Time-varying joint distribution through copulas

M. Concepcion Ausin*
Department of Mathematics
University of A Coruña, Spain

Hedibert F. Lopes
Graduate School of Business
University of Chicago, USA

Abstract

This paper deals with the analysis of temporal dependence in multivariate time series. The dependence structure between the marginal series is modelled through the use of copulas which, unlike the correlation matrix, give a complete description of the joint distribution. The parameters of the copula function vary through time following certain evolution equations depending on their previous values and the historical data. The marginal time series follow standard univariate GARCH models. We develop full Bayesian inference where the whole set of model parameters is estimated simultaneously. This represents an essential difference with previous approaches in the literature where the marginal and the copula parameters are estimated separately in two consecutive steps. Moreover, we propose a Bayesian procedure for the estimation of the Value-at-Risk (VaR) of a portfolio of assets, providing point estimates and predictive intervals. The proposed copula model allows us to capture the dependence structure between the individual assets which strongly influences the portfolio VaR. Finally, we also address the problem of optimal portfolio selection based on the estimation of mean-VaR efficient frontiers. The proposed approach is illustrated with simulated and real financial time series.

Keywords: Bayesian estimation; Copulas; Multivariate GARH; Kendall's tau; Value at Risk; Portfolio allocation.

*Corresponding author. M. Concepcion Ausin, Departamento de Matemáticas, Facultade de Informática, Campus de Elviña, Universidade da Coruña, 15071 A Coruña, Spain. Tel.: +34 981 167 000 (ext. 1318); fax: + 34 981 167 160. E-mail address: mausin@udc.es

1 Introduction

There have been much work on the extension to the multivariate case of the successful univariate autoregressive conditional heteroscedasticity (ARCH) and the generalized ARCH (GARCH) models in order to describe temporal dependence in financial data, see e.g. Bauwens et al. (2006) for a survey on multivariate GARCH models. The usual assumption in most cases is that the conditional joint distribution of the returns follows a multivariate normal or multivariate t-distribution. However, it is well known that these elliptical distribution models require a very strong symmetry of the data and might not be appropriate in many circumstances.

A recent alternative approach to the study of dependence in financial time series is the use of copulas, see e.g. Nelsen (2006). The main advantage of this approach is that the individual marginal densities of the returns can be defined separately from their dependence structure. Then, the models for the marginal series can be firstly specified using the required univariate characteristics and then, the dependence between the returns can be completely modeled by selecting an appropriate copula function. Using this approach, many non-elliptical and flexible multivariate distributions can be obtained. In this context, most of researchers have considered copula GARCH models where the marginal series follow univariate GARCH processes and the dependence structure between them is specified by a copula function, see e.g. Dias and Embrechts (2004), Rodriguez (2003), Hu (2006), Bartram et al. (2006), Patton (2006), Jondeau and Rockinger (2006).

Most of previous works using copula GARCH models does not account for the parameter uncertainty simultaneously. Proposed methods are generally based on a two-stage approach where, in the first step, the marginal series are estimated assuming independence and then, in a second step, these estimations are plugged in the copula function in order to estimate the copula parameters. Although two-steps approaches are usually straightforward to implement for copula models, they may produce inadequate measures of uncertainty. In fact, it can be shown that two-stage maximum likelihood estimation approaches lead to consistent but not efficient estimators, see Patton (2006).

One of the main interest in modelling multivariate financial time series is portfolio management. The Value-at-Risk (VaR) has become an important and widely used measure of the risk inherent in asset portfolios, see Jorion (2000). Clearly, the VaR of a portfolio depends on the behaviour of the individual assets in the portfolio and also on the dependence structure between them. In particular, the dependence in the tails of the distribution strongly influences the VaR calculation, see e.g. Kiesel and Kleinow (2002). Thus, the correlation coefficient, which is not adequate to measure the dependence in the tails, may lead to inaccurate estimations of VaR. Alternatively, copulas provide a useful tool to model tail dependence and obtain precise VaR estimations, see Embrechts et al. (1999).

Besides the applications of VaR in risk measurement, it also provides a useful tool in optimal portfolio selection. Investors are mainly interested in how to allocate their investments between different assets so as to minimize the overall risk for a given expected return. The classical portfolio optimization approach follows the pioneering work of Markowitz (1959) based on the variance-covariance matrix. As pointed out in Embrechts et al. (1999), Markowitz optimization makes sense for elliptical distributions, as the variance-covariance matrix is only valid to measure linear dependence, but may be insufficient to capture other dependence structures between portfolio assets. The use of the VaR in portfolio optimization is a relatively novel alternative, see e.g. Gaivoronski and Pflug (2005). Using this approach, the optimal portfolio is the one which minimizes the VaR subject to achieving a specified level of expected return. The analogous to the classical mean-variance efficient frontiers are the mean-VaR efficient frontiers which gives the minimum VaR for each expected return. Note that for the case of elliptical distributions, the portfolio minimizing VaR coincides with the Markowitz variance minimizing portfolio.

In this paper, we propose a Bayesian methodology to make inference and prediction in copula GARCH models. We develop a one-step Bayesian procedure where all the parameters are estimated simultaneously using the whole likelihood function. Copulas are modelled to be time-varying in order to capture the time evolution in the dependence structure. This allows us to identify contagion effects or changing dependence structures during periods of financial instability, see e.g. Rodriguez (2003) and Arakelian and Dellaportas (2006). Moreover, we describe how to obtain Bayesian estimations of portfolio VaR and predictive intervals which are much more informative than simple point forecasts. We also show how to obtain optimal portfolios based on VaR and obtain confidence regions for the mean-VaR efficient frontiers.

The rest of the paper is organized as follows. Section 2 briefly reviews the definition, main properties and some examples of copulas. Section 3 presents the time-varying copula GARCH models considered in this article. Section 4 describes how to carry out Bayesian inference and prediction for these models making use of the Markov chain Monte Carlo (MCMC) methodology. The estimation of volatilities, rank correlation and tail dependence is also addressed. Section 5 is devoted to portfolio management. Predictive mean and intervals for the portfolio VaR are obtained and used for optimal portfolio allocation. Bayesian prediction of mean-VaR efficient frontiers is also addressed. Section 6 illustrates the methodology with simulated and real financial time series. Section 7 concludes with some discussion and extensions.

2 Copulas

A p-dimensional copula $C(u_1, ..., u_p)$, is a multivariate distribution function in the unit hypercube $[0,1]^p$, with uniform $U(0,1)$ marginal distributions. It can be shown, see e.g. Schweizer and Sklar (1983), that every joint distribution, $F(x_1, ..., x_p)$, whose marginals are given by $F_1(x_1), \ldots, F_p(x_p)$, can be written as,

$$F(x_1, ..., x_p) = C(F_1(x_1), ..., F_p(x_p)), \tag{1}$$

for a function C that is called a *copula* of F. Furthermore, if the marginal distributions are continuous, then there is a unique copula associated to the joint distribution, F, that can be obtained from,

$$C(u_1, ..., u_p) = F\left(F_1^{-1}(u_1), ..., F_p^{-1}(u_p)\right). \tag{2}$$

Conversely, given a p-dimensional copula, $C(u_1, ..., u_p)$, and p univariate distributions, $F_1(x_1), ..., F_p(x_p)$, the function (1) is a p-variate distribution function with margins $F_1, ..., F_p$, whose corresponding density function is given by,

$$f(x_1, ..., x_p) = c(F_1(x_1), ..., F_p(x_p)) \prod_{i=1}^{p} f_i(x_i), \tag{3}$$

where f_i represents the marginal density functions and c is the density function of the copula which is derived from (2) and is given by,

$$c(u_1, ..., u_p) = \frac{f\left(F_1^{-1}(u_1), ..., F_p^{-1}(u_p)\right)}{\prod_{i=1}^{p} f_i\left(F_i^{-1}(u_i)\right)}.$$

There is a large number of parametric families of copulas in the literature, see e.g. Nelsen (2006). The basic example is the Gaussian copula, which is obtained from the multivariate normal distribution with correlation matrix, R, and is given by,

$$C_R^{\text{Ga}}(u_1, ..., u_p) = \int_{-\infty}^{\Phi^{-1}(u_1)} \cdots \int_{-\infty}^{\Phi^{-1}(u_p)} \frac{1}{\sqrt{(2\pi)^p |R|}} \exp\left\{\frac{-\mathbf{u}' R^{-1}\mathbf{u}}{2}\right\} d\mathbf{u},$$

where $\mathbf{u} = (u_1, ..., u_p)$ and Φ is the cumulative distribution function of the univariate standard normal distribution. The normal copula assumes that there is no dependence in the tails of the distribution and then, in financial economics, it is often more useful to consider the t-copula, which is obtained from the multivariate t-distribution with η degrees of freedom and correlation matrix, R, and is given by,

$$C_{\eta,R}^{\text{t}}(u_1, \ldots, u_p) = \int_{-\infty}^{t_\eta^{-1}(u_1)} \cdots \int_{-\infty}^{t_\eta^{-1}(u_p)} \frac{\Gamma\left(\frac{\eta+p}{2}\right)\left(1 + \frac{\mathbf{u}' R^{-1}\mathbf{u}}{\eta}\right)^{-\frac{\eta+p}{2}}}{\Gamma\left(\frac{\eta}{2}\right)\sqrt{(\pi\eta)^p |R|}} d\mathbf{u}, \tag{4}$$

where t_η denotes the cumulative distribution function of the standard univariate Student-t distribution with η degrees of freedom. Note that the Gaussian copula is obtained as a special case of the t-copula when η goes to infinity.

Copulas have the property that various dependence measures between two random variables depend only their copula function. For example, an important measure of dependence is the Kendall's tau rank correlation, which is defined by,

$$\tau = E\left[sign\left(X_1 - X_1'\right)\left(X_2 - X_2'\right)\right],$$

where (X_1, X_2) and (X_1', X_2') are two independent and equally distributed pairs of random variables.

The Kendall's tau is a very useful alternative to the linear correlation coefficient because it does not depend on the marginal distributions of X_1 and X_2. In fact, the Kendall's tau only depends on the copula function and it can be shown that,

$$\tau = 4\int_0^1 \int_0^1 C\left(u_1, u_2\right) c\left(u_1, u_2\right) du_1 du_2 - 1.$$

The Kendall's tau takes the same form for the bivariate Gaussian copula and the t-copula with correlation coefficient ρ, and is given by,

$$\tau = \frac{2}{\pi} \arcsin \rho. \tag{5}$$

Other useful dependence measures between two variables are the coefficients of upper tail dependence, λ_u, and lower tail dependence, λ_l, which are defined by,

$$\lambda_u = \lim_{q \to 1} P\left(X_2 > F_{X_2}^{-1}\left(q\right) \mid X_1 > F_{X_1}^{-1}\left(q\right)\right), \qquad \lambda_l = \lim_{q \to 0} P\left(X_2 \leq F_{X_2}^{-1}\left(q\right) \mid X_1 \leq F_{X_1}^{-1}\left(q\right)\right),$$

and can be expressed in terms of the copula as follows,

$$\lambda_u = \lim_{q \to 1} \frac{1 - 2q + C\left(q, q\right)}{1 - q}, \qquad \lambda_l = \lim_{q \to 0} \frac{C\left(q, q\right)}{q}.$$

As commented above, the Gaussian copula is characterized by zero tail dependence. The t-copula exhibits tail dependence which is determined by,

$$\lambda_u = \lambda_l = 2t_{\eta+1}\left(\frac{-\sqrt{\eta+1}\sqrt{1-\rho}}{\sqrt{1+\rho}}\right). \tag{6}$$

3 Copula GARCH models

A p-dimensional vector of financial time series, $\mathbf{y}_t = (y_{1t}, \ldots, y_{pt})$, follows a copula GARCH model if the joint cumulative distribution function is given by,

$$F(\mathbf{y}_t \mid \boldsymbol{\mu}, \mathbf{h}_t) = C\left(F_1(y_{1t} \mid \mu_1, h_{1t}), \ldots, F_p(y_{pt} \mid \mu_p, h_{pt})\right),$$

where C is a p-dimensional copula, F_i is the conditional distribution function of the marginal series y_{it}, for $i = 1, \ldots, p$, and y_{it} follows a standard univariate GARCH model,

$$\begin{aligned} y_{it} &= \mu_i + \sqrt{h_{it}}\epsilon_{it}, \\ h_{it} &= \omega_i + \alpha_i (y_{i,t-1} - \mu_i)^2 + \beta_i h_{i,t-1}, \end{aligned}$$

where h_{it} is the conditional variance of y_{it} given the previous information $I_{i,t-1} = \{y_{i,t-1}, y_{i,t-2}, ...\}$, ϵ_{it} are independent and identically distributed random variables with zero mean and $\omega_i, \alpha_i, \beta_i > 0$ and $\alpha_i + \beta_i < 1$ to ensure positivity of h_{it} and covariance stationarity, respectively.

We assume that the innovations follow the standard Student t-distribution, $\epsilon_{it} \sim t_{\nu_i}$, with ν_i degrees of freedom, zero mean and variance $\nu_i/(\nu_i - 2)$, for $i = 1, \ldots, p$, which is the usual choice to model fat tails in univariate time series, see e.g. Bollerslev (1987). Then, the conditional distribution function of each marginal series is $F_i(y_{it} \mid \mu_i, h_{it}) = t_{\nu_i}\left((y_{it} - \mu_i) h_{it}^{-1/2}\right)$, for $i = 1, \ldots, p$.

We also assume that the dependence structure between the marginal series is described by a time-varying t-copula function with η degrees of freedom, as defined in (4), whose density for each time, t, is given by,

$$c^{\mathrm{t}}_{\eta,R_t}(u_{1t}, \ldots, u_{pt}) = \frac{f^{\mathrm{t}}_{\eta,R_t}\left(t^{-1}_{\eta}(u_{1t}), \ldots, t^{-1}_{\eta}(u_{pt})\right)}{\prod_{i=1}^{p} f^{\mathrm{t}}_{\eta}\left(t^{-1}_{\eta}(u_{it})\right)},$$

where $u_{it} = F_i(y_{it} \mid \mu_i, h_{it})$, for $i = 1, \ldots, p$; $f^{\mathrm{t}}_{\eta,R_t}$ is the joint density of the standard multivariate Student-t distribution with η degrees of freedom and correlation matrix R_t and f^{t}_{η} is the density of the standard univariate t-distribution with η degrees of freedom. Note that the resulting joint distribution of the multivariate series is only elliptically contoured if the degrees of freedom η of the t-copula and the degrees of freedom ν_i of the marginals coincide, in which case the joint distribution correspond to the multivariate t-distribution.

Finally, we assume that the parameter matrix, R_t, of the t-copula varies through time according to the following equation,

$$R_t = (1 - a - b)\, R + a\Psi_{t-1} + bR_{t-1}, \tag{7}$$

where a and b are nonnegative parameters, R is a time-invariant $p \times p$ positive definite parameter matrix with unit diagonal elements and Ψ_{t-1} is a $p \times p$ matrix whose (i, j)-th element is given by,

$$\Psi_{ij,t-1} = \frac{\sum_{h=1}^{m} x_{it-h} x_{jt-h}}{\sqrt{\sum_{h=1}^{m} x_{it-h}^2 \sum_{h=1}^{m} x_{jt-h}^2}},$$

which gives the sample correlation of $\{\mathbf{x}_{t-1}, ..., \mathbf{x}_{t-m}\}$, with $m \geq 2$, where,

$$\mathbf{x}_t = (x_{1t}, ..., x_{pt}) = \left(t_\eta^{-1}\left(t_{\nu_1}\left(\epsilon_{1t}\right)\right), ..., t_\eta^{-1}\left(t_{\nu_p}\left(\epsilon_{pt}\right)\right)\right). \tag{8}$$

Note that $\mathbf{x}_t$ follows a standard multivariate Student t-distribution with η degrees of freedom. For stationarity to be guaranteed, we impose the constraints $0 \leq a, b \leq 1$, $a + b \leq 1$ and $-1 \leq r_{ij} \leq 1$, where r_{ij} is the (i, j)-th element of the parameter matrix R. The equation (7) is based in the dynamics for the correlation matrix proposed by Tse and Tsui (2002) in a multivariate GARCH model. A similar equation has been also used in Jondeau and Rockinger (2006) in a bivariate copula GARCH model. Note that the time-varying equation (7) has the advantage that the parameter matrix, R_t, is a well-defined correlation matrices i.e., positive definite with unit diagonal elements. Then, there is no need of using any transformation such as the logistic function considered e.g. in Patton (2006) and Dias and Embrechts (2004) to keep the correlation parameter of the bivariate t-copula inside the interval $[-1, 1]$.

Thus, using (3), the joint density function of the multivariate time series can be computed by,

$$\begin{aligned} f\left(\mathbf{y}_t \mid \boldsymbol{\mu}, \mathbf{h}_t\right) &= c^{\mathrm{t}}_{\eta,R_t}\left(t_{\nu_1}\left(\frac{y_{1t}-\mu_1}{h_{1t}^{1/2}}\right),\ldots,t_{\nu_p}\left(\frac{y_{pt}-\mu_p}{h_{pt}^{1/2}}\right)\right)\prod_{i=1}^{p} f^{\mathrm{t}}_{\nu_i}\left(\frac{y_{it}-\mu_i}{h_{it}^{1/2}}\right)\frac{1}{h_{it}^{1/2}} \\ &= \frac{f^{\mathrm{t}}_{\eta,R_t}\left(x_{1t},\ldots,x_{pt}\right)}{\prod_{i=1}^{p} f^{\mathrm{t}}_{\eta}\left(x_{it}\right)}\prod_{i=1}^{p} f^{\mathrm{t}}_{\nu_i}\left(\frac{y_{it}-\mu_i}{h_{it}^{-1/2}}\right)\frac{1}{h_{it}^{-1/2}} \end{aligned} \tag{9}$$

where x_{it}, for $i = 1, \ldots, p$, is given in (8).

4 Bayesian inference and prediction

We want to make inference for the model parameters, $\boldsymbol{\theta} = \{(\mu_i, \omega_i, \alpha_i, \beta_i, \nu_i)_{i=1}^{p}, (a, b, R, \eta)\}$, constituted by the parameters of the dynamic copula function, (η, a, b, R), the parameters of the conditional variances' equations, $(\mu_i, \omega_i, \alpha_i, \beta_i)$, for $i = 1, \ldots, p$, and the degrees of freedom of each marginal series, ν_i. Firstly, we define prior distributions for $\boldsymbol{\theta}$. For each one of the parameters $(\mu_i, \omega_i, \alpha_i, \beta_i)$, we assume a uniform prior over their respective domains imposing the stationary condition, $\alpha_i + \beta_i < 1$. For the degrees of freedom parameters, we assume a half-right side Cauchy prior,

$$\pi\left(\nu_i\right) \propto \frac{1}{1+\nu_i^2}, \qquad \nu_i > 0, \tag{10}$$

for $i = 1, \ldots, p$. Note that a flat prior on these parameters would lead to an improper posterior distribution, as shown in Bauwens and Lubrano (1998). For the time-varying copula parameters (a, b, R), we assume a uniform prior distribution restricted to $0 \leq a, b \leq 1$, $a + b \leq 1$ and $-1 \leq r_{ij} \leq 1$, where r_{ij} is the (i, j)-th element of R. And finally, we assume a half-right side Cauchy distribution as the given in (10) for the degrees of freedom η of the t-copula.

Given an observed multivariate series, $\mathbf{y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_T\}$, and the priors specified above, the evaluation of the joint posterior distribution $\pi\left(\theta \mid \mathbf{y}\right)$ is analitycally intractable. Then, we make use of the MCMC sampling strategies in order to obtain a sample from the joint posterior distribution which allows us to develop Bayesian inference. Initially, we propose a Gibbs sampling scheme which is carried out by cycling repeatedly through draws of each parameter conditional on the remaining parameters, see Tierney (1994). Given the prior distributions and the likelihood function obtained from (9), it can be shown that the conditional posterior distributions of the model parameters are given by,

$$\pi\left(\phi_i \mid \cdot\right) \propto \prod_{t=1}^{T} \frac{\left(1+\frac{\mathbf{x}' R_t^{-1}\mathbf{x}}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}\left(1+\frac{x_{it}^2}{\eta}\right)^{-\frac{\eta+1}{2}}} \frac{\left(1+\frac{(y_{it}-\mu_i)^2}{\nu_i h_{it}}\right)^{-\frac{\nu_i+1}{2}}}{\sqrt{h_{it}}}, \tag{11}$$

for $\phi_i = \mu_i, \omega_i, \alpha_i, \beta_i$, for $i = 1, \ldots p$,

$$\pi\left(\nu_i \mid \cdot\right) \propto \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)^T \nu_i^{-T/2}}{\Gamma\left(\frac{\nu_i}{2}\right)^T (1+\nu_i)^2} \prod_{t=1}^{T} \frac{\left(1+\frac{\mathbf{x}' R_t^{-1}\mathbf{x}}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}\left(1+\dfrac{x_{it}^2}{\eta}\right)^{-\frac{\eta+1}{2}}\left(1+\frac{(y_{it}-\mu_i)^2}{\nu_i h_{it}}\right)^{\frac{\nu_i+1}{2}}}, \tag{12}$$

for $i = 1, 2$,

$$\pi\left(\eta \mid \cdot\right) \propto \frac{\Gamma\left(\frac{\eta+p}{2}\right)^T \Gamma\left(\frac{\eta}{2}\right)^{T(p-1)}}{\Gamma\left(\frac{\eta+1}{2}\right)^{Tp} (1+\eta)^2} \prod_{t=1}^{T} \frac{\left(1+\frac{\mathbf{x}' R_t^{-1}\mathbf{x}}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}\prod_{i=1}^{p}\left(1+\dfrac{x_{it}^2}{\eta}\right)^{-\frac{\eta+1}{2}}}, \tag{13}$$

and,

$$\pi\left(d \mid \cdots\right) \propto \prod_{t=1}^{T} \frac{1}{\sqrt{|R_t|}} \left(1+\frac{\mathbf{x}' R_t^{-1}\mathbf{x}}{\eta}\right)^{-\frac{\eta+p}{2}}, \tag{14}$$

or $d = a, b, r_{ij}$. Then, using these conditional posterior distributions, the simplest sampling approach that we propose here is to update each model parameter separately in the MCMC algorithm. For example, we can consider a simple one-dimensional random walk Metropolis for each parameter using normal candidate distributions whose variances can be calibrated to obtain good acceptance rates. However, we have observed in practice that the mixing in the MCMC algorithm can be significantly improved and the computational cost can be drastically reduced by using simultaneous updating of highly correlated subvectors of the model parameters $\boldsymbol{\theta}$, as suggested in e.g. Vrontos et al. (2001). More specifically, we propose to update simultaneously the subset of parameters $\phi_i = (\mu_i, \omega_i, \alpha_i, \beta_i)$ using a multivariate Metropolis step for each $i = 1, \ldots, p$. We generate a candidate vector from a multivariate normal distribution $N\left(\phi_i^{(n)}, c\Sigma\right)$, where $\phi_i^{(n)}$ denotes the current value of the parameter subvector, Σ is a estimation of the variance covariance matrix associated to this subvector and c is a constant to calibrate the acceptance rate. The matrix Σ can be obtained for example from a moderate number of iterations from the one-dimensional random walk Metropolis algorithm considered previously. Analogously, we can update simultaneously the subset of parameters (a, b, R) of the copula evolution equations.

In financial time series, it is frequent to observe changes in the temporal dependence during periods of high volatility. This effect is known as financial contagion, see e.g. Rodriguez (2003) and Arakelian and Dellaportas (2006). Thus, in multivariate GARCH models, it is important the estimation of dependence measures and volatilities as a function of t. Given the MCMC output, we can obtain samples from the posterior distribution of the individual volatilities, h_{it}, by evaluating their values $h_{it}^{(n)}$ for each draw $\boldsymbol{\theta}^{(n)}$ of the model parameters in the MCMC sample. Then, we can approximate their posterior means using,

$$E\left[h_{it} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} h_{it}^{(n)}.$$

Also, we can approximate the posterior median and 95% credible intervals for h_{it} by just calculating the median and the 0.025 and 0.975 quantiles, respectively, of the posterior sample of h_{it}.

Analogously, we can obtain samples from the posterior distribution of the individual dependence parameters, r_{ijt}, of the matrix R_t by evaluating their values $r_{ijt}^{(n)}$ for each draw $\boldsymbol{\theta}^{(n)}$ of the MCMC sample. For the particular case of having a bivariate times series, we can easily estimate the posterior mean of the Kendall's tau correlation, given in (5), for each time t using,

$$E\left[\tau_t \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} \frac{2}{\pi} \arcsin \rho_t^{(n)}, \tag{15}$$

where $\rho_t^{(n)} = r_{12t}^{(n)}$, which is the off-diagonal element of the time-varying matrix given in (7) evaluated for each draw $\boldsymbol{\theta}^{(n)}$ of the MCMC sample of size N. Also, we can approximate the posterior median of τ_t and credible intervals as before. Analogously, we can estimate the posterior mean, median and credible intervals for the coefficient of tail dependence λ_t as a function of t using (6).

Finally, note that using this approach we can also estimate the predictive distribution and intervals for the one-step ahead volatilities, $h_{i,T+1}$, the one-step ahead Kendall's tau, τ_{T+1}, and the one-step ahead coefficient of tail dependence, λ_{T+1}, which are of particular interest for prediction purposes.

5 Value-at-Risk estimation and portfolio allocation

The VaR of a portfolio is defined as a low order quantile of the portfolio return in a given period of time. As the losses should exceed VaR only a small percentage of time, it can be thought as the worst case outcome of the portfolio performance. More specifically, given a portfolio obtained from a multivariate log return series,

$$\sum_{i=1}^{p} \delta_i y_{it}, \qquad 0 \leq \delta_i \leq 1, \qquad \sum_{i=1}^{p} \delta_i = 1,$$

the t-period $\pi\%$ VaR is given by,

$$\pi = \Pr\left(\sum_{i=1}^{p} \delta_i y_{it} \leq -\mathrm{VaR}_t\right),$$

where π is supposed to be a small probability such as 0.01 or 0.05.

Given the observed series, $\mathbf{y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_T\}$, it is particularly interesting the estimation of the one-step ahead VaR, that is, the $(T+1)$-period VaR. A consistent estimator of the one-step ahead VaR is given by,

$$E\left[\mathrm{VaR}_{T+1} \mid \mathbf{y}\right] \approx \frac{1}{N} \sum_{n=1}^{N} \mathrm{VaR}_{T+1}^{(n)}, \tag{16}$$

where $\mathrm{VaR}_{T+1}^{(n)}$ is the one-step-ahead Value-at-Risk given the model parameters $\boldsymbol{\theta}^{(n)}$ of the n-th MCMC iteration. Although the conditional distribution of the multivariate series $(y_{1,T+1}, \ldots, y_{p,T+1})$ is explicit given the model parameters, it is not straightforward to derive the distribution of the one-step ahead portfolio $\sum_{i=1}^{p} \delta_i y_{iT+1}$. Therefore, it is complicated to obtain an analytic expression for $\mathrm{VaR}_{T+1}^{(n)}$ given $\boldsymbol{\theta}^{(n)}$. However, it can be easily approximated by generating values from our copula GARCH model as follows. For each value of the parameters $\boldsymbol{\theta}^{(n)}$, the value of $(h_{1,T+1}^{(n)}, \ldots, h_{p,T+1}^{(n)})$

is known and then, we can generate M replicas $\{(y_{1,T+1}^{(n,m)}, \ldots, y_{p,T+1}^{(n,m)})\}_{m=1}^{M}$ from the one-step ahead density of the multivariate series using the following two steps.

For each $m = 1, \ldots, M$:

1. Simulate $\left(x_{1,T+1}^{(n,m)}, \ldots, x_{p,T+1}^{(n,m)}\right)$ from a multivariate-t with parameters $\eta^{(n)}$ and $R_{T+1}^{(n)}$.
2. Set,

$$y_{i,T+1}^{(n,m)} = t_{\nu_i}^{-1}\left(t_\eta\left(x_{i,T+1}^{(n,m)}\right)\right)\sqrt{h_{i,T+1}^{(n)}} + \mu_i, \qquad \text{for } i = 1, \ldots, p.$$

Then, the value of $\text{VaR}_{T+1}^{(n)}$ can be approximated by the empirical π-quantile of the sample of portfolios $\left\{\sum_{i=1}^{p} \delta_i y_{i,T+1}^{(n,m)}\right\}_{m=1}^{M}$. Now, we can estimate the posterior mean (16) and obtain 95% predictive intervals for VaR using the 0.025 and 0.975 quantiles of the posterior sample of $\text{VaR}^{(n)}$, for $n = 1, \ldots, N$. A similar simulation procedure is considered in Ausin and Galeano (2007) to obtain predictive intervals for the VaR in univariate GARCH models.

We have shown how to estimate the one-step ahead VaR of a given portfolio. A different problem is how to choose the optimal portfolio which minimizes the VaR. Unfortunately, given the model parameters, it is not easy to obtain a closed expression for the optimal weights which minimize the VaR. However, using the simulation procedure described above, we can approximate the optimal weights by evaluating (16) in a grid of values for $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_p)$, and choosing the vector of weights which minimizes the predictive mean of the portfolio VaR.

Finally, assume that we are interested in the portfolio which minimizes the VaR subject to achieving at least some specified expected gain. The predictive mean of the portfolio expected gain is given by,

$$g = E\left[\sum_{i=1}^{p} \delta_i \mu_i \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N}\sum_{i=1}^{p} \delta_i \mu_i^{(n)}.$$

Thus, for a given value of g, we can find the set of values for $\boldsymbol{\delta}$ which lead to that expected gain. Observe that for each vector of weights, $\boldsymbol{\delta}$, and for each set of model parameters, $\boldsymbol{\theta}^{(n)}$, we can approximate the portfolio VaR as before. Then, we choose the vector of weights, $\boldsymbol{\delta}^{(n)}$, which minimizes the portfolio VaR, denoted by $\text{VaR}_{T+1}^{(n)}$. Thus, we obtain a posterior sample of the portfolio VaR, $\{\text{VaR}_{T+1}^{(n)}\}_{n=1}^{N}$, for each given expected gain, g, which allows for the construction of predictive intervals. Repeating this procedure for a number of values of the expected gain, g, we can approximate the mean-VaR efficient frontier with associated predictive regions of credibility.

6 Illustration

6.1 Simulated data

In this section, we illustrate the proposed methodology with one of the many artificial time series analysis that we have performed to examine our procedure. We simulate a bivariate time series of size $T = 1000$ from the time-varying copula GARCH model described in Section 4 with the following univariate GARCH models,

$$\begin{aligned} y_{1t} &= 0.05 + \sqrt{h_{1t}}\epsilon_{1t}, \quad \epsilon_{1t} \sim t_6\left(0,1\right), & y_{2t} &= 0.08 + \sqrt{h_{2t}}\epsilon_{2t}, \quad \epsilon_{2t} \sim t_3\left(0,1\right), \\ h_{1t} &= 0.05 + 0.15\left(y_{i,t-1} - 0.05\right)^2 + 0.75 h_{i,t-1}, & h_{2t} &= 0.01 + 0.10\left(y_{2,t-1} - 0.08\right)^2 + 0.8 h_{1,t-1}, \end{aligned}$$

and the following pattern for the time-varying copula parameter, $\rho_t = r_{12t}$, which is the off-diagonal of the matrix, R_t, defined in (7),

$$\rho_t = (1 - 0.5 - 0.2) \times 0.9 + 0.5\psi_{t-1} + 0.2\rho_{t-1},$$

where,

$$\psi_{t-1} = \frac{\sum_{h=1}^{5} x_{1t-h} x_{2t-h}}{\sqrt{\sum_{h=1}^{5} x_{1t-h}^2 \sum_{h=1}^{m} x_{2t-h}^2}},$$

where $x_{it} = t_\eta^{-1}(t_{\nu_i}(\epsilon_{it-h}))$, and $\eta = 4$. Note that the conditional distribution of the bivariate series is not elliptical as the degrees of freedom parameters ν_1, ν_2 and η take different values.

The proposed MCMC algorithm is run for 20000 iterations discarding the first 10000 as burn-in iterations. As described in Section 4, we consider simultaneous updating for the vectors of parameters $(\mu_i, \omega_i, \alpha_i, \beta_i)$, for $i = 1, 2$, and (a, b, r_{12}) using multivariate normal candidate distributions, whose variance-covariance matrices are estimated with the last 1000 iterations of a previously run single-updating MCMC algorithm with 2000 iterations. The whole algorithm is programmed in MATLAB (The MathWorks, Inc.) using the internal Gaussian and uniform random number gener-

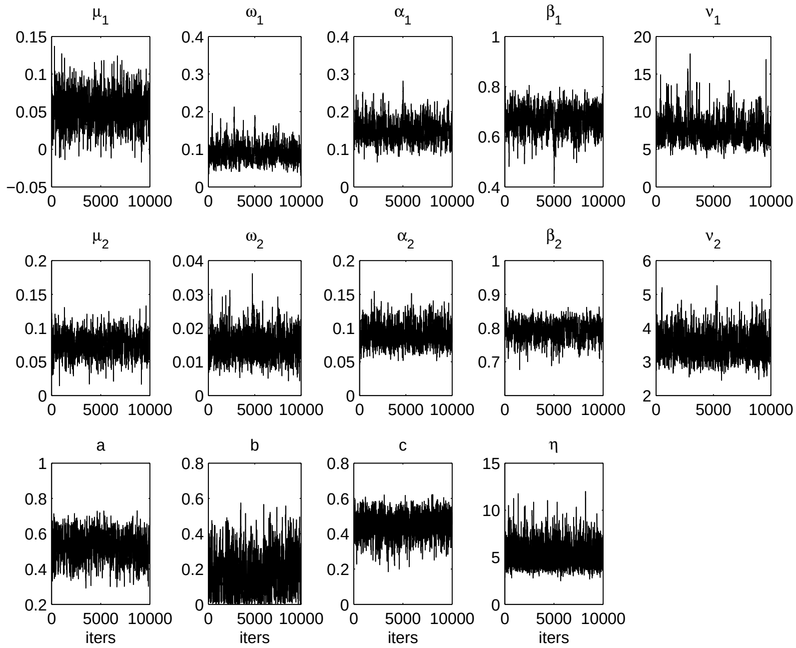

Figure 1: Convergence diagrams of the posterior samples of each parameter for the simulated series.

ators. The convergence behaviour of the algorithm can be examined in Figure 1, where the traces of the posterior samples of each model parameter are shown indicating a good mixing performance.

Table 1 presents the posterior mean and standard deviations for each model parameters obtained from the MCMC output. These are compared with the true parameter values in order to show the accuracy of the estimations.

Figure 2 shows the true values and posterior means of the marginal volatilities, h_{it}, for $i = 1, 2$, and the Kendall's tau, τ_t, for the last 100 observations, $t = 900, \ldots, 1000$, together with their 95% credible intervals. Observe the accuracy of the estimations and that the Bayesian credible intervals always include the true values of h_{it} and τ_t for all time periods. Predictions for the one-step ahead volatilities, $h_{i,T+1}$, and for the one-step ahead Kendall's tau, τ_{T+1}, where $T + 1 = 1001$, are also shown Figure 2. Note that the respective predictive intervals also include the true values for $h_{i,T+1}$, for $i = 1, 2$, and τ_{T+1}.

Table 2 shows the Bayesian estimations and 95% predictive intervals for the VaR of the one-step ahead portfolio, $\delta y_{1T+1} + (1 - \delta)\, y_{2T+1}$, for different values of δ and π. These are obtained using

Table 1: Parameter estimation results for the simulated series.

Parameter	True value	Posterior Mean	Posterior Std.
μ_1	0.05	0.0529	0.0214
μ_2	0.08	0.0761	0.0154
ω_1	0.05	0.0910	0.0233
ω_2	0.01	0.0151	0.0039
α_1	0.15	0.1494	0.0290
α_2	0.10	0.0908	0.0154
β_1	0.75	0.6680	0.0496
β_2	0.80	0.7914	0.0250
ν_1	6.00	7.4503	1.5892
ν_2	3.00	3.4931	0.3695
η	4.00	5.1496	1.1712
a	0.50	0.5235	0.0755
b	0.20	0.1820	0.1106
r_{12}	0.50	0.4459	0.0641

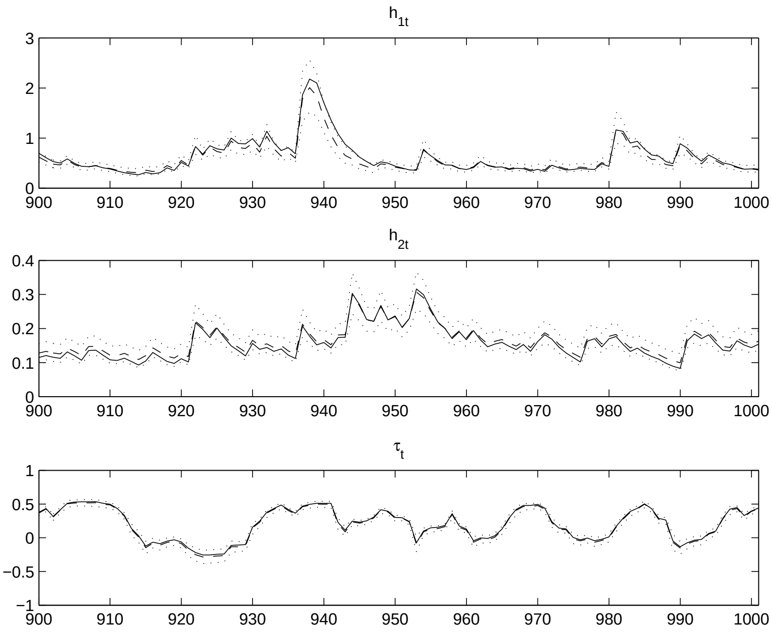

Figure 2: True (solid lines), predictive mean (dashed lines) and 95% credible intervals (dotted lines) for the volatilities, h_{1t} and h_{2t}, and for the Kendall's tau coefficients, τ_t, for $t = 900, \ldots, 1001$, for the simulated series.

Table 2: True values, predictive means and 95% predictive intervals for the one-step ahead VaR of different portfolios, for the simulated data.

	$\pi = 0.01$			$\pi = 0.05$		
	True	Mean	95% Interval	True	Mean	95% Interval
$\delta = 0$	1.7056	1.7482	$(0.8881, 3.6909)$	0.8464	0.8346	$(0.5383, 1.2389)$
$\delta = .25$	1.5982	1.6352	$(0.8904, 3.2575)$	0.8396	0.8252	$(0.5427, 1.1963)$
$\delta = .5$	1.5965	1.6075	$(0.9336, 2.9834)$	0.8895	0.8744	$(0.5933, 1.2328)$
$\delta = .75$	1.6725	1.6748	$(1.0329, 2.8430)$	0.9903	0.9746	$(0.6803, 1.3456)$
$\delta = 1$	1.8299	1.8308	$(1.1665, 2.9177)$	1.1288	1.1127	$(0.7909, 1.5165)$

the simulation procedure described in Section 5 with $M = 100$ replicas for each MCMC iteration. These estimations are compared with the true VaR values which are obtained for the true model parameters. Note that the estimated values of the VaR are very close to the true values, which are always inside the predictive intervals.

Finally, Figure 3 illustrates the Bayesian estimation of the mean-VaR efficient frontier with the corresponding 95% predictive region, obtained as described in Section 5. It is compared with the theoretical mean-VaR efficient frontier obtained for the true model parameters. Observe that the estimated curve is very similar to the theoretical curve, which lies inside the confidence region.

6.2 Real data

In this section, we apply our Bayesian procedure to the daily closing prices of the Dow Jones Industrial Average and DAX indices for the period 07/Sep/1998 to 07/Sep/2004. The log return bivariate series, whose sample size is $T = 1543$, is plotted in Figure 4. This sample has been analyzed previously in Arakelian and Dellaportas (2006) using a copula threshold model which changes discretely over time. Their model predicts four structural breaks in the dependence structure of the series. Alternatively, we show here that our approach can capture the temporal dependence of the series using the time-varying copula GARCH model described in Section 4, which change continuously across time.

As in the previous section, we run the proposed MCMC method for 20000 iterations discarding the first 10000 as burn-in iterations. Table 3 shows the posterior means and standard deviations

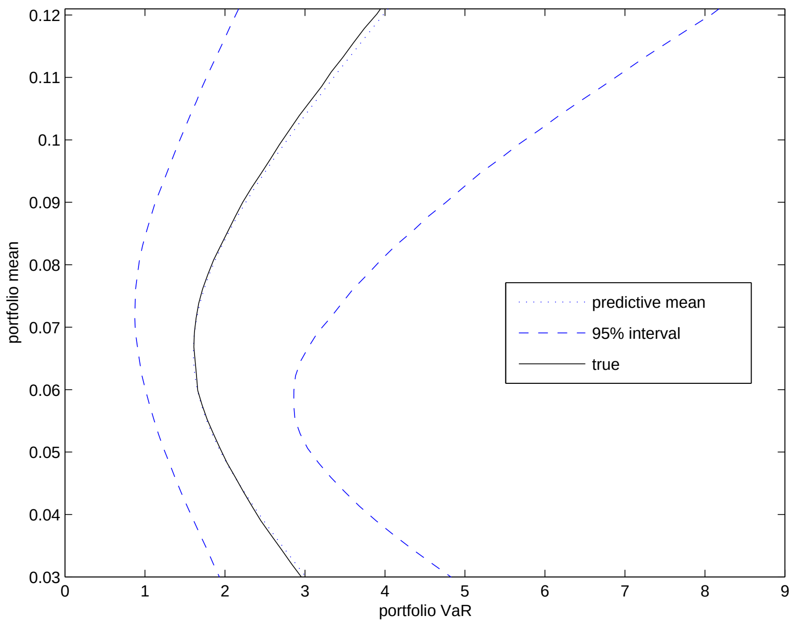

Figure 3: Bayesian estimation, 95% predictive region and true value for the mean-VaR efficient frontier for the simulated series.

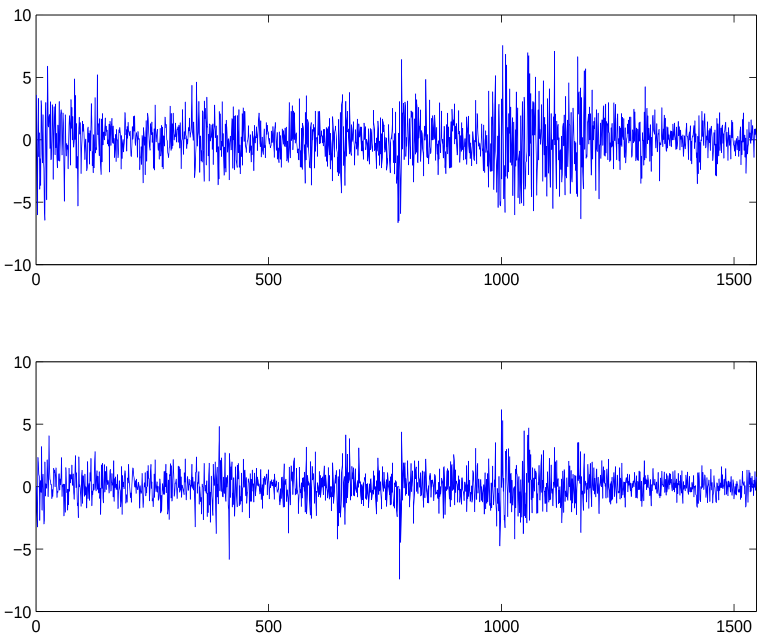

Figure 4: Daily returns of the DAX (top) and Dow Jones Industrial Average (bottom) indices.

Table 3: Parameter estimation results for the Dow Jones and DAX indices.

Parameter	Posterior Mean	Posterior Std.
μ_1	0.0236	0.0364
μ_2	0.0374	0.0249
ω_1	0.0441	0.0143
ω_2	0.0204	0.0071
α_1	0.0724	0.0112
α_2	0.0508	0.0099
β_1	0.9061	0.0136
β_2	0.9181	0.0146
ν_1	157.12	1501.2
ν_2	10.549	2.4754
η	8.6929	2.5068
a	0.0225	0.0079
b	0.9684	0.0135
r_{12}	0.5806	0.1467

obtained from this algorithm. Observe that the large value for the posterior mean and standard deviation of the degrees of freedom, ν_1, of the DAX returns indicates that the innovations, ϵ_{1t}, of this marginal series may be normally distributed, while the relatively small value for ν_2 indicates that the tails of the innovations of the Dow Jones returns are longer and the normal distribution is not appropriate in this case. Then, it is clear that the multivariate normal or multivariate t-distributions would not be adequate in this case to describe the multivariate distribution of the innovation process.

Now, we examine the temporal dependence between the two time series. Figure 5 illustrates the posterior mean of the Kendall's tau, τ_t, as a function of time, t, which has been obtained as described in Section 4, see (15). In order to assess the quality of these estimations, following Arakelian and Dellaportas (2006), we compare these posterior means with the sample estimates of the Kendall's tau obtained with two different rolling windows of sizes 100 and 250 observations. Observe that the dynamics of the Bayesian posterior means are very similar to the sample estimations.

Figure 6 shows the estimated VaR of the one-step ahead portfolio, $\delta \times DAX_{T+1} + (1-\delta) \times DowJ_{T+1}$, for different values of the weight δ, ranging in the interval $(0,1)$, for $\pi = 0.01$. We compare these estimations, obtained with our time-varying copula GARCH model, with those

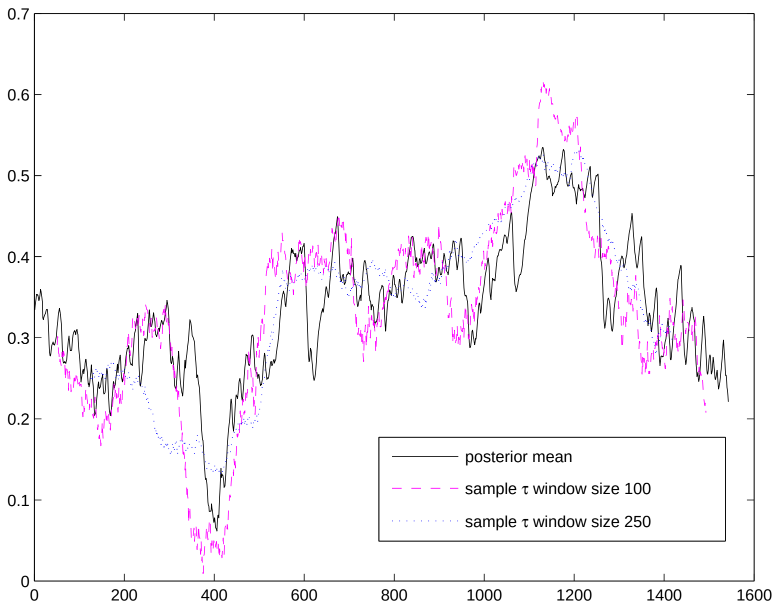

Figure 5: Posterior means of the Kendall's τ_t as a function of t compared with two window-based sample estimates, for the DAX and Dow Jones indices.

obtained assuming a constant copula model, where a and b are assumed to be equal to cero in (7) and the resulting posterior mean and standard deviation for r_{12} are 0.5048 and 0.0217, respectively. Observe that the predictive means of the one-step ahead VaR are quite different if we impose that the copula is constant rather than assuming a time-varying copula function. Note also that, as commented in Section 5, we can make use of the VaR estimations for portfolio selection, as we may be interested in that portfolio which minimizes the VaR. Thus, our time-varying copula model predicts that the optimal portfolio should assign a weight of approximately $\delta = 0.34$, while the constant copula model predicts that $\delta = 0.28$.

Finally, Figure 7 illustrates the Bayesian estimations for the mean-VaR efficient frontiers, assuming a constant and a time-varying copula model. Observe that again the predictive curves are quite different if we impose that the copula is constant rather than assuming a time-varying copula function.

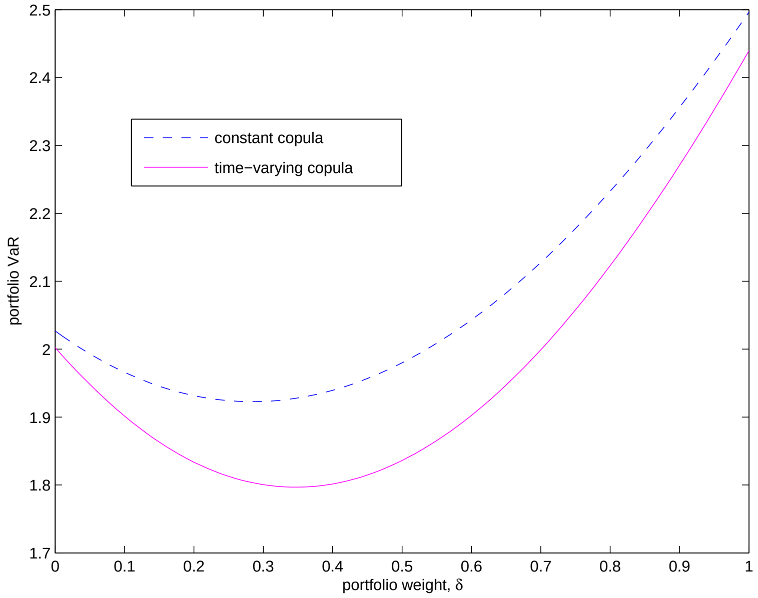

Figure 6: Bayesian estimation of the one-step ahead VaR for the portfolio, $\delta \times DAX_{T+1} + (1-\delta) \times DowJ_{T+1}$, as a function of δ, for $\pi = 0.01$, using a constant and a time-varying copula function.

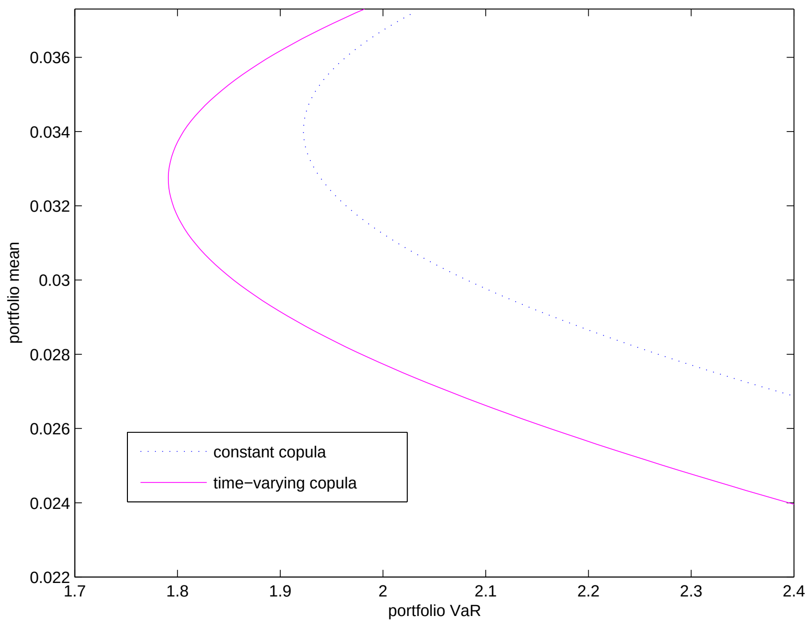

Figure 7: Bayesian estimations for the mean-VaR efficient frontiers, assuming a constant and a time-varying copula model.

7 Conclusions and extensions

In this paper, we have proposed a Bayesian procedure for the analysis of multivariate time series. Time-varying copula models have been considered to describe the structure of temporal dependence in the joint distributions. Our approach allows for the simultaneous estimation of the marginal and copula parameters, which is in contrast with the classical two-stage estimation procedures. We have described how to make Bayesian prediction of volatilities, various dependence measures, portfolio VaR and mean-VaR efficient frontiers for portfolio selection. The method has been illustrated with simulated and real financial time series.

Although, in this article, we have considered multivariate t-copula-GARCH models with Student's t-distributed marginals, the same approach can be straightforwardly extended to other alternative or more general models. For example, we can assume that the innovations follow a mixture of two Gaussian distributions instead of a t-distribution, as in Ausin and Galeano (2007). This mixture model for the marginals is known to be statistically more stable and avoids the use of informative priors for the degrees of freedom parameter as in (10). The disadvantage is that the number of parameters is increased in one unit for each marginal variable.

Other alternative models can be constructed and compared using the plethora of copula functions existing in the literature. For example, the t-copula model assumes a symmetric dependence structure and inference could be improved in certain circumstances if it is replaced by the Clayton copula, see e.g. Nelsen (2006). Furthermore, we could construct more flexible models based on mixtures of copula functions which could capture most of the tail dependence considered in the literature. Inference for these general families could be performed by defining an MCMC algorithm that visits all the copula functions included in the mixture, "chooses" the best one or a subset of best models and provide a coherent way of combining results with different copulas. These ideas are related to the Bayesian selection method for copulas proposed in Huard et al. (2006). Mixtures of copulas have been considered for dynamic models in Rodriguez (2003) and Hu (2006), using classical techniques.

References

Arakelian, V., Dellaportas, P., (2006). Contagion tests via copula threshold models. Discussion paper.

Ausin, M. C., Galeano, P., (2007). Bayesian estimation of the Gaussian mixture GARCH model. *Computational Statistics and Data Analysis,* **51**, 2636-2652.

Bartram, S., Taylor, S., and Wang, Y. H. (2007). The euro and European financial market integration. *Journal of Banking and Finance*, **31**, 1461-1481.

Bauwens, L. and Lubrano, M. (1998). Bayesian inference on GARCH models using the Gibbs sampler. *Econometrics Journal*, **1**, C23-C46.

Bauwens, L., Laurent, S., and Rombouts, J., (2006). Multivariate GARCH Models: A Survey. *Journal of Applied Econometrics*, **21**, 79-109.

Bollerslev, T., (1987). A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return. *Review of Economics and Statistics*, 69, 542-547.

Dias, A., and Embrechts, P. (2004). Dynamic copula models for multivariate high-frequency data in finance. Discussion paper.

Embrechts, P., McNeil, A., and Straumann, D., (1999). Correlation: pitfalls and alternatives. *Risk*, 5, 69-71.

Gaivoronski, A. A., Pflug, G., (2005). Value-at-Risk in Portfolio Optimization: Properties and Computational Approach. *Journal of Risk*, **7**, 1-31.

Hu, L. (2006). Dependence patterns across financial markets: A mixed copula approach. *Applied Financial Economics*, **10**, 717-729.

Huard, D., Évin, G., and Favre, A. C., (2006). Bayesian copula selection, *Computational Statistics and Data Analysis,* **51**, 809-822.

Jondeau, E., and Rockinger, M., (2006). The Copula-GARCH model of conditional dependencies: An international stock-market application. *Journal of International Money and Finance*, **25**, 827-853

Jorion, P., (2000). Value at Risk: The New Benchmark for Managing Financial Risk. second ed. McGraw-Hill, NewYork.

Kiesel, R., and Kleinow, T., (2002). Sensitivity analysis of credit portfolio models. In W. Härdle, T. Kleinow, G. Stahl (eds.), *Applied Quantitative Finance*, Springer.

Markowitz, H., (1952). Portfolio selection, *Journal of Finance*, **8**, 77-91.

Nelsen, R. B., (2006). *An introduction to copulas*, Springer, New York, 2nd ed.

Patton, A. (2006). Estimation of multivariate models for time series of possibly different lengths. *Journal of Applied Econometrics*, **21,** 147-173.

Rodriguez, J. C. (2003). Measuring financial contagion: A copula approach. Tech. rep., EURANDOM.

Schweizer, B., and Sklar, A., (1983). *Probabilistic Metric Spaces*, Elsevier Science, New York.

Tierney, L., (1994). Markov chains for exploring posterior distributions. *Annals of Statistics*, **22**, 1701-1762.

Tse, Y. K., and Tsui, A. K. C., (2002). A Multivariate Generalized Autoregressive Conditional Heteroscedasticity Model with Time-Varying Correlations. *Journal of Business and Economics Statistics*, **20**, 351-362.

Vrontos, I.D., Dellaportas, P. and Politis, D. N., (2001). Full bayesian inference for GARCH and EGARCH models. *Journal of Business and Economic Statistics*, **18**, 187-198.