussian, Student t(df=2), Student t(df=4), Student t(df=8), eGumbel, eClayton}\}$. The estimated dynamic vine copula model is analyzed in more detail in the following.

The first tree of the selected vine tree structure is shown in Figure 8. We see that some currencies that are connected by an edge are from countries of the same region. For example, the currencies GBP/USD (British Pound to USD) and DKK/USD (Danish Krone to USD) are connected to the EUR/USD (Euro to USD). Since the vine structure is selected as the maximum spanning tree, where the absolute value of Kendall's τ serves as weight, this indicates high dependence among those currencies. Further, we see that the selected vine structure is neither a C-vine nor a D-vine structure. The generalization of C-vine and D-vine structures to R-vine structures seems to be necessary.

In Table 4 we show the selected types of dependence per tree level. Above tree nine, all selected copulas are equal to the independence copula, i.e. the type of dependence is estimated to be zero. Further we see that only few static copulas were selected. The number of dynamic copulas selected decreases as we move up to higher tree levels. In total $\frac{20 \cdot 21}{2} = 210$ pair copulas are estimated, of which 150 were set equal to the independence copula. Our proposed procedure is able to detect sparse structures. Note that the level of sparsity can be increased by adjusting the selection of the type of dependence accordingly. As mentioned in Section 2.3, we decide for the more complex type of dependence if the WAIC of the more complex model is at least 2 standard errors smaller. By increasing 2 to for example 4 standard errors, we achieve more sparsity. This might be interesting in higher dimensional settings. Since for most pair copulas in the first trees the dynamic type of dependence is selected, a vine copula model with static dependence might not be appropriate for those selected pair copula terms with time-varying dependence.




Figure 8: The first tree of the vine tree structure selected for the 21-dimensional exchange rates data set. Nodes which belong to the same region have the same color (Europe: grey, Asia: white, America: green, Australasia: orange, Africa: yellow).

Tree	Dynamic	Static	Zero
1	18	2	0
2	16	1	2
3	7	0	11
4	6	1	10
5	2	0	14
6	3	0	12
7	2	0	12
8	1	0	12
9	1	0	11

Table 4: We show how often the different types of dependence (dynamic, static, zero) were selected per tree level for the first nine trees.

Figure 9 shows how the dynamic Kendall's τ values evolve over time. We see that the dependence between the exchange rates AUD/USD (Australian Dollar to USD) and ZAR/USD (South African Rand to USD) varies more in 2007 and 2008, during the financial crisis, and remains almost constant after that period. Further we observe that the Kendall's τ between SGD/USD (Singapore Dollar to USD) and THB/USD (Thai Baht to USD) is close to zero in 2007 and then starts to increase after 2007. The dependence between DKK/USD (Danish Krone to USD) and CHF/USD (Swiss Franc to USD) is rather high but decreases in 2010 and reaches its lowest point in 2011. This might be the effect of the introduction of the cap on the Swiss Franc on 6 September 2011 by the Swiss National Bank. The minimum exchange rate was set at 1.2 CHF (Swiss Franc) per EUR (Euro). The second row of Figure 9 shows fitted conditional Kendall's τ values. For example, we see how the Kendall's τ of the exchange rates DKK/USD (Danish Krone to USD) and JPY/USD (Japanese Yen to USD) conditional on CHF/USD (Swiss Franc to USD) evolves over time. The conditional dependence varies between -0.5 and 0 . We also provide quantification of uncertainty of the Kendall's τ values through credible intervals. This is an advantage of our Bayesian approach compared to the frequentist approach of Almeida et al (2016) for dynamic D-vine copulas, where uncertainty quantification was not provided.




Figure 9: Visualization of the dynamic Kendall's τ for some chosen pair copulas of the dynamic vine copula model estimated for the 21-dimensional exchange rates data set. Rows 1 and 2 of this plot correspond to pair copulas in Trees 1 and 2, respectively. The black line shows marginal posterior mode estimates of Kendall's τ plotted against time t . A 90% credible region constructed from the estimated 5% and 95% posterior quantiles is added in grey.

As already mentioned, the proposed dynamic vine copula model can be seen as a generalization of static vine copulas and as a generalization of the dynamic C-vine and D-vine copula models of Goel and Mehra (2019) and Almeida et al (2016). We would like to further support the hypothesis that our proposed dynamic vine copula model is a needed generalization and is able to describe the dependence structure more appropriate than its competitor models: a dynamic C-vine, a dynamic D-vine and a static vine copula. Therefore we compare these models with respect to their predictive accuracy. To estimate the dynamic C-vine and D-vine copula model, we adjust our structure selection procedure in Section 3.2 accordingly. The families are selected from the same set \mathcal{M} that we used for the dynamic vine copula. The static vine copula model is estimated with the algorithm of Dissmann et al (2013), as implemented in the R-package `rvinecopulib` (Nagler and Vatter (2018)). Here we allow for all parametric copula families that are implemented in the `rvinecopulib` package. For all models we use skew Student t stochastic volatility models for the margins. For the static vine copula model the pseudo copula data is obtained by fixing the parameters of the skew Student t stochastic volatility models at marginal posterior mode estimates. The competitor models are also estimated based on the first 1500 observations of our data. For all models we obtain one day ahead predictive scores for the other 1630 days in our data set. We proceed as in Kreuzer and Czado (2019). Instead of refitting the models 1629 times, we keep static model parameters fixed at point estimates (marginal posterior mode estimates for the dynamic models, maximum likelihood estimates for the static vine copula) and only update dynamic parameters. This reduces computation time a lot. To further reduce computation time, we evaluate the log predictive likelihoods at point estimates instead of averaging over all posterior draws. Similar to Kastner (2019), we call this pseudo log predictive scores (plps). The plps at time $t > 1500$ has the following structure

$$\text{plps}_t(y_1, \dots, y_{21}) = \ln(\hat{c}_t(\hat{u}_1, \dots, \hat{u}_{21})) + \sum_{j=1}^{21} \ln \left(\text{sst} \left(y_j \exp \left(-\frac{\hat{s}_{tj}^{st}}{2} \right) \middle| \hat{\alpha}_j^{st}, \hat{d}f_j^{st} \right) \right) - \frac{\hat{s}_{tj}^{st}}{2}, \quad (38)$$

where \hat{c}_t is the estimated copula density of one of the four considered models obtained by fixing the corresponding parameters at point estimates (marginal posterior mode estimates for the dynamic models, maximum likelihood estimates for the static vine copula). Further $\hat{u}_j = \text{SST} \left(y_j \exp \left(-\frac{\hat{s}_{tj}^{st}}{2} \right) \middle| \hat{\alpha}_j^{st}, \hat{d}f_j^{st} \right)$ with marginal posterior mode estimates $\hat{s}_{tj}^{st}, \hat{\alpha}_j^{st}, \hat{d}f_j^{st}$ for $j = 1, \dots, 21$. The marginal contribution $\sum_{j=1}^{21} \ln \left(\text{sst} \left(y_j \exp \left(-\frac{\hat{s}_{tj}^{st}}{2} \right) \middle| \hat{\alpha}_j^{st}, \hat{d}f_j^{st} \right) \right) - \frac{\hat{s}_{tj}^{st}}{2}$ is the same for all considered models. So we compare the models with respect to the copula contributions

$\ln(\hat{c}(\hat{u}_1, \dots, \hat{u}_{21}))$ to which we refer as copula plps. Note that a higher (copula) plps is an indication for better forecasting accuracy.

Table 5 shows the cumulative copula plps, i.e. the sum over all 1630 plps. We see that the two vine copula models with flexible tree structure outperform the dynamic C-vine and D-vine copulas. Further the dynamic vine copula, for which the selected structure deviates clearly from a C-vine and a D-vine structure and for which many pair copulas have a dynamic type of dependence provides the most accurate forecasts. Our conclusion is that the dynamic vine copula model provides a useful generalization of static vine copula models as well as of dynamic C-vine and D-vine copula models.

	Dynamic vine	Dynamic C-vine	Dynamic D-vine	Static vine
copula plps	11643	11132	11126	11267

Table 5: Cumulative one day ahead copula plps for the four considered models: Dynamic vine, dynamic C-vine, dynamic D-vine and static vine copula.

5 Conclusion and future research

We introduced a class of dynamic vine copula models and provided a novel Bayesian estimation procedure based on an approximation of the posterior distribution allowing for simplification of the sampler. Here we allowed for the selection of the pair copula family, the selection of the type of dependence for each pair copula term and the sequential selection of a static (time-constant) vine structure. The application showed that the dynamic vine copula model is a useful extension of static vine copulas and of dynamic C-vine and D-vine copulas.

Our estimation procedure propagates uncertainty of copula parameter estimation from lower to higher trees. But the type of dependence is selected with WAIC and then fixed before we move to the next tree. Instead of using the WAIC or any other information criteria, shrinkage priors as proposed by Bitto and Frühwirth-Schnatter (2019) that allow to shrink dynamic parameters to static ones might be an interesting alternative to be studied in future research.

One restriction of our approach is that both the vine tree structure as well as the pair copula families are assumed not to change over time. In the future we are interested in overcoming these restrictions.

Further, it would be interesting to study in more detail how the dynamic vine copula model performs in situations, where appropriate dependence modeling is crucial, for example in financial risk management. The dynamic vine copula model might lead to more accurate value at risk predictions than those obtained from static vine copula models. Another example is pairs trading. Stübinger et al (2018) showed that profitable trading strategies can be constructed with static vine copula models. These strategies might be improved by allowing for dynamic dependencies.

Lastly, we think that the ability of the vine copula framework to scale bivariate copula models to copula models of arbitrary dimensions has not been fully exploited yet. We are sure that there is a variety of useful extensions of static vine copula models that build on sophisticated bivariate copula models. For example, one could allow for bivariate copula families with more than one parameter such as the BB1 family or for bivariate dynamic mixture copulas as studied by Kreuzer and Czado (2019), where both mixture components share the same dynamic on Kendall's τ . Alternatively one could also study bivariate mixture copula models with one dynamic and one static component.

A Additional material for parameter sharing (Section 2.1)

Our procedure in Section 2.1 shares parameters among different copula families. This is motivated by the fact that the parameter s_t , the Fisher's Z transform of Kendall's τ , is similar for different copula families. To support this statement, we conduct the following experiment: We simulate 100 bivariate data sets, each containing 1000 observations, from the bivariate Student t copula with 4 degrees of freedom and copula parameter ρ_{true} . The corresponding Kendall's τ

C Exchange rates (to the US Dollar) data set

Ticker	Currency
BRL	Brazilian Real
CAD	Canadian Dollar
CNY	Chinese Yuan
DKK	Danish Krone
HKD	Hong Kong Dollar
INR	Indian Rupees
JPY	Japanese Yen
KRW	South Korean Won
MYR	Malaysian Ringgit
MXN	Mexican New Pesos
NOK	Norwegian Krone
SEK	Swedish Krona
ZAR	South African Rand
SGD	Singapore Dollar
CHF	Swiss Franc
NTD	New Taiwan Dollar
THB	Thai Baht
AUD	Australian Dollar
EUR	Euro
NZD	New Zealand Dollar
GBP	British Pound

Table 6: The 21 currencies with corresponding ticker symbols used in the application in Section 4.

Acknowledgements

The second author is supported by the German Research Foundation (DFG grant CZ 86/4-1). Computations were performed on a Linux cluster supported by DFG grant INST 95/919-1 FUGG.

References

- Aas K (2016) Pair-copula constructions for financial applications: A review. *Econometrics* 4(4):43
- Aas K, Czado C, Frigessi A, Bakken H (2009) Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics* 44(2):182–198
- Acar EF, Czado C, Lysy M (2019) Flexible dynamic vine copula models for multivariate time series data. *Econometrics and Statistics* 12:181–197
- Almeida C, Czado C (2012) Efficient Bayesian inference for stochastic time-varying copula models. *Computational Statistics & Data Analysis* 56(6):1511–1527
- Almeida C, Czado C, Manner H (2016) Modeling high-dimensional time-varying dependence using dynamic D-vine models. *Applied Stochastic Models in Business and Industry* 32(5):621–638
- Baele L, Bekaert G, Inghelbrecht K (2010) The determinants of stock and bond return comovements. *The Review of Financial Studies* 23(6):2374–2428
- Barthel N, Geerdens C, Czado C, Janssen P (2018) Dependence modeling for recurrent event times subject to right-censoring with D-vine copulas. *Biometrics* 75:439–451

- Bedford T, Cooke RM (2001) Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence* 32(1-4):245–268
- Bitto A, Frühwirth-Schnatter S (2019) Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics* 210(1):75–97
- Brechmann EC, Czado C (2013) Risk management with high-dimensional vine copulas: An analysis of the Euro Stoxx 50. *Statistics & Risk Modeling* 30(4):307–342
- Brechmann EC, Czado C, Aas K (2012) Truncated regular vines in high dimensions with application to financial data. *Canadian Journal of Statistics* 40(1):68–85
- Czado C (2019) *Analyzing Dependent Data with Vine Copulas*. Lecture Notes in Statistics, Springer
- Dissmann J, Brechmann EC, Czado C, Kurowicka D (2013) Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis* 59:52–69
- Embrechts P, McNeil A, Straumann D (2002) Correlation and dependence in risk management: properties and pitfalls. *Risk Management: Value at Risk and Beyond* 1:176–223
- Engle R (2002) Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20(3):339–350
- Erhardt TM, Czado C, Schepsmeier U (2015) R-vine models for spatial time series with an application to daily mean temperature. *Biometrics* 71(2):323–332
- Garthwaite PH, Fan Y, Sisson SA (2016) Adaptive optimal scaling of Metropolis–Hastings algorithms using the Robbins–Monro process. *Communications in Statistics-Theory and Methods* 45(17):5098–5111
- Goel A, Mehra A (2019) Analyzing Contagion Effect in Markets During Financial Crisis Using Stochastic Autoregressive Canonical Vine Model. *Computational Economics* 53(3):921–950
- Green PJ (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82(4):711–732
- Gruber LF, Czado C (2015) Sequential Bayesian model selection of regular vine copulas. *Bayesian Analysis* 10(4):937–963
- Gruber LF, Czado C (2018) Bayesian model selection of regular vine copulas. *Bayesian Analysis* 13(4):1107–1131
- Haff IH, Aas K, Frigessi A (2010) On the simplified pair-copula construction—simply useful or too simplistic? *Journal of Multivariate Analysis* 101(5):1296–1310
- Hafner CM, Manner H (2012) Dynamic stochastic copula models: Estimation, inference and applications. *Journal of Applied Econometrics* 27(2):269–295
- Harvey A, Ruiz E, Shephard N (1994) Multivariate stochastic variance models. *The Review of Economic Studies* 61(2):247–264
- Joe H (2014) *Dependence modeling with copulas*. CRC Press
- Kastner G (2016) Dealing with stochastic volatility in time series using the R package stochvol. *Journal of Statistical Software* 69(5):1–30
- Kastner G (2019) Sparse Bayesian time-varying covariance estimation in many dimensions. *Journal of Econometrics* 210(1):98–115

- Kastner G, Frühwirth-Schnatter S, Lopes HF (2017) Efficient Bayesian inference for multivariate factor stochastic volatility models. *Journal of Computational and Graphical Statistics* 26(4):905–917
- Kreuzer A, Czado C (2019) Efficient Bayesian inference for univariate and multivariate non linear state space models with univariate autoregressive state equation. arXiv e-prints arXiv:1902.10412
- Liu JS (1994) The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association* 89(427):958–966
- Min A, Czado C (2011) Bayesian model selection for D-vine pair-copula constructions. *Canadian Journal of Statistics* 39(2):239–258
- Möller A, Spazzini L, Kraus D, Nagler T, Czado C (2018) Vine copula based post-processing of ensemble forecasts for temperature. arXiv preprint arXiv:181102255
- Morales-Nápoles O (2010) Counting vines. In: *Dependence modeling: Vine copula handbook*, World Scientific, pp 189–218
- Murray I, Adams RP, MacKay DJ (2010) Elliptical slice sampling. *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)* 9:541–548
- Nagler T, Vatter T (2018) rvinecopulib: High performance algorithms for vine copula modeling. R package version 02 8(0)
- Oh DH, Patton AJ (2018) Time-varying systemic risk: Evidence from a dynamic copula model of cds spreads. *Journal of Business & Economic Statistics* 36(2):181–195
- Pitt M, Shephard N (1999) Time varying covariances: a factor stochastic volatility approach. *Bayesian Statistics* 6:547–570
- Richard JF, Zhang W (2007) Efficient high-dimensional importance sampling. *Journal of Econometrics* 141(2):1385–1411
- Sklar M (1959) Fonctions de repartition an dimensions et leurs marges. *Publ inst statist univ Paris* 8:229–231
- Stoeber J, Joe H, Czado C (2013) Simplified pair copula constructions—limitations and extensions. *Journal of Multivariate Analysis* 119:101–118
- Stübinger J, Mangold B, Krauss C (2018) Statistical arbitrage with vine copulas. *Quantitative Finance* 18(11):1831–1849
- Tan BK, Panagiotelis A, Athanasopoulos G (2019) Bayesian inference for the one-factor copula model. *Journal of Computational and Graphical Statistics* 28(1):155–173
- Van Dyk DA, Jiao X (2015) Metropolis-Hastings within partially collapsed Gibbs samplers. *Journal of Computational and Graphical Statistics* 24(2):301–327
- Vatter T, Chavez-Demoulin V (2015) Generalized additive models for conditional dependence structures. *Journal of Multivariate Analysis* 141:147–167
- Vatter T, Nagler T (2018) Generalized additive models for pair-copula constructions. *Journal of Computational and Graphical Statistics* 27(4):715–727
- Vehtari A, Gelman A, Gabry J (2017) Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing* 27(5):1413–1432
- Wainwright MJ, Jordan MI, et al (2008) Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning* 1(1–2):1–305
- Watanabe S (2010) Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research* 11(Dec):3571–3594

Yu Y, Meng XL (2011) To center or not to center: That is not the question—an Ancillarity–Sufficiency Interweaving Strategy (ASIS) for boosting MCMC efficiency. *Journal of Computational and Graphical Statistics* 20(3):531–570