

Efficient Bayesian inference for nonlinear state space models with univariate autoregressive state equation

Alexander Kreuzer* Claudia Czado

Zentrum Mathematik, Technische Universität München

October 31, 2019

Abstract

Latent autoregressive processes are a popular choice to model time varying parameters. These models can be formulated as nonlinear state space models for which inference is not straightforward due to the high number of parameters. Therefore maximum likelihood methods are often infeasible and researchers rely on alternative techniques, such as Gibbs sampling. But conventional Gibbs samplers are often tailored to specific situations and suffer from high autocorrelation among repeated draws. We present a Gibbs sampler for general nonlinear state space models with an univariate autoregressive state equation. For this we employ an interweaving strategy and elliptical slice sampling to exploit the dependence implied by the autoregressive process. Within a simulation study we demonstrate the efficiency of the proposed sampler for bivariate dynamic copula models. Further we are interested in modeling the volatility return relationship. Therefore we use the proposed sampler to estimate the parameters of stochastic volatility models with skew Student t errors and the parameters of a novel bivariate dynamic mixture copula model. This model allows for dynamic asymmetric tail dependence. Comparison to relevant benchmark models, such as the DCC-GARCH or a Student t copula model, with respect to predictive accuracy shows the superior performance of the proposed approach. `R` scripts, an `R` package, data and additional results are provided as supplementary material.

1 Introduction

There are many situations where statistical models with constant parameters are no longer sufficient to appropriately represent certain aspects of the economy. For example it is well known that volatility of financial assets changes over time (Schwert (1989)). This is why many models that allow for variation in the parameter have been proposed. There are time varying vector autoregressive models (Primiceri (2005), Nakajima et al (2011)), stochastic volatility models (Kim et al (1998)), GAM copula models (Vatter and Chavez-Demoulin (2015), Vatter and Nagler (2018)) and many more. Stochastic volatility models and the bivariate dynamic copula model of Almeida and Czado (2012) assume that the parameter follows a latent autoregressive process of order 1 (AR(1) process). These two models belong to the class of models that we will study.

In a general time varying parameter framework we consider a d dimensional random variable at time t , $\mathbf{Y}_t \in \mathbb{R}^d$, which is generated from a d dimensional density $f(\cdot|s_t)$. We are interested in models, where the density $f(\cdot|s_t)$ has a univariate dynamic parameter $s_t \in \mathbb{R}$ following an AR(1) process. These models can be formulated as state space models with observation equation

$$\mathbf{Y}_t|s_t \sim f(\mathbf{y}_t|s_t) \text{ independently,} \quad (1)$$

for $t = 1, \dots, T$. The state equation describes an AR(1) process with mean parameter $\mu \in \mathbb{R}$, persistence parameter $\phi \in (-1, 1)$ and standard deviation parameter $\sigma \in (0, \infty)$ and is given

*Corresponding author: E-mail: a.kreuzer@tum.de

by

$$s_t = \mu + \phi(s_{t-1} - \mu) + \sigma\epsilon_t, \quad (2)$$

where $\epsilon_t \sim N(0, 1)$ iid for $t = 1, \dots, T$ and $s_0 | \mu, \phi, \sigma \sim N\left(\mu, \frac{\sigma^2}{1 - \phi^2}\right)$. In the state equation we assume Gaussian innovations ϵ_t but in the observation equation we do not put any restrictions on the density f . Thus we allow for nonlinear and non Gaussian state space models. Several established models can be analyzed within this framework.

By choosing $f(\cdot | s_t)$ as the univariate normal density with mean 0 and variance $\exp(s_t)$, denoted by $\varphi(\cdot | 0, \exp(s_t))$, we obtain the stochastic volatility model (Kim et al (1998)) given by

$$\begin{aligned} Y_t | s_t &\sim \varphi(y_t | 0, \exp(s_t)) \text{ independently,} \\ s_t &= \mu + \phi(s_{t-1} - \mu) + \sigma\epsilon_t, \end{aligned} \quad (3)$$

for $t = 1, \dots, T$. By modeling the log variance as a latent AR(1) process this model allows for time varying volatility. Time varying volatility is a stylized fact of financial time series. In this context, the stochastic volatility model has also shown better performance than the frequently used GARCH models (Engle (1982), Bollerslev (1986)) for several data sets (Yu (2002), Chan and Grant (2016)). To allow for heavy tails and skewness, other distributions have been considered in the observation equation. One example is the stochastic volatility model with skew Student t errors (Abanto-Valle et al (2015)), which can also be analyzed within our framework.

Dependence modeling is another research area, where models that allow for time varying parameters have been introduced. Vine copulas (Bedford and Cooke (2001), Aas et al (2009), Czado (2019)) are widely used models to capture complex dependence structures. To name a few, Brechmann and Czado (2013) and Nagler et al (2019) employ vine copulas for forecasting the value at risk of a portfolio, Aas (2016) gives an overview of applications of vine copulas in finance including asset pricing, credit risk management and portfolio optimization and Barthel et al (2018) model the association pattern between gap times with D-vine copulas to study asthma attacks. A vine copula model is made up of different bivariate copulas with corresponding dependence parameters. Since dependencies may change over time, extensions that allow for variation in the dependence parameter have been proposed. Vatter and Chavez-Demoulin (2015) introduce a bivariate copula model, where the dependence parameter follows a generalized additive model. Another approach is the dynamic bivariate copula model proposed by Almeida and Czado (2012) and Hafner and Manner (2012), which we will analyze within our state space framework. For this model, we consider single parameter copula families for which there is a one-to-one correspondence between the copula parameter, denoted by θ , and Kendall's τ . So we can express Kendall's τ as a function of the copula parameter θ and we write $\tau(\theta)$. The restriction of Kendall's τ to the interval $(-1, 1)$ is removed by applying the Fisher' Z transformation $F_Z(x) = \frac{1}{2} \log\left(\frac{1+x}{1-x}\right)$. This transformed time varying Kendall's τ is then modeled by an AR(1) process. More precisely, we consider T bivariate random vectors, $(U_{t1}, U_{t2})_{t=1, \dots, T} \in [0, 1]^{T \times 2}$, corresponding to T time points. We assume for $t = 1, \dots, T$ that

$$\begin{aligned} (U_{t1}, U_{t2}) | \theta_t &\sim c(u_{t1}, u_{t2}; \theta_t) \text{ independently} \\ s_t &= \mu + \phi(s_{t-1} - \mu) + \sigma\epsilon_t, \text{ for } s_t = F_Z(\tau(\theta_t)), \end{aligned} \quad (4)$$

where $c(u_{t1}, u_{t2}; \theta_t)$ is a bivariate copula density with parameter θ_t .

Nonlinear state space models as specified with (1) and (2) are typically difficult to estimate, since there is a large number of parameters and likelihood evaluation requires high dimensional integration. This often makes maximum likelihood approaches infeasible. Gibbs sampling (Geman and Geman (1984)) is a frequently used Bayesian approach to infer parameters of such nonlinear state space models (Carlin et al (1992)). But the posterior samples, resulting from conventional Gibbs sampling, often suffer from high autocorrelation. Furthermore, the availability of the full conditional distributions or at least an efficient MCMC approach to sample from them is required. This is often tailored to specific situations. We present a Gibbs sampling approach that is designed to handle models with a latent AR(1) process and general likelihood functions as specified by the state space formulation in Equations (1) and (2). To sample from

the associated posterior distribution we rely on elliptical slice sampling (Murray et al (2010)) and on an ancillarity-sufficiency interweaving strategy (Yu and Meng (2011)). Elliptical slice sampling is used to sample the latent states. This allows us to exploit the Gaussian dependence structure, that is implied by the AR(1) process. But even if we provide efficient methods to sample from the full conditionals, the sampler may still suffer from the dependence among the parameters in the posterior distribution. Additionally, its performance may vary for different model parameterizations (Frühwirth-Schnatter and Sögner (2003), Strickland et al (2008)). This problem is tackled with the ancillarity-sufficiency interweaving strategy, where the parameters of the latent AR(1) process are sampled from two different parameterizations. The decision between two parameterizations is avoided by using both. This approach has already shown good results for several models, including univariate and multivariate stochastic volatility models (Kastner and Frühwirth-Schnatter (2014), Kastner et al (2017)). The efficiency of our proposed sampler is illustrated with a simulation study.

The second part of this work has a more applied focus and deals with modeling the volatility return relationship, i.e. the dependence between an index and the corresponding volatility index. More precisely, we investigate the American index S&P500 and its volatility index the VIX as well as the German index DAX and the VDAX. It is important to provide appropriate models for this relationship, since it has influence on hedging and risk management decisions (Allen et al (2012)). For our analysis we make use of a two step approach commonly used in copula modeling (Joe and Xu (1996)) motivated by Sklar’s Theorem (Sklar (1959)). We first model the marginal distribution with a univariate skew Student t stochastic volatility model. In the second step we model the dependence for which we propose a dynamic copula model allowing for asymmetric tail dependence. This model is a dynamic mixture of a Gumbel and a Student t copula and can be seen as an alternative to the symmetrized Joe-Clayton copula of Patton (2006). Estimation is carried out through a two step approach, where we first estimate the marginal stochastic volatility models, fix their parameters at point estimates and then estimate the dynamic mixture copula. At both steps, estimation is straightforward with the proposed sampler. Our model is able to capture several characteristics of the joint distribution of volatility and return. With respect to the marginal distribution we observe positive skewness for volatility indices compared to slight negative skewness for the return indices. In the dependence structure we identify asymmetry and time variation. Finally, we compare the proposed model to several restricted models with constant or symmetric dependence and to a bivariate DCC-GARCH model (Engle (2002)). Model comparison with respect to predictive accuracy shows the superiority of our approach.

To summarize, the main contribution of this paper is an approach to efficiently sample from the posterior distribution of general nonlinear state space models as specified by Equations (1) and (2). In addition, we propose a dynamic mixture copula for time varying asymmetric tail dependence. We discuss Bayesian inference for this model class and demonstrate how it can be utilized to model the volatility return relationship.

The outline of the paper is as follows: After the introduction, we discuss the proposed MCMC approach in Section 2. In Section 3 we investigate the efficiency of the sampler for bivariate dynamic copula models through an extensive simulation study. Section 4 deals with modeling the volatility return relationship and Section 5 concludes.

2 Bayesian inference

In the following, we denote an observation of the d -dimensional random vector \mathbf{Y}_t by \mathbf{y}_t and the associated data matrix is given by $Y = (\mathbf{y}_1, \dots, \mathbf{y}_T)^\top \in \mathbb{R}^{T \times d}$. To subset vectors and matrices we make use of the following notation: For sets of indices A and B we set $\mathbf{x}_A = (x_i)_{i \in A}$ for a vector \mathbf{x} and $X_{A;B} = (x_{ij})_{i \in A, j \in B}$ for a matrix X . We use a capital letter to refer to a matrix and small letters to refer to its components. The set $\{n, \dots, k\}$ of integers will be abbreviated by $n : k$.

We consider the state space model as specified by Equations (1) and (2). To obtain a fully specified Bayesian model we equip the parameters μ , ϕ and σ with prior distributions. We follow Kastner and Frühwirth-Schnatter (2014), who propose the following prior distributions

for the latent AR(1) process of the stochastic volatility model:

$$\mu \sim N(0, \sigma_\mu^2), \quad \frac{\phi + 1}{2} \sim \text{Beta}(a_\phi, b_\phi), \quad \sigma^2 \sim \text{Gamma}\left(\frac{1}{2}, \frac{1}{2B_\sigma}\right), \quad (5)$$

where $\sigma_\mu, a_\phi, b_\phi, B_\sigma > 0$. Our standard choice for the prior hyperparameters is $\sigma_\mu = 100$, $a_\phi = 5$, $b_\phi = 1.5$ and $B_\sigma = 1$ as in Kastner (2016). With these prior distributions our Bayesian model is complete. For sampling from the posterior distribution of this model we should take into account that sampling efficiency may highly depend on the model parameterization (Frühwirth-Schnatter and Sögner (2003), Strickland et al (2008)). Yu and Meng (2011) differentiate between two parameterizations: A sufficient augmentation and an ancillary augmentation scheme. In our case a sufficient augmentation is characterized by an observation equation that is free of the parameters μ, ϕ and σ and only depends on the latent states $\mathbf{s}_{1:T}$. In this case $\mathbf{s}_{1:T}$ is a sufficient statistics for the parameters μ, ϕ and σ . In an ancillary augmentation the state equation is independent of the parameters μ, ϕ and σ , then $\mathbf{s}_{1:T}$ is an ancillary statistics for the parameters μ, ϕ and σ . The standard parameterization of our model is already a sufficient augmentation and we refer to this parameterization as given by Equations (1) and (2) as (SA).

$$(SA) : \quad \begin{aligned} & \mathbf{Y}_t | s_t \sim f(\mathbf{y}_t | s_t) \text{ independently} \\ & s_t = \mu + \phi(s_{t-1} - \mu) + \sigma \epsilon_t. \end{aligned}$$

An ancillary augmentation is obtained by the following parameterization

$$\tilde{s}_t = \frac{s_t - \mu - \phi(s_{t-1} - \mu)}{\sigma}, \quad \text{with inverse } s_t = \mu + \phi(s_{t-1} - \mu) + \sigma \tilde{s}_t, \quad (6)$$

for $t = 1, \dots, T$. This reparameterization is obtained by solving Equation (2) for ϵ_t and implies that the state space model is given by

$$(AA) : \quad \begin{aligned} & \mathbf{Y}_t | \tilde{\mathbf{s}}_{1:T}, s_0, \mu, \phi, \sigma \sim f(\mathbf{y}_t | s_t(\tilde{\mathbf{s}}_{1:T}, s_0, \mu, \phi, \sigma)) \text{ independently} \\ & \tilde{s}_t \sim N(0, 1) \text{ independently,} \end{aligned}$$

where $s_t(\tilde{\mathbf{s}}_{1:T}, s_0, \mu, \phi, \sigma)$ is the function that calculates s_t according to (6). We refer to this model representation as (AA). Instead of deciding between (SA) and (AA), we combine them in an ancillarity-sufficiency interweaving strategy (Yu and Meng (2011)), given by

- a) Sample $\mathbf{s}_{0:T}$ in (SA) from $\mathbf{s}_{0:T} | Y, \mu, \phi, \sigma$.
- b) Sample (μ, ϕ, σ) in (SA) from $\mu, \phi, \sigma | Y, \mathbf{s}_{0:T}$.
- c) Move to (AA) via $\tilde{s}_t = \frac{s_t - \mu - \phi(s_{t-1} - \mu)}{\sigma}$, for $t = 1, \dots, T$.
- d) Sample (μ, ϕ, σ) in (AA) from $\mu, \phi, \sigma | Y, s_0, \tilde{\mathbf{s}}_{1:T}$.
- e) Move back to (SA) via the recursion $s_t = \mu + \phi(s_{t-1} - \mu) + \sigma \tilde{s}_t$ for $t = 1, \dots, T$.

Kastner and Frühwirth-Schnatter (2014) employed interweaving for the stochastic volatility model and showed its superior performance with an extensive simulation study. For the stochastic volatility model, Kastner and Frühwirth-Schnatter (2014) propose to move between a sufficient augmentation and a reparameterization of the latent states $\mathbf{s}_{1:T}$ given by $s_t^K = \frac{s_t - \mu}{\sigma}$. Within this reparameterization parameters can be sampled conveniently from its full conditional distribution by recognizing a linear regression model. This is possible for the standard stochastic volatility model, but not in our case since our sampler is designed to handle more general likelihood functions. Therefore we have chosen the reparameterization of (SA) such that it is optimal in the sense of Yu and Meng (2011), i.e. we move between a sufficient and a ancillary augmentation.

Note that reducing the sampler to the first two steps a) and b) results in a standard Gibbs sampler in (SA). This sampler typically suffers from the dependence among the parameters μ, ϕ, σ and the latent states $\mathbf{s}_{1:T}$ in the posterior distribution.

Step a: Sampling of the latent states in the sufficient augmentation

To sample the latent states $\mathbf{s}_{0:T}$ from its full conditional in (SA) we make use of elliptical slice sampling as proposed by Murray et al (2010). It was developed for models, where dependencies are generated through a latent multivariate normal distribution. In (SA), the AR(1) structure implies, that the vector $\mathbf{s}_{0:T}|\mu, \phi, \sigma$ has a $(T+1)$ dimensional multivariate normal distribution with mean vector $\boldsymbol{\mu}^{AR}$ and covariance matrix Σ^{AR} given by

$$\boldsymbol{\mu}^{AR} = \begin{pmatrix} \mu \\ \mu \\ \vdots \\ \mu \end{pmatrix} \in \mathbb{R}^{T+1}, \quad \Sigma^{AR} = \frac{\sigma^2}{1-\phi^2} \begin{pmatrix} 1 & \phi & \phi^2 & \dots & \phi^T \\ \phi & 1 & \phi & \dots & \phi^{T-1} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \phi^T & \phi^{T-1} & \phi^{T-2} & \dots & 1 \end{pmatrix} \in \mathbb{R}^{(T+1) \times (T+1)}. \quad (7)$$

(see e.g. Brockwell et al (2002), Chapter 2.2). The posterior density is proportional to

$$\left(\prod_{t=1}^T f(\mathbf{y}_t | s_t) \right) \varphi(\mathbf{s}_{0:T} | \boldsymbol{\mu}^{AR}, \Sigma^{AR}) \pi(\mu) \pi(\phi) \pi(\sigma),$$

where $\boldsymbol{\mu}^{AR}$ and Σ^{AR} are given in (7) and $\pi(\cdot)$ denotes the corresponding prior density as specified in (5). The initial state can be sampled from its full conditional density given by

$$f(s_0 | \mathbf{s}_{1:T}, \mu, \phi, \sigma) = \varphi(s_0 | \mu + \phi(s_1 - \mu), \sigma^2).$$

The full conditional density of the latent states $\mathbf{s}_{1:T}$ is given by

$$f(\mathbf{s}_{1:T} | Y, s_0, \mu, \phi, \sigma) \propto \left(\prod_{t=1}^T f(\mathbf{y}_t | s_t) \right) \varphi(\mathbf{s}_{1:T} | \boldsymbol{\mu}_{1:T|0}, \Sigma_{1:T|0}),$$

with a corresponding mean vector $\boldsymbol{\mu}_{1:T|0}$ and covariance matrix $\Sigma_{1:T|0}$. The mean vector and the covariance matrix of the conditional distribution are derived in a more general way in Appendix A. By reparameterizing the model with $\mathbf{s}'_{1:T} = \mathbf{s}_{1:T} - \boldsymbol{\mu}_{1:T|0}$, we impose a multivariate normal prior with zero mean. We obtain the situation elliptical slice sampling was designed for. However updating the whole T dimensional vector $\mathbf{s}_{1:T}$ with elliptical slice sampling at once will lead to high autocorrelation in the posterior draws. This is illustrated in Section 3 and was also observed by Hahn et al (2019), where elliptical slice sampling was used for linear regression models. Hahn et al (2019) circumvent this problem by partitioning the vector $\mathbf{s}_{1:T}$ into smaller blocks. We follow this approach and partition the set $\{1, \dots, T\}$ into m different blocks $B_1, \dots, B_m \subset \{1, \dots, T\}$. Let $a_i = \min_{s \in B_i} s$ denote the minimal and $b_i = \max_{s \in B_i} s$ denote the maximal index in the i -th block. The blocks are chosen such that $B_i = \{t \in \{1, \dots, T\} : a_i \leq t \leq b_i\}$, for $i = 1, \dots, m$. The full conditional density for the i -th block can be expressed as $f(\mathbf{s}_{B_i} | Y, s_0, \mathbf{s}_{-B_i}, \mu, \phi, \sigma) \propto \left(\prod_{t \in B_i} f(\mathbf{y}_t | s_t) \right) f(\mathbf{s}_{B_i} | s_0, \mathbf{s}_{-B_i}, \mu, \phi, \sigma)$, where $-B_i = \{1, \dots, T\} \setminus B_i$. The vector $\mathbf{s}_{B_i} | s_0, \mathbf{s}_{-B_i}, \mu, \phi, \sigma$ is multivariate normal distributed with mean denoted by $\boldsymbol{\mu}_{B_i|}$ and covariance matrix $\Sigma_{B_i|}$ (see Appendix A) and therefore the full conditional density can be written as

$$f(\mathbf{s}_{B_i} | Y, s_0, \mathbf{s}_{-B_i}, \mu, \phi, \sigma) \propto \left(\prod_{t \in B_i} f(\mathbf{y}_t | s_t) \right) \varphi(\mathbf{s}_{B_i} | \boldsymbol{\mu}_{B_i|}, \Sigma_{B_i|}).$$

To sample the latent states of the i -th block \mathbf{s}_{B_i} from its full conditional we proceed as follows

- Set $\mathbf{s}'_{B_i} = \mathbf{s}_{B_i} - \boldsymbol{\mu}_{B_i|}$
- Draw \mathbf{s}'_{B_i} from the density

$$f(\mathbf{s}'_{B_i} | Y, s_0, \mathbf{s}_{-B_i}, \mu, \phi, \sigma) \propto \left(\prod_{t \in B_i} f(\mathbf{y}_t | s_t) \right) \varphi(\mathbf{s}'_{B_i} | \mathbf{0}, \Sigma_{B_i|}),$$

using elliptical slice sampling, where $\varphi(\mathbf{s}'_{B_i} | \mathbf{0}, \Sigma_{B_i|})$ is interpreted as the prior density for \mathbf{s}'_{B_i} .

- Set $\mathbf{s}_{B_i} = \mathbf{s}'_{B_i} + \boldsymbol{\mu}_{B_i|}$

Step b: Sampling of the constant parameters in the sufficient augmentation

In (SA) the observation equation only depends on $\mathbf{s}_{1:T}$ and is independent of the parameters μ , ϕ and σ . The parameters μ , ϕ and σ only depend on $\mathbf{s}_{0:T}$. This allows to use the same approach as in Kastner and Frühwirth-Schnatter (2014) to sample the parameters μ , ϕ and σ in (SA). We reparameterize the model such that proposals can be found using Bayesian linear regression. We define $\gamma = \mu(1 - \phi)$ and the state equation is given by

$$s_t = \gamma + \phi s_{t-1} + \sigma \eta_t,$$

where $\eta_t \sim N(0, 1)$. For fixed $\mathbf{s}_{0:T}$, this is a linear regression model with regression parameters γ , ϕ and variance σ^2 . Proposals for (μ, ϕ, σ) are found and accepted or rejected as described in Kastner and Frühwirth-Schnatter (2014) Section 2.4 (two block sampler).

Step d: Sampling of the constant parameters in the ancillary augmentation

To sample μ , ϕ and σ in (AA) we deploy a random walk Metropolis-Hastings scheme with Gaussian proposal, where the proposal variance or covariance matrix is adapted during the burn-in period. For the adaptations we use the Robbins Monro process (Robbins and Monro (1985)) as suggested by Garthwaite et al (2016). More details are given in Appendix A.

Implementation

For the implementation of the sampler we use `Rcpp` (Eddelbuettel et al (2011)) which allows to embed C++ code into R. In addition we make use of `rvinecopulib` (Nagler and Vatter (2018)) to evaluate copula densities and of `RcppEigen` (Bates et al (2013)). For sampling (μ, ϕ, σ) in (SA) we use corresponding parts of the implementation of the R package `stochvol` (Kastner (2016)). The R package `coda` (Plummer et al (2008)) is used to compute effective sample sizes in the following section.

3 Illustration of the proposed sampler for bivariate dynamic copula models

We illustrate the MCMC sampler we proposed in the previous section for the bivariate dynamic copula model of Almeida and Czado (2012). Kastner and Frühwirth-Schnatter (2014) have already shown that interweaving improves sampling efficiency a lot for the stochastic volatility model. We investigate if this is also the case for the bivariate dynamic copula model. Further we study how the sampling efficiency is affected by the chosen block size and by the data generating process (DGP). Therefore we perform an extensive simulation study. We consider different modifications of the sampler. A sampler is specified by a vector (b,i) which indicates its blocksize (b) and if interweaving is used (i=I) or not (i=NI). We consider ten different sampler specifications $(b,i) \in \{1, 5, 20, 100, T\} \times \{I, NI\}$, where T is the length of the time series. By using blocks of size T , we obtain the sampler which updates the parameters $\mathbf{s}_{1:T}$ jointly with elliptical slice sampling. If we turn off interweaving (i=NI) we obtain a standard Gibbs sampler updating parameters in the sufficient augmentation. These samplers are run for different simulated data sets. A data set is simulated from the bivariate dynamic copula model (see (4)) with parameters: Family, T, μ, ϕ, σ . The parameters are chosen from the following grid $(\text{Family}, T, \mu, \phi, \sigma) \in \{\text{Gauss}, \text{eClayton}\} \times \{500, 1000, 1500\} \times \{0, 1\} \times \{0, 0.1, 0.5, 0.9, 0.99\} \times \{0.05, 0.1, 0.2\}$. Here, eClayton denotes the extended Clayton copula, which extends the Clayton copula to allow for negative Kendall's τ values. More precisely, the extended Clayton copula has the following density

$$c(u_1, u_2; \theta) = \begin{cases} c_{\text{Clayton}}(u_1, u_2; \theta) & \text{if } \theta \geq 0 \\ c_{\text{Clayton}}(1 - u_1, u_2; -\theta) & \text{if } \theta < 0, \end{cases}$$

where $c_{\text{Clayton}}(\cdot, \cdot; \theta)$ is the density of the bivariate Clayton copula with parameter θ (see Joe (2014), Chapter 4). So the extended Clayton copula density is equal to the Clayton copula density for non negative Kendall's τ and equal to a 90 degree rotation of the Clayton copula density for negative Kendall's τ . Among the different DGPs, most distinct values are considered

for ϕ . We expect its choice to be influential since it controls the dependence among the latent variables. With this grid we obtain 180 different DGPs. For each of the different DGPs we generate 100 simulated data sets and for each data set we run the 10 different samplers with the correctly specified copula family for 25000 iterations and discard the first 5000 iterations for burn-in. So, in total we obtain 18000 simulated data sets and each of the 10 samplers is run 18000 times.

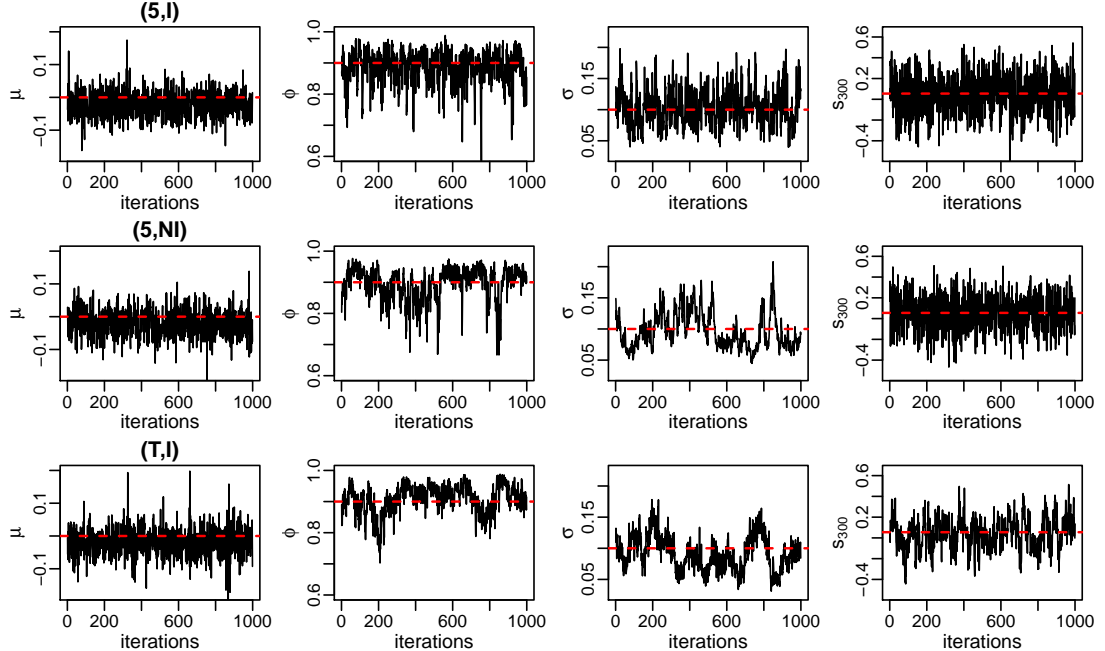


Figure 1: Trace plots of 1000 MCMC draws based on a total of 25000 iterations, where the first 5000 draws are discarded for burn-in and the remaining 20000 draws are thinned with factor 20. The trace plots are shown for the parameters μ , ϕ , σ and s_{300} for three different sampler specifications: (5,I) (top row), (5,NI) (middle row), (T,I) (bottom row). The corresponding data was generated from the following DGP: $T = 1000$, Family=eClayton, $\mu = 0$, $\phi = 0.9$, $\sigma = 0.1$. True values are added in red (dashed).

Figure 1 shows trace plots of different parameters (μ , ϕ , σ , s_{300}) based on one simulated data set for three different sampler specifications. We consider specification (5,I), and the same specification with interweaving turned off, i.e. (5,NI) and the specification (T,I). We observe that all three samplers produce posterior samples covering the true values. Further, the trace plots suggest that the (5,I) sampler achieves better mixing than the other two considered samplers.

The runtime of the sampler is mainly affected by the choice of T and the sampler specification. From Table 1 we see that the runtime is increasing in T and that interweaving adds considerable additional runtime.

	(1,I)	(5,I)	(20,I)	(100,I)	(T,I)	(1,NI)	(5,NI)	(20,NI)	(100,NI)	(T,NI)
$T = 500$	0.6	0.5	0.5	0.7	0.8	0.3	0.3	0.3	0.4	0.5
$T = 1000$	1.1	1.0	1.1	1.3	1.6	0.7	0.5	0.6	0.8	1.1
$T = 1500$	1.7	1.4	1.6	1.9	2.6	1.0	0.7	0.9	1.2	1.8

Table 1: Average runtime in minutes for 25000 draws. We consider averages for different sampler specifications and different values of T . The sampler was run on a Linux cluster with CPU Intel Xeon E5-2697 v3.

To measure efficiency of the samplers we consider the effective sample size per minute which we call effective sampling rate, similar to Hosszejni and Kastner (2019). We average the effective sampling rate of 100 runs, where the same sampler and the same DGP was used. Then we obtain 180 average effective sampling rates (AESR) per sampler. Since the AESR decreases for higher

values of T for every sampler, we compare the ten different samplers among DGPs with the same value for T . For a fixed T we have 60 AESR values per sampler. For each sampler, the minimum of these 60 values (mAESR) is given in Table 2. We consider the minimum since we are interested in samplers that are reliable for all DGPs. We see that for the parameters μ , ϕ and σ sampler specifications with interweaving have higher mAESR values, while for the latent states specifications without interweaving perform better. Although interweaving adds additional runtime, it still increases the mAESR of μ , ϕ and σ considerably. Further we observe that choosing the blocksize too big results in very low mAESR values for the latent states. In this case many parameters are updated jointly with elliptical slice sampling which results in high autocorrelation among consecutive draws. The performance of samplers with blocksize T is especially poor. For these samplers the length of the time series ($T = 500, 1000, 1500$) also has strong effects. For example the mAESR for μ for specification (T,NI) decreases by 85 % from 78 to 11 when the length of the time series is increased from 500 to 1500. The most inefficient sampler is (T,NI), the sampler without interweaving and with the largest blocksize. In our opinion the best results are obtained for sampler specification (5,I). It provides the highest mAESR values for μ , ϕ and σ for all choices of T and also provides rather high mAESR values for the latent states.

	(1,I)	(5,I)	(20,I)	(100,I)	(T,I)	(1,NI)	(5,NI)	(20,NI)	(100,NI)	(T,NI)
$T = 500$										
μ	2829	3178	1813	642	381	437	838	780	277	78
ϕ	471	749	556	191	60	266	326	232	78	24
σ	272	411	263	58	33	53	74	56	23	6
$s(a)$	938	2982	2484	199	36	1347	4942	3908	273	49
$s(m)$	346	1092	441	35	5	496	1011	627	41	5
$T = 1000$										
μ	1167	1365	840	309	180	172	330	371	121	21
ϕ	202	269	212	73	17	81	97	76	27	6
σ	97	175	133	25	9	18	25	21	12	1
$s(a)$	298	1011	1312	99	12	400	1911	1982	133	16
$s(m)$	115	417	200	15	1	161	559	296	17	2
$T = 1500$										
μ	711	833	578	200	111	89	218	231	75	11
ϕ	120	145	107	43	9	38	49	42	14	3
σ	63	109	84	16	3	10	14	13	8	1
$s(a)$	162	583	879	66	7	247	1168	1359	90	10
$s(m)$	61	239	112	9	1	95	426	176	11	1

Table 2: For different lengths of the time series ($T = 500, 1000, 1500$) the minimum of 60 AESR values (mAESR) corresponding to 60 different DGPs is shown for different parameters and ten different sampler specifications. In the $s(a)$ row we calculate the mAESR based on the average of the effective sampling rates of s_0, \dots, s_T , while in the $s(m)$ row the mAESR values are calculated based on the minimum of the effective sampling rates of s_0, \dots, s_T .

In addition to the previous analysis, we investigate how different DGPs affect the samplers. Therefore we consider the best sampler according to Table 2, i.e. sampler specification (5,I). In addition we have a look at the same sampler specification without interweaving (5,NI) and the same sampler with joint updates of the latent states, i.e. (T,I). We consider the AESR for σ , which is usually the parameter which causes most problems. Table 3 shows for each of these three samplers the DGPs with $T = 1000$ which resulted in the lowest and the highest AESR for σ . For the (5,I) specification satisfactory AESR values, ranging from 175 to 456, are obtained for all DGPs. In the (T,I) specification the latent states are updated jointly. In scenarios with strong dependence among the latent states ($\phi = 0.99$) this sampler performs poorly whereas the best performance of the sampler was seen for a DGP with low dependence among the latent states ($\phi=0$). The (5,NI) specification is a standard Gibbs sampler in the sufficient augmentation. We see that for this specification the AESR values vary a lot, ranging from 25 to 558. The best performance of this sampler specification was seen for a DGP with high persistence ($\phi = 0.99$). Kastner and Frühwirth-Schnatter (2014) studied different sampling schemes for the stochastic volatility model and have also seen that a standard Gibbs sampler in the sufficient augmentation performs well for scenarios with high persistence.

Lastly we compare our sampler to the coarse grid sampler employed by Almeida and Czado

AESR(σ)		(5,I)		(T,I)		(5,NI)	
		min	max	min	max	min	max
		175	456	9	121	25	558
DGP	family	Gauss	eClayton	eClayton	eClayton	Gauss	eClayton
	μ	0	1	0	0	0	0
	ϕ	0.1	0.9	0.99	0	0.1	0.99
	σ	0.2	0.2	0.2	0.05	0.05	0.2

Table 3: For three sampler specifications ((5,I), (T,I), (5,NI)) we show the DGPs with $T = 1000$ which resulted in the highest and in the lowest AESR values for σ , respectively. The corresponding AESR is also shown.

(2012). We already covered six DGPs that were also analyzed by Almeida and Czado (2012). Instead of running their sampler we make the comparison with respect to these six cases. There are several points which make the comparison slightly less reliable. First Almeida and Czado (2012) did not report exact computation times but they note that 100 000 iterations of their sampler take about 15 minutes. We use this number to calculate the AESR values from the effective sample sizes they report in their paper. Second their calculations were performed on a different computer and third they use different prior distributions for ϕ and σ^2 . But we think that this comparison should still give us a rough idea of how the sampling efficiencies compare to each other. From Table 4 we see that the (5,I) specification considerably outperforms the coarse grid sampler (CG). For every parameter (μ , ϕ , σ) we obtain way higher AESR values.

family	DGP			AESR(μ)		AESR(ϕ)		AESR(σ)	
	μ	ϕ	σ	(5,I)	CG	(5,I)	CG	(5,I)	CG
Gauss	1	0.9	0.1	8187	2130	527	81	393	50
Gauss	1	0.1	0.2	1637	305	364	83	437	69
Gauss	0	0.9	0.2	9935	4150	549	150	316	85
eClayton	1	0.9	0.1	8382	2088	539	74	397	46
eClayton	1	0.1	0.2	1778	279	378	71	416	52
eClayton	0	0.9	0.2	10079	4375	574	175	323	97

Table 4: Comparison of sampler specification (5,I) and the coarse grid sampler (CG) of Almeida and Czado (2012) for six different DGPs with $T = 1000$ with respect to the AESR of μ , ϕ and σ .

4 Application: Modeling the volatility return relationship

We investigate the volatility return relationship through the bivariate joint distribution of a stock index and the corresponding volatility index. The joint distribution of return and volatility incorporates all the marginal information as well as information about the dependence, which are both relevant for hedging and risk management (Allen et al (2012)). Since there has already been evidence for asymmetry in the joint distribution of volatility and return (Allen et al (2012), Fink et al (2017)) models which are able to handle these characteristics are necessary.

The two step copula modeling approach motivated by Sklar's theorem (Sklar (1959)) provides a very flexible method for the construction of multivariate distributions. We can combine arbitrary marginal distributions with any copula. Here, we propose a bivariate model that combines the skew Student t stochastic volatility model (Abanto-Valle et al (2015)) for the margins with a novel dynamic mixture copula. This model allows for asymmetry and heavy tails in the marginal distribution as well as for time varying asymmetric tail dependence in the dependence structure. Both, the marginal as well as the copula model can be estimated with the proposed sampler.

Marginal model

The stochastic volatility model with skew Student t errors is obtained by replacing the normal distribution of the stochastic volatility model (Kim et al (1998)) by a skew Student t distribution. This allows for heavy tails and skewness. The stochastic volatility model with skew

Student t errors as considered by Abanto-Valle et al (2015) is given by

$$\begin{aligned} Y_t &= \exp\left(\frac{s_t}{2}\right)\epsilon_t \\ s_t &= \mu + \phi(s_{t-1} - \mu) + \sigma\epsilon_t, \end{aligned} \quad (8)$$

where $\epsilon_t|\alpha, df \sim sst(\epsilon_t|\alpha, df)$ independently for $t = 1, \dots, T$. We denote by $sst(\epsilon_t|\alpha, df)$ the density of the standardized skew Student t distribution with parameters $\alpha \in \mathbb{R}$ and $df > 2$ (cf. Appendix B). Compared to our framework this model has two additional parameters α and df . For these additional parameters we choose the following prior distributions

$$\alpha \sim N(0, 100), \quad df \sim N_{>2}(5, 25), \quad (9)$$

where $N_{>2}$ denotes the normal distribution truncated to $(2, \infty)$. We need to ensure that $df > 2$ since the standardized skew Student t distribution would not be well defined otherwise.

Conditional on α and df our sampler can be applied directly to sample $(\mu, \phi, \sigma, \mathbf{s}_{0:T})$ from its full conditional. Another approach, which lead to better mixing, is to include the parameters α and df in the interweaving strategy. The sampler is slightly modified in the following way:

- a) Sample $\mathbf{s}_{0:T}$ from $\mathbf{s}_{0:T}|Y, \mu, \phi, \sigma, \alpha, df$.
- b) Sample $(\mu, \phi, \sigma, \alpha, df)$ in (SA) from $\mu, \phi, \sigma, \alpha, df|Y, \mathbf{s}_{0:T}$.
- c) Move to (AA) via $\tilde{s}_t = \frac{s_t - \mu - \phi(s_{t-1} - \mu)}{\sigma}$, for $t = 1, \dots, T$.
- d) Sample $(\mu, \phi, \sigma, \alpha, df)$ in (AA) from $\mu, \phi, \sigma, \alpha, df|Y, s_0, \tilde{\mathbf{s}}_{1:T}$.
- e) Move back to (SA) via the recursion $s_t = \mu + \phi(s_{t-1} - \mu) + \sigma\tilde{s}_t$ for $t = 1, \dots, T$.

For step a) we proceed as described in Section 2. For step b) we draw α and df from its univariate full conditional distributions using Metropolis-Hastings, similar to Step d) in Section 2. The parameters (μ, ϕ, σ) are drawn from its full conditional as described in Section 2. For step d) we investigated different blocking strategies for the parameters $(\mu, \phi, \sigma, \alpha, df)$. We compared the different strategies with respect to effective sample sizes and decided to use the following three blocks: (μ, df) , (ϕ, σ) and α . Each block is updated using Metropolis-Hastings as in Step d) in Section 2.

Dependence model

Dependence among financial assets is often modeled with a Student t copula. This copula allows for tail dependence symmetric in the upper and lower tail. Evidence against the assumption of symmetric tail dependence has been provided and models to handle this characteristic have become necessary (Patton (2006), Nikoloulopoulos et al (2012), Jondeau (2016)). Patton (2006) proposes the symmetrized Joe-Clayton copula. This is a modification of the BB7 copula (Joe (2014), Chapter 4) that is symmetric if upper and lower tail dependence coincide, which he describes as a desirable property. In the application of Nikoloulopoulos et al (2012) the Student t copula provides the best fit in terms of the likelihood. But they argue that if the focus is on the tails a BB1 or BB7 copula might be more appropriate. The BB1 and BB7 copulas have two parameters which might not be enough, if we want to model three characteristics in a flexible way: upper tail dependence, lower tail dependence and overall dependence as measured with Kendall's τ . We provide another approach to relax the symmetric tail dependence assumption. We propose a mixture of a Student t and a extended Gumbel copula with parameters $\tau \in (-1, 1)$, $\nu > 2$ and $p \in [0, 1]$ given by

$$C^M(u_1, u_2; \tau, \nu, p) = pC^t(u_1, u_2; \tau, \nu) + (1 - p)C^G(u_1, u_2; \tau), \quad (10)$$

where C^t is the bivariate Student t copula specified by Kendall's τ and the degree of freedom ν and C^G is the bivariate extended Gumbel copula specified by Kendall's τ . The extended Gumbel copula is defined similarly to the extended Clayton copula in Section 2, i.e. its density is equal to the Gumbel copula density for positive values of Kendall's τ and equal to a 90 degree rotation of the Gumbel copula density for negative Kendall's τ values. Both copulas C^t and C^G

share the dependence parameter τ and we expect the mixture copula to have a similar strength of dependence. The corresponding Kendall's τ of the mixture copula is given by

$$\begin{aligned}\tau^M &= \int_{(0,1)^2} C^M(u_1, u_2; \tau, \nu, p) c^M(u_1, u_2; \tau, \nu, p) du_1 du_2 = (p^2 + (1-p)^2)\tau + \\ &+ \int_{(0,1)^2} p(1-p) (C^G(u_1, u_2; \tau))(c^t(u_1, u_2; \tau, \nu) + C^t(u_1, u_2; \tau, \nu)c^G(u_1, u_2; \tau)) du_1 du_2.\end{aligned}\tag{11}$$

We evaluated the integral in (11) numerically for different values of τ , p and ν and observed only negligible difference between τ and τ^M . The upper and lower tail dependence coefficients λ_M^L and λ_M^U of the mixture copula can, for $\tau > 0$, be obtained as

$$\begin{aligned}\lambda_M^L(\tau, p, \nu) &= \lim_{u \rightarrow 0} \frac{C^M(u, u)}{u} = \lim_{u \rightarrow 0} \frac{pC^t(u_1, u_2; \tau, \nu) + (1-p)C^G(u_1, u_2; \tau)}{u} \\ &= p2T_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\sin(\frac{\pi}{2}))}{1+\sin(\frac{\pi}{2})}}\right) + 0 \\ \lambda_M^U(\tau, p, \nu) &= p2T_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\sin(\frac{\pi}{2}))}{1+\sin(\frac{\pi}{2})}}\right) + (1-p)(2-2^{1-\tau}),\end{aligned}$$

where we used the well known formulas for the tail dependence coefficients of the Student t and the Gumbel copula (Joe (2014), Chapter 4). Whereas the upper and lower tail dependence coefficients measure dependence in the upper right and lower left corner, we are also interested in the dependence in the upper left and the lower right corner when $\tau < 0$. We consider the following tail dependence coefficients in the upper left corner λ_M^{UL} and in the lower right corner λ_M^{LR} if $\tau < 0$

$$\lambda_M^{LR} = \lambda_M^L(-\tau, p, \nu), \quad \lambda_M^{UL} = \lambda_M^U(-\tau, p, \nu),$$

analogous to the definition of quarter tail dependence in Fink et al (2017).

The tail dependence coefficient of the mixture copula is a linear combination of the tail dependence coefficients of its two components, the Student t and the Gumbel copula. The Student t copula has symmetric tail dependence, whereas the Gumbel copula has upper but no lower tail dependence. So we expect upper tail dependence to be higher than lower tail dependence in the mixture copula. The amount of asymmetry in the tails is controlled by p , whereas the copula is symmetric in the tails for $p = 1$ and the level of asymmetry increases as we decrease p . So this copula allows for great flexibility: The overall dependence can be described by Kendall's τ , the degrees of freedom parameter controls the upper and lower tail dependence coefficient and p controls the difference between upper and lower tail dependence. This is visualized in Figure 4 in Appendix C. Note that the desirable property according to Patton (2006) of symmetry in case of coinciding upper and lower tail dependence is here fulfilled. If we expected higher lower than upper tail dependence we can replace the Gumbel copula by a survival Gumbel copula which has the density $c^{SG}(u_1, u_2) = c^G(1-u_1, 1-u_2)$. To allow for time variation we use the mixture copula C^M of (10) within the dynamic bivariate copula model of Almeida and Czado (2012). A nonlinear state space model for T bivariate random vectors $(U_{t1}, U_{t2})_{t=1, \dots, T} \in [0, 1]^{T \times 2}$, corresponding to T time points, is given by

$$\begin{aligned}(U_{t1}, U_{t2}) &\sim c^M(u_{t1}, u_{t2}; \tau_t, \nu, p) \text{ independently} \\ s_t &= \mu + \phi(s_{t-1} - \mu) + \sigma \epsilon_t, \quad s_t = F_Z(\tau_t)\end{aligned}\tag{12}$$

for $t = 1, \dots, T$. We assign a uniform prior on $[0, 1]$ for p , a normal prior with mean 5 and standard deviation 20 truncated to the interval $(2, \infty]$ for ν and the same priors as in (5) for the remaining parameters. Sampling is done in the following way.

- Draw $\log\left(\frac{p}{1-p}\right)$ and $\log(\nu - 2)$ from its univariate full conditionals with random walk Metropolis-Hastings with Gaussian proposal (proposal standard deviation: 0.3).
- Draw $\mu, \phi, \sigma, \mathbf{s}_{0:T}$ conditioned on p and ν as in Section 2.

Two step estimation

We consider the S&P500 (SPX) and its volatility index the VIX as well as the DAX and its volatility index the VDAX. The daily log returns from 2006 to 2013 of these indices are obtained from Yahoo finance (<https://finance.yahoo.com>). With approximately 250 trading days per year this results in 2063 observations, visualized in Figure 6 in Appendix D. The corresponding data matrix with 2063 rows and 4 columns is denoted by Y .

Combining the marginal and the dependence model we obtain that for T bivariate random vectors $(Y_{t1}, Y_{t2})_{t=1, \dots, T} \in \mathbb{R}^{T \times 2}$ the following holds

$$(Y_{t1}, Y_{t2}) \sim C^M \left(ssT \left(\frac{y_{t1}}{\exp(s_{t1}^{st}/2)} \middle| \alpha_1^{st}, df_1^{st} \right), ssT \left(\frac{y_{t2}}{\exp(s_{t2}^{st}/2)} \middle| \alpha_2^{st}, df_2^{st} \right); F_Z^{-1}(s_t^{cop}), \nu^{cop}, p^{cop} \right) \quad (13)$$

independently, where

$$\begin{aligned} s_{tj}^{st} &= \mu_j^{st} + \phi_j^{st}(s_{t-1;j}^{st} - \mu_j^{st}) + \sigma_j^{st} \epsilon_{tj}^{st} \\ s_t^{cop} &= \mu^{cop} + \phi^{cop}(s_{t-1}^{cop} - \mu^{cop}) + \sigma^{cop} \epsilon_t^{cop} \end{aligned}$$

and $\epsilon_{tj}^{st}, \epsilon_t^{cop} \sim N(0, 1)$ iid, α_j^{st}, df_j^{st} as in (9), ν^{cop}, p^{cop} as in (12) and $\mu_j^{st}, \mu^{cop}, \phi_j^{st}, \phi^{cop}, \sigma_j^{st}, \sigma^{cop}, s_{0j}^{st}, s_0^{cop}$ as in (2) and (5) for $j = 1, 2$ and $t = 1, \dots, T$. Here, ssT denotes the distribution function of the standardized skew Student t distribution (cf. Appendix B). We refer to the probability integral transforms $ssT(y_{t1} \exp(-s_{t1}^{st}/2) | \alpha_1^{st}, df_1^{st})$ and $ssT(y_{t2} \exp(-s_{t2}^{st}/2) | \alpha_2^{st}, df_2^{st})$ for $t = 1, \dots, T$ as copula data.

For inference we rely on a two step approach. We first estimate marginal distributions and based on the resulting estimated copula data we estimate the copula parameters. This approach is also called inference for margins (Joe and Xu (1996)) and is commonly used in (Bayesian) copula modeling (Min and Czado (2011), Almeida and Czado (2012), Smith (2015), Gruber et al (2015), Loaiza-Maya et al (2018)).

First we fit a skew Student t stochastic volatility model for each of the indices. For each index we run the sampler (5,I) for 31000 iterations and discard the first 1000 draws as burn-in. As it is typical for financial data all indices show a high persistence parameter ϕ (Posterior mode estimates for ϕ : SPX: 0.99, VIX: 0.90, DAX: 0.99, VDAX: 0.96). A notable difference is that for stock indices we observe negative skewness, whereas for the volatility indices positive skewness is observed (Posterior mode estimates for α : SPX: -0.51 , VIX: 1.33, DAX: -0.48 , VDAX: 0.96). Evidence for negative skewness has also been observed for the log returns of other stock indices, as e.g., for the NASDAQ by Abanto-Valle et al (2015). Posterior mode estimates, posterior quantiles and effective samples sizes for several parameters of the four marginal models are summarized in Table 7 in Appendix D. The estimated daily log variances are shown in Figure 2. In the end of 2009, the estimated variances are high for all indices due to the financial crisis.

In the next step we obtain data on the $[0,1]$ scale by applying the probability integral transform using the posterior mode estimates of the marginal parameters. We refer to this data as pseudo copula data and it is obtained as

$$\hat{u}_{tj} = ssT \left(y_{tj} \exp \left(-\frac{\hat{s}_{tj}^{st}}{2} \right); \hat{\alpha}_j^{st}, \hat{df}_j^{st} \right),$$

where $\hat{s}_{tj}^{st}, \hat{\alpha}_j^{st}, \hat{df}_j^{st}$ are the posterior mode estimates of the corresponding marginal skew Student t stochastic volatility model for $t = 1, \dots, T, j = 1, \dots, 4$. In the copula data marginal characteristics are removed and what is left is information about the dependence structure. Based on the pseudo copula data two dynamic mixture copula models are fitted, one corresponding to the pair (SPX,VIX) and one corresponding to the pair (DAX,VDAX). For each pair we obtain 31000 iterations with the sampler specification (5,I) and discard the first 1000 draws as burn-in. The posterior mode estimates for p are 0.29 for the model for (SPX,VIX) and 0.66 for the model for (DAX,VDAX), respectively. (Further, posterior statistics for the

model parameters μ, ϕ, σ, p and ν are shown in Table 8 in Appendix D). So both fitted models allow for asymmetric tail dependence, whereas the asymmetry is stronger for the (SPX,VIX) model. For these models tail dependence in the upper left corner λ_M^{UL} is stronger than the one in the lower right corner λ_M^{LR} . This means that joint extreme comovements, where the stock index decreases and the volatility index increases are more likely to occur than vice versa, which agrees with the statement that the market reacts more extreme in bad market situations (Sun and Wu (2018)).




Figure 2: Posterior mode estimates of the daily log variances of the four skew Student t stochastic volatility models for the indices SPX, VIX, DAX, VDAX from 2006 to 2013 plotted against time. The 90% credible region, added in grey, is constructed from the 5% and 95% posterior quantiles.

The time varying estimates for Kendall's τ and the tail dependence coefficients are shown in Figure 3. The figure visualizes the asymmetry in tail dependence. We also observe changes in tail dependence as time evolves: for (SPX,VIX), λ_M^{UL} ranges from 0.41 to 0.71 and λ_M^{LR} from 0.004 to 0.10. For the pair (DAX,VDAX), λ_M^{UL} ranges from 0.12 to 0.67 and λ_M^{LR} from 0.003 to 0.35. This variation over time in tail dependence goes hand in hand with variation in Kendall's τ . For (SPX,VIX), Kendall's τ ranges from -0.76 to -0.48 and for (DAX,VDAX) from -0.80 to -0.30 .




Figure 3: In these plots the top row corresponds to the pair (SPX,VIX), the bottom row to the pair (DAX,VDAX). The first column shows posterior mode estimates of Kendall's τ , where the Kendall's τ obtained from rolling window estimates (Kendall's τ at time t is estimated as empirical Kendall's τ based on the 50 observations before and after time t) is added in red. The middle column shows posterior mode estimates of λ_M^{UL} and the right column shows estimates of λ_M^{LR} . Credible regions, constructed from the 5% and 95% posterior quantiles, are added in grey.

Out of sample predictions

We aim to further support the findings we obtained through the dynamic copula model, i.e. that the dependence structure is asymmetric and varies over time. Therefore we consider several restrictions with respect to the dependence structure. Giving up time variation leads a constant mixture copula, giving up asymmetry leads to a dynamic Student t copula and giving up time variation and asymmetry leads to a constant Student t copula. In addition, we compare our model to the frequently used dynamic conditional correlation (DCC) GARCH model of Engle (2002). The DCC-GARCH allows for time varying symmetric dependence. So we take five different models into consideration. These models are summarized in Table 5. The model \mathcal{M}_{dyn}^{mix} is the model given in (13), which was also used for the previous analysis.

We compare the models with respect to cumulative pseudo log predictive scores (Kastner (2019)), which are obtained by evaluating the corresponding density at point estimates, instead of averaging over all posterior draws. In comparison to other multivariate scoring rules, such as the energy score or the variogram score (Scheuerer and Hamill (2015)), pseudo log predictive scores have here the advantage that they can be computed fast since they require only one evaluation of the density per observation. We consider $T + K$ observations of dimension two, stored in the data matrix $Y_{1:(T+K);1:2}$, where the first T observations are used to train the model and the last K are used for testing.

model	specification	dependence		
		margin asymmetric	asymmetric	dynamic
\mathcal{M}_{dyn}^{mix}	sstSV + dynamic mixture copula	Yes	Yes	Yes
$\mathcal{M}_{const}^{mix}$	sstSV + constant mixture copula	Yes	Yes	No
\mathcal{M}_{dyn}^t	sstSV + dynamic Student t copula	Yes	No	Yes
\mathcal{M}_{const}^t	sstSV + constant Student t copula	Yes	No	No
\mathcal{M}_{DCC}	DCC(1,1)-GARCH(1,1)	No	No	Yes

Table 5: Different models considered for comparison. Models are specified by: marginal model + copula model. The skew Student t stochastic volatility model as given in (8) is denoted by sstSV. The mixture copula is defined in (10). If a copula is dynamic, the corresponding copula is considered within the dynamic bivariate copula model framework of Almeida and Czado (2012) given in (4).

The DCC-GARCH model is estimated based on T data points in the training period. Based on this model, we obtain rolling one-day ahead estimates of the covariance matrix for each day in the test set. The pseudo log predictive scores are obtained by evaluating the corresponding multivariate normal log densities at the observations. To estimate the DCC-GARCH models we used the R package `rmgarch` of Ghalanos (2012).

Similarly, the model \mathcal{M}_{dyn}^{mix} is estimated with the training data. Instead of computing daily updates for all model parameters we fix the constant parameters at their posterior mode estimates to save computation time. The dynamic parameters are updated daily and the one-day ahead forecasts are obtained by evolving the AR(1) process. To obtain the pseudo log predictive scores we evaluate the density implied by (13) at the corresponding observations. Appendix D contains a detailed description of this procedure. The pseudo log predictive scores for $\mathcal{M}_{const}^{mix}$, \mathcal{M}_{dyn}^t and \mathcal{M}_{const}^t are obtained similarly.

This procedure for calculating the pseudo log predictive scores is applied for both data sets corresponding to the pairs (SPX,VIX) and (DAX,VDAX). Using the last two years (2012 - 2013) of our data set as test data yields $K = 517$. As training period we use $T = 1000$ which corresponds to a training period of approximately four years.

Table 6 summarizes the cumulative pseudo log predictive scores. In both cases, for the (SPX,VIX) as well as for the (DAX,VDAX) data, the best model is provided by the dynamic mixture copula model \mathcal{M}_{dyn}^{mix} . Furthermore, we see that in both cases the constant mixture copula model $\mathcal{M}_{const}^{mix}$ is preferred over the constant and dynamic Student t copula models \mathcal{M}_{const}^t and \mathcal{M}_{dyn}^t . For the (DAX,VDAX) data, the second best model is provided by the constant mixture copula model $\mathcal{M}_{const}^{mix}$. For this data the rolling window estimates of Kendall's τ in Figure 3 vary less than for the (SPX,VIX) data. For the (SPX,VIX) data, the DCC-

GARCH \mathcal{M}_{DCC} yields the second best cumulative pseudo log predictive scores.

	\mathcal{M}_{dyn}^{mix}	$\mathcal{M}_{const}^{mix}$	\mathcal{M}_{dyn}^t	\mathcal{M}_{const}^t	\mathcal{M}_{DCC}
(SPX,VIX)	2740.7	2733.5	2732.0	2725.9	2737.2
(DAX,VDAX)	2817.7	2814.1	2810.1	2809.9	2789.0

Table 6: Cumulative pseudo log predictive scores for the models \mathcal{M}_{dyn}^{mix} , $\mathcal{M}_{const}^{mix}$, \mathcal{M}_{dyn}^t , \mathcal{M}_{const}^t and \mathcal{M}_{DCC} .

5 Conclusion

We propose a sampler, applicable to general nonlinear state space models with univariate autoregressive state equation. Sampling efficiency is demonstrated for bivariate dynamic copula models within a simulation study. Furthermore we use the sampler to estimate the parameters of a dynamic bivariate mixture copula model. This mixture copula model turns out to be a good candidate to model the volatility return relationship, since in our application it produces more accurate forecasts than a bivariate DCC-GARCH model or a Student t copula model.

In this work, there are two objectives that might be extended: The sampler and the bivariate mixture copula model. The sampler could be extended to allow for a broader class of models. For example we might consider autoregressive processes of higher order in the state equation. In this case we can still rely on elliptical slice sampling and on an interweaving strategy. Another extension could relax the assumption of a Gaussian dependence structure in the state equation by replacing the autoregressive process by a D-vine copula model. In this case elliptical slice sampling can no longer be applied to sample the latent states and an alternative sampling method is required.

The bivariate dynamic mixture copula could serve as a building block for regular vine copula models. Thus we could extend the bivariate model to arbitrary dimensions. This is interesting if we study not only the bivariate volatility return relationship, but for example the dependence structure among several exchange rates.

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
- Abanto-Valle C, Lachos V, Dey DK (2015) Bayesian estimation of a skew-student-t stochastic volatility model. *Methodology and Computing in Applied Probability* 17(3):721–738
- Allen D, Singh A, Powell R, McAleer M, Taylor J (2012) The Volatility-Return Relationship: Insights from Linear and Non-Linear Quantile Regressions. School of Accounting. Finance and Economics & FEMARC Working Paper Series
- Almeida C, Czado C (2012) Efficient Bayesian inference for stochastic time-varying copula models. *Computational Statistics & Data Analysis* 56(6):1511–1527
- Azzalini A, Capitanio A (2003) Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65(2):367–389
- 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
- Bates D, Eddelbuettel D, et al (2013) Fast and elegant numerical linear algebra using the RcppEigen package. *Journal of Statistical Software* 52(5):1–24

- 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
- Bennett J, Grout R, Pébay P, Roe D, Thompson D (2009) Numerically stable, single-pass, parallel statistics algorithms. In: 2009 IEEE International Conference on Cluster Computing and Workshops, IEEE, pp 1–8
- Bollerslev T (1986) Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3):307–327
- 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
- Brockwell PJ, Davis RA, Calder MV (2002) Introduction to time series and forecasting, vol 2. Springer
- Carlin BP, Polson NG, Stoffer DS (1992) A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *Journal of the American Statistical Association* 87(418):493–500
- Chan JC, Grant AL (2016) Modeling energy price dynamics: GARCH versus stochastic volatility. *Energy Economics* 54:182–189
- Czado C (2019) Analyzing Dependent Data with Vine Copulas. *Lecture Notes in Statistics*, Springer
- Eddelbuettel D, François R, Allaire J, Ushey K, Kou Q, Russel N, Chambers J, Bates D (2011) Rcpp: Seamless R and C++ integration. *Journal of Statistical Software* 40(8):1–18
- 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
- Engle RF (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society* 50:987–1007
- Engle RF, Sheppard K (2001) Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. Tech. rep., National Bureau of Economic Research
- Fink H, Klimova Y, Czado C, Stöber J (2017) Regime switching vine copula models for global equity and volatility indices. *Econometrics* 5(1):3
- Frühwirth-Schnatter S, Sögner L (2003) Bayesian estimation of the Heston stochastic volatility model. In: *Operations Research Proceedings 2002*, Springer, pp 480–485
- 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
- Geman S, Geman D (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6(6):721–741
- Ghalanos A (2012) rmgarch: Multivariate GARCH models. R package version 098
- Gruber L, Czado C, et al (2015) Sequential Bayesian model selection of regular vine copulas. *Bayesian Analysis* 10(4):937–963
- Hafner CM, Manner H (2012) Dynamic stochastic copula models: Estimation, inference and applications. *Journal of Applied Econometrics* 27(2):269–295
- Hahn PR, He J, Lopes HF (2019) Efficient sampling for Gaussian linear regression with arbitrary priors. *Journal of Computational and Graphical Statistics* 28(1):142–154
- Hosszejni D, Kastner G (2019) Approaches Toward the Bayesian Estimation of the Stochastic Volatility Model with Leverage. arXiv preprint arXiv:190111491

- Joe H (2014) *Dependence modeling with copulas*. CRC Press
- Joe H, Xu JJ (1996) The estimation method of inference functions for margins for multivariate models. Technical Report 166, Department of Statistics, University of British Columbia
- Jondeau E (2016) Asymmetry in tail dependence in equity portfolios. *Computational Statistics & Data Analysis* 100:351–368
- 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 (2014) Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis* 76:408–423
- 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
- Kim S, Shephard N, Chib S (1998) Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies* 65(3):361–393
- Loaiza-Maya R, Smith MS, Maneesoonthorn W (2018) Time series copulas for heteroskedastic data. *Journal of Applied Econometrics* 33(3):332–354
- Min A, Czado C (2011) Bayesian model selection for D-vine pair-copula constructions. *Canadian Journal of Statistics* 39(2):239–258
- 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 5(0)
- Nagler T, Bumann C, Czado C (2019) Model selection in sparse high-dimensional vine copula models with an application to portfolio risk. *Journal of Multivariate Analysis* 172:180–192
- Nakajima J, Kasuya M, Watanabe T (2011) Bayesian analysis of time-varying parameter vector autoregressive model for the Japanese economy and monetary policy. *Journal of the Japanese and International Economies* 25(3):225–245
- Neal RM (1998) Regression and classification using Gaussian process priors. *Bayesian statistics* 6:475
- Neal RM, et al (2003) Slice sampling. *The Annals of Statistics* 31(3):705–767
- Nikoloulopoulos AK, Joe H, Li H (2012) Vine copulas with asymmetric tail dependence and applications to financial return data. *Computational Statistics & Data Analysis* 56(11):3659–3673
- Patton AJ (2006) Modelling asymmetric exchange rate dependence. *International Economic Review* 47(2):527–556
- Plummer M, Best N, Cowles K, Vines K (2008) *coda*: Output analysis and diagnostics for MCMC. R package version 013-3
- Primiceri GE (2005) Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3):821–852
- Robbins H, Monro S (1985) A stochastic approximation method. In: *Herbert Robbins Selected Papers*, Springer, pp 102–109

- Roberts GO, Gelman A, Gilks WR, et al (1997) Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability* 7(1):110–120
- Roberts GO, Rosenthal JS, et al (2001) Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science* 16(4):351–367
- Scheuerer M, Hamill TM (2015) Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Monthly Weather Review* 143(4):1321–1334
- Schwert GW (1989) Why does stock market volatility change over time? *The Journal of Finance* 44(5):1115–1153
- Sklar A (1959) Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris* 8:229–231
- Smith MS (2015) Copula modelling of dependence in multivariate time series. *International Journal of Forecasting* 31(3):815–833
- Strickland CM, Martin GM, Forbes CS (2008) Parameterisation and efficient MCMC estimation of non-Gaussian state space models. *Computational Statistics & Data Analysis* 52(6):2911–2930
- Sun Y, Wu X (2018) Leverage and Volatility Feedback Effects and Conditional Dependence Index: A Nonparametric Study. *Journal of Risk and Financial Management* 11(2):29
- 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
- Yu J (2002) Forecasting volatility in the New Zealand stock market. *Applied Financial Economics* 12(3):193–202
- 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

Appendix A. Details on the sampling procedure

Sampling of the latent states in the sufficient augmentation

Here we derive $\boldsymbol{\mu}_{\mathbf{B}_i|}$ and $\Sigma_{\mathbf{B}_i|}$. By the conditional independence assumptions of the AR(1) process and the way we defined the blocks B_1, \dots, B_m we obtain

$$\begin{aligned} f(\mathbf{s}_{\mathbf{B}_i}|s_0, \mathbf{s}_{-\mathbf{B}_i}, \mu, \phi, \sigma) &= f(\mathbf{s}_{\mathbf{B}_i}|s_{a_i-1}, s_{b_i+1}, \mu, \phi, \sigma), \text{ for } i = 1, \dots, m-1, \text{ and} \\ f(\mathbf{s}_{\mathbf{B}_m}|s_0, \mathbf{s}_{-\mathbf{B}_m}, \mu, \phi, \sigma) &= f(\mathbf{s}_{\mathbf{B}_m}|s_{a_m-1}, \mu, \phi, \sigma). \end{aligned}$$

Conditional on μ, ϕ and σ , the vector $(s_{a_i-1}, \mathbf{s}_{\mathbf{B}_i}, s_{b_i+1})$ is multivariate normal distributed with mean vector $\boldsymbol{\mu}_{(\mathbf{a}_i-1, \mathbf{B}_i, \mathbf{b}_i+1)}^{\text{AR}} \in \mathbb{R}^{c_i+2}$ and covariance matrix $\Sigma_{(\mathbf{a}_i-1, \mathbf{B}_i, \mathbf{b}_i+1); (\mathbf{a}_i-1, \mathbf{B}_i, \mathbf{b}_i+1)}^{\text{AR}} \in \mathbb{R}^{(c_i+2) \times (c_i+2)}$, where c_i is the cardinality of B_i . Thus the vector $\mathbf{s}_{\mathbf{B}_i}|s_{a_i-1}, s_{b_i+1}, \mu, \phi, \sigma$ follows a multivariate normal distribution with mean vector $\boldsymbol{\mu}_{\mathbf{B}_i|}$ and covariance matrix $\Sigma_{\mathbf{B}_i|}$ given by

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{B}_i|} &= \boldsymbol{\mu}_{\mathbf{B}_i}^{\text{AR}} + \Sigma_{\mathbf{B}_i; (\mathbf{a}_i-1, \mathbf{b}_i+1)}^{\text{AR}} \frac{1 - \phi^2}{(1 - \phi^{2(c_i+1)})\sigma^2} \begin{pmatrix} 1 & -\phi^{c_i+1} \\ -\phi^{c_i+1} & 1 \end{pmatrix} \begin{pmatrix} s_{a_i-1} - \mu \\ s_{b_i+1} - \mu \end{pmatrix}, \\ \Sigma_{\mathbf{B}_i|} &= \Sigma_{\mathbf{B}_i; \mathbf{B}_i}^{\text{AR}} - \Sigma_{\mathbf{B}_i; (\mathbf{a}_i-1, \mathbf{b}_i+1)}^{\text{AR}} \frac{1 - \phi^2}{(1 - \phi^{2(c_i+1)})\sigma^2} \begin{pmatrix} 1 & -\phi^{c_i+1} \\ -\phi^{c_i+1} & 1 \end{pmatrix} \Sigma_{(\mathbf{a}_i-1, \mathbf{b}_i+1); \mathbf{B}_i}^{\text{AR}}. \end{aligned} \quad (14)$$

The vector $\mathbf{s}_{\mathbf{B}_m}|s_{a_m-1}, \mu, \phi, \sigma$ corresponding to the last block is multivariate normal distributed with mean vector $\boldsymbol{\mu}_{\mathbf{B}_m|}$ and covariance matrix $\Sigma_{\mathbf{B}_m|}$ obtained as

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{B}_m|} &= \boldsymbol{\mu}_{\mathbf{B}_m}^{\text{AR}} + \Sigma_{\mathbf{B}_m; a_m-1}^{\text{AR}} \frac{1 - \phi^2}{\sigma^2} (s_{a_m-1} - \mu), \\ \Sigma_{\mathbf{B}_m|} &= \Sigma_{\mathbf{B}_m; \mathbf{B}_m}^{\text{AR}} - \Sigma_{\mathbf{B}_m; a_m-1}^{\text{AR}} \frac{1 - \phi^2}{\sigma^2} \Sigma_{a_m-1; \mathbf{B}_m}^{\text{AR}}. \end{aligned}$$

We need to sample from $N(\mathbf{0}, \Sigma_{\mathbf{B}_i|})$ several times during elliptical slice sampling. Instead of working with the $c_i \times c_i$ covariance matrix $\Sigma_{\mathbf{B}_i|}$ we can more efficiently sample from the c_i dimensional normal distribution by using the conditional independence assumptions of the AR(1) process. It holds that

$$\begin{aligned} f(\mathbf{s}_{\mathbf{B}_i}|s_{a_i-1}, s_{b_i+1}, \mu, \phi, \sigma) &= \prod_{t=0}^{c_i-1} f(s_{a_i+t}|s_{a_i-1}, s_{b_i+1}, \mu, \phi, \sigma) \\ &= \prod_{t=0}^{c_i-1} f(s_{a_i+t}|s_{a_i+t-1}, s_{b_i+1}, \mu, \phi, \sigma), \end{aligned}$$

where $f(s_{a_i+t}|s_{a_i+t-1}, s_{b_i+1}, \mu, \phi, \sigma)$ is the univariate normal density with mean

$$\mu_{a_i+t|a_i+t-1, b_i+1} = \frac{1}{1 - \phi^{2(c_i+1-t)}} \left((\phi - \phi^{2c_i+1-2t})(s_{a_i+t-1} - \mu) + (\phi^{c_i-t} - \phi^{c_i+2-t})(s_{b_i+1} - \mu) \right),$$

and variance

$$\sigma_{a_i+t|a_i+t-1, b_i+1}^2 = \frac{\sigma^2}{1 - \phi^2} \left(1 - \frac{1}{1 - \phi^{2(c_i+1-t)}} \left(\phi^2 - 2\phi^{2(c_i-t+1)} + \phi^{2(c_i-t)} \right) \right).$$

So we can sample $\mathbf{s}_{\mathbf{B}_i} = (s_{a_i+t})_{t=0, \dots, c_i-1}$ conditioned on $s_{a_i-1}, s_{b_i+1}, \mu, \phi, \sigma$ recursively by

$$s_{a_i+t} \sim N(\mu_{a_i+t|a_i+t-1, b_i+1}, \sigma_{a_i+t|a_i+t-1, b_i+1}^2),$$

for $t = 0, \dots, c_i - 1$ and then $\mathbf{s}_{\mathbf{B}_i} - \boldsymbol{\mu}_{\mathbf{B}_i|}$ is a sample from $N(\mathbf{0}, \Sigma_{\mathbf{B}_i|})$.

Sampling of the constant parameters in the ancillary augmentation

To sample μ , ϕ and σ in (AA) we deploy an adaptive random walk Metropolis-Hastings scheme as suggested by Garthwaite et al (2016), where tuning parameters are selected automatically using the Robbins Monro process (Robbins and Monro (1985)). For sampling, it is convenient to move to unconstrained parameter spaces which is achieved by the following transformations

$$\psi = \ln(\sigma), \quad \xi = F_Z(\phi).$$

Here $F_Z(x) = \frac{1}{2} \log\left(\frac{1+x}{1-x}\right)$ is Fisher's Z transformation. The from (5) implied log prior densities for ξ and ψ are given by

$$\begin{aligned} \ln(\pi(\xi)) &= (a_\phi - 1) \ln(F_Z^{-1}(\xi) + 1) + (b_\phi - 1) \ln(1 - F_Z^{-1}(\xi)) + \ln(1 - (F_Z^{-1}(\xi))^2) + c_1 \\ \ln(\pi(\psi)) &= -\psi - \frac{1}{2B_\sigma} \exp(2\psi) + 2\psi + c_2. \end{aligned}$$

where $c_1 \in \mathbb{R}$ and $c_2 \in \mathbb{R}$ are constants. The log posterior density in (AA) is obtained as

$$\begin{aligned} l_{p(AA)}(\mu, \xi, \psi, s_0, \tilde{\mathbf{s}}_{1:T} | Y) &= \sum_{t=1}^T \ln(f(\mathbf{y}_t | s_t(\tilde{\mathbf{s}}_{1:T}, \mu, F_Z^{-1}(\xi), \exp(\psi)))) - \frac{1}{2} \sum_{t=1}^T \tilde{s}_t^2 \\ &+ \ln\left(\varphi\left(s_0 | \mu, \frac{\exp(\psi)^2}{1 - F_Z^{-1}(\xi)^2}\right)\right) + \ln(\pi(\mu)) + \ln(\pi(\xi)) + \ln(\pi(\psi)) + c_3, \end{aligned}$$

where $c_3 \in \mathbb{R}$ is a constant. We sample (μ, ϕ, σ) in two blocks, one block for μ and one block for (ϕ, σ) .

Update for μ

To sample the mean parameter μ from its full conditional we propose a new state μ_{prop} in the r -th iteration of the MCMC procedure by

$$\mu_{prop} \sim N(\mu_{cur}, \sigma_{MH,\mu}^{r-1}),$$

where μ_{cur} is the current value for μ . The proposal μ_{prop} is accepted with probability

$$R = \exp(l_{p(AA)}(\mu_{prop}, \xi, \psi, s_0, \tilde{\mathbf{s}}_{1:T} | Y) - l_{p(AA)}(\mu_{cur}, \xi, \psi, s_0, \tilde{\mathbf{s}}_{1:T} | Y))$$

and the scaling parameter $\sigma_{MH,\mu}^r$ is updated according to Garthwaite et al (2016) by

$$\ln(\sigma_{MH,\mu}^r) = \ln(\sigma_{MH,\mu}^{r-1}) + 4.058 \frac{(R - 0.44)}{r - 1}.$$

The scaling parameter is reduced, if the acceptance probability is larger than 0.44 and increased if the acceptance probability is smaller than 0.44. We target an average acceptance probability of 0.44, as recommended by Roberts et al (2001) for univariate random walk Metropolis-Hastings. The constant 4.058 controls the step size and is chosen as suggested by Garthwaite et al (2016).

Joint update for ϕ and σ

In the r -th iteration, a two dimensional proposal $(\xi_{prop}, \psi_{prop})$ for (ξ, ψ) is obtained by

$$(\xi_{prop}, \psi_{prop})^\top \sim N_2((\xi_{cur}, \psi_{cur})^\top, \Sigma_{MH,\xi,\psi}^{r-1}),$$

where (ξ_{cur}, ψ_{cur}) are the current values. The proposal is accepted with probability

$$R = \exp(l_{p(AA)}(\mu, \xi_{prop}, \psi_{prop}, s_0, \tilde{\mathbf{s}}_{1:T} | Y) - l_{p(AA)}(\mu, \xi_{cur}, \psi_{cur}, s_0, \tilde{\mathbf{s}}_{1:T} | Y)).$$

For adapting the covariance matrix we follow a suggestion of Garthwaite et al (2016). Let I_n denote the n -dimensional identity matrix. We set $\Sigma_{MH,\xi,s}^r = I_2$ if $r < 100$ and

$$\Sigma_{MH,\xi,s}^r = (\sigma_{MH,\xi,s}^r)^2 \left(\hat{\Sigma}^r + \frac{(\sigma_{MH,\xi,s}^r)^2}{r} I_2 \right) \quad \text{if } r \geq 100.$$

Here $\hat{\Sigma}^r$ is the empirical covariance matrix of $(\xi^i, \psi^i)_{i=1, \dots, r}$, the first r samples for (ξ, ψ) , and

$$\ln(\sigma_{MH, \xi, s}^r) = \ln(\sigma_{MH, \xi, s}^{r-1}) + 6.534 \frac{(R - 0.234)}{r - 1}.$$

The matrix $\hat{\Sigma}^r + \frac{(\sigma_{MH, \xi, \psi}^r)^2}{r} I_2$ is a positive definite estimate of the covariance matrix. This covariance estimate is scaled by $(\sigma_{MH, \xi, \psi}^r)^2$ to obtain the covariance matrix for the proposal in the next iteration. The scaling $(\sigma_{MH, \xi, \psi}^r)^2$ is tuned to achieve an average acceptance probability of 0.234 as suggested by Roberts et al (1997) for multivariate random walk Metropolis-Hastings. To reduce computational cost the empirical covariance matrix $\hat{\Sigma}^r$ can be updated in every step by the following recursion (see e.g. Bennett et al (2009))

$$\hat{\Sigma}^r = \frac{r-2}{r-1} \hat{\Sigma}^{r-1} + \frac{1}{r} ((\xi^r, \psi^r)^\top - \hat{\mu}^{r-1}) ((\xi^r, \psi^r)^\top - \hat{\mu}^{r-1})^\top,$$

where $\hat{\mu}^{r-1}$ is the sample mean of $(\xi^i, \psi^i)_{i=1, \dots, r-1}$. We also update the sample mean recursively by

$$\hat{\mu}^r = \frac{1}{r} ((r-1) \hat{\mu}^{r-1} + (\xi^r, \psi^r)^\top).$$

We have seen that the adaptations for the μ and the (ϕ, σ) updates tend to be very small after burn-in and therefore we only adapt during the burn-in period. This also ensures a correct sampling procedure without the need to verify the validity of an adaptive MCMC scheme.

Appendix B. The standardized skew Student t distribution

According to Azzalini and Capitanio (2003), the density of the univariate skew Student t distribution with parameters $\xi \in \mathbb{R}, \omega \in (0, \infty), \alpha \in \mathbb{R}$ and $df \in (0, \infty)$ is given by

$$st(x|\xi, \omega, \alpha, df) = \frac{2}{\omega} t(x|df) T \left(\alpha \frac{x - \xi}{\omega} \sqrt{\frac{df + 1}{\left(\frac{x - \xi}{\omega}\right)^2 + df}} \middle| df + 1 \right),$$

where $t(\cdot|df)$ is the density function of the univariate Student t distribution with df degrees of freedom and $T(\cdot|df)$ the corresponding distribution function. The expectation and variance of a random variable X following a skew Student t distribution with parameters ξ, ω, α as above and $df > 2$ are given by

$$E(X) = \xi + \omega b_{df} \delta, \text{ and } Var(X) = \omega^2 \left(\frac{df}{df - 2} - b_{df}^2 \delta^2 \right),$$

where $\delta = \frac{\alpha^2}{\sqrt{1 + \alpha^2}}$ and $b_{df} = \sqrt{\frac{df}{\pi} \frac{\Gamma(\frac{df-1}{2})}{\Gamma(\frac{df}{2})}}$. If we set

$$\omega = \sqrt{\frac{1}{\left(\frac{df}{df-2} - b_{df}^2 \delta^2\right)}} \text{ and } \xi = -\omega b_{df} \delta = -\sqrt{\frac{1}{\left(\frac{df}{df-2} - b_{df}^2 \delta^2\right)}} b_{df} \delta,$$

only the parameters α and df remain unknown and the random variable has zero mean and a variance of one. We refer to the corresponding distribution as the standardized skew Student t distribution. Its density is denoted by sst and is obtained as

$$sst(x|\alpha, df) = st \left(x \middle| -\sqrt{\frac{1}{\left(\frac{df}{df-2} - b_{df}^2 \delta^2\right)}} b_{df} \delta, \sqrt{\frac{1}{\left(\frac{df}{df-2} - b_{df}^2 \delta^2\right)}}, \alpha, df \right). \quad (15)$$

Appendix C. Additional material for the bivariate dynamic mixture copula (Section 4)

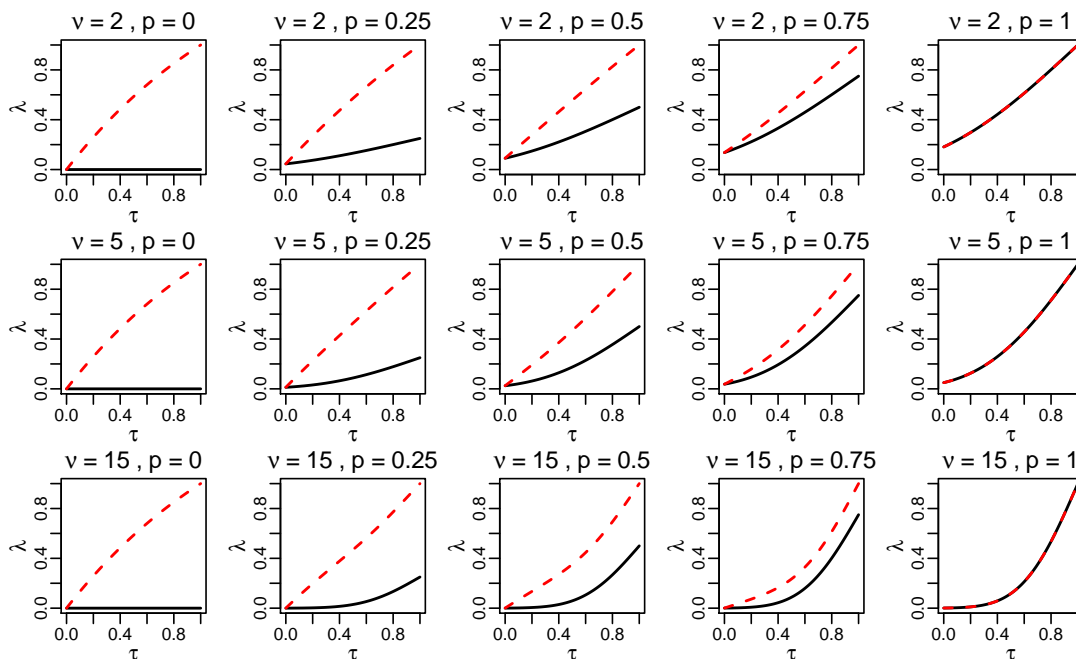


Figure 4: Upper (red, dashed) and lower (black) tail dependence coefficient of the mixture copula defined in (10) plotted against Kendall's τ for different values of ν and p .

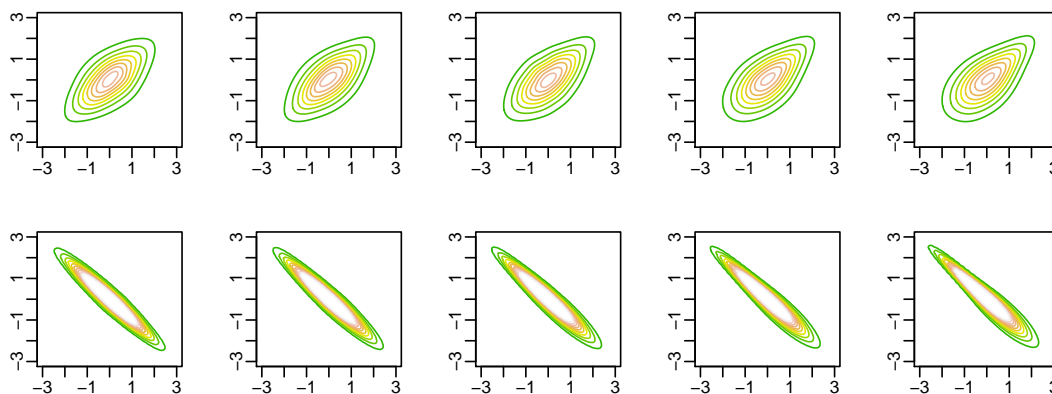


Figure 5: Normalized contour plots for the mixture copula model in (10) with $\tau = 0.4$ (top row), $\tau = -0.8$ (bottom row), $\nu = 5$ and $p = 1, 0.75, 0.5, 0.25, 0$ (from left to right).

Appendix D. Additional material for the application (Section 4)

Daily log returns




Figure 6: Daily log returns of the four indices SPX, VIX, DAX, VDAX from 2006 to 2013 plotted against time.

Posterior statistics

	mode	5% quantile	95% quantile	effective sample size
SPX				
μ	-9.32	-9.94	-8.70	13896.29
ϕ	0.99	0.98	1.00	565.10
σ	0.15	0.13	0.19	208.31
α	-0.51	-0.80	-0.22	4293.41
df	6.84	5.46	10.40	1821.93
VIX				
μ	-5.65	-5.80	-5.50	3586.38
ϕ	0.90	0.84	0.93	362.48
σ	0.36	0.28	0.48	311.76
α	1.33	0.97	1.73	1376.24
df	9.30	6.50	15.14	1251.32
DAX				
μ	-8.89	-9.30	-8.50	19573.46
ϕ	0.99	0.97	0.99	598.27
σ	0.15	0.12	0.19	249.07
α	-0.48	-0.80	-0.06	5662.26
df	9.74	7.31	15.11	2483.44
VDAX				
μ	-6.06	-6.21	-5.89	8086.52
ϕ	0.96	0.92	0.97	416.02
σ	0.18	0.13	0.24	293.01
α	0.96	0.66	1.27	3305.73
df	8.35	6.44	12.88	1386.65

Table 7: Posterior mode estimates, posterior quantiles and effective sample sizes for the univariate skew Student t stochastic volatility models for the four indices SPX, VIX, DAX, VDAX.

	mode	5% quantile	95% quantile	effective sample size
(SPX,VIX)				
μ	-0.74	-0.77	-0.71	1862.77
ϕ	0.94	0.85	0.97	306.33
σ	0.05	0.03	0.08	215.92
p	0.29	0.13	0.44	1436.78
ν	9.03	5.29	41.14	1039.38
(DAX,VDAX)				
μ	-0.81	-0.84	-0.78	1785.30
ϕ	0.86	0.73	0.92	285.29
σ	0.10	0.06	0.13	207.29
p	0.66	0.50	0.81	1406.11
ν	8.30	5.91	34.32	756.50

Table 8: Posterior mode estimates, posterior quantiles and effective sample sizes for the dynamic mixture copula models for the pairs (SPX,VIX) and (DAX,VDAX).

Calculating the log predictive score

We describe in detail how we proceed for model \mathcal{M}_{dyn}^{mix} . We consider $T + K$ observations of dimension two, stored in the data matrix $Y_{1:(T+K);1:2}$, where the first T observations are used to train the model and the last K are used for evaluation.

Step 1: (Model fitting based on the training period)

- We fit two marginal skew Student t stochastic volatility models to $\mathbf{y}_{1:T;1}$ and $\mathbf{y}_{1:T;2}$. This yields R_{train} draws of the parameters denoted by $\mathbf{s}_{1:T;j}^{st,r}$, $\mu_j^{st,r}$, $\phi_j^{st,r}$, $\sigma_j^{st,r}$, $\alpha_j^{st,r}$ and $df_j^{st,r}$, $r = 1, \dots, R_{train}$ and corresponding posterior mode estimates $\hat{\mathbf{s}}_{1:T;j}^{st}$, $\hat{\mu}_j^{st}$, $\hat{\phi}_j^{st}$, $\hat{\sigma}_j^{st}$, $\hat{\alpha}_j^{st}$ and \hat{df}_j^{st} for $j = 1, 2$.
- We estimate the copula data

$$\hat{u}_{tj} = ssT \left(y_{tj} \exp\left(-\frac{\hat{s}_{tj}^{st}}{2}\right) \middle| \hat{\alpha}_j^{st}, \hat{df}_j^{st} \right)$$

for $t = 1, \dots, T, j = 1, 2$.

- We fit the dynamic bivariate mixture copula model introduced in (12) based on the pseudo copula data $\hat{U}_{1:T;1:2}$ and obtain posterior draws $\mathbf{s}_{1:T}^{cop,r}$, $\mu^{cop,r}$, $\phi^{cop,r}$, $\sigma^{cop,r}$, $\nu^{cop,r}$, $p^{cop,r}$ for $r = 1, \dots, R_{train}$ and corresponding posterior mode estimates $\hat{\mathbf{s}}_{1:T}^{cop}$, $\hat{\mu}^{cop}$, $\hat{\phi}^{cop}$, $\hat{\sigma}^{cop}$, $\hat{\nu}^{cop}$, \hat{p}^{cop} .

Step 2: (The one-day ahead predictive density)

Estimating the one-day ahead predictive density at time $T + k, 1 \leq k \leq K$ would usually require to fit daily models with observations up to time $T + k - 1$ for $k = 1, \dots, K$. In order to save computational resources we use another approach where we only update the dynamic parameters, i.e. the log variances and Kendall's τ . For the constant parameters we use the estimates from the training period $1, \dots, T$. In this case we found that it is enough to only consider a time horizon of 100 time points, i.e. to estimate a dynamic parameter at time $T + k$ we consider data in the period $T + k - 100, \dots, T + k - 1$. We proceed as follows to obtain the one-day ahead predictive density at time point $T + k$ with $1 \leq k \leq K$.

- We consider a skew Student t stochastic volatility model as in (8), where we keep the parameters $\mu, \phi, \sigma, \alpha$ and df fixed and only update the latent log variances. Therefore we draw the latent log variances $\mathbf{s}_{(T+k-100):(T+k-1);j}^{st}$ conditional on $\mathbf{y}_{(T+k-100):(T+k-1);j}$, $\hat{\mu}_j^{st}$, $\hat{\phi}_j^{st}$, $\hat{\sigma}_j^{st}$, \hat{df}_j^{st} and $\hat{\alpha}_j^{st}$ for $j = 1, 2$. We denote the draws by $\mathbf{s}_{(T+k-100):(T+k-1);j}^{st,r}$, $r = 1, \dots, R_{test}$ for $j = 1, 2$. Corresponding posterior mode estimates are denoted by $\hat{\mathbf{s}}_{(T+k-100):(T+k-1);j}^{st}$, $j = 1, 2$.

- We estimate the copula data via the probability integral transform, i.e. for $j = 1, 2$ and $t = T + k - 100, \dots, T + k - 1$ we calculate

$$\hat{u}_{tj} = sst \left(y_{tj} \exp \left(-\frac{\hat{s}_{tj}^{st}}{2} \right) \middle| \hat{\alpha}_j^{st}, \hat{d}f_j^{st} \right).$$

- We fit the dynamic mixture copula model to the pseudo copula data $\hat{U}_{(T+k-100):(T+k-1):(1:2)}$ where we keep the constant parameters fixed. We only update $\mathbf{s}_{(T+k-100):(T+k-1)}^{cop}$ conditional on $\hat{U}_{(T+k-100):(T+k-1):(1:2)}$, $\hat{\mu}^{cop}$, $\hat{\phi}^{cop}$, $\hat{\sigma}^{cop}$, $\hat{\nu}^{cop}$, \hat{p}^{cop} . The corresponding draws are denoted by $\mathbf{s}_{(T+k-100):(T+k-1)}^{cop,r}$, $r = 1, \dots, R_{test}$ and the posterior mode estimates by $\hat{\mathbf{s}}_{(T+k-100):(T+k-1)}^{cop}$.
- For $j = 1, 2$, we obtain an estimate for the log variance at time point $T + k$ as $\hat{s}_{T+k;j}^{st} = \hat{\mu}_j^{st} + \hat{\phi}_j^{st}(\hat{s}_{T+k-1;j}^{st} - \hat{\mu}_j^{st})$.
- We obtain an estimate for Fisher's Z transform of Kendall's τ at time point $T + k$, as $\hat{s}_{T+k}^{cop} = \hat{\mu}^{cop} + \hat{\phi}^{cop}(\hat{s}_{T+k-1}^{cop} - \hat{\mu}^{cop})$.
- The predictive density evaluated at (y_1, y_2) is given by

$$f_{T+k}^p(y_1, y_2) = c_{T+k}^p(y_1, y_2) g_{T+k}^p(y_1, y_2),$$

with

$$c_{T+k}^p(y_1, y_2) = c^M \left(sst \left(x_1 \middle| \hat{\alpha}_1^{st}, \hat{d}f_1^{st} \right), sst \left(x_2 \middle| \hat{\alpha}_2^{st}, \hat{d}f_2^{st} \right); F_Z^{-1}(\hat{s}_{T+k}^{cop}), \hat{\nu}^{cop}, \hat{p}^{cop} \right),$$

where c^M is the density of the mixture copula defined in (10) and

$$g_{T+k}^p(y_1, y_2) = sst \left(x_1 \middle| \hat{\alpha}_1^{st}, \hat{d}f_1^{st} \right) sst \left(x_2 \middle| \hat{\alpha}_2^{st}, \hat{d}f_2^{st} \right) \exp \left(-\frac{\hat{s}_{T+k;1}^{st}}{2} \right) \exp \left(-\frac{\hat{s}_{T+k;2}^{st}}{2} \right),$$

with $x_j = y_j \exp \left(-\frac{\hat{s}_{T+k;j}^{st}}{2} \right)$ for $j = 1, 2$.

Step 3: (The cumulative pseudo log predictive score)

The cumulative pseudo log predictive score is obtained as

$$LP = \sum_{k=1}^K \log(f_{T+k}^p(y_{T+k;1}, y_{T+k;2})).$$

During the training period we run $R_{train} = 31000$ iterations with a burn-in of 1000, while for updating only the dynamic parameters 11000 iterations with a burn-in of 1000 is enough, i.e. we use $R_{test} = 11000$.

Supplementary material

1 Elliptical slice sampling

We assume that the posterior density for a parameter vector $\boldsymbol{\theta}$ given data D is proportional to

$$f(\boldsymbol{\theta}|D) \propto \ell(\boldsymbol{\theta}|D)\varphi(\boldsymbol{\theta}|\mathbf{0}, \Sigma), \quad (16)$$

where $\ell(\boldsymbol{\theta}|D)$ is the likelihood function and $\varphi(\boldsymbol{\theta}|\mathbf{0}, \Sigma)$ is the multivariate normal density with zero mean and covariance matrix Σ . Murray et al (2010) consider the Metropolis-Hastings sampler of Neal (1998) where a proposal $\boldsymbol{\theta}'$ is obtained from the following stochastic representation

$$\boldsymbol{\theta}' = \sqrt{1 - \alpha^2}\boldsymbol{\theta} + \alpha\mathbf{v}, \mathbf{v} \sim N(\mathbf{0}, \Sigma). \quad (17)$$

Here $\alpha \in [-1, 1]$ is a fixed step size parameter. The proposal is accepted with probability

$$\min\left(1, \frac{\ell(\boldsymbol{\theta}')}{\ell(\boldsymbol{\theta})}\right). \quad (18)$$

Elliptical slice sampling adapts the step size parameter α during sampling. This eliminates the need to select the parameter before sampling and it may be a better approach for situations where good choices of the step size parameter depend on the region of the state space. Murray et al (2010) first suggest alternatively to propose a new state by

$$\boldsymbol{\theta}' = \cos(\omega)\boldsymbol{\theta} + \sin(\omega)\mathbf{v}, \mathbf{v} \sim N(\mathbf{0}, \Sigma). \quad (19)$$

Here the angle ω corresponds to the step size. As we move ω towards zero the proposal gets closer to the initial value $\boldsymbol{\theta}$. Murray et al (2010) argue that (19) provides a more flexible choice for the proposals compared to (17), if the parameter ω is also updated, which is here the case. In elliptical slice sampling we first draw an angle ω from the uniform distribution on $[0, 2\pi]$ and obtain a proposal as outlined in (19). This proposal is accepted according to (18). If the proposal is not accepted a new angle is selected with a slice sampling approach (Neal et al (2003)) such that the angle approaches zero as more samples are rejected. This ensures that at some point the proposal will be accepted. The approach is outlined in Algorithm 1. Murray et al (2010) show that this samples from a Markov chain, where (16) is the corresponding stationary distribution.

Algorithm 1 Elliptical slice sampling

```
1:  $\mathbf{v} \sim N(\mathbf{0}, \Sigma)$ 
2:  $u \sim \text{uniform}(0, 1)$ 
3:  $\omega \sim \text{uniform}(0, 2\pi)$ 
4:  $\omega_{min} = \omega - 2\pi, \omega_{max} = \omega$ 
5:  $\boldsymbol{\theta}' = \cos(\omega)\boldsymbol{\theta} + \sin(\omega)\mathbf{v}$ 
6: while  $\frac{\ell(\boldsymbol{\theta}')}{\ell(\boldsymbol{\theta})} \leq u$  do
7:   if  $\omega < 0$  then
8:      $\omega_{min} = \omega$ 
9:   else
10:     $\omega_{max} = \omega$ 
11:   end if
12:    $\omega \sim \text{uniform}(\omega_{min}, \omega_{max})$ 
13:    $\boldsymbol{\theta}' = \cos(\omega)\boldsymbol{\theta} + \sin(\omega)\mathbf{v}$ 
14: end while
```

2 The DCC-GARCH model

The DCC-GARCH model was introduced by Engle (2002). Mathematical properties were developed by Engle and Sheppard (2001). In the DCC-GARCH model it is assumed that a d

dimensional random vector $\boldsymbol{\epsilon}_t \in \mathbb{R}^d$ at time t is multivariate normal distributed with zero mean and dynamic covariance matrix $H_t \in \mathbb{R}^{d \times d}$, i.e.

$$\boldsymbol{\epsilon}_t | H_t \sim N(\mathbf{0}, H_t),$$

for $t = 1, \dots, T$. The covariance matrix can be written as

$$H_t = D_t C_t D_t,$$

where C_t is a $d \times d$ correlation matrix and D_t is a $d \times d$ diagonal matrix. The diagonal matrix D_t contains the marginal standard deviations. We assume GARCH(P_j, Q_j) innovations for the j -th diagonal entry of D_t , i.e.

$$d_{t,j}^2 = \omega_j + \sum_{p=1}^{P_j} \alpha_{pj} \epsilon_{t-p,j}^2 + \sum_{q=1}^{Q_j} \beta_{qj} d_{t-q,j}^2,$$

where

- $\omega_j > 0$,
- $d_{t,0} > 0$,
- $\sum_{p=1}^{P_i} \alpha_{jp} + \sum_{q=1}^{Q_i} \beta_{qj} < 1$ and the roots of $1 - \sum_{p=1}^{P_i} \alpha_{jp} Z^p + \sum_{q=1}^{Q_i} \beta_{qj} Z^q$ lie outside the unit circle,
- α_{jp} for all $p \in \{1, \dots, P_j\}$ and β_{jq} for all $q \in \{1, \dots, Q_j\}$ are such that $d_{t,j}^2$ is positive.

The correlation matrix is decomposed as follows

$$C_t = \text{diag}(L_t)^{-1/2} L_t \text{diag}(L_t)^{-1/2},$$

where L_t is a positive definite matrix. Further we denote by

$$\mathbf{r}_t = D_t^{-1} \boldsymbol{\epsilon}_t$$

the standardized return and obtain \bar{L} as

$$\bar{L} = \frac{1}{T} \sum_{t=1}^T \mathbf{r}_t \mathbf{r}_t^\top.$$

For L_t we assume the following dynamic structure

$$L_t = \left(1 - \sum_{m=1}^M a_m - \sum_{n=1}^N b_n \right) \bar{L} + \sum_{m=1}^M a_m \mathbf{r}_{t-m} \mathbf{r}_{t-m}^\top + \sum_{n=1}^N b_n L_{t-n},$$

where

- $\alpha_m \geq 0$ for all $m \in \{1, \dots, M\}$ and $b_n \geq 0$ for all $n \in \{1, \dots, N\}$,
- $\sum_{m=1}^M a_m + \sum_{n=1}^N b_n < 1$,
- L_0 is positive definite.

Engle and Sheppard (2001) show that under the above conditions H_t is a proper covariance matrix. The DCC(M,N)-GARCH(P,Q) model is obtained by setting $P_j = P$ and $Q_j = Q$ for $j = 1, \dots, d$.

3 Further results for the application

Results for the marginal models

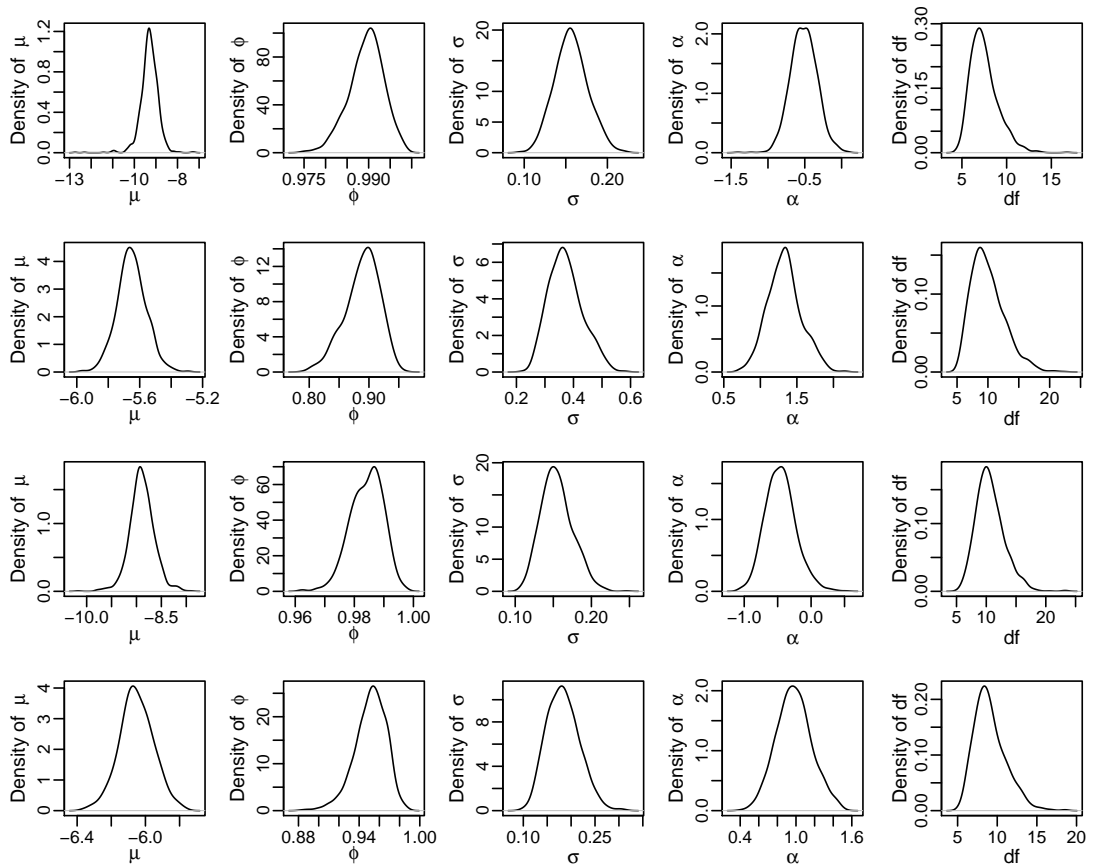


Figure 7: Estimated posterior densities based on 30000 MCMC iterations after a burn-in of 1000 for the parameters of the univariate skew Student t stochastic volatility models for SPX, VIX, DAX and VDAX (from top to bottom row).




Figure 8: Trace plots of 1000 MCMC draws based on a total of 31000 iterations, where the first 1000 draws are discarded for burn-in and the remaining 30000 draws are thinned with factor 30. The trace plots are shown for parameters of the univariate skew Student t stochastic volatility models for SPX, VIX, DAX and VDAX (from top to bottom row).

Results for the dependence models

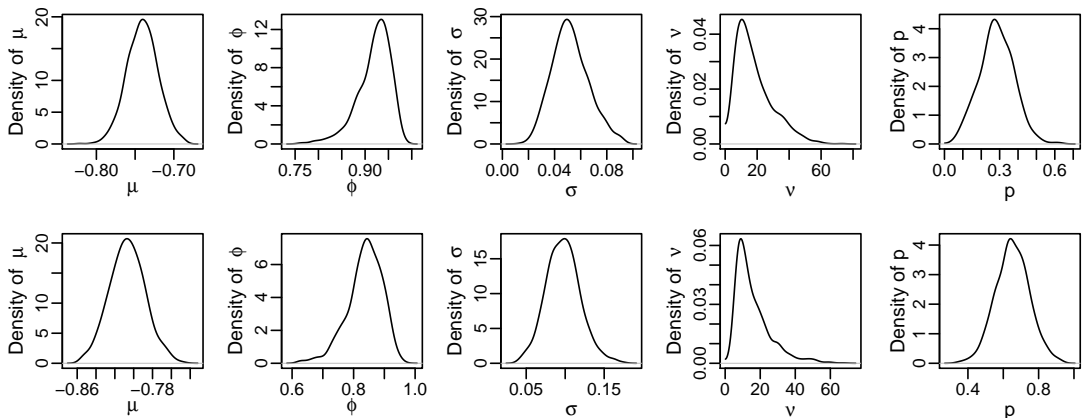


Figure 9: Estimated posterior densities based on 30000 MCMC iterations after a burn-in of 1000 for parameters of the dynamic mixture copula model for (SPX,VIX) in the top row and for (DAX,VDAX) in the bottom row.

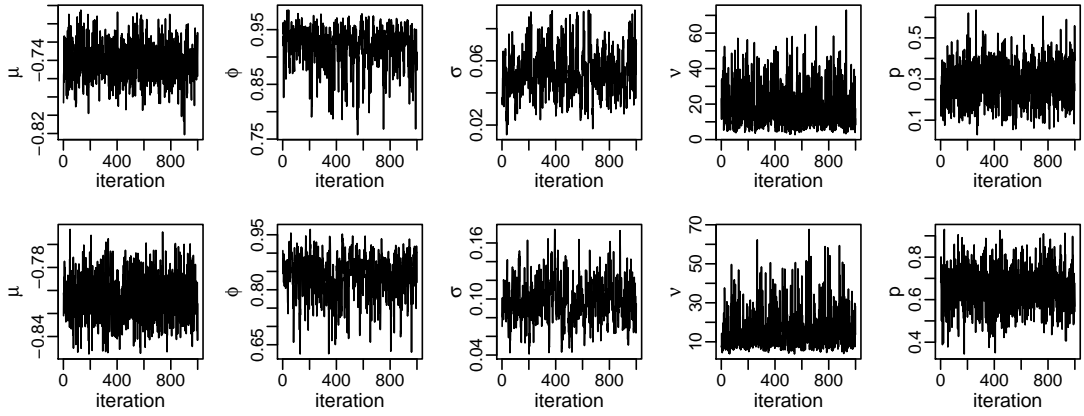


Figure 10: Trace plots of 1000 MCMC draws based on a total of 31000 iterations, where the first 1000 draws are discarded for burn-in and the remaining 30000 draws are thinned with factor 30. The trace plots are shown for parameters of the dynamic mixture copula model for (SPX,VIX) in the top row and for (DAX,VDAX) in the bottom row.