Computational Statistics and Data Analysis 54 (2010) 2383–2399

Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

Time-varying joint distribution through copulas

M. Concepcion Ausin [a,*], Hedibert F. Lopes [b]

[a] Department of Statistics and Operations Research, Universidad Complutense de Madrid, Spain
[b] Booth School of Business, University of Chicago, USA

ARTICLE INFO

Article history:
Available online 21 March 2009

ABSTRACT

The analysis of temporal dependence in multivariate time series is considered. 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. Full Bayesian inference is developed where the whole set of model parameters is estimated simultaneously. This represents an essential difference from previous approaches in the literature where the marginal and the copula parameters are estimated separately in two consecutive steps. Moreover, a Bayesian procedure is proposed for the estimation of several measures of risk, such as the variance, Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) of a portfolio of assets, providing point estimates and predictive intervals. The proposed copula model enables to capture the dependence structure between the individual assets which strongly influences these risk measures. Finally, the problem of optimal portfolio selection based on the estimation of mean–variance, mean–VaR and mean–CVaR efficient frontiers is also addressed. The proposed approach is illustrated with simulated and real financial time series.

1. Introduction

There has 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 time 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 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. (2007), Patton (2006), Jondeau and Rockinger (2006), Roch and Alegre (2006), Liu and Luger (2009).

* Corresponding address: Department of Statistics and Operations Research, Universidad Complutense de Madrid, Facultad de Matemáticas, Plaza de Ciencias, 3, Ciudad Universitaria, 28040 Madrid, Spain Tel.: +34 91 394 4377; fax: +34 91 394 4406.
E-mail addresses: mcausin@mat.ucm.es (M.C. Ausin), hlopes@chicagobooth.edu (H.F. Lopes).

0167-9473/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.csda.2009.03.008

Most of previous works using copula GARCH models do 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). Recently, Liu and Luger (2009) have proposed an algorithm to improve efficiency in copula GARCH models using maximum-likelihood estimation.

One of the main interests in modelling multivariate financial time series is portfolio management. The Value-at-Risk (VaR, see e.g. Jorion (2000)) and the Conditional Value-at Risk (CVaR, see e.g. Artzner et al. (1997)) have become important and widely used measures of the risk inherent in asset portfolios. The CVaR, which is also known as Mean Excess Loss, Mean Shortfall, or Tail VaR, is a coherent risk measure, while VaR is coherent only when it is based on normal distributions, e.g. Artzner et al. (1997). Clearly, the VaR and the CVaR of a portfolio depends on the behavior 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 and CVaR calculation, see e.g. Embrechts et al. (1999a,b) and 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 and CVaR. Alternatively, copulas provide a useful tool to model tail dependence and obtain precise VaR and CVaR estimations.

Besides the applications of VaR and CVaR in risk measurement, they also provide useful tools 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 (1952) based on the variance–covariance matrix. As pointed out in Embrechts et al. (1999a), 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 and CVaR in portfolio optimization is a relatively novel alternative, see e.g. Gaivoronski and Pflug (2005), Winker and Maringer (2007) and Rockafellar and Uryasev (2000). Using this approach, the optimal portfolio is the one which minimizes the VaR or the CVaR, respectively, subject to achieving a specified level of expected return. The analogous to the classical mean–variance efficient frontiers are the mean–VaR and mean–CVaR efficient frontiers which give the minimum VaR and minimum CVaR, respectively, for given expected return. Note that for the case of elliptical distributions, the portfolio minimizing VaR and CVaR 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 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 (2008). Moreover, we describe how to obtain Bayesian estimation of portfolio variance, VaR and CVaR and predictive intervals, which are much more informative than simple point forecasts. We also show how to obtain optimal portfolios based on minimizing the variance, the VaR and the CVaR and how to obtain confidence regions for the mean–variance, mean–VaR and mean–CVaR efficient frontiers.

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 variance, VaR and CVaR are obtained and used for optimal portfolio allocation. Bayesian prediction of 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 $C\left(u_1, \ldots, u_p\right)$, 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, $F\left(x_1, \ldots, x_p\right)$, whose marginals are given by $F_1(x_1), \ldots, F_p\left(x_p\right)$, can be written as,

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

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,

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

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

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

provided the density exists, 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

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

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

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

where $\mathbf{u}=\left(u_1,\ldots,u_p\right)$ and Φ^{-1} is the inverse of 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. Therefore, 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

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

where t_η^{-1} denotes the inverse of 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.

It is important to note that various dependence measures between two random variables depend only on their copula function. For example, an important measure of dependence is 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. 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, Kendall's tau only depends on the copula function and it can be shown that,

$$\tau=4\int_0^1\int_0^1 C\left(u_1,u_2\right)c\left(u_1,u_2\right)\mathrm{d}u_1\mathrm{d}u_2-1.$$

Kendall's tau admits the same form for the bivariate Gaussian copula and for the bivariate 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,

$$\lambda_u=\lim_{q\to1}P\left(X_2>F_{X_2}^{-1}(q)\mid X_1>F_{X_1}^{-1}(q)\right),\qquad\lambda_l=\lim_{q\to0}P\left(X_2\le F_{X_2}^{-1}(q)\mid X_1\le F_{X_1}^{-1}(q)\right),$$

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

$$\lambda_u=\lim_{q\to1}\frac{1-2q+C(q,q)}{1-q},\qquad\lambda_l=\lim_{q\to0}\frac{C(q,q)}{q}.$$

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

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

where $t_{\eta+1}$ denotes the cumulative distribution function of the standard univariate Student-t distribution with $\eta+1$ degrees of freedom.

3. Copula GARCH models

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

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

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(1,1)-model,

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

where h_{it} is the conditional variance of y_{it} given the previous information $I_{i,t-1} = \{y_{i,t-1}, y_{i,t-2}, \ldots\}$, 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 $i = 1, \ldots, p$, given by

$$f\left(\epsilon_{it}\right) = \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)}{\Gamma\left(\frac{\nu_i}{2}\right)\sqrt{\nu_i \pi}} \left(1 + \frac{\epsilon_{it}^2}{\nu_i}\right)^{-\frac{\nu_i+1}{2}},$$

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 $F_i\left(y_{it} \mid \mu_i, h_{it}\right) = t_{\nu_i}\left(\left(y_{it} - \mu_i\right) h_{it}^{-1/2}\right)$, 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,

$$c_{\eta,R_t}^t\left(u_{1t}, \ldots, u_{pt}\right) = \frac{f_{\eta,R_t}^t\left(t_\eta^{-1}\left(u_{1t}\right), \ldots, t_\eta^{-1}\left(u_{pt}\right)\right)}{\prod_{i=1}^p f_\eta^t\left(t_\eta^{-1}\left(u_{it}\right)\right)},$$

where $u_{it} = F_i\left(y_{it} \mid \mu_i, h_{it}\right)$, for $i = 1, \ldots, p$; f_{η,R_t}^t is the joint density of the standard multivariate Student-t distribution with η degrees of freedom and correlation matrix R_t and f_η^t 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 corresponds 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 $\{\mathbf{x}_{t-1}, \ldots, \mathbf{x}_{t-m}\}$, with $m \geq 2$, where,

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

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 Eq. (7) is based on 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 matrix 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 time series can be computed by,

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

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

$$\begin{aligned} l(\boldsymbol{\theta} \mid \mathbf{y}_t) &= \prod_{t=1}^{T} f(\mathbf{y}_t \mid \boldsymbol{\mu}, \mathbf{h}_t) \\ &= \prod_{t=1}^{T} \frac{\Gamma\left(\frac{\eta+p}{2}\right)\Gamma\left(\frac{\eta}{2}\right)^{p-1}}{\Gamma\left(\frac{\eta+1}{2}\right)^{p}} \frac{\left(1+\frac{\mathbf{x}_t' R_t^{-1}\mathbf{x}_t}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}} \prod_{i=1}^{p}\left(1+\frac{x_{it}^2}{\eta}\right)^{\frac{\eta+1}{2}} \\ &\quad \times \prod_{i=1}^{p} \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)}{\Gamma\left(\frac{\nu_i}{2}\right)\sqrt{\pi\nu_i h_{it}}}\left(1+\frac{(y_{it}-\mu_i)^2}{\nu_i h_{it}}\right)^{-\frac{\nu_i+1}{2}} \end{aligned} \tag{10}$$

where $\boldsymbol{\theta} = \{(\mu_i, \omega_i, \alpha_i, \beta_i, \nu_i)_{i=1}^{p}, (a, b, R, \eta)\}$ and $\mathbf{x}_t = (x_{1t}, \ldots, x_{pt})$.

4. Bayesian inference and prediction

We want to make inference for the model parameters, $\boldsymbol{\theta}$, 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,

$$f(\nu_i) \propto \frac{1}{1+\nu_i^2}, \quad \nu_i > 0, \tag{11}$$

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 (11) for the degrees of freedom η of the t-copula.

Given an observed series, $\mathbf{y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_T\}$, and the priors specified above, the evaluation of the joint posterior distribution $f(\boldsymbol{\theta} \mid \mathbf{y})$ is analytically intractable. Therefore, 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 given in (10), it is straightforward to see that the conditional posterior distribution of the model parameters are given by,

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

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

$$f(\nu_i \mid \cdot) \propto \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)^{T} \nu_i^{-T/2}}{\Gamma\left(\frac{\nu_i}{2}\right)^{T}(1+\nu_i)^2} \prod_{t=1}^{T} \frac{\left(1+\frac{\mathbf{x}_t' R_t^{-1}\mathbf{x}_t}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}\left(1+\frac{x_{it}^2}{\eta}\right)^{-\frac{\eta+1}{2}}\left(1+\frac{(y_{it}-\mu_i)^2}{\nu_i h_{it}}\right)^{\frac{\nu_i+1}{2}}}, \tag{13}$$

for $i = 1, 2$,

$$f(\eta \mid \cdot) \propto \frac{\Gamma\left(\frac{\eta+p}{2}\right)^{T}\Gamma\left(\frac{\eta}{2}\right)^{T(p-1)}}{\Gamma\left(\frac{\eta+1}{2}\right)^{Tp}(1+\eta)^2} \prod_{t=1}^{T} \frac{\left(1+\frac{\mathbf{x}_t' R_t^{-1}\mathbf{x}_t}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}\prod_{i=1}^{p}\left(1+\frac{x_{it}^2}{\eta}\right)^{-\frac{\eta+1}{2}}}, \tag{14}$$

and,

$$f(d \mid \cdots) \propto \prod_{t=1}^{T} \frac{1}{\sqrt{|R_t|}}\left(1+\frac{\mathbf{x}_t' R_t^{-1}\mathbf{x}_t}{\eta}\right)^{-\frac{\eta+p}{2}}, \tag{15}$$

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 $\boldsymbol{\phi}_i = (\mu_i, \omega_i, \alpha_i, \beta_i)$ using a multivariate Metropolis step for each $i = 1, \ldots, p$. We generate a candidate vector from a multivariate normal distribution $N\left(\boldsymbol{\phi}_i^{(n)}, c\,\Sigma\right)$, where $\boldsymbol{\phi}_i^{(n)}$ 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.

4.1. Two-step Bayesian approach

In order to compare our approach with an analogous two-stage estimation procedure, we design here a two-step Bayesian algorithm. Firstly, we estimate each marginal series independently considering a univariate GARCH model with Student-t innovations. For each marginal series, we have five parameters to estimate $\boldsymbol{\theta}_i = (\mu_i, \omega_i, \alpha_i, \beta_i, \nu_i)$, for $i = 1, 2$, and the likelihood is given by,

$$l(\boldsymbol{\theta}_i \mid \mathbf{y}_t) \propto \prod_{t=1}^{T} \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)}{\Gamma\left(\frac{\nu_i}{2}\right)\sqrt{\nu_i h_{it}}} \left(1 + \frac{(y_{it}-\mu_i)^2}{\nu_i h_{it}}\right)^{-\frac{\nu_i+1}{2}}.$$

Then, it is straightforward to define an MCMC algorithm to sample from the joint posterior of $\boldsymbol{\theta}_i$ for each series with a Gibbs sampling scheme, where each parameter is updated using a Metropolis–Hastings step as described before.

Note that for each element of the Monte Carlo sample of size N, we can obtain a set of residuals,

$$\mu_i^{(n)}, \omega_i^{(n)}, \alpha_i^{(n)}, \beta_i^{(n)}, \nu_i^{(n)} \Longrightarrow \epsilon_{it}^{(n)} = \frac{y_{it} - \mu_i^{(n)}}{\sqrt{h_{it}^{(n)}}}, \quad \text{for } t = 1, \ldots, T$$

for $n = 1, \ldots, N$. Thus, we can estimate the residual for each time t and the degrees of freedom of the marginal t-distribution for each series as follows,

$$\hat{\epsilon}_{it} = \frac{1}{N}\sum_{n=1}^{N} \epsilon_{it}^{(n)}, \quad \text{and} \quad \hat{\nu}_i = \frac{1}{N}\sum_{n=1}^{N} \nu_i^{(n)},$$

for $i = 1, 2$. In order to estimate the copula parameters, $\boldsymbol{\theta}_c = \left(a, b, r_{ij}, \eta\right)$, we plug in these estimations in the likelihood (10) by approximating

$$\hat{x}_{it} = t_{\eta}^{-1}\left(t_{\hat{\nu}_i}\left(\hat{\epsilon}_{it}\right)\right),$$

and obtaining the following pseudo-likelihood function for $\boldsymbol{\theta}_c$,

$$l(\boldsymbol{\theta}_c \mid \mathbf{y}_t) = \prod_{t=1}^{T} \frac{\Gamma\left(\frac{\eta+p}{2}\right)\Gamma\left(\frac{\eta}{2}\right)^{p-1}}{\Gamma\left(\frac{\eta+1}{2}\right)^{p}} \frac{\left(1 + \frac{\hat{\mathbf{x}}_t' R_t^{-1} \hat{\mathbf{x}}_t}{\eta}\right)^{-\frac{\eta+p}{2}}}{\sqrt{|R_t|}} \prod_{i=1}^{p}\left(1 + \frac{\hat{x}_{it}^2}{\eta}\right)^{\frac{\eta+1}{2}}, \tag{16}$$

where $\hat{\mathbf{x}}_t = (\hat{x}_{1t}, \ldots, \hat{x}_{pt})$. Now, we can construct another Markov chain to sample from the posterior distribution of $\boldsymbol{\theta}_c$ using analogous Metropolis–Hastings steps as the previously considered for updating a, b, r_{ij} and η but now based on the pseudo-likelihood (16).

4.2. Prediction of volatilities and dependence measures

In financial time series, one frequently observes 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 (2008). Thus, in multivariate GARCH models, it is important to estimate the 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% Bayesian confidence 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 elements, r_{ijt}, of the parameter 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 Kendall's tau correlation, given in (5), for each time t using,

$$E\left[\tau_t \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N}\frac{2}{\pi}\arcsin\rho_t^{(n)}, \tag{17}$$

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 model parameters in the MCMC sample of size N. Also, we can approximate the posterior median of τ_t and Bayesian confidence intervals as before. Analogously, we can estimate the posterior mean, median and Bayesian confidence intervals for the coefficient of tail dependence λ_t as a function of t using (6) as follows,

$$E\left[\lambda_t \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} 2t_{\eta^{(n)}+1}\left(\frac{-\sqrt{\eta^{(n)}+1}\sqrt{1-\rho_t^{(n)}}}{\sqrt{1+\rho_t^{(n)}}}\right). \tag{18}$$

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. Estimation of risk measures and portfolio allocation

In this section, we consider three of the most usual measures of the risk of a portfolio, namely, the portfolio variance, VaR and CVaR.

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 log return series,

$$\sum_{i=1}^{p}\delta_i y_{it}, \quad \text{where } \sum_{i=1}^{p}\delta_i = 1, \tag{19}$$

the t-period q% VaR is given by,

$$q = \Pr\left(\sum_{i=1}^{p}\delta_i y_{it} \le -\mathrm{VaR}_t\right), \tag{20}$$

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

As an alternative measure of risk, the CVaR is known to have better properties than VaR, see Artzner et al. (1997). By definition, with respect to a specified probability level, q, the CVaR is the conditional expectation of losses above the q% VaR. More specifically, given the portfolio (19), the t-period q% CVaR is given by,

$$\mathrm{CVaR}_t = -E\left[\sum_{i=1}^{p}\delta_i y_{it} \,\middle|\, \sum_{i=1}^{p}\delta_i y_{it} \le -\mathrm{VaR}_t\right],$$

where VaR_t is the q% VaR at time t given in (20).

5.1. Bayesian procedure for the estimation of risk measures

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 and CVaR, that is, the $(T+1)$-period VaR and CVaR. We can obtain consistent estimators of these quantities using,

$$E\left[\mathrm{VaR}_{T+1} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N}\mathrm{VaR}_{T+1}^{(n)}, \tag{21}$$

and

$$E\left[\mathrm{CVaR}_{T+1} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N}\mathrm{CVaR}_{T+1}^{(n)}, \tag{22}$$

where $\mathrm{VaR}_{T+1}^{(n)}$ and $\mathrm{CVaR}_{T+1}^{(n)}$ are the one-step-ahead VaR and CVaR, respectively, given the model parameters $\boldsymbol{\theta}^{(n)}$ of the n-th MCMC iteration. Although the conditional distribution of the multivariate series $\left(y_{1,T+1}, \ldots, y_{p,T+1}\right)$ is explicit given the model parameters, it is not straightforward to derive the distribution of the portfolio (19). Therefore, it is complicated to

obtain an analytic expression for $\text{VaR}_{T+1}^{(n)}$ and $\text{CVaR}_{T+1}^{(n)}$ given $\boldsymbol{\theta}^{(n)}$. However, these can be easily approximated by generating values from our copula GARCH model as follows. For each value of the parameters $\boldsymbol{\theta}^{(n)}$, the values of $(h_{1,T+1}^{(n)}, \ldots, h_{p,T+1}^{(n)})$ and $R_{T+1}^{(n)}$ are known and then, we can generate M replications $\{(y_{1,T+1}^{(n,m)}, \ldots, y_{p,T+1}^{(n,m)})\}_{m=1}^{M}$ from the one-step-ahead density of the 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,

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

Then, the value of $\text{VaR}_{T+1}^{(n)}$ can be approximated by the negative value of the empirical q-quantile of the sample of portfolios $\left\{\sum_{i=1}^{p} \delta_i y_{i,T+1}^{(n,m)}\right\}_{m=1}^{M}$. Also, the value of $\text{CVaR}_{T+1}^{(n)}$ can be approximated by the negative value of the empirical mean of the following conditional sample of portfolios,

$$\left\{\sum_{i=1}^{p} \delta_i y_{i,T+1}^{(n,m)} \text{ such that } \sum_{i=1}^{p} \delta_i y_{i,T+1}^{(n,m)} \leq -\text{VaR}_{T+1}^{(n)}\right\}_{m=1}^{M}.$$

Now, we can estimate the posterior means (21) and (22) and obtain 95% predictive intervals for VaR and CVaR using the 0.025 and 0.975 quantiles of the posterior sample of $\text{VaR}^{(n)}$ and $\text{CVaR}^{(n)}$, respectively, 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. Finally, note that by using a similar approach we can also approximate the variance of the one-step-ahead portfolio using the empirical variance of the sample of portfolios $\left\{\sum_{i=1}^{p} \delta_i y_{i,T+1}^{(n,m)}\right\}_{m=1}^{M}$, for $n = 1, \ldots, N$.

Suppose now that we are interested in the prediction of multi-step-ahead variance, VaR and CVaR estimation. We cannot obtain samples directly from the multiple-step-ahead portfolio using the same procedure as for the one-step-ahead portfolio because the values of $R_{T+s}^{(n)}$ and $h_{i,T+s}^{(n)}$ given $\boldsymbol{\theta}^{(n)}$ are unknown for $s > 1$. However, we can use a simulation procedure similar to the proposed in Ausin and Galeano (2007) as follows. For each $\boldsymbol{\theta}^{(n)}$, the values for $R_{T+1}^{(n)}$ and $h_{i,T+1}^{(n)}$ are known and a sample $y_{i,T+1}^{(n,m)}$ can be generated as before. Sequentially, given $y_{i,T+s-1}^{(n,m)}$, the values for $R_{T+s}^{(n,m)}$ and $h_{i,T+s}^{(n,m)}$ are known and a sample $y_{i,T+s}^{(n,m)}$ can be generated leading to the sample of portfolios $\left\{\sum_{i=1}^{p} \delta_i y_{i,T+s}^{(n,m)}\right\}_{m=1}^{M}$, for $n = 1, \ldots, N$, which allows us to estimate the s-step-ahead variance, VaR and CVaR as before.

5.2. Bayesian procedure for portfolio allocation

We have shown how to estimate the one-step-ahead variance, VaR and CVaR of a given portfolio. A different problem is how to choose the optimal portfolio which minimizes the one-step-ahead variance, VaR or CVaR. Unfortunately, given the model parameters, it is not easy to obtain a closed expression for the optimal weights,

$$\delta_{opt,V} = \arg\min_{\delta}\left\{V\left(\sum_{i=1}^{p} \delta_i y_{i,T+1}\right) : \sum_{i=1}^{p} \delta_i = 1\right\}, \tag{23}$$

$$\delta_{opt,\text{VaR}} = \arg\min_{\delta}\left\{\text{VaR}_{T+1} : \sum_{i=1}^{p} \delta_i = 1\right\}, \tag{24}$$

$$\delta_{opt,\text{CVaR}} = \arg\min_{\delta}\left\{\text{CVaR}_{T+1} : \sum_{i=1}^{p} \delta_i = 1\right\}. \tag{25}$$

In fact, it is well known that the solutions to problems (24) and (25) are not analytically tractable even under the Gaussianity assumption. However, we can make use of numerical optimization procedures to approximate the optimal weights, $\delta_{opt,V}^{(n)}$, $\delta_{opt,\text{VaR}}^{(n)}$ and $\delta_{opt,\text{CVaR}}^{(n)}$ for each set of model parameters, $\boldsymbol{\theta}^{(n)}$, and then, obtain consistent estimators of the posterior mean of the optimal weights using,

$$E\left[\delta_{opt,V} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} \delta_{opt,V}^{(n)},$$

$$E\left[\delta_{opt,\text{VaR}} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} \delta_{opt,\text{VaR}}^{(n)},$$

and

$$E\left[\delta_{opt,\mathrm{CVaR}} \mid \mathbf{y}\right] \approx \frac{1}{N}\sum_{n=1}^{N} \delta_{opt,\mathrm{CVaR}}^{(n)}.$$

Also, 95% Bayesian confidence intervals can be obtained by just calculating .025 and .975 quantiles of the posterior samples of the optimal weights.

Observe that in the definition of portfolio, given in (19), we have not imposed that $0 \leq \delta_i$, $i = 1, \ldots, p$. However, if one wanted to impose that the portfolio weights are positive in order to prevent investors from short-selling, these restrictions should be included in (23), (24) and (25), and imposed in the numerical optimization procedure to approximate the optimal weights, $\delta_{opt,V}^{(n)}$, $\delta_{opt,\mathrm{VaR}}^{(n)}$ and $\delta_{opt,\mathrm{CVaR}}^{(n)}$ for each set of model parameters, $\boldsymbol{\theta}^{(n)}$.

Finally, assume that we are interested in the portfolio which minimizes the variance, the VaR or the CVaR 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 $(\delta_1, \ldots, \delta_p)$ which lead to that expected gain. Then, for each set of weights and for each set of model parameters, $\boldsymbol{\theta}^{(n)}$, we can approximate the $V_{T+1}^{(n)}$, $\mathrm{VaR}_{T+1}^{(n)}$ and $\mathrm{CVaR}_{T+1}^{(n)}$ as before, and choose the ones which minimizes the portfolio variance, VaR and CVaR. Repeating this procedure for a number of values for the expected gain, g, we can approximate the mean–variance, mean–VaR and mean–CVaR 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 are also interested in comparing the proposed one-step Bayesian approach with the two-stage Bayesian method described in Section 4.1. 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,

$$y_{1t} = 0.08 + \sqrt{h_{1t}}\epsilon_{1t}, \quad \epsilon_{1t} \sim t_3(0,1), \qquad y_{2t} = 0.05 + \sqrt{h_{2t}}\epsilon_{2t}, \quad \epsilon_{2t} \sim t_4(0,1),$$
$$h_{1t} = 0.01 + 0.04\left(y_{i,t-1} - 0.08\right)^2 + 0.91h_{i,t-1}, \qquad h_{2t} = 0.01 + 0.05\left(y_{2,t-1} - 0.05\right)^2 + 0.9h_{1,t-1},$$

and the following pattern for the time-varying copula parameter,

$$\rho_t = (1 - 0.02 - 0.96) \times 0.5 + 0.02\xi_{t-1} + 0.96\rho_{t-1},$$

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 20 000 iterations discarding the first 10 000 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, c) 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 two-stage MCMC algorithm described in Section 4.1. is also run using the same number of iterations. Both algorithms are programmed in MATLAB (The MathWorks, Inc.) using the internal Gaussian and uniform random number generators. The resulting samples are checked for convergence using the test proposed by Geweke (1992). Table 1 shows the values of the Geweke's statistic for each parameter obtained for the two algorithms. These indicate that convergence has been achieved in both cases because if the samples are drawn from the stationary distribution of the chain, the Geweke's statistic has an asymptotically standard normal distribution. Also, Fig. 1 shows the traces of the posterior samples of each model parameter obtained with the one-step Bayesian method. These indicate a good mixing performance of the Markov chain as it moves fluidly through all possible states. The trace plots obtained with the two-step method also indicate good mixing but are not reported to save space.

Table 2 presents the posterior mean and standard deviations for each model parameter obtained from the one-step and the two-step MCMC outputs. These are compared with the true parameter values in order to show the accuracy of the estimations. Observe that, although the posterior means obtained with the two methods are similar, the posterior standard deviations are in general smaller with the one-step approach than with the two-stage procedure, indicating that the full Bayesian approach leads to more precise estimations. We have also observed that the posterior distributions have shorter tails using the one-step method rather than the two-step procedure. This is illustrated in Fig. 2(top), where the histograms of the posterior samples of the parameter a obtained with the two methods are shown. Furthermore, there are also essential differences in the joint posterior distribution obtained with the one and the two-step procedures. This is also illustrated in Fig. 2(bottom), where the scatter plots of the posterior samples of the parameters μ_1 and μ_2 are shown for the one and the two stage approaches. Note that the posterior relationship existing between these two parameters, which is observed with

Table 1
Values of the Geweke's statistic for each parameter obtained for the one- and the two-step algorithms.

	μ_1	μ_2	ω_1	ω_2	α_1	α_2	β_1	β_2
One-step	0.2778	0.4943	0.3195	−0.4189	1.544	1.983	−1.173	−1.803
Two-step	0.3716	0.4838	0.2243	1.251	−0.064	0.2064	0.1634	−0.868
	ν_1	ν_2	a	b	c	η		
One-step	0.083	−0.758	1.201	−1.837	−0.4322	−0.5461		
Two-step	1.398	−0.091	0.222	0.622	−0.3058	1.076		

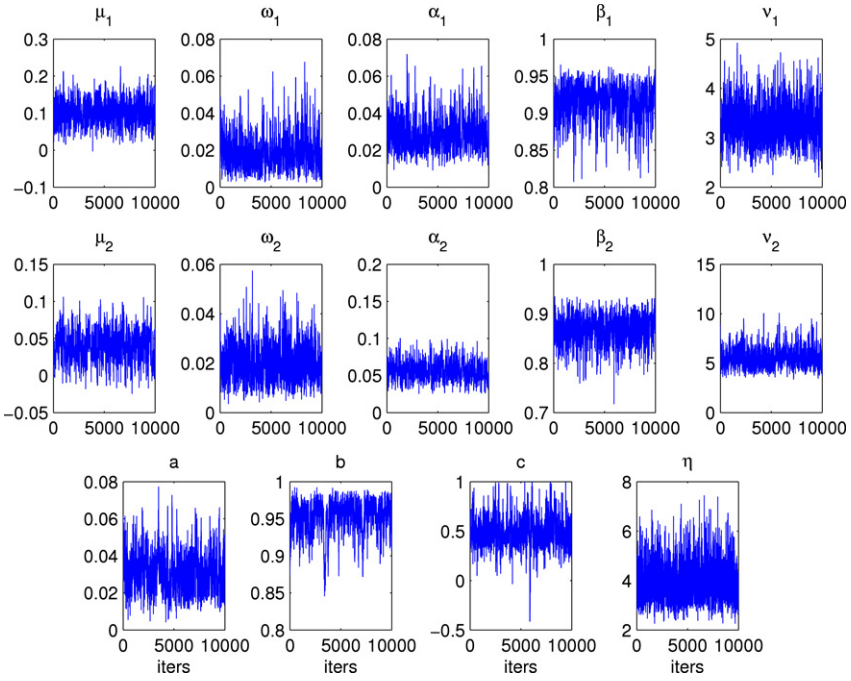

Fig. 1. Convergence diagrams of the posterior samples of each parameter for the simulated series.

Table 2
Parameter estimation results for the simulated series using the proposed one- and two-step Bayesian approaches.

Parameter	True value	One-step		Two-step	
		Posterior mean	Posterior std.	Posterior mean	Posterior std.
μ_1	0.08	0.1024	0.0309	0.1100	0.0309
μ_2	0.05	0.0398	0.0205	0.0474	0.0213
ω_1	0.01	0.0187	0.0090	0.0138	0.0094
ω_2	0.01	0.0198	0.0065	0.0263	0.0117
α_1	0.04	0.0284	0.0085	0.0246	0.0086
α_2	0.05	0.0562	0.0125	0.0562	0.0161
β_1	0.91	0.9144	0.0242	0.9265	0.0257
β_2	0.90	0.8691	0.0262	0.8488	0.0458
ν_1	3.00	3.2779	0.3598	3.1437	0.3404
ν_2	4.00	5.5228	0.8732	5.3716	0.9235
η	4.00	3.8728	0.6787	3.8935	0.6239
a	0.02	0.0305	0.0103	0.0340	0.0211
b	0.96	0.9527	0.0227	0.9298	0.1108
c	0.50	0.5010	0.1629	0.4933	0.1487

the one-step Bayesian approach, disappears with the two-step Bayesian method. Note that using a two-step procedure, the marginal series are assumed to be independent and this assumption clearly affects the shape of the joint posterior. Thus, Table 2 and Fig. 2 suggest that the one-step approach is in a sense more "efficient" than the two-step one. A similar result was found by Liu and Luger (2009) in a classical context.

We now consider Bayesian prediction for the volatilities and dependence measures. Fig. 3 shows the true values, posterior means and 95% Bayesian confidence intervals for the marginal volatilities, h_{it}, for $i = 1, 2$, Kendall's tau, τ_t, the copula parameter, ρ_t, and the tail dependence, λ_t, for the last 100 observations, $t = 900, \ldots, 1000$, obtained as described in

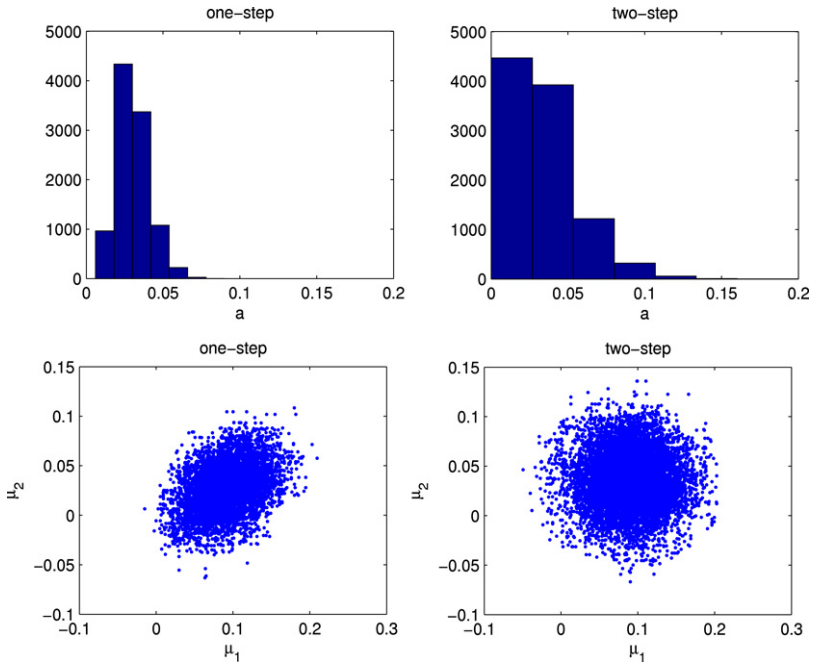

Fig. 2. Top: Histograms of the posterior samples of the parameter a obtained with the one- and the two-step methods. Bottom: Scatter plots of the the posterior samples of the parameters μ_1 and μ_2 for the one- and the two-step methods.

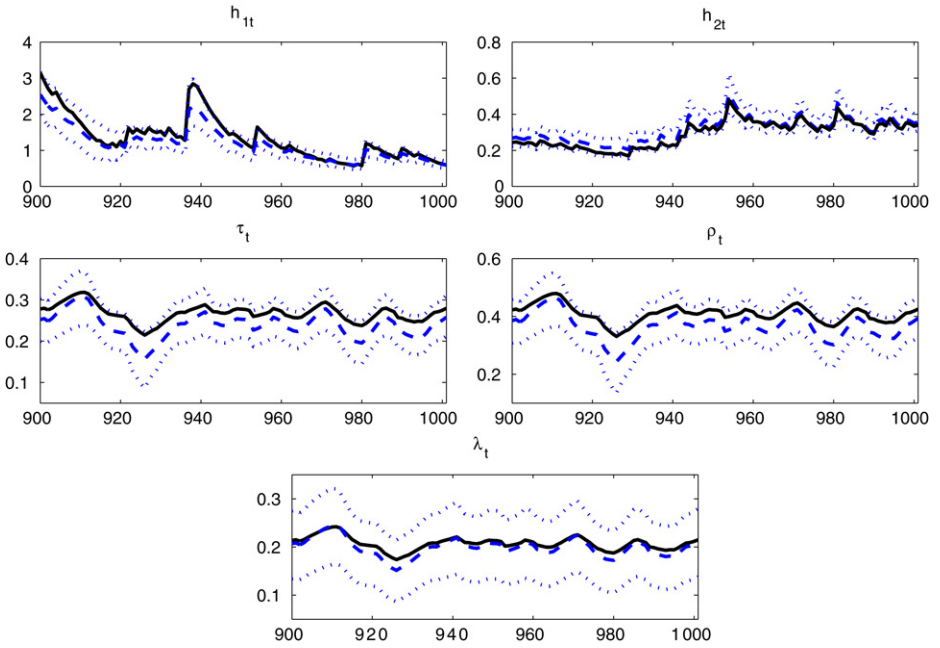

Fig. 3. True (solid lines), predictive mean (dashed lines) and 95% Bayesian confidence intervals (dotted lines) for the volatilities, h_{1t} and h_{2t}, Kendall's tau, τ_t, copula parameter, ρ_t, and tail dependence, λ_t, for $t = 900, \ldots, 1001$, for the simulated series.

Section 4.2. Observe that the Bayesian confidence intervals always include the true values of the volatilities and dependence measures for all time periods. Predictions for the one-step-ahead volatilities, $h_{i,T+1}$, and for the one-step-ahead dependence measures, τ_{T+1}, ρ_{T+1} and λ_t, where $T + 1 = 1001$, are also shown in Fig. 3. Note that their respective predictive intervals also include the corresponding true values.

Next, we perform Bayesian prediction of portfolio risk measures. Fig. 4 shows the Bayesian estimations and 95% predictive intervals for the CVar, VaR and variance of the one-step-ahead portfolio, $\delta y_{1T+1} + (1 - \delta)\, y_{2T+1}$, for different values of δ in the interval $(-1, 1)$ and $q = 0.05$ and $q = 0.01$. These are obtained using the simulation procedure described in Section 5

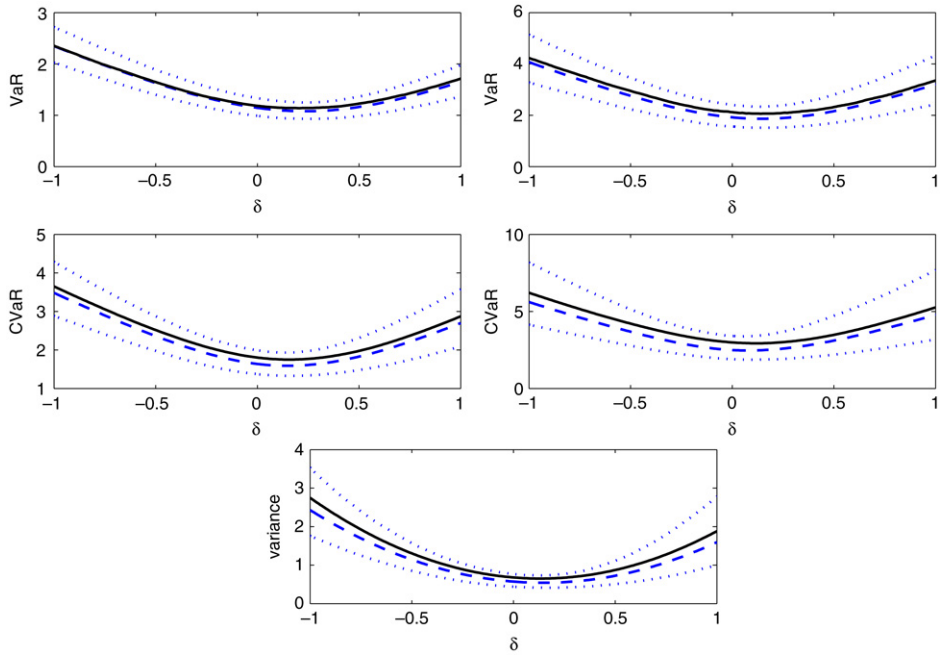

Fig. 4. True (solid lines), predictive mean (dashed lines) and 95% Bayesian confidence intervals (dotted lines) for the VaR, CVaR and variance of the one-step-ahead portfolio, $\delta y_{1T+1} + (1-\delta)\, y_{2T+1}$, for different values of δ and $q = 0.05$ (left) and $q = 0.01$ (right) for the simulated series.

Table 3
True values, predictive means and 95% predictive intervals for the optimal weight, δ_{opt}, which minimizes the VaR, CVaR and variance of the one-step-ahead portfolio $\delta y_{1T+1} + (1-\delta)\, y_{2T+1}$ for the simulated data.

	$q = 0.05$			$q = 0.01$		
	True	Mean	95% interval	True	Mean	95% interval
VaR	0.2219	0.2396	(0.0971, 0.4054)	0.1583	0.1510	(−0.0477, 0.4009)
CVaR	0.1559	0.1617	(0.0038, 0.3388)	0.0917	0.0997	(−0.1296, 0.3632)
Variance	0.1506	0.1599	(0.0151, 0.2976)			

with $M = 1000$ replications for each MCMC iteration. These estimations are compared with the true one-step-ahead VaR, CVaR and variances which are obtained using the true model parameters. Note that the estimated values are very close to the true values, which are always inside the predictive intervals.

Table 3 shows the posterior means and 95% Bayesian confidence intervals of the optimal weight, δ_{opt}, which minimizes the CVar, VaR and variance of the one-step-ahead portfolio $\delta y_{1T+1} + (1-\delta)\, y_{2T+1}$, using the initial value $\delta_0 = 0.5$ in the numerical optimization method (other initial values lead to similar results). Observe that the obtained optimal weights are coherent with the plots shown in Fig. 4. Also note that the posterior means are close to the true values and that the true values are always inside the predictive intervals. Finally, observe that there are certain differences between the optimal weights obtained with the three different optimizing criteria based on the VaR, the CVaR and the variance, respectively. However, these differences do not seem to be statistically significant.

Finally, Fig. 5 illustrates the Bayesian estimations of the mean–VaR, mean–CVaR and mean–variance efficient frontiers with the corresponding 95% predictive region for $q = 0.05$ and $q = 0.01$, obtained as described in Section 5. These are compared with the theoretical efficient frontiers which are obtained using the true model parameters. Observe that the estimated curves are very similar to the theoretical ones, which always lie inside the confidence regions.

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 Fig. 6. Table 4 shows some summary statistics. This sample has been analyzed previously in Arakelian and Dellaportas (2008) 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 time-varying copula GARCH models, as described in Section 4, which change continuously across time and that we believe are more natural in economic terms than a copula model which varies discretely over time.

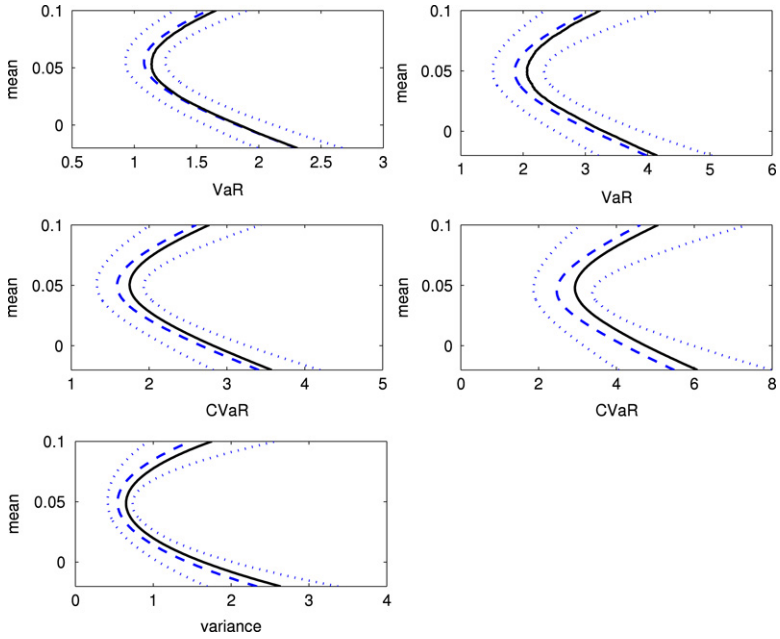

Fig. 5. True (solid lines), predictive mean (dashed lines) and 95% Bayesian confidence intervals (dotted lines) for the mean–VaR, mean–CVaR and mean–variance efficient frontiers for $q = 0.05$ (left) and $q = 0.01$ (right) for the simulated series.

Table 4
Summary statistics of the daily returns of the DAX and Dow Jones Industrial Average indices.

	DAX	Dow Jones
Mean	−0.0152	0.0164
Std. dev.	1.8255	1.2177
Skewness	−0.0300	−0.0369
Kurtosis	4.5419	5.6049
Linear correlation	0.5300	

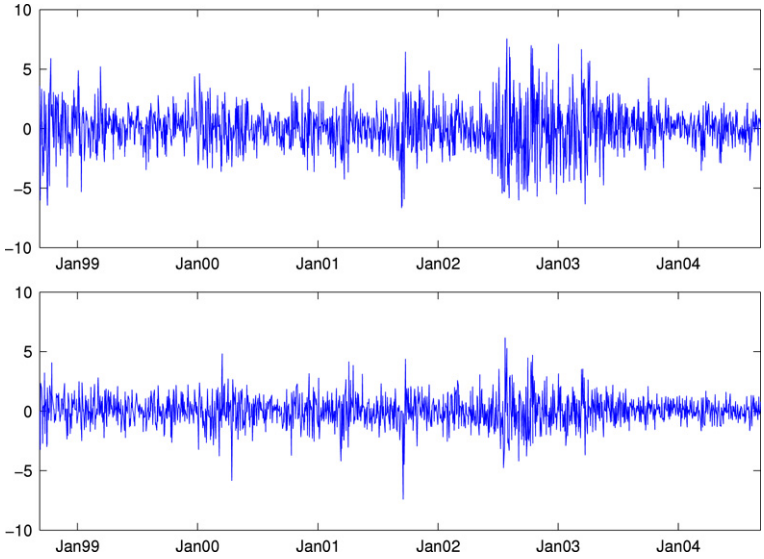

Fig. 6. Daily returns of the DAX (top) and Dow Jones Industrial Average (bottom) indices.

As in the previous section, we run the proposed MCMC method for 20 000 iterations, discarding the first 10 000 as burn-in iterations. Table 5 shows the values of the Geweke's statistic for each parameter, indicating that convergence has been achieved. Given the sample of residuals, we also check whether a GARCH(1,1) setting is appropriate for our data set or

Table 5
Values of the Geweke's statistic for each parameter obtained for the DAX and Dow Jones indices.

μ_1	μ_2	ω_1	ω_2	α_1	α_2	β_1	β_2
1.637	1.798	1.422	−1.080	1.716	0.5189	−1.430	0.3344
ν_1	ν_2	a	b	c	η		
−0.7018	0.718	0.1210	0.05079	−0.2559	1.065		

Table 6
Parameter estimation results for the DAX and Dow Jones 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
c	0.5806	0.1467

whether lags (p, q) of higher order would be necessary. To do that, we perform the hypothesis test proposed by Engle (1982) which examines the presence of ARCH/GARCH effects by regressing the squared residuals on a constant and the lagged values of the previous $M = p+q$ squared residuals. The Engle's LM statistics for lags $M = 30$ and $M = 40$ are 18.81 and 21.55 with p-values 0.94 and 0.99, respectively, for the Dow Jones index, and 44.27 and 51.58 with p-values 0.045 and 0.103, respectively, for the residuals of the DAX index, concluding that the GARCH(1,1) setting seems appropriate. Moreover, we check if the choice of a t-copula model with univariate GARCH volatilities is adequate by performing the Kolmogorov–Smirnov goodness of fit test for $t_{\hat{\eta}}(t_{\hat{\nu}_i}(\hat{\epsilon_{it}}))$, for $i = 1, 2$, where $\hat{\eta}$ and $\hat{\nu}_i$ denote the posterior means of η and ν_i, respectively, and $\hat{\epsilon}_{it}$ are the standardized residuals. The p-values of testing if these samples follow univariate Student-t distributions with $\hat{\eta}$ degrees of freedom are 0.2375 and 0.2306 for the residuals of the DAX and Dow Jones indices, respectively.

Table 6 shows the posterior means and standard deviations obtained from this algorithm. We have chosen $m = 5$ in the computation of $\Psi_{ij,t-1}$ in order to approximate the short-term correlation over one week, as suggested in Jondeau and Rockinger (2006). However, we have observed that similar results are obtained using different values for m such that $m = 10$ and $m = 20$. Also, observe in Table 6 the large value for the posterior mean and standard deviation of the degrees of freedom, ν_1, of the DAX returns, which 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 than the tails of the normal distribution, which 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. Fig. 7 illustrates the posterior means of Kendall's τ_t, the copula coefficient ρ_t and the tail dependence λ_t, as functions of time, which have been obtained as described in Section 4.2., see (17) and (18). In order to assess the quality of these estimations, following Arakelian and Dellaportas (2008), we compare in Fig. 8 the posterior means of τ_t with sample estimates of Kendall's τ_t evaluated in two rolling windows of sizes 100 and 250 observations. Observe that the dynamics of the Bayesian posterior means are very similar to the sample estimations. However, as opposed to our dynamic model, the moving window is not capable of forecasting ahead the quantities of interest.

Next, we obtain Bayesian estimations of portfolio risk measures. Fig. 9(left) shows the predictive means and 95% Bayesian confidence intervals for the VaR, CVaR and variance of the one-step-ahead portfolio, $\delta \times DAX_{T+1} + (1 - \delta) \times DowJ_{T+1}$, for different values of δ and $q = 0.05$. Also, Fig. 9(right) shows the Bayesian estimations and 95% predictive intervals for the mean–VaR, mean–CVaR and mean–variance efficient frontiers for $q = 0.05$ for the DAX and Dow Jones indices. These predictions are obtained using the simulation procedure described in Section 5 with $M = 1000$ replications for each MCMC iteration. Finally, Table 7 shows the Bayesian estimations and 95% predictive intervals of the optimal weight, δ_{opt}, which minimizes the CVar, VaR and variance of the one-step-ahead portfolio $\delta y_{1T+1} + (1 - \delta) y_{2T+1}$. Observe that the obtained optimal weights are coherent with the plots shown in Fig. 9(left). As for the simulated data, observe that there are certain differences between the optimal weights obtained with the three different optimizing criteria but they are not statistically significant.

In order to analyze the effect of assuming a time-varying copula model instead of a constant copula, Fig. 10(left) compares 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,

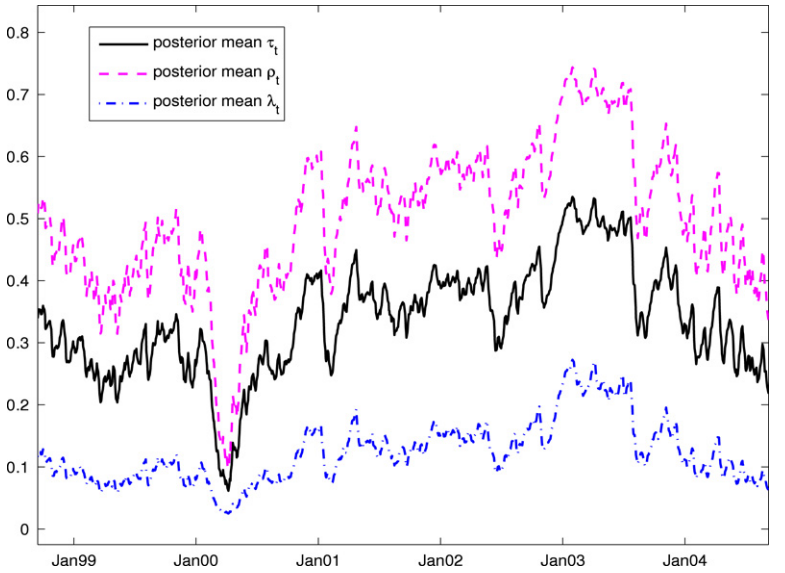

Fig. 7. Posterior means of Kendall's τ_t, the copula coefficient ρ_t and the tail dependence λ_t, as functions of time, t, for the DAX and Dow Jones indices.

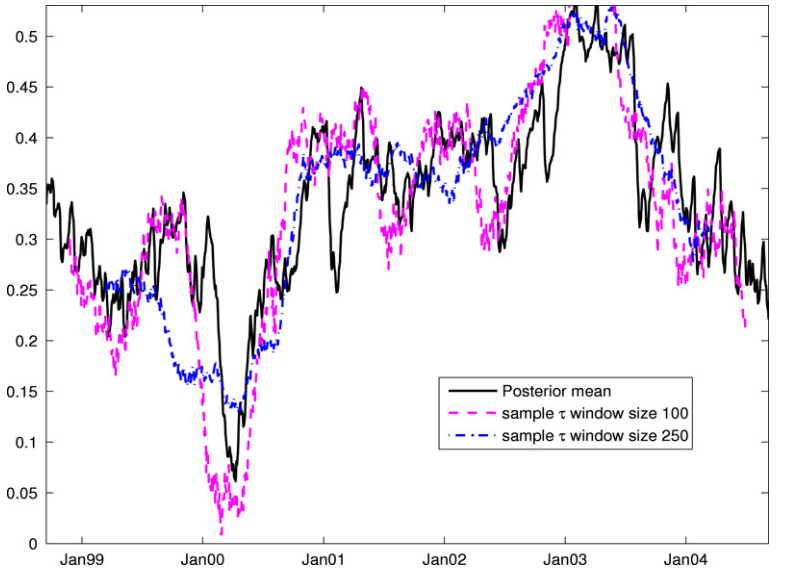

Fig. 8. Posterior means of Kendall's τ_t as a function of t compared with two window-based sample estimates, for the DAX and Dow Jones indices.

Table 7
Predictive means and 95% predictive intervals for the optimal weight, δ_{opt}, which minimizes the VaR, CVaR and variance of the one-step-ahead portfolio $\delta \times DAX_{T+1} + (1-\delta) \times Dowj_{T+1}$.

	$\pi = 0.05$		$\pi = 0.01$	
	Mean	95% interval	Mean	95% interval
VaR	0.3192	(0.1984, 0.4407)	0.3473	(0.1855, 0.5224)
CVaR	0.3405	(0.2147, 0.4650)	0.3736	(0.1844, 0.5737)
Variance	0.3223	(0.2395, 0.4036)		

obtained with our time-varying copula GARCH model with those obtained assuming a constant copula model, where a and b are assumed to be equal to zero in (7) and the resulting posterior mean and standard deviation for R 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 our time-varying copula model predicts that the optimal portfolio should assign a weight of approximately $\delta_{opt} = 0.32$ (see also Table 7) while the constant copula model predicts that $\delta_{opt} = 0.28$. Fig. 10(right) illustrates the Bayesian estimations for the mean–VaR efficient frontiers, assuming

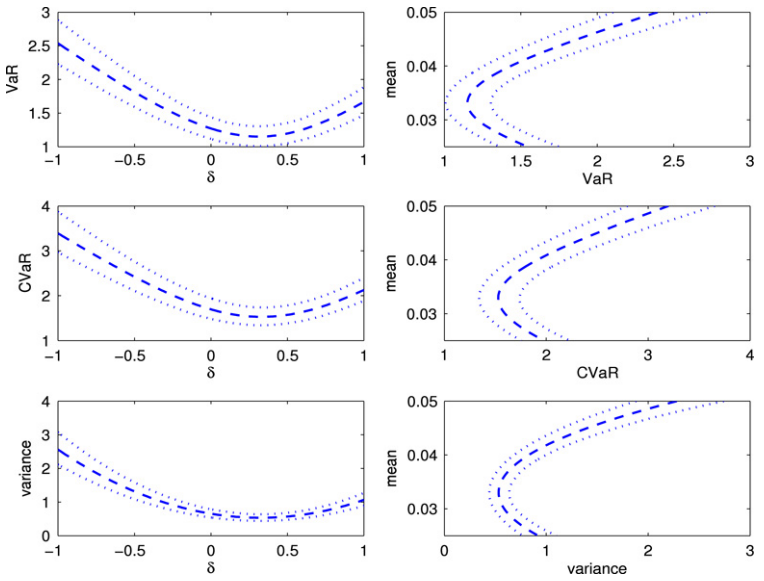

Fig. 9. Predictive mean (dashed lines) and 95% Bayesian confidence intervals (dotted lines) for the VaR, CVaR and variance of the one-step-ahead portfolio, $\delta \times DAX_{T+1} + (1-\delta) \times Dow J_{T+1}$, for different values of δ and $q = 0.05$ (left) and for the mean–VaR, mean–CVaR and mean–variance efficient frontiers for $q = 0.05$ for the DAX and Dow Jones indices (right).

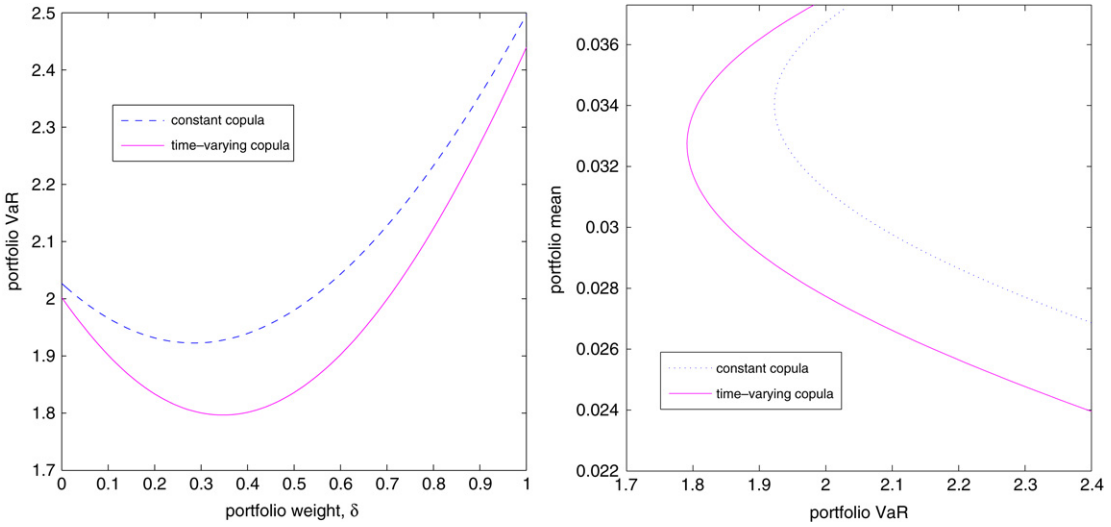

Fig. 10. Bayesian estimations of the one-step-ahead VaR for the portfolio, $\delta \times DAX_{T+1} + (1-\delta) \times Dow J_{T+1}$, as a function of δ, for $q = 0.01$, (left) and for the mean–VaR efficient frontiers (right), assuming a constant and a time-varying copula model.

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.

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 also addressed the problem of Bayesian prediction of volatilities, various dependence measures, portfolio VaR and CVaR and efficient frontiers for portfolio selection. The method has been illustrated with simulated and real financial time series.

Although, in this article, we have considered a t-copula–GARCH model 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 (11). The only inconvenience 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 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 model selection criteria for copulas considered in Huard et al. (2006) and Silva and Lopes (2008). Mixtures of copulas have been considered for dynamic models in Rodriguez (2003) and Hu (2006), using classical techniques. Finally, it would be interesting to examine the influence of the copula model in the Bayesian estimation of VaR and CVaR, as considered by Fantazzini (2009) in a classical context.

References

Arakelian, V., Dellaportas, P., 2008. Contagion tests via copula threshold models. Discussion paper.
Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1997. Thinking Coherently. Risk 10, 68–71.
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., Wang, Y.H., 2007. The euro and European financial market integration. Journal of Banking and Finance 31, 1461–1481.
Bauwens, L., Lubrano, M., 1998. Bayesian inference on GARCH models using the Gibbs sampler. Econometrics Journal 1, C23–C46.
Bauwens, L., Laurent, S., 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., Embrechts, P., 2004. Dynamic copula models for multivariate high-frequency data in finance. Discussion paper.
Embrechts, P., McNeil, A., Straumann, D., 1999a. Correlation: Pitfalls and alternatives. Risk 5, 69–71.
Embrechts, P., Resnick, S., Samorodnitsky, G., 1999b. Extreme value theory as a risk management tool. North American Actuarial Journal 3, 30–41.
Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50, 987–1007.
Fantazzini, D., 2009. The effects of misspecified marginals and copulas on computing the value at risk: A Monte Carlo study. Computational Statistics and Data Analysis 53, 2168–2188.
Gaivoronski, A.A., Pflug, G., 2005. Value-at-Risk in portfolio optimization: Properties and computational approach. Journal of Risk 7, 1–31.
Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In: Bernardo, J.M., Berger, J.O., Dawid, A.P., Smith, A.F.M. (Eds.), Bayesian Statistics, vol 4. Clarendon Press, Oxford.
Hu, L., 2006. Dependence patterns across financial markets: A mixed copula approach. Applied Financial Economics 10, 717–729.
Huard, D., Évin, G., Favre, A.C., 2006. Bayesian copula selection. Computational Statistics & Data Analysis 51, 809–822.
Jondeau, E., 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, Kleinow, 2002. Sensitivity analysis of credit portfolio models. In: Härdle, W., Kleinow, T., Stahl, G. (Eds.), Applied Quantitative Finance. Springer.
Liu, Y., Luger, R., 2009. Efficient estimation of copula-GARCH models. Computational Statistics & Data Analysis 53, 2284–2297.
Markowitz, H., 1952. Portfolio selection. Journal of Finance 8, 77–91.
Nelsen, R.B., 2006. An Introduction to Copulas, 2nd ed. Springer, New York.
Patton, A., 2006. Estimation of multivariate models for time series of possibly different lengths. Journal of Applied Econometrics 21, 147–173.
Roch, O., Alegre, A., 2006. Testing the bivariate distribution of daily equity returns using copulas. An application to the Spanish stock market. Computational Statistics & Data Analysis 51, 1312–1329.
Rockafellar, R.T., Uryasev, S., 2000. Optimization of conditional value-at-risk. Journal of Risk 2, 21–41.
Rodriguez, J.C., 2003. Measuring financial contagion: A copula approach. Journal of Empirical Finance 14, 401–423.
Schweizer, B., Sklar, A., 1983. Probabilistic Metric Spaces. Elsevier Science, New York.
Silva, R., Lopes, H.F., 2008. Copula, marginal distributions and model selection: A Bayesian note. Statistics and Computing 18, 313–320.
Tierney, L., 1994. Markov chains for exploring posterior distributions. Annals of Statistics 22, 1701–1762.
Tse, Y.K., 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., Politis, D.N., 2001. Full Bayesian inference for GARCH and EGARCH models. Journal of Business & Economic Statistics 18, 187–198.
Winker, P., Maringer, D., 2007. The hidden risks of optimizing bond portfolios under VAR. Journal of Risk 9, 1–7.