



TI 2008-105/4

Tinbergen Institute Discussion Paper

Out-of-sample Comparison of Copula Specifications in Multivariate Density Forecasts

Cees Diks^{1,4}

Valentyn Panchenko²

Dick van Dijk^{3,4}

¹ CeNDEF, Amsterdam School of Economics, University of Amsterdam;

² School of Economics, University of New South Wales;

³ Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam;

⁴ Tinbergen Institute.

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Out-of-sample comparison of copula specifications in multivariate density forecasts

Cees Diks*

*CeNDEF, Amsterdam School of Economics
University of Amsterdam*

Valentyn Panchenko[†]

*School of Economics
University of New South Wales*

Dick van Dijk[‡]

*Econometric Institute
Erasmus University Rotterdam*

November 1, 2008

Abstract

We introduce a statistical test for comparing the predictive accuracy of competing copula specifications in multivariate density forecasts, based on the Kullback-Leibler Information Criterion (KLIC). The test is valid under general conditions: in particular it allows for parameter estimation uncertainty and for the copulas to be nested or non-nested. Monte Carlo simulations demonstrate that the proposed test has satisfactory size and power properties in finite samples. Applying the test to daily exchange rate returns of several major currencies against the US dollar we find that the Student's t copula is favored over Gaussian, Gumbel and Clayton copulas. This suggests that these exchange rate returns are characterized by symmetric tail dependence.

Keywords: Copula-based density forecast; semiparametric statistics; out-of-sample forecast evaluation; Kullback-Leibler Information Criterion; empirical copula

JEL Classification: C12; C14; C32; C52; C53

*Corresponding author: Center for Nonlinear Dynamics in Economics and Finance, Faculty of Economics and Business, University of Amsterdam, Roetersstraat 11, NL-1018 WB Amsterdam, The Netherlands. E-mail: C.G.H.Diks@uva.nl

[†]School of Economics, Faculty of Business, University of New South Wales, Sydney, NSW 2052, Australia. E-mail: v.panchenko@unsw.edu.au

[‡]Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands. E-mail: djvandijk@few.eur.nl

1 Introduction

Copulas have become an important tool for modeling multivariate distributions, with applications in economics and particularly finance expanding rapidly, see Cherubini *et al.* (2004) and Patton (2008) for recent surveys. The main attractive property of copulas is that they allow for modeling the marginal distributions and the dependence structure between the variables of interest separately. Many parametric copula families are available, which share the advantage of having only a low-dimensional parameter vector, so that generally the copula estimation problem is essentially low-dimensional. On the other hand, popular parametric copulas such as the Gaussian, Student's t , Clayton, and Gumbel copulas have rather different dependence properties. An important issue in empirical applications therefore is the choice of an appropriate copula specification. Not surprisingly then, considerable interest in goodness-of-fit testing for copulas has arisen recently. Several different types of copula specification tests have been put forward. The alternatives that appear most promising have in common that they reduce the evaluation of the multivariate copula to a univariate problem, and then apply some univariate test. Examples include goodness-of-fit tests based on the probability integral transform (PIT) of Rosenblatt (1952) by Breymann *et al.* (2003), Malevergne and Sornette (2003), Fermanian (2005), and Berg and Bakken (2007); based on the empirical copula (DeHeuvels, 1979) and the copula distribution function by Genest and Rémillard (in press) and Genest *et al.* (in press); based on the empirical copula and the CDF of the copula function by Genest *et al.* (2006); and based on non-parametric distance measures by Panchenko (2005). We refer the interested reader to Berg (2007) for a detailed review of these approaches and a power comparison based on extensive simulation experiments.

The related problem of selecting a copula specification from two or more parametric families has mostly been approached in an *indirect* way. Typically, several goodness-of-fit tests are applied to the competing specifications, and the copula that performs best on these statistics is selected; see Kole *et al.* (2007) for an empirical example. A *direct* comparison of two alternative copulas has only been considered by Chen and Fan (2006) and Patton (2006), adopting the approach based on pseudo likelihood ratio (PLR) tests

for model selection originally developed by Vuong (1989) and Rivers and Vuong (2002). These tests compare the candidate copula specifications in terms of their Kullback-Leibler Information Criterion (KLIC), which measures the distance from the true (but unknown) copula. Similar to the goodness-of-fit tests, these PLR tests are based on the in-sample fit of competing copulas.

In this paper we approach the copula selection problem from an out-of-sample forecasting perspective. In particular, we extend the PLR testing approach of Chen and Fan (2006) and Patton (2006) to compare the predictive accuracy of alternative copula specifications. We adopt the framework of semiparametric copula-based multivariate dynamic (SCOMDY) models developed in Chen and Fan (2005, 2006), which has parametric specifications for the conditional mean and conditional variance together with a semiparametric specification for the distribution of the (standardized) innovations, consisting of a (time-invariant) parametric copula with nonparametric univariate marginal distributions.¹ The crucial difference with Chen and Fan (2006) is that our test is based on out-of-sample log-likelihood scores obtained using copula density forecasts rather than in-sample scores. An important motivation for considering the (relative) predictive accuracy of copulas is that multivariate density forecasting is one of the main purposes of copula applications in economics and finance, in particular in risk management.

Comparison of KLIC for assessing relative predictive accuracy has recently become popular for evaluation of univariate density forecasts, see Mitchell and Hall (2005), Amisano and Giacomini (2007) and Bao *et al.* (2007). Amisano and Giacomini (2007) provide an interesting interpretation of the KLIC-based comparison in terms of scoring rules, which are loss functions depending on the density forecast and the actually observed data. In particular, it is shown that the difference between the log-likelihood scoring rule for two competing density forecasts corresponds exactly to their relative KLIC values. The same interpretation continues to hold for multivariate density forecasts as considered in this paper.

¹The set-up in Patton (2006) is similar, except that the marginal distributions need to be specified parametrically, while the copula can be time-varying. Our test of equal predictive accuracy of alternative copulas can be readily extended to this context of *conditional* copulas.

Our test of equal predictive accuracy is valid under general conditions, which is achieved by adopting the framework of Giacomini and White (2006). This assumes that the unknown model parameters are estimated on the basis of a moving window of fixed size. The finite estimation window essentially allows us to treat competing density forecasts based on different copula specifications, including the time-varying estimated model parameters, as two competing *forecast methods*. Comparing scores for forecast methods rather than for models simplifies the resulting test procedures considerably, the reason being that parameter estimation uncertainty does not play any role as it is considered to be part of the respective competing forecasting methods. In addition, the asymptotic distribution of our test statistic in this case does not depend on whether or not the competing copulas are nested.

We examine the size and power properties of our copula predictive accuracy test in small samples via Monte Carlo simulations. One aspect that is of particular interest is that we aim to compare copulas using the log-likelihood score, without assuming that the (conditional) marginals are known. As a result, the copula vector cannot be observed directly, and therefore in practice is replaced by the empirical copula. Especially for small estimation windows this might affect the actual size (and power) of our test. Reassuringly, our simulation results show that this aspect does not greatly affect the size and power properties of the copula predictive accuracy test, which generally are satisfactory in realistic sample sizes.

We consider an empirical application to daily exchange rate returns of the Canadian dollar, Swiss franc, Euro, British pound, and Japanese yen against the US dollar over the period January 1980 - July 2008. Based on the relative predictive accuracy of one-step ahead density forecasts we find that the Student's t copula is favored over Gaussian, Gumbel and Clayton copulas. This suggests that these exchange rate returns are characterized by symmetric tail dependence.

The paper is organized as follows. In Section 2 we briefly review the class of SCOMDY models introduced by Chen and Fan (2006). We develop our predictive accuracy test for copulas based on out-of-sample log-likelihood scores in Section 3. In Section 4 we investigate its size and power properties by means of Monte Carlo simulations. In section 5

we illustrate our test with an application to daily exchange rate returns for several major currencies. We summarize and conclude in Section 6.

2 SCOMDY models

We briefly review the class of SCOMDY models introduced by Chen and Fan (2006). Let $\{(\mathbf{Y}'_t, \mathbf{X}'_t)\}$ denote a vector stochastic process, in which \mathbf{Y}_t represents a d -dimensional vector, and \mathbf{X}_t is a vector of exogenous or pre-determined variables. Furthermore, let \mathcal{F}_{t-1} denote the sigma-field generated by $\{\mathbf{Y}_{t-1}, \mathbf{Y}_{t-2}, \dots; \mathbf{X}_t, \mathbf{X}_{t-1}, \dots\}$. The class of SCOMDY models is then specified as

$$\mathbf{Y}_t = \boldsymbol{\mu}_t(\boldsymbol{\theta}_1) + \sqrt{\mathbf{H}_t(\boldsymbol{\theta})}\boldsymbol{\varepsilon}_t, \quad (1)$$

where

$$\boldsymbol{\mu}_t(\boldsymbol{\theta}_1) = (\mu_{1,t}(\boldsymbol{\theta}_1), \dots, \mu_{d,t}(\boldsymbol{\theta}_1))' = \mathbb{E}[\mathbf{Y}_t | \mathcal{F}_{t-1}]$$

is a specification of the conditional mean, parameterized by a finite dimensional vector of parameters $\boldsymbol{\theta}_1$, and

$$\mathbf{H}_t(\boldsymbol{\theta}) = \text{diag}(h_{1,t}(\boldsymbol{\theta}), \dots, h_{d,t}(\boldsymbol{\theta})),$$

where

$$h_{j,t}(\boldsymbol{\theta}) = h_{j,t}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = \mathbb{E}[(Y_{j,t} - \mu_{j,t}(\boldsymbol{\theta}_1))^2 | \mathcal{F}_{t-1}], \quad j = 1, \dots, d,$$

is the conditional variance of $Y_{j,t}$ given \mathcal{F}_{t-1} , parameterized by a finite dimensional vector of parameters $\boldsymbol{\theta}_2$, where $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ do not have common elements. The innovations $\boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \dots, \varepsilon_{d,t})'$ are independent of \mathcal{F}_{t-1} and independent and identically distributed (i.i.d.) with $\mathbb{E}(\varepsilon_{j,t}) = 0$ and $\mathbb{E}(\varepsilon_{j,t}^2) = 1$ for $j = 1, \dots, d$. The joint distribution function of $\boldsymbol{\varepsilon}_t$ is denoted as $F(\boldsymbol{\varepsilon})$. We define the copula vector associated with $\boldsymbol{\varepsilon}_t$ as the vector whose elements are equal to the respective marginal probability integral transforms (PITs) of ε_t , that is

$$\mathbf{U}_t = (U_{1,t}, \dots, U_{d,t})' = (F_1(\varepsilon_{1,t}), \dots, F_d(\varepsilon_{d,t}))',$$

where $F_j(\cdot)$ is the marginal distribution function of $\varepsilon_{j,t}$, $j = 1, \dots, d$, which is assumed to be continuous. It then follows from Sklar's (1959) theorem that the distribution function

$F(\boldsymbol{\varepsilon})$ can be written as

$$F(\boldsymbol{\varepsilon}) = C(F(\varepsilon_1), \dots, F(\varepsilon_d); \boldsymbol{\alpha}) \equiv C(u_1, \dots, u_d; \boldsymbol{\alpha}), \quad (2)$$

where $C(u_1, \dots, u_d; \boldsymbol{\alpha}): [0, 1]^d \rightarrow [0, 1]$ is a member of a parametric family of copula functions with finite dimensional parameter vector $\boldsymbol{\alpha}$. Popular copula families are the normal or Gaussian copula, the Student's t copula, the Gumbel copula and the Clayton copula, see Joe (1997) and Nelsen (2006) for overviews.

Although many other examples of SCOMDY models can be given by considering different specifications of the conditional means $\mu_{j,t}(\boldsymbol{\theta}_1)$ and conditional variances $h_{j,t}(\boldsymbol{\theta})$, the following two specific cases are popular in financial applications.

Model Type 1 *Multivariate i.i.d. time series with parametric copula*

In the simplest case the conditional means and variances are constant over time, and we have

$$Y_{j,t} = \delta_j + \sigma_j \varepsilon_{j,t}, \quad \text{for } j = 1, \dots, d. \quad (3)$$

In terms of the general SCOMDY specification in (1), this corresponds to $\mu_{j,t}(\boldsymbol{\theta}_1) = \delta_j$ and $h_{j,t}(\boldsymbol{\theta}) = \sigma_j^2$, $j = 1, \dots, d$. The parameter vectors $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ in this case equal $\boldsymbol{\theta}_1 = (\delta_1, \dots, \delta_d)'$ and $\boldsymbol{\theta}_2 = (\sigma_1, \dots, \sigma_d)'$. Note that $\{\mathbf{Y}_t\}$ is an i.i.d. sequence of random vectors, each of whose d elements can have any continuous marginal distribution function. Moreover, if we require, without loss of generality, that $\sigma_j > 0$, $j = 1, \dots, d$, it can be shown that the copula of \mathbf{Y}_t is the same as that of $\boldsymbol{\varepsilon}_t$, that is

$$F_{\mathbf{Y}_t}(y_1, \dots, y_d) = C(F_{Y_{1,t}}(y_1), \dots, F_{Y_{d,t}}(y_d); \boldsymbol{\alpha}).$$

□

Model Type 2 *GARCH(1,1) model with parametric copula*

It is straightforward to extend the previous example with a time-varying conditional mean and a GARCH(1,1) structure for the conditional variance, to obtain

$$\begin{aligned} Y_{j,t} &= X'_{j,t} \delta_j + \sqrt{h_{j,t}} \varepsilon_{j,t} \\ h_{j,t} &= \kappa_j + \gamma_j (Y_{j,t-1} - X'_{j,t-1} \delta_j)^2 + \beta_j h_{j,t-1}, \end{aligned}$$

where $\kappa_j > 0$, $\beta_j \geq 0$, $\gamma_j > 0$ and $\beta_j + \gamma_j < 1$, for $j = 1, \dots, d$. This case corresponds to $\boldsymbol{\theta}_1 = (\delta'_1, \dots, \delta'_d)'$, $\boldsymbol{\theta}_2 = (\kappa_1, \dots, \kappa_d; \beta_1, \dots, \beta_d; \gamma_1, \dots, \gamma_d)'$. \square

An important characteristic of SCOMDY models is that the univariate marginal densities $F_j(\cdot)$, $j = 1, \dots, d$ are not specified parametrically (up to an unknown parameter vector) but are estimated nonparametrically. Chen and Fan (2005, 2006) suggest the following three-stage procedure to estimate the parameters in the SCOMDY model, based on a sample of observations $\{(\mathbf{Y}'_t, \mathbf{X}'_t)\}_{t=1}^T$. First, estimate the parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ in the conditional mean and conditional variance specifications using quasi-maximum likelihood (QMLE), that is, under the assumptions of normality of the standardized innovations $\varepsilon_{j,t}$. Second, estimate the marginal distributions $F_j(\cdot)$ by means of the empirical CDF transformation

$$u_{j,t} = \frac{\text{rank of } \hat{\varepsilon}_{j,t} \text{ among the residuals } \hat{\varepsilon}_{j,1}, \dots, \hat{\varepsilon}_{j,T}}{T + 1} \quad \text{for } j = 1, \dots, d.$$

Finally, estimate the parameters $\boldsymbol{\alpha}$ of a given copula specification by maximizing the corresponding copula log-likelihood function. (Chen and Fan, 2006) study the large sample properties of the resulting parameter estimates. The parameters $\boldsymbol{\theta}$ in the conditional mean and variance specifications as well as the marginal distributions are estimated consistently at root- T , with T being the sample size. The estimator of the copula parameters $\boldsymbol{\alpha}$ converges to the pseudo-true value that minimizes the KLIC between the candidate copula and the true unknown copula. The limiting distribution of this estimator is not affected by the estimation of the parameters $\boldsymbol{\theta}$, but it does depend on the estimation of the unknown marginal distributions $F_j(\cdot)$.

3 Equal predictive accuracy test for copulas

As mentioned in the introduction, the Kullback-Leibler information criterion (KLIC), measuring the divergence between the true probability density and a candidate density, has become a popular measure for evaluating the (relative) predictive performance of univariate density forecasts. In this section we extend the KLIC-based tests of equal predictive accuracy to the context of multivariate density forecasts based on SCOMDY models with alternative parametric copulas. Recall that the (pseudo-) likelihood ratio test of Chen

and Fan (2006) can be interpreted as a copula selection procedure based on in-sample KLICs. As will become clear below, our test statistic is the out-of-sample analogue of their approach.

In order to develop our test, we first make precise what we mean by copula KLIC. Given any copula density $c_A(\cdot)$, the copula KLIC is defined as

$$\text{KLIC}(c_A) = \int_{[0,1]^d} c(\mathbf{u}) \log \left(\frac{c_A(\mathbf{u})}{c(\mathbf{u})} \right) d\mathbf{u} = \mathbb{E}(\log c(\mathbf{U}_t) - \log c_A(\mathbf{U}_t)), \quad (4)$$

where $c(\cdot)$ denotes the true copula density corresponding with the copula defined in (2), and the dependence of $c(\cdot)$ and $c_A(\cdot)$ on parameters α and α_A is suppressed for notational convenience. In words, the copula KLIC measures the distance between a given copula density $c_A(\cdot)$ and the true copula density $c(\cdot)$. Copulas which achieve a smaller KLIC value are preferred. In fact, a crucial observation is that by taking the difference of the KLICs of two competing copula densities, c_A and c_B , say, the term $\mathbb{E}(\log c(\mathbf{U}_t))$ in (4) drops out. This means that it is not necessary to know the true copula $c(\cdot)$ in order to test whether two competing copulas have significantly different KLIC values.

For a given copula density $c_A(\cdot)$, we define the copula score based on a copula vector \mathbf{U}_t as the part of the SCOMDY log-likelihood associated with the copula specification, that is

$$S_t^{c_A} = \log c_A(\mathbf{U}_t).$$

The difference between $\text{KLIC}(c_A)$ and $\text{KLIC}(c_B)$ for two competing copula densities then is identical to the difference between their expected copula scores $\mathbb{E}(S_t^{c_A})$ and $\mathbb{E}(S_t^{c_B})$. The preferred copula with the smallest KLIC value has the highest expected copula score.

In order to implement the KLIC-based comparison of two different parametric copulas c_A and c_B in the SCOMDY model (1), Chen and Fan (2006) obtain estimates \hat{c}_A and \hat{c}_B based on observations $\{(\mathbf{Y}'_t, \mathbf{X}'_t)\}_{t=1}^T$, which are used to compute a PLR test statistic based on in-sample scores $\{S_t^{\hat{c}_A}, S_t^{\hat{c}_B}\}_{t=1}^T$. Here we consider an alternative testing procedure based on out-of-sample scores, constructed as follows.

Consider the case of a SCOMDY Model Type 2 specified above. Suppose, moreover, that there are two competing parametric copula densities c_A and c_B , say, available. Let \hat{c}_A^0 and \hat{c}_B^0 denote the corresponding copula pseudo parameters. The null hypothesis of

interest is that the average copula score for one-step-ahead density forecasts based on the two copula specifications are the same, that is²

$$H_0 : \mathbb{E}(S_{t+1}^{\hat{c}_A^0}) \equiv \mathbb{E}(\log \hat{c}_A^0(\mathbf{U}_{t+1})) = \mathbb{E}(S_{t+1}^{\hat{c}_B^0}) \equiv \mathbb{E}(\log \hat{c}_B^0(\mathbf{U}_{t+1})).$$

Note that in practice neither the pseudo-parameters \hat{c}_A^0 and \hat{c}_B^0 are known exactly, nor the actual innovations $\varepsilon_{j,t+1}$ and their marginals F_j , so that also \mathbf{U}_t is not known exactly. We deal with this by estimating the pseudo parameters as well as the copula vectors in a moving window framework for the SCOMDY model, as follows. Fix the length of the in-sample period at R observations, and at time t use the observations $\mathbf{Y}_{t-R+1}, \dots, \mathbf{Y}_t$ to estimate the unknown parameters. Recall that these consist of the finite dimensional parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ for the conditional mean and variance specifications, the nonparametric (infinite dimensional) innovation marginals, and the finite dimensional copula parameters $\boldsymbol{\alpha}_A$ and $\boldsymbol{\alpha}_B$. Let us denote the estimated copulas at time t from copula families A and B by $\hat{c}_{A,t}$ and $\hat{c}_{B,t}$, respectively, where the dependence on the size of the moving window, R , is suppressed for notational convenience. Given a point estimate $\hat{\boldsymbol{\theta}}_t$ of the parameter vector, we compute the sequence of in-sample residuals $\{\hat{\varepsilon}_s\}_{s=t-R+1}^t$ as $\hat{\varepsilon}_{j,s} = (Y_{j,s} - \mu_{j,s}(\hat{\boldsymbol{\theta}}_{1,t})) / \sqrt{h_{j,s}(\hat{\boldsymbol{\theta}}_t)}$. In addition, we obtain the one-step ahead forecast error

$$\hat{\varepsilon}_{j,t+1|t} = \frac{Y_{j,t+1} - \mu_{j,t+1}(\hat{\boldsymbol{\theta}}_{1,t})}{\sqrt{h_{j,t+1}(\hat{\boldsymbol{\theta}}_t)}}.$$

We define the empirical copula vector $\hat{\mathbf{U}}_{t+1}$ associated with the forecast errors $\hat{\varepsilon}_{t+1|t}$ as $\hat{\mathbf{U}}_{t+1} = (\hat{U}_{1,t+1}, \dots, \hat{U}_{d,t+1})'$ where

$$\hat{U}_{j,t+1} = \frac{\text{rank of } \hat{\varepsilon}_{j,t+1|t} \text{ among } \{\hat{\varepsilon}_{j,t-R+1}, \dots, \hat{\varepsilon}_{j,t}\}}{R+1}. \quad (5)$$

Let $d_{t+1} = \log \hat{c}_{A,t}(\hat{\mathbf{U}}_{t+1}) - \log \hat{c}_{B,t}(\hat{\mathbf{U}}_{t+1})$ denote the difference in log scores for the two competing copulas, and assume that P observations are available for $t = R, \dots, T =$

²Chen and Fan (2005, 2006) test the null hypothesis that scores for a given copula are *at least as high as* for a competing copula. Although it is possible to consider this hypothesis also within our framework, we focus on the null of equal average copula scores, because this appears relevant to a symmetric treatment of the competing copula families (i.e. appropriate if no benchmark model is available) and has the additional advantage that the asymptotic distribution of the resulting test is free from nuisance parameters. (Specifically, the null distribution is asymptotically standard normal, regardless of the actual processes and competing copula families satisfying the null.)

$R + P$. The test for equal log scores is then obtained as

$$Q_{R,P} = \sqrt{P} \frac{\bar{d}_{R,P}}{\hat{\sigma}_{R,P}},$$

where $\bar{d}_{R,P} = \frac{1}{P} \sum_{t=R}^{T-1} d_{t+1}$ is the sample mean of the log score difference, and $\hat{\sigma}_{R,P}$ is an autocovariance and heteroskedasticity consistent (HAC) estimate of the asymptotic standard deviation of $\sqrt{P}\bar{d}_{R,P}$. The HAC-estimator used for the asymptotic variance of $\sqrt{P}\bar{d}_{R,P}$ is $\hat{\sigma}^2 = \hat{\gamma}_0 + 2 \sum_{k=1}^{K-1} a_k \hat{\gamma}_k$, where $\hat{\gamma}_k$ denotes the lag- k sample covariance of the sequence $\{d_{t+1}\}_{t=R}^{T-1}$ and a_k are the Bartlett weights $a_k = 1 - k/K$ with $K = \lfloor P^{-1/4} \rfloor$.

Taken together, although the null hypothesis H_0 specified above is of our main interest, the actual null hypothesis that we test in practice using $Q_{R,P}$ is given by

$$H'_0 : \quad \mathbb{E}(S_{t+1}^{\hat{c}_{A,t}}) \equiv \mathbb{E}(\log \hat{c}_{A,t}(\hat{\mathbf{U}}_{t+1})) = \mathbb{E}(S_{t+1}^{\hat{c}_{B,t}}) \equiv \mathbb{E}(\log \hat{c}_{B,t}(\hat{\mathbf{U}}_{t+1})).$$

We have the following results for the $Q_{R,P}$ test statistic.

Theorem 1 *Under H'_0 , for an in-sample estimation window of fixed length R , suppose (i) the score difference sequence $\{d_{t+1}\} = \{S_{t+1}^{\hat{c}_{A,t}} - S_{t+1}^{\hat{c}_{B,t}}\}$ is strictly stationary and ϕ -mixing with index $-r/(2r - 1)$, $r \geq 1$, or α -mixing with index $-r/(r - 1)$, $r > 1$, and (ii) $\mathbb{E}|d_{t+1}|^{2(r+\delta)}$ for some $r, \delta > 0$. Then, $Q_{R,P}$ is asymptotically standard normally distributed.*

PROOF: This is an intermediate result in the proof of Giacomini and White's (2006) Theorem 1. The same mixing- and moment conditions apply. \square

This theorem demonstrates the asymptotic behaviour of the $Q_{R,P}$ statistic. Nevertheless, the scores used in practice are an approximation to the true scores involving the null hypothesis H_0 , which is of main interest, which may affect the small sample properties of the test statistic. For that reason we investigate the effect of replacing the copula scores by the empirical copula scores numerically, in the next section, by Monte Carlo simulation.

4 Monte Carlo simulation

In this section we examine the behavior of our copula predictive accuracy test in small samples. We conduct our simulation experiments in the context of the SCOMDY model

(1), for the bivariate case with $d = 2$. For brevity, we only report results for data generating processes (DGPs) of Model Type 1 as given in Section 2, where the conditional means and conditional variances are constant. Throughout (but without loss of generality) we set $\delta_j = 0$ and $\sigma_j = 1$ for $j = 1, 2$ in (3). The correct specifications of the conditional means and variances are assumed to be known. Both marginal distributions of $\varepsilon_{j,t}$ are specified as i.i.d. standard normal.³

In the size and power experiments reported below, we limit ourselves to two popular copulas: the Gaussian copula and the Student's t copula. Both copulas can be obtained using the so-called inversion method, that is

$$C(u_1, u_2, \dots, u_d) = F(F_1^{-1}(u_1), F_2^{-1}(u_2), \dots, F_d^{-1}(u_d)), \quad (6)$$

where F is the joint CDF and $F_i^{-1}(u) = \min\{x | u \leq F_i(x)\}$ is the (quasi)-inverse of the corresponding marginal CDF F_i .

The Gaussian copula is obtained from (6) by taking F to be a multivariate normal distribution with mean zero, unit variances, and correlations ρ_{ij} , $i, j = 1, \dots, d$, and standard normal marginals F_i . The corresponding copula density is given by

$$c(\mathbf{u}; \Sigma) = |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}(\Phi^{-1}(\mathbf{u}))'(\Sigma^{-1} - \mathbf{I}_d)\Phi^{-1}(\mathbf{u})\right), \quad (7)$$

where \mathbf{I}_d is the d -dimensional identity matrix, Σ is the correlation matrix, and $\Phi^{-1}(\mathbf{u}) = (\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d))'$ with $\Phi^{-1}(\cdot)$ the inverse of the standard normal CDF. In the bivariate case $d = 2$, the correlation coefficient $\rho \equiv \rho_{12}$ is the only parameter of the Gaussian copula.

The Student's t copula is obtained similarly, but using a t distribution instead of the Gaussian. The t copula density is given by

$$c(\mathbf{u}; \Sigma, \nu) = |\Sigma|^{-1/2} \frac{\Gamma([\nu + d]/2)\Gamma^{d-1}(\nu/2)}{\Gamma^d((\nu + 1)/2)} \frac{\left(1 + \frac{T_\nu^{-1}(\mathbf{u})'\Sigma^{-1}T_\nu^{-1}(\mathbf{u})}{\nu}\right)^{-(\nu+d)/2}}{\prod_{i=1}^d \left(1 + \frac{(T_\nu^{-1}(u_i))^2}{\nu}\right)^{-(\nu+1)/2}}, \quad (8)$$

³We repeat all experiments reported in this section for DGPs involving time-varying structures in the specification of the conditional means and variances. Specifically, we use Model Type 2 as given in Section 2 with an AR(1) model for the conditional mean and a GARCH(1,1) model for the conditional variance. The results of these experiments are qualitatively and quantitatively similar to results reported here for the case of constant conditional means and variances. Further details are available upon request.

where $T_\nu^{-1}(\mathbf{u}) = (T_\nu^{-1}(u_1), \dots, T_\nu^{-1}(u_d))'$, and $T_\nu(\cdot)$ is the inverse of univariate Student's t CDF, Σ is the correlation matrix and ν is the number of degrees of freedom. In the bivariate case the Student's t copula has two parameters, the degrees of freedom ν and the correlation coefficient $\rho \equiv \rho_{12}$. Note that the Student's t copula nests the Gaussian copula when $\nu = \infty$.

The major difference between the Gaussian copula and the Student's t copula is their ability to capture tail dependence, which may be important for financial applications. For two random variables X and Y with CDFs F_X and F_Y , respectively, the coefficients of lower and upper tail dependence are given by

$$\lambda_l = \lim_{\alpha \rightarrow 0} P(X \leq F_X^{-1}(\alpha) | Y \leq F_Y^{-1}(\alpha)) \quad \text{and} \quad \lambda_u = \lim_{\alpha \rightarrow 1} P(X \geq F_X^{-1}(\alpha) | Y \geq F_Y^{-1}(\alpha)).$$

For the Gaussian copula both tail dependence coefficients are equal to zero, while for the Student's t copula the tail dependence is symmetric⁴ and positive. Specifically, the tail dependence coefficients are given by

$$\lambda_l = \lambda_u = 2T_{\nu+1} \left(-\sqrt{(\nu+1)(1-\rho)/(1+\rho)} \right),$$

which is increasing in the correlation coefficient ρ and decreasing in the degrees of freedom ν .

We evaluate the properties of the copula predictive accuracy test for different combinations of the length of the moving in-sample window $R \in \{100, 1000\}$, and the number of out-of-sample evaluations $P \in \{100, 1000, 5000\}$. For a given window ending at time t , the three-stage estimation procedure outlined in Section 2 is used to estimate the model parameters. That is, having estimated the parameters θ_1 and θ_2 using QMLE, the in-sample copula vectors $\hat{\mathbf{u}}_s$, $s = t - R + 1, \dots, t$ are obtained from the empirical CDF transformation $u_{j,s} = \frac{R_{j,s}}{R+1}$ where $R_{j,s}$ is the rank of $\hat{\varepsilon}_{j,s}$ among the residuals $\hat{\varepsilon}_{j,t-R+1}, \dots, \hat{\varepsilon}_{j,t}$, for $j = 1, \dots, d$. The parameters of both the Gaussian copula and t copula then are estimated using maximum likelihood. Finally, the empirical copula vector $\hat{\mathbf{U}}_{t+1}$ corresponding to the one-step ahead forecast errors $\hat{\varepsilon}_{j,t+1|t}$ is obtained from (5). The number of replications in each experiment is set equal to 1000.

⁴In Section 5 we consider two other copulas that have asymmetric tail dependence, that is, the Clayton copula with only lower tail dependence and the Gumbel copula with only upper tail dependence.

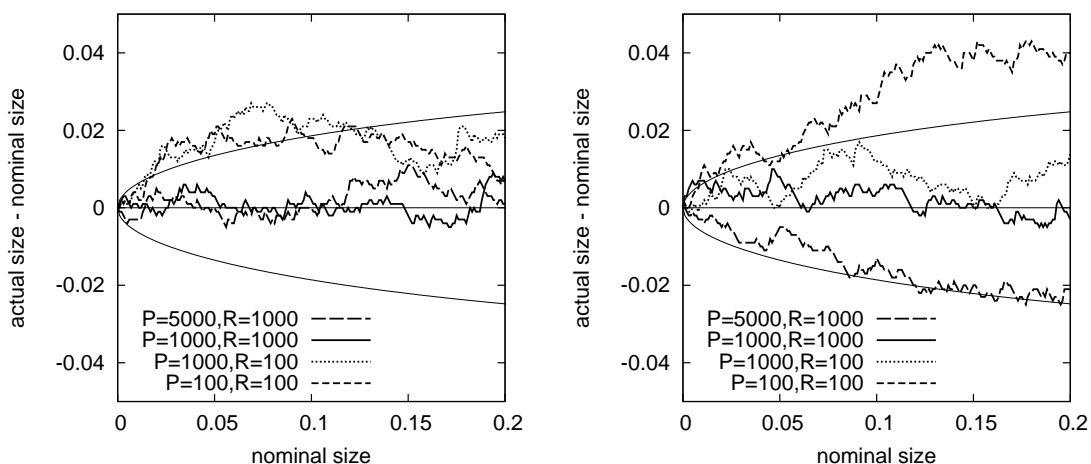


Figure 1: Size-discrepancy plots. The figure displays size-discrepancy curves, showing the difference between the actual two-sided rejection rate and the nominal size (on the vertical axis) for varying nominal significance levels (on the horizontal axis). The thin lines indicate 95% (point-wise) confidence bounds based on 1000 replications. The DGP is the SCOMDY model (3) with standard normal marginal distributions and a bivariate product (independence) copula. The test of equal predictive accuracy compares a t copula with $\rho = 0.3$ against a t copula with $\rho = -0.3$, where ν is fixed at 5 for both copulas in the left panel and ν is estimated for both copulas in the right panel. R denotes the number of observations R in the moving in-sample window and P denotes the number of out-of-sample evaluations.

4.1 Size

In order to assess the size properties of the test a case is required with two competing copulas that are both ‘equally (in)correct’. We achieve this in two different ways, which we call ‘non-nested null’ and ‘nested null’. Under the ‘non-nested null’ the copula in the DGP is specified as the product (or independence) copula. We then test the null hypothesis of equal predictive accuracy of two Student’s t copulas with the correlation parameter fixed at either $\rho = 0.3$ or $\rho = -0.3$. We consider two possibilities for the degrees of freedom parameter: (i) ν is also fixed for both copulas and set equal to 5, and (ii) ν is estimated at each point in time based on a moving window of length R . In either case the two copulas considered in the test are equally distant from the true copula. By considering these two cases we aim to investigate the effect of copula parameter estimation uncertainty on the performance of the $Q_{R,P}$ test statistic.

The discrepancy between the actual size (or observed rejection rate) and the nominal

size of the test is shown on Figure 1, where the null hypothesis is tested against the two-sided alternative that the averages scores of the two copulas are not equal. The left panel corresponds to case (i) where the degrees of freedom parameter ν is fixed at 5, while the right panel corresponds to case (ii) with ν being estimated. In the absence of any estimation uncertainty in the copula parameters, we observe that the empirical rejection rate is close to the nominal significance level when both the length of the moving in-sample window and the number of out-of-sample evaluations are relatively large, that is, when $R = 1000$ and $P = 1000, 5000$. The test is somewhat liberal and rejects the correct null hypothesis too frequently when $R = 100$ and $P = 100$ or 1000 . This dependence of the size properties on the length of the moving in-sample window R may appear surprising at first sight, but recall that the empirical CDF transformation is used to obtain the empirical copula vectors, which still entails some estimation uncertainty. Adding parameter uncertainty due to the estimation of the copula parameter ν results in larger deviations from the nominal size, as shown by the right panel.

Next, we investigate the performance of the test under the ‘nested null’. The DGP in this case is taken to be a Gaussian copula with correlation coefficient $\rho = 0.7$. We test for equal performance between the Gaussian copula and the Student’s t copula.⁵ For both copulas we estimate the correlation parameter ρ , while the degrees of freedom ν in the t copula is estimated as well. The actual/nominal size discrepancy is shown on Figure 2. Similar to the previous experiment, when the in-sample window is of reasonable length ($R = 1000$), the actual size of the test is within the 95% (pointwise) confidence bounds of the nominal significance level. For small estimation windows, $R = 100$, the test is conservative and under-rejects the null hypothesis.

4.2 Power

We evaluate the power of the test of equal predictive accuracy by performing two simulation experiments where one of the competing copula specifications is correct, in the sense that it corresponds with the underlying copula DGP. Again we consider the Gaussian copula and Student’s t copula, hence, we focus on the question whether the $Q_{R,P}$ test statistic

⁵Recall that the Student’s t copula nests the Gaussian copula when $\nu = \infty$.

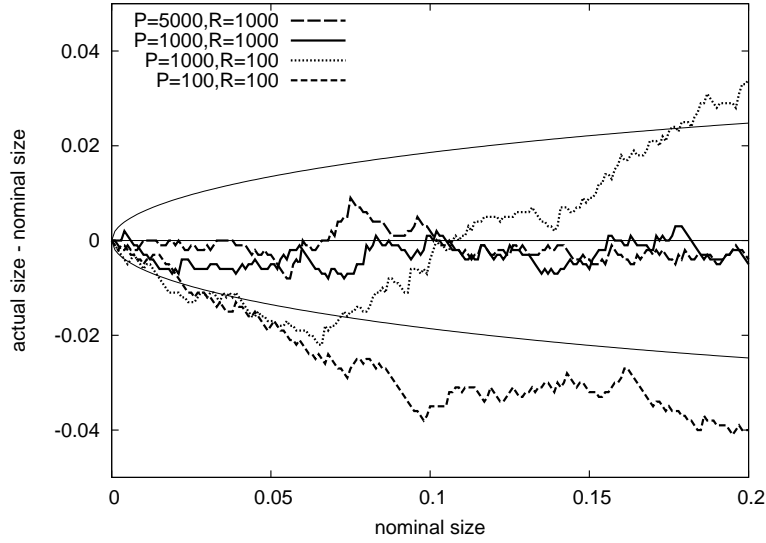


Figure 2: Size-discrepancy plots. The figure displays size-discrepancy curves, showing the difference between the actual two-sided rejection rate and the nominal size (on the vertical axis) for varying nominal significance levels (on the horizontal axis). The thin lines indicate 95% (point-wise) confidence bounds based on 1000 replications. The DGP is the SCOMDY model (3) with standard normal marginal distributions and a bivariate Gaussian copula with $\rho = 0.7$. The test of equal predictive accuracy compares a t copula with both parameters ρ estimated, against a Gaussian copula with parameter ρ estimated. R denotes the number of observations R in the moving in-sample window and P denotes the number of out-of-sample evaluations. The graphs are based on 1000 replications.

can distinguish between copulas with and without tail dependence.

We start with the DGP specified by a Student's t copula with correlation coefficient $\rho = 0.7$ and consider two different values of the degrees of freedom parameter, $\nu = 5$ and $\nu = 10$. We test the t copula specification (with both parameters estimated) against the Gaussian copula specification (with the correlation parameter ρ estimated). Intuitively, as the t copula approaches the Gaussian copula as ν increases (such that the tail dependence disappears with the coefficients λ_l and λ_u converging to zero), the higher the value of ν in the DGP, the more difficult it is to distinguish between these two copulas.

The results are shown in Figure 3 in the form of size-power curves, showing the observed rejection rate for different nominal significance levels. The left and right panels relate to the t copulas with $\nu = 5$ and $\nu = 10$, respectively. Here we consider a one-sided test, against the correct alternative hypothesis that the true copula has a higher average score. Note that these graphs also can be used to infer the rejection rates against the incor-

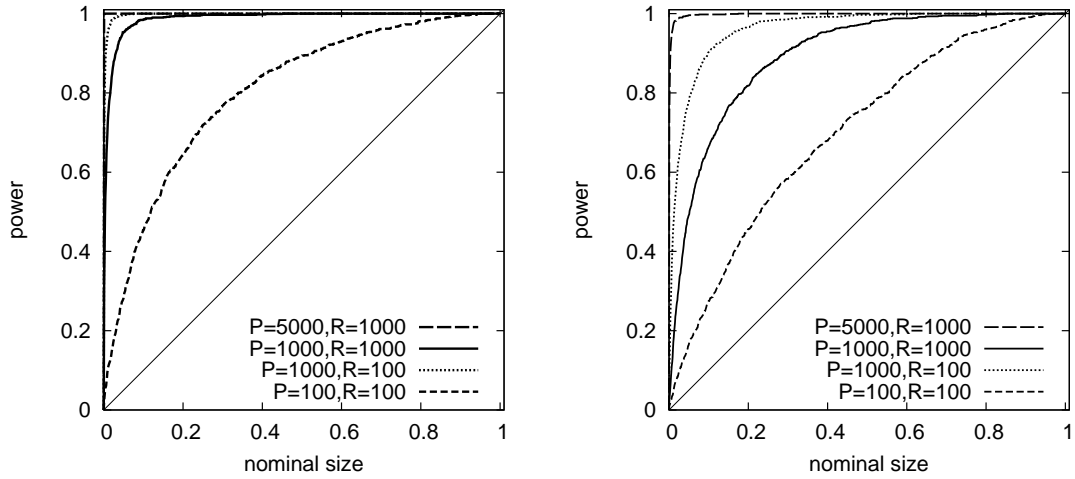


Figure 3: Size-power curves. The figure displays size-power curves, showing the rejection rate of a one-sided test against the alternative hypothesis that the true copula has higher average score (on the vertical axis) for varying nominal significance levels (on the horizontal axis). The DGP is the SCOMDY model (3) with standard normal marginal distributions and a bivariate Student's t copula with $\rho = 0.7$ and $\nu = 5$ in the left panel and $\nu = 10$ in the right panel. The test of equal predictive accuracy compares a t copula with both parameters ρ and ν estimated against a Gaussian copula with parameter ρ estimated. R denotes the number of observations R in the moving in-sample window and P denotes the number of out-of-sample evaluations. Results are based on 1000 replications.

rect alternative hypothesis that the other copula has a higher average score. This 'spurious power' at a certain nominal significance level α is equal to one minus the 'true power' (i.e. the rejection rate against the correct alternative) at nominal significance level $1 - \alpha$.

Figure 3 shows that, as expected, the test has higher power when $\nu = 5$, that is, when the distance between the two copula specifications being compared is larger. The 'true' power of the test quickly goes to unity when $P = 5000$. In case $R = 100$ and $P = 100$ the power is relatively low. The size-power curves are completely above the 45 degree line, showing that 'spurious' power of the $Q_{R,P}$ test statistic is always below the nominal significance level.

In our final experiment we consider a DGP without tail dependence by using a Gaussian copula with correlation coefficient $\rho = 0.7$. We examine whether the test can distinguish between the correct Gaussian copula and a Student's t copula with the degrees of freedom parameter fixed at $\nu = 5$ or $\nu = 10$, which imposes positive tail dependence. The

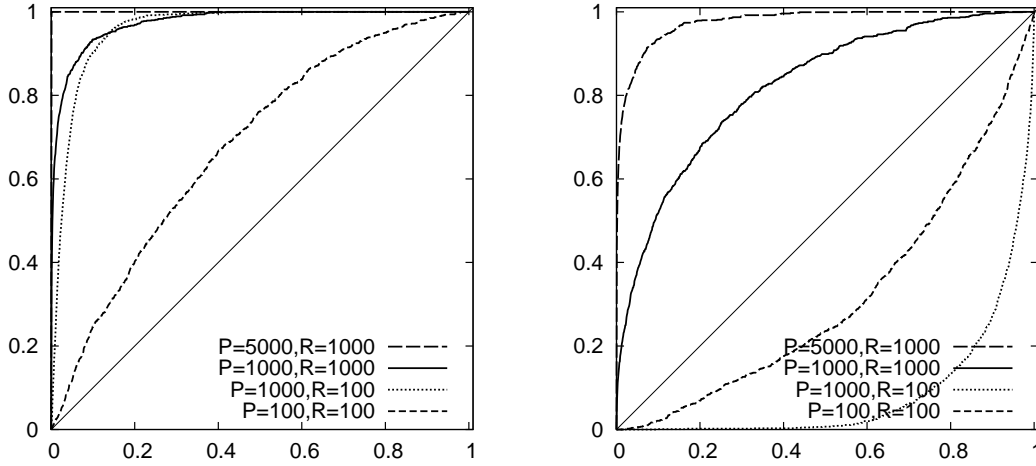


Figure 4: Size-power curves. The figure displays size-power curves, showing the rejection rate of a one-sided test against the alternative hypothesis that the true copula has higher average score (on the vertical axis) for varying nominal significance levels (on the horizontal axis). The DGP is the SCOMDY model (3) with standard normal marginal distributions and a bivariate Gaussian copula with $\rho = 0.7$. The test of equal predictive accuracy compares a t copula with parameter ρ estimated and ν fixed at 5 in the left panel and at 10 in the right panel, against a Gaussian copula with parameter ρ estimated. R denotes the number of observations R in the moving in-sample window and P denotes the number of out-of-sample evaluations. Results are based on 1000 replications.

parameter ρ is estimated for both the Gaussian and t copulas. Figure 4 shows size-power curves for this experiment, with $\nu = 5$ and 10 in the left and right panels, respectively. We observe that the power of the test is higher when $\nu = 5$, again due to the fact that the distance of the t copula from the true Gaussian copula is larger in this case. The power of the test is close to one when $P = 5000$ even for very low significance levels. Also for $\nu = 5$ we observe that power is still reasonable when the moving estimation window is small ($R = 100$) as long as the number of out-of-sample forecasts is reasonably large ($P = 1000$). For $P = 100$ power declines considerably, although we still do not observe rejection rates against the incorrect alternative exceeding the nominal significance level. This is rather different for the comparison of the Gaussian copula with a t copula with $\nu = 10$ in the right panel of Figure 4. Here for $R = 100$ we observe substantial ‘spurious’ power, in the sense that the rejection rates against the incorrect alternative are considerable. Of course in this case the distance between the true and alternative copulas is not

large to begin with, and due to the estimation uncertainty the $Q_{R,P}$ test statistic apparently is misguided to such an extent that it often incorrectly selects the Student's t copula.

From the above simulation experiments we conclude that the test of equal predictive accuracy has good finite sample properties when the number of out-of-sample evaluations is sufficiently large ($P = 1000, 5000$).

5 Empirical application

We illustrate the use of the predictive accuracy test for comparing alternative copula specifications with an empirical application. We consider daily returns on US dollar exchange rates of the Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP) and Japanese yen (JPY) over the period from January 1, 1980 until June 25, 2008 (7160 observations). Up to December 31, 1998, the euro series actually concerns the exchange rate of the German Deutschmark, while the euro is used as of January 1, 1999. The data are obtained from the Federal Reserve Bank of New York and concerns noon buying rates in New York. We work with two groups of three exchange rates each, that is, GBP-CHF-EUR, and CAD-JPY-EUR. The first, inter-European group presumably has a relatively high level of dependence compared to the second, inter-continental group. This is confirmed by the unconditional correlations between the daily returns series (not shown here to save space).

We use SCOMDY models as discussed in Section 2 to model the exchange rate returns and their dependence. The conditional means and the conditional variances for the three series in each group are specified by an AR(5)-GARCH(1,1) model given by

$$Y_{j,t} = c_j + \sum_{l=1}^5 \phi_{j,l} Y_{j,t-l} + \sqrt{h_{j,t}} \varepsilon_{j,t}$$

$$h_{j,t} = \kappa_j + \gamma_j \left(Y_{j,t-1} - c_j - \sum_{l=1}^5 \phi_{j,l} Y_{j,t-1-l} \right)^2 + \beta_j h_{j,t-1},$$

where $\kappa_j > 0$, $\beta_j \geq 0$, $\gamma_j > 0$ and $\beta_j + \gamma_j < 1$, for $j = 1, \dots, 3$. The marginal distributions $F_j(\cdot)$ of the standardized innovations $\varepsilon_{j,t}$ are specified nonparametrically.

The dependence between the marginals is captured by four alternative copula specifications, which we compare in terms of their relative performance in out-of-sample den-

sity forecasting. In particular, we consider the Gaussian copula, the Student's t copula, the Clayton copula and the Gumbel copula. The Gaussian and Student's t copula were introduced earlier in Section 4. The Clayton and Gumbel copulas are from the family of Archimedean copulas (see Nelsen (2006) for details) and are briefly described below.

The d -dimensional Clayton copula is given by

$$C(u_1, u_2, \dots, u_d; \alpha) = \left(\sum_{i=1}^d u_i^{-\alpha} - d + 1 \right)^{-1/\alpha}, \quad \text{with } \alpha > 0.$$

In contrast to the Gaussian and Student's t copulas, the Clayton copula is able to capture asymmetric tail dependence. In fact, it can only exhibit lower tail dependence, while upper tail dependence is absent. In the bivariate case the lower tail dependence coefficient for the Clayton copula is $\lambda_l = 2^{-1/\alpha}$, which is increasing in the parameter α . The density function of the Clayton copula is

$$c(\mathbf{u}, \alpha) = \left(\prod_{i=1}^d (1 + (i-1)\alpha) \right) \left(\prod_{i=1}^d u_i^{-(\alpha+1)} \right) \left(\sum_{i=1}^d u_i^{-\alpha} - d + 1 \right)^{-1/\alpha+d}.$$

The Gumbel copula is specified as

$$C(u_1, u_2, \dots, u_d; \alpha) = \exp \left(- \left[\sum_{i=1}^d (-\ln u_i)^\alpha \right]^{1/\alpha} \right), \quad \text{with } \alpha > 1.$$

The Gumbel copula also exhibits asymmetric tail dependence, but in contrast to the Clayton copula it only allows upper tail dependence, while lower tail dependence is absent. In the bivariate case the upper tail dependence coefficient for the Gumbel copula is $\lambda_u = 2 - 2^{1/\alpha}$, which is again increasing in the parameter α . The expression for the multivariate Gumbel copula density is rather cumbersome and we do not show it here. Note that neither the Clayton copula nor the Gumbel copula can capture negative dependence and, thus, may not always be suitable for modeling.

We compare the one-step ahead density forecasting performance of the four copulas using a rolling window scheme. We estimate the SCOMDY model parameters using the three-stage procedure described in Section 2. The length of the in-sample estimation window is set to $R = 2000$ observations, which leaves $P = 5160$ observations for out-of-sample evaluation, covering the period January 1988 - July 2008. Patton (2006)

documents a structural break in exchange rate relationships after the introduction of the euro on January 1, 1999. We therefore partition our evaluation period at this date into two sub-periods, called pre-euro and post-euro, to explore the implications of the structural break on the forecasting performance of the copula specifications. The number of observations differs slightly for the pre- and post-euro periods, being equal to $P = 2772$ and $P = 2388$, respectively.

Table 1 reports the values of the pairwise $Q_{R,P}$ test statistic for the full sample, as well as the pre- and post-euro sub-periods, for the two groups of three exchange rates. In most cases the test makes a clear judgment about the two copulas being compared, in the sense that the null of equal predictive accuracy is generally rejected at the 1% significance level or better. The Clayton copula shows the worst performance in comparison to any other copula. The Gumbel copula is dominated by the Gaussian copula in all cases except for the CAD-JPY-EUR group during the post-euro subperiod, where the equal performance of these two copulas is not rejected at the 10% level. The Gaussian copula is, in turn, outperformed by the Student's t copula in all cases except for the GBP-CHF-EUR group during the post-euro subperiod, where again there is not enough evidence to reject the null of equal performance of these two copulas. Thus, overall we conclude that the Student's t copula performs best in terms of out-of-sample multivariate density forecasts. A similar conclusion is reached by Chen and Fan (2006) based on the in-sample versions of our tests.

To provide more insight into these results, in Figure 5 we show the rolling window parameter estimates of the Student's t copula for the GBP-CHF-EUR group in the left panel and for the CAD-JPY-EUR group in the right panel. The dates on the horizontal axis correspond with the end of the estimation window, that is, the moment these estimates are used for constructing the density forecast. For presentation purposes the parameter estimates are reported for every 100th day in the evaluation period. Three aspects are worth discussing in more detail. First, the estimates of the degrees of freedom parameter ν exhibit substantial variation over time, ranging between 6 and 18 for the CAD-JPY-EUR group and between 4 and 10 for the GBP-CHF-EUR group. Recall that lower values of ν indicate stronger tail dependence; this is found to be the major reason for the rejection of

Table 1: Pair-wise tests for out-of-sample performance of copulas

Copula B	Copula A		
	Gumbel	Student's t	Gaussian
GBP-CHF-EUR			
<i>Full sample</i>			
Clayton	−8.38***	−21.66***	−21.50***
Gumbel	−	−19.59***	−19.71***
Student's t	−	−	2.51**
<i>Pre-euro</i>			
Clayton	−7.72***	−15.95***	−15.67***
Gumbel	−	−14.42***	−13.91***
Student's t	−	−	3.17***
<i>Post-euro</i>			
Clayton	−3.80***	−16.32***	−16.06***
Gumbel	−	−14.16***	−14.79***
Student's t	−	−	0.55
CAD-JPY-EUR			
<i>Full sample</i>			
Clayton	−16.90***	−19.89***	−19.14***
Gumbel	−	−8.11***	−6.65***
Student's t	−	−	4.04***
<i>Pre-euro</i>			
Clayton	−15.88***	−18.72***	−18.15***
Gumbel	−	−7.61***	−6.91***
Student's t	−	−	1.78*
<i>Post-euro</i>			
Clayton	−8.76***	−10.35***	−9.79***
Gumbel	−	−3.66***	−1.57
Student's t	−	−	4.54***

Note: The table present values of the $Q_{R,P}$ test statistic of the null hypothesis of equal predictive accuracy of two alternative copula specifications. Positive (Negative) values indicate better performance of copula A (B). The asterisks *, ** and *** indicate significance at (two-sided) 10%, 5% and 1% levels respectively. The length of the moving in-sample estimation window is equal to $R = 2000$ in all cases. The number of observations in the out-of-sample period is equal to $P = 5160$ for the full sample, $P = 2772$ for the pre-euro period and $P = 2388$ for the post-euro period.

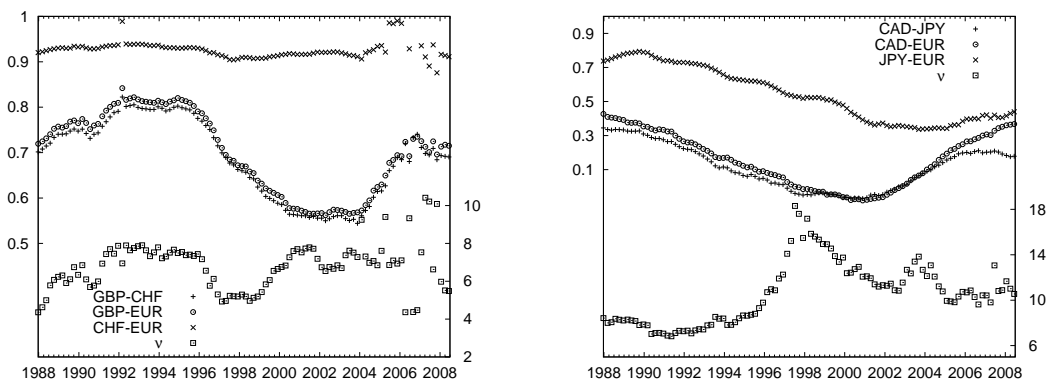


Figure 5: Parameter estimates of the Student's t copula over time for the two groups of exchange rates: GBP-CHF-EUR in the left panel and CAD-JPY-EUR in the right panel. The parameter estimates are reported for every 100th day in the evaluation period. Year labels indicate the beginning of the corresponding year.

the Gaussian copula against the t copula.

Second, the correlation estimates clearly differ among the exchange rate pairs within each group. This suggests that the one-parameter Clayton and Gumbel copulas are not suitable as they are not able to reflect these differences. This explains the inferior forecasting performance of these two copulas relative to the Gaussian and t specifications.

Third, we observe slowly changing correlations over time. There is some evidence for a structural break due to the euro introduction on January 1, 1999. The post-euro period is characterized by slowly increasing increasing correlations for CAD-EUR, CAD-JPY and JPY-EUR, while they were slowly decreasing before. The same is true for the exchange rate pairs in the other group, that is the GBP-CHF and GBP-EUR correlations also start increasing after the introduction of the euro. The CHF-EUR correlations are close to unity over the whole sample period, with a slight tendency to increase during the post-euro period. This creates some challenges for the maximization procedure and results in somewhat unstable estimates around 2006-2008. It is also more difficult to distinguish the Gaussian copula from the Student's t copula when the correlation is close to one. In the limiting case of unit correlation the Gaussian copula also has tail dependence.⁶ All these factors contribute to the inability of the test to reject the equal performance between the

⁶The correlations estimates obtained for the Gaussian copula specifications are very similar to the values shown in Figure 5 for the Student's t copula.

Gaussian copula and the Student's t during this subperiod.

6 Summary and discussion

In this paper we have introduced a new statistical test for comparing alternative parametric copula specifications in the context of multivariate density forecasts. The test is based on the Kullback-Leibler Information Criterion (KLIC), which is a measure of distance of the candidate copula specifications to the true copula. Although the true copula is unknown, score differences can be used to compare copula specifications pair-wise; the copula with the smallest average KLIC is considered to be superior. The main difference with other tests for competing copula families is that we use an out-of-sample measure for the suitability of the copulas.

Following Giacomini and White (2006), the potential problem that parameter estimation uncertainty might affect the distribution of test statistics is avoided by considering this aspect as an integral part of the forecasting method. The resulting test is valid under general conditions, allowing for parameter estimation uncertainty and for comparing either nested or non-nested copula families.

In practice, the test would be used to compare two competing copula families by testing for equal average scores against the two-sided alternative that one of the families receives a higher average score. However, to distinguish between true power and spurious power (rejections of the null hypothesis in favor of the wrong copula family) we presented Monte Carlo simulation results in the form of one-sided size-power plots. The Monte Carlo results suggest that the proposed test has satisfactory size and power properties in finite samples. Spurious power is observed, but only in cases of very small estimation windows and alternative copula specifications that are not very different from each other in terms of their dependence properties.

In an application to daily exchange rate returns of several major currencies against the US dollar we found that the Student's t copula is clearly favored over Gaussian, Gumbel and Clayton specifications. This suggests that these exchange rate returns are characterized by symmetric tail dependence.

References

- Amisano, G. and Giacomini, R. (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business and Economic Statistics*, **25**, 177–190.
- Bao, Y., Lee, T.-H. and Saltoğlu, B. (2007). Comparing density forecast models. *Journal of Forecasting*, **26**, 203–225.
- Berg, D. (2007). Copula goodness-of-fit testing: An overview and power comparison. Working paper, Norwegian Computing Center.
- Berg, D. and Bakken, H. (2007). A copula goodness-of-fit approach based on the conditional probability integral transform. Working paper, Norwegian Computing Center.
- Breymann, W., Dias, A. and Embrechts, P. (2003). Dependence structures for multivariate high-frequency data in finance. *Quantitative Finance*, **3**, number 1, 1–14.
- Chen, X. and Fan, Y. (2005). Pseudo-likelihood ratio tests for semiparametric multivariate copula model selection. *Canadian Journal of Statistics*, **33**, 389–414.
- Chen, X. and Fan, Y. (2006). Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification. *Journal of Econometrics*, **135**, 125–154.
- Cherubini, U., Luciano, E. and Vecchiato, E. (2004). *Copula Methods in Finance*. Wiley Finance Series. Wiley, Chichester.
- DeHeuvels, P. (1979). La fonction de dépendance empirique et ses propriétés. Un test non paramétrique d'indépendance. *Académie Royale de Belgique. Bulletin de la Classe de Sciences, 5ème Serie* **65**, 274–292.
- Fermanian, J.-D. (2005). Goodness-of-fit tests for copulas. *Journal of Multivariate Analysis*, **95**, 119–152.
- Genest, C., Quessy, J.-F. and Remillard, B. (2006). Goodness-of-fit procedures for copula models based on the probability integral transformation. *Scandinavian Journal of Statistics*, **33**, 337–366.

- Genest, C. and Rémillard, B. (in press). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Annales de l'Institut Henri Poincaré - Probabilités et Statistiques*.
- Genest, C., Rémillard, B. and Beaudoin, D. (in press). Omnibus goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*.
- Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, **74**, 1545–1578.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. London: Chapman & Hall.
- Kole, E., Koedijk, K. and Verbeek, M. (2007). Selecting copulas for risk management. *Journal of Banking and Finance*, **31**, 2405–2423.
- Malevergne, Y. and Sornette, D. (2003). Testing the Gaussian copula hypothesis for financial assets dependence. *Quantitative Finance*, **3**, 231–250.
- Mitchell, J. and Hall, S.G. (2005). Evaluating, comparing and combining density forecasts using the KLIC with an application to the Bank of England and NIESR ‘fan’ charts of inflation. *Oxford Bulletin of Economics and Statistics*, **67**, 995–1033.
- Nelsen, R.B. (2006). *An Introduction to Copulas*, Second edn. New York: Springer Verlag.
- Panchenko, V. (2005). Goodness-of-fit tests for copulas. *Physica A*, **355**, 176–182.
- Patton, A.J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, **47**, 527–556.
- Patton, A.J. (2008). Copula-based models for financial time series. In *Handbook of Financial Time Series* (eds T.G. Andersen, R.A. Davis, J.-P. Kreiss and T. Mikosch). Berlin: Springer-Verlag.
- Rivers, D. and Vuong, Q. (2002). Model selection tests for non-linear dynamic models. *Econometrics Journal*, **5**, 1–39.

- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Annals of Mathematical Statistics*, **23**, 470–472.
- Sklar, A. (1959). Fonction de répartition à n dimensions et leurs marges. *Publications de l'Institut Statistique de l'Université de Paris*, **8**, 229–231.
- Vuong, Q. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, **57**, 307–333.