

Copula-GARCH Model Selection: A Bayesian Approach

João Luiz Rossi^a, Ricardo S. Ehlers^{a*}, Marinho G. Andrade Filho^a

^a*Universidade de São Paulo, São Carlos, Brazil*

Abstract

We study multivariate time series models where the dependence structure between the series is modelled through copulas. The advantage of this approach is that copulas provide a full description of the joint distribution. In terms of inference we adopt the Bayesian approach and computationally intensive methods based on Markov chain Monte Carlo (MCMC) simulations are used due to the complexities present in these models. In the context of copula-GARCH models, a simulation study is carried out to investigate the impact of the sample sizes, the choice of copula function and marginal distributions in the model selection procedure.

Keywords: Copula-GARCH models, Markov chain Monte Carlo, Metropolis-Hastings, skewed distributions.

1 Introduction

Recently, there has been a considerable interest in modelling the temporal dependence of financial returns through copula functions (see for example Patton 2009 and Arakelian and Dellaportas 2010). In this approach, the joint distribution of the returns is split into the marginal density functions and the

*Corresponding author. Email: ehlers@icmc.usp.br

dependence structure. So, one can specify marginal distributions exploring well known stylised facts for financial returns and then in the second stage choose the most appropriate copula function so as to fully characterise the dependence structure.

There is now empirical evidence in the literature that using copulas allows for a more realistic description of the dependence that goes beyond the linear correlation common to usual elliptical models like the multivariate normal or multivariate t distributions (Embrechts et al. 2002; Embrechts et al. 2003). In this context, many authors have considered copula-GARCH models, i.e. where the marginal time series follow a usual univariate GARCH process and the dependence structure between them is specified by a copula function. Interesting non-elliptical and flexible multivariate distributions can then be obtained. See for example, Dias and Embrechts (2004), Rodriguez (2003), Hu (2006), Jondeau and Rockinger (2006), Liu and Luger (2009) and Patton (2006).

Proposed estimation methods for this class of models are generally based on a two-stage approach. Univariate GARCH models are fitted to each of the return series separately, assuming independence between them, and in a second step the point estimates of GARCH parameters are plugged in the copula function in order to estimate the copula parameters. Of course this approach does not take into account parameter uncertainty simultaneously but is relatively straightforward to implement. More recently, Silva and Lopes (2008) proposed a fully Bayesian approach where the whole set of model parameters is estimated simultaneously. Also, Ausin and Lopes (2010) proposed to use a standard multivariate t distribution with a correlation matrix that varies through time according to a dynamics proposed in Tse and Tsui (2002) for multivariate GARCH models.

In this paper, we focus on the bivariate case and use some of the most well known copulas, namely the Normal, Student- t , Gumbel, Heavy Tail and Clayton. Our main objective is to provide a Monte Carlo study to compare copula-GARCH models with different marginal distributions (both skewed and symmetric) and dependence structures. We also compare model performances using one and two-step estimation methods.

The paper is organised as follows. In Section 2 we briefly review the main

definitions concerning copulas and embed this in the context of multivariate GARCH models. The Monte Carlo study which is the bulk of our contribution is detailed in Section 3. We conclude with a discussion in Section 4.

2 Copulas and Dependence Structure

A copula is a multivariate cumulative distribution function whose marginal distributions are uniform. In the bivariate case for example, the joint cumulative distribution function of two random variables X and Y can be represented by a copula C which in turn is a function of the marginal distribution functions,

$$F_{X,Y}(x, y) = C(F_X(x), F_Y(y)) = C(u, v).$$

Therefore, the joint distribution function of (X, Y) is fully determined by the marginal distribution functions and the copula. In particular, the copula is unique if $F_X(x)$ and $F_Y(y)$ are continuous in which case we have the following expression for the copula,

$$C(u, v) = F(F_X^{-1}(u), F_Y^{-1}(v)).$$

The joint density function can be obtained as

$$f_{X,Y}(x, y) = c(u, v)f_X(x)f_Y(y),$$

where $c(\cdot)$ is the joint density function of the copula. It is clear then that one can define different joint distributions by either changing the copula while keeping the same marginal distributions or changing the marginal distributions for the same copula function.

The extension to more than 2 dimensions is straightforward. A p -dimensional copula $C(u_1, \dots, u_p)$ is a multivariate distribution function in the unitary hypercube $[0, 1]^p$ with uniform $(0,1)$ marginal distributions. In this case, it can be shown (Schweizer and Sklar 1983) that a joint distribution function $F(x_1, \dots, x_p)$ with marginal distributions $F_1(x_1), \dots, F_p(x_p)$ can be written as,

$$F(x_1, \dots, x_p) = C(F_1(x_1), \dots, F_p(x_p)) = C(u_1, \dots, u_p).$$

Then, given a p -dimensional copula and the p univariate distributions it follows that the joint density function is,

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

where $c(\cdot)$ is the density function of $C(\cdot)$. There are several copula functions available in the literature which differ basically in terms of the dependence they represent (see for example, Nelsen 2006). For easy of presentation we focus on the bivariate case.

The Gaussian copula is perhaps the most widely applied copula function and is obtained from a bivariate normal distribution with correlation matrix R and is given by,

$$C(u, v) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi|R|^{1/2}} \exp\{-(u, v)'R^{-1}(u, v)/2\} dudv$$

where Φ^{-1} is the inverse of the cumulative distribution function of the univariate standard normal distribution. Pitt et al. (2006) describe a Bayesian approach for these copula models. The main drawback with this copula is that it assumes no dependence in the tails of the distribution. In financial applications, it is often more useful to consider at least a t -copula which is obtained from a multivariate standardised t distribution with η degrees of freedom and correlation matrix R . This is given by,

$$C(u, v) = \int_{-\infty}^{t_{\eta}^{-1}(u)} \int_{-\infty}^{t_{\eta}^{-1}(v)} \frac{\Gamma(\frac{\eta+2}{2})}{\Gamma(\frac{\eta}{2})\sqrt{(\pi\eta)^2|R|}} \left[1 + \frac{(u, v)'R^{-1}(u, v)}{\eta} \right]^{-\frac{\eta+2}{2}} dudv,$$

where t_{η}^{-1} denotes the inverse of the cumulative distribution function of the standard univariate t distribution with η degrees of freedom.

Other frequently used copulas that will be studied in this paper are, the Clayton copula,

$$C(u, v) = (u^{-\theta} + v^{-\theta} - 1)^{-1/\theta}, \quad \theta > 0.$$

the Gumbel copula,

$$C(u, v) = \exp \left\{ -[(-\log u)^{\theta} + (-\log v)^{\theta}]^{1/\theta} \right\}, \quad \theta \geq 1.$$

and the Heavy Tail copula,

$$C(u, v) = u + v - 1 + ((u - 1)^{-1/\theta} + (v - 1)^{-1/\theta} - 1)^{-\theta}, \quad \theta > 0.$$

The above copula models were chosen so as to cover different possible characteristics in the data in terms of dependence. For example, Gumbel and Heavy Tail copulas show upper tail dependence while Clayton copulas show lower tail dependence. Elliptical copulas (normal and Student) are characterised by symmetric tail dependencies.

2.1 Copula GARCH Models

For a multivariate time series $\mathbf{y}_t = (y_{1t}, \dots, y_{pt})'$ the multivariate GARCH model can be written as $\mathbf{y}_t = \mathbf{H}_t^{1/2} \boldsymbol{\epsilon}_t$ where $\mathbf{H}_t^{1/2}$ is any $p \times p$ positive definite matrix such that the conditional variance of \mathbf{y}_t is \mathbf{H}_t . There are several possible specifications for \mathbf{H}_t proposed in the literature (see e.g. Bauwens et al. 2006 for a review on the many existing proposals). Here, the dependence structure between the individual series is specified in terms of marginal distributions for the error terms ϵ_{it} combined with a copula function. We assume that each series y_{it} follows a standard univariate GARCH(1,1) model defined by

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

with $\omega_i > 0$, $\alpha_i \geq 0$, $\beta_i \geq 0$ and $\alpha_i + \beta_i < 1$, $i = 1, \dots, p$, so that the marginal distribution function for each y_{it} is given by $u_{it} = F_i(y_{it}) = F_{\epsilon_i}(y_{it}/\sqrt{h_{it}})$ where $F_{\epsilon_i}(\cdot)$ denotes the univariate distribution function of ϵ_{it} . Therefore, the joint density of \mathbf{y}_t is then given by,

$$f(y_{1t}, \dots, y_{pt}) = c(u_{1t}, \dots, u_{pt}) \prod_{i=1}^p f_i(y_{it}) = c(u_{1t}, \dots, u_{pt}) \prod_{i=1}^p \frac{1}{\sqrt{h_{it}}} f_{\epsilon_i} \left(\frac{y_{it}}{\sqrt{h_{it}}} \right).$$

where $f_{\epsilon_i}(\cdot)$ is the marginal density function of each ϵ_{it} . In this paper we assume that the errors follow either a standardised Student- t distribution with ν_i degrees of freedom or its skewed version as proposed in Fernandez

and Steel (1998). In this skewed version the density function of each ϵ_{it} is given by $\frac{2\sigma_\gamma}{\gamma_i+1/\gamma_i}f_i(\epsilon_{it}^*/\gamma_i)$ if $\epsilon_{it}^* \geq 0$ and by $\frac{2\sigma_\gamma}{\gamma_i+1/\gamma_i}f_i(\epsilon_{it}^*\gamma_i)$ if $\epsilon_{it}^* < 0$, where $\epsilon_{it}^* = \sigma_\gamma\epsilon_{it} + \mu_\gamma$, $f_i(\cdot)$ is the symmetric density function of the standardised Student- t distribution, and μ_γ and σ_γ are the mean and standard deviation of the non-standardised skew t distribution (see Fernandez and Steel 1998 for details). We note that, the symmetric version of the standardised t is obtained by simply setting $\gamma_i = 1$.

Now, given a multivariate density function $f(\cdot)$ with joint distribution function $F(\cdot)$ and corresponding marginal densities $f_i(\cdot)$ the copula density is obtained using (1) and then,

$$f(y_{1t}, \dots, y_{pt}) = \frac{f(F^{-1}(u_{1t}), \dots, F^{-1}(u_{pt}))}{\prod_{i=1}^p f_i(F^{-1}(u_{it}))} \prod_{i=1}^p \frac{1}{\sqrt{h_{it}}} f_{\epsilon_i} \left(\frac{y_{it}}{\sqrt{h_{it}}} \right).$$

Prior Distributions

In the Bayesian approach we need to specify prior distributions for the parameters which define the marginal GARCH models, i.e. ω_i , α_i , β_i , ν_i and γ_i , $i = 1, 2$ plus the parameters in the copula function, i.e. ρ in the Gaussian copula, ρ and η in the t copula, and θ in the Clayton, Gumbel and Heavy Tail copulas. The choices made here are intended to avoid introducing strong prior information. All the parameters are assumed to be independent and normally distributed a priori truncated to the interval that defines each one. For the GARCH(1,1) coefficients we adopted the same prior distributions as proposed in Ardia (2006), $\omega_i \sim N(\mu_{\omega_i}, \sigma_{\omega_i}^2)I_{(\omega_i > 0)}$, $\alpha_i \sim N(\mu_{\alpha_i}, \sigma_{\alpha_i}^2)I_{(0 < \alpha_i < 1)}$ and $\beta_i \sim N(\mu_{\beta_i}, \sigma_{\beta_i}^2)I_{(0 < \beta_i < 1)}$ while for the correlation parameters in the copulas we set $\rho \sim N(\mu_\rho, \sigma_\rho^2)I_{(-1 < \rho < 1)}$ and $\theta \sim N(\mu_\theta, \sigma_\theta^2)I_{(\theta > k)}$ where $k = 0, 1$. The hyperparameters were specified as $\mu_{\omega_i} = \mu_{\alpha_i} = \mu_{\beta_i} = \mu_\rho = \mu_\theta = 0$ and $\sigma_{\omega_i}^2 = \sigma_{\alpha_i}^2 = \sigma_{\beta_i}^2 = \sigma_\rho^2 = \sigma_\theta^2 = 100$.

The prior distributions for the degrees of freedom parameters ν_i in the marginal t distributions and η in the t copula are assigned as $\nu_i \sim N(\mu_\nu, \sigma_\nu^2)I_{(\nu > 2)}$ and $\eta \sim N(\mu_\eta, \sigma_\eta^2)I_{(\eta > 2)}$ with hyperparameters $\mu_\nu = \mu_\eta = 8$ and $\sigma_\nu^2 = \sigma_\eta^2 = 30$.

Finally, for the skewness parameters γ_i when using Skew Student marginal distributions we assigned a $N(0, 0.64^{-1})$ distribution truncated at $\gamma_i > 0$.

This last prior distribution is the one proposed in Fernandez and Steel (1998) and is centred around the symmetric version of the skewed distribution giving approximately equal weights to left and right skewness. We shall adopt these prior choices in the simulation studies of Section 3.

2.2 Estimation

For all combinations of copula functions and marginal distributions adopted, combining likelihood functions and prior densities via Bayes theorem lead to analytically intractable posterior densities. We then adopt MCMC sampling strategies for obtaining samples from the joint posterior distributions. In particular, we use a Metropolis-Hastings algorithm to yield the required realisations. For both estimation approaches described below, we adopt a random walk Metropolis algorithm where the parameters are updated as a block and at each iteration we generate new values from a multivariate normal distribution centred around the current value with a variance-covariance proposal matrix which was calculated from a pilot tuning procedure. This pilot tuning was carried out by running one-dimensional random walk Metropolis updates with univariate normal candidate distributions whose variances were calibrated to obtain good acceptance rates.

In the two-step Bayesian approach, point estimates for each marginal GARCH model are obtained using the estimated posterior mean of the parameters ω_i , α_i and β_i and the residuals $e_{it} = y_{it}/\sqrt{h_{it}}$ are computed. These residuals plus the point estimates of ν_i and γ_i are then used to obtain $u_{it} = F_{\epsilon_i}(e_{it})$ where $F_{\epsilon_i}(\cdot)$ is the cumulative distribution function of ϵ_{it} . In the one step approach, the random walk Metropolis algorithm described above is applied to update the whole set of parameters jointly.

In our simulation study, the models were compared and selected according to the expected value of the Akaike information criterion (EAIC), the expected value of the Bayesian information criterion (EBIC) and the deviance information criterion (DIC). These are given by $EAIC = E[D(\boldsymbol{\theta}_M)] + 2 np_M$, $EBIC = E[D(\boldsymbol{\theta}_M)] + \log(n) np_M$ and $DIC = 2E[D(\boldsymbol{\theta}_M)] - D(E[\boldsymbol{\theta}_M])$ respectively, where np_M represents the number of parameters in model M , $\boldsymbol{\theta}_M$ is the set of parameters in model M and $D(\cdot)$ is the deviance function de-

defined as minus twice the log-likelihood function. These measures share the important feature of invariance under monotone increasing transformations of the marginal distributions and are straightforward to approximate given a sample from the posterior distribution of θ_M (see Spiegelhalter et al. 2002 for details about these measures).

3 Simulations

In this section, we report the results of a simulation study to check whether the selection of copula-GARCH models is influenced by the estimation method (joint or two-steps), the size of the series and the values of parameters in the copula function.

We studied 10 bivariate models which were obtained by combining the dependence structure in each of five copula functions plus two marginal GARCH(1,1) models with Student and Skew Student error terms. We generated 500 replications of time series with sizes 200, 500 and 1000 for each of the GARCH-copula models obtained from the Gaussian, Student- t , Clayton, Gumbel and Heavy Tail copulas. In order to define the copula parameters we used the Kendall's τ dependence measure. This measure is an alternative to the correlation coefficient and is analytically available for the copulas studied here. We have that, $\tau = (2/\pi) \arcsin(\rho)$ for the Gaussian and t copulas and $\tau = \theta/(\theta + 2)$, $\tau = 1 - 1/\theta$ and $\tau = \theta^{-1}/(\theta^{-1} + 2)$ for the Clayton, Gumbel and Heavy Tail copulas respectively. The copula parameters were then chosen so that Kendall's τ equal 1/3 and 2/3. For each replication and sample size, the 5 copula-GARCH models were fitted with parameters being estimated separately (first the marginal GARCH parameters and then the copula parameter) and jointly (all parameters simultaneously). In any case, Monte Carlo approximations of the posterior expectations were obtained after running the MCMC sampler for 30,000 iterations and discarding the first 20,000 as burn-in. The information criteria EAIC, EBIC and DIC were then calculated from the MCMC output and verified whether the selected model was the same used to generate the series. In what follows, we report results based on the DIC only since those for the EAIC and EBIC were quite similar. The results obtained are summarised in Table 1 (two stage estimation) and

Table 2 (joint estimation) where the proportions of true copula distributions correctly identified according to DIC are reported.

Table 1: Correct model selection rates according to DIC using a two-stage estimation procedure with marginal distributions Skew Student (SST) and Symmetric Student (ST).

| copula | marginals | $\tau = 1/3$ | | | $\tau = 2/3$ | | |
|------------|-----------|--------------|-----------|------------|--------------|-----------|------------|
| | | $n = 200$ | $n = 500$ | $n = 1000$ | $n = 200$ | $n = 500$ | $n = 1000$ |
| Normal | SST | 0.810 | 0.950 | 0.992 | 0.838 | 0.938 | 0.956 |
| Student-t | SST | 0.808 | 0.976 | 1.000 | 0.874 | 0.988 | 1.000 |
| Clayton | SST | 0.946 | 0.998 | 1.000 | 0.990 | 1.000 | 1.000 |
| Gumbel | SST | 0.572 | 0.852 | 0.968 | 0.828 | 0.968 | 0.998 |
| Heavy-Tail | SST | 0.750 | 0.862 | 0.958 | 0.776 | 0.970 | 1.000 |
| Normal | ST | 0.846 | 0.968 | 0.996 | 0.888 | 0.936 | 0.978 |
| Student-t | ST | 0.850 | 0.980 | 0.998 | 0.856 | 0.994 | 1.000 |
| Clayton | ST | 0.936 | 0.998 | 1.000 | 0.988 | 1.000 | 1.000 |
| Gumbel | ST | 0.648 | 0.894 | 0.990 | 0.900 | 0.992 | 1.000 |
| Heavy-Tail | ST | 0.796 | 0.894 | 0.970 | 0.828 | 0.978 | 1.000 |

We first look at the proportions of correct selection as a whole, i.e. for both tables. For $\tau = 1/3$ and $n = 200$, we note that the proportions of correct selection were above 0.95 in only 10% of the cases while increasing the sample sizes to $n = 500$ and $n = 1000$ leads these proportions to increase to 45% and 100% of the cases respectively. Increasing the correlation ($\tau = 2/3$), the proportions of correct selection above 0.95 were 25%, 80% and 95% for the sample sizes 200, 500 and 1000 respectively. We can conclude that, in the context of this work, series of size 1000 are enough for the correct identification of the copula model.

Now, comparing the rows of Tables 1 and 2, it is worth noting that in some cases the proportions of correct selection were smaller when joint estimation was adopted. This was the case for normal copulas with both skew and symmetric Student- t marginal distributions, for all sample sizes

Table 2: Correct model selection rates according to DIC using a joint estimation procedure with marginal distributions Skew Student (SST) and Symmetric Student (ST).

| copula | marginals | $\tau = 1/3$ | | | $\tau = 2/3$ | | |
|------------|-----------|--------------|-----------|------------|--------------|-----------|------------|
| | | $n = 200$ | $n = 500$ | $n = 1000$ | $n = 200$ | $n = 500$ | $n = 1000$ |
| Normal | SST | 0.732 | 0.924 | 0.988 | 0.790 | 0.918 | 0.946 |
| Student-t | SST | 0.838 | 0.980 | 1.000 | 0.840 | 0.988 | 1.000 |
| Clayton | SST | 0.944 | 0.992 | 1.000 | 0.990 | 1.000 | 1.000 |
| Gumbel | SST | 0.460 | 0.832 | 0.962 | 0.820 | 0.968 | 0.994 |
| Heavy-Tail | SST | 0.774 | 0.884 | 0.960 | 0.934 | 1.000 | 1.000 |
| Normal | ST | 0.768 | 0.942 | 0.990 | 0.816 | 0.928 | 0.966 |
| Student-t | ST | 0.866 | 0.984 | 0.998 | 0.896 | 0.992 | 1.000 |
| Clayton | ST | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 |
| Gumbel | ST | 0.616 | 0.862 | 0.986 | 0.944 | 0.996 | 1.000 |
| Heavy-Tail | ST | 0.828 | 0.908 | 0.970 | 0.984 | 1.000 | 1.000 |

and correlations. It occurred also for the Gumbel copula with skew Student- t margins while for symmetric Student- t margins the proportions of correct selection decreased for low correlation ($\tau = 1/3$) only.

Figure 1 provides a visual summary of our simulation study in the case of Skew Student marginal error distributions in the GARCH models with the one step estimation procedure. Each column represents a copula model used to generate the data and the legends indicate the selected model based on DIC. Such a figure is useful to verify whether data generated by a particular copula model is being consistently misclassified by another one. We can see for example that normal and Student copulas misclassified each other for small sample sizes ($n = 200$) which would indicate that copulas with similar properties might be difficult to discriminate (a feature previously detected in Silva and Lopes (2008) for non time series data). For larger sample sizes however ($n = 500, 1000$), Figure 1 corroborates the information in Table 2 and we would not expect a great deal of misclassification in practice. Similar

figures were obtained for symmetric Student marginal error distributions and two steps estimation but are not reported here to save space.

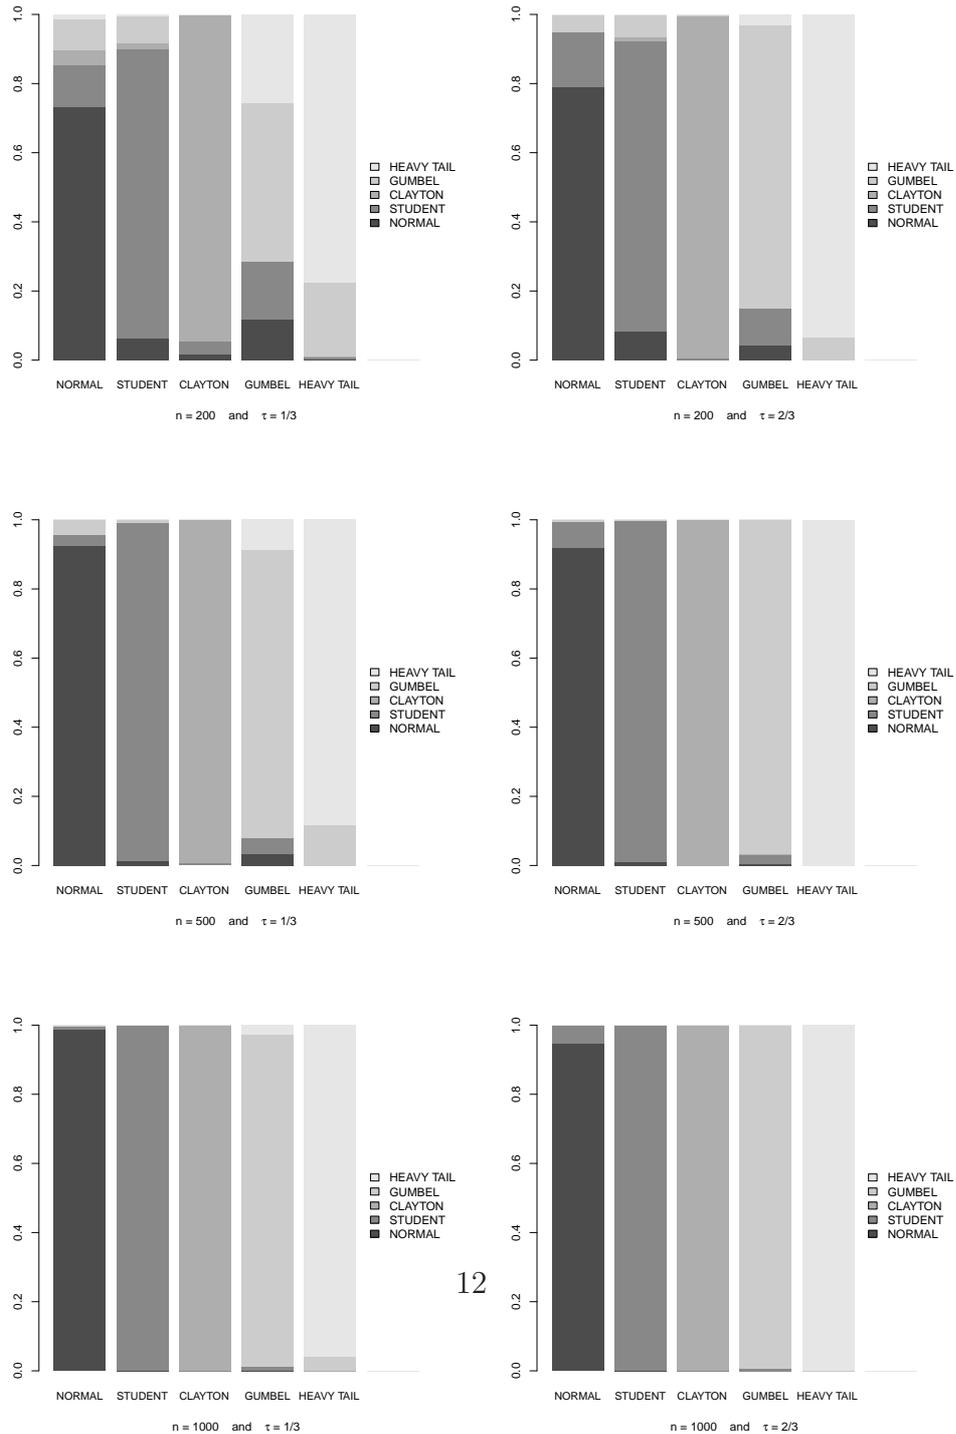
Computationally, the two step procedure is less demanding and we would recommend it if the interest is mainly in point estimates. However, as pointed out in Silva and Lopes (2008), from a Bayesian perspective jointly estimating all unknown parameters also takes into account parameter dependence. It should be noted however that, from a practical point of view, in the context of copula-GARCH models with higher dimensions the one step procedure is likely to be quite time consuming.

4 Discussion

In this paper we adopted a Bayesian approach to estimate and compare copula models traditionally utilised to analyse multivariate financial time series. Our main contribution was to provide a Monte Carlo study to compare multivariate GARCH models where the dependence structure between the series is modelled through copulas. Our results suggest that the DIC may be used as an appropriate selection criterion for reasonable sample sizes typically encountered in financial applications. To the best of our knowledge, such a study has not yet appeared in the multivariate GARCH literature. Of course, as in any Monte Carlo study, our results are limited to our particular selection of sample sizes, degrees of dependence, marginal distributions and copula functions. However, we note that these are typical choices in most financial application and we hope that our findings are useful to the practitioners.

As a possible extension, the parameters in the copula function could vary along time, possibly changing at distinct points which define different contagion regimes. In this case, the number of regimes might also be estimated rather than being specified a priori using a reversible jump MCMC scheme. This is the approach proposed in Arakelian and Dellaportas (2010).

Figure 1: Proportions of correct model selection according to DIC for $\tau = 1/3$ (first column), $\tau = 2/3$ (second column) and three sample sizes ($n= 200, 500, 1000$). Marginal error distributions in the GARCH models are Skew Student with one step estimation procedure.



Acknowledgements

The work of the first author was funded by CAPES - Brazil. The work of the second author was supported by FAPESP - Brazil, under grant number 2011/22317-0.

References

- Arakelian, V. and P. Dellaportas (2010). Contagion tests via copula threshold models. *Quantitative Finance*. To appear.
- Ardia, D. (2006). Bayesian estimation of the GARCH(1,1) model with normal innovations. *Student* 5(3–4), 283–298. URL <http://ssrn.com/abstract=1543409>.
- Ausin, M. C. and H. F. Lopes (2010). Time-varying joint distribution through copulas. *Computational Statistics and Data Analysis* 54(11), 2383–2399.
- Bauwens, L., S. Laurent, and J. V. K. Rombouts (2006). Multivariate GARCH models: a survey. *Journal of Applied Econometrics* 21(1), 79–109.
- Dias, A. and P. Embrechts (2004). Change-point analysis for dependence structures in finance and insurance. In G. Szegoe (Ed.), *Risk Measures for 21st Century*, pp. 321–335. Wiley Finance Series.
- Embrechts, P., F. Lidskog, and A. J. McNeal (2003). Modelling dependence with copulas and applications to risk management. In S. Rachev (Ed.), *Handbook of Heavy Tailed Distributions in Finance*, pp. 329–384. Elsevier.
- Embrechts, P., A. J. McNeal, and D. Straumann (2002). Correlation and dependence in risk management: properties and pitfalls. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*, pp. 176–223. Cambridge University Press.
- Fernandez, C. and M. Steel (1998). On Bayesian modelling of fat tails and skewness. *Journal of the American Statistical Association* 93, 359–371.

- Hu, L. (2006). Dependence patterns across financial markets: A mixed copula approach. *Applied Financial Economics* 10, 717–729.
- Jondeau, E. and M. Rockinger (2006). The copula-GARCH model of conditional dependencies: An international stock-market application. *Journal of International Money and Finance* 25, 827–853.
- Liu, Y. and R. Luger (2009). Efficient estimation of copula-GARCH models. *Computational Statistics & Data Analysis* 53, 2284–2297.
- Nelsen, R. B. (2006). *An Introduction to Copulas*. Springer, New York.
- Patton, A. (2006). Estimation of multivariate models for time series of possibly different lengths. *Journal of Applied Econometrics* 21, 147–173.
- Patton, A. J. (2009). Copula-based models for financial time series. In T. G. Andersen, R. A. Davis, J. Kreiss, & T. Mikosch (Eds.), *Handbook of Financial Time Series*. Springer Berlin Heidelberg.
- Pitt, M., D. Chan, and R. Kohn (2006). Efficient Bayesian inference for Gaussian copula regression models. *Biometrika* 93(3), 537–554.
- Rodriguez, J. C. (2003). Measuring financial contagion: a copula approach. *Journal of Empirical Finance* 14, 401–423.
- Schweizer, B. and A. Sklar (1983). *Probabilistic Metric Spaces*. Elsevier Science, New York.
- Silva, R. S. and H. F. Lopes (2008). Copula, marginal distributions and model selection: a Bayesian note. *Statistics and Computing* 18, 313–320.
- Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. van der Linde (2002). Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B* 64, 1–34.
- Tse, Y. K. and A. K. C. Tsui (2002). A multivariate GARCH model with time-varying correlations. *Journal of Business and Economic Statistics* 20, 351–362.