UvA-DARE (Digital Academic Repository)

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

Diks, C.; Panchenko, V.; van Dijk, D.

DOI

[10.1016/j.jedc.2010.06.021](https://doi.org/10.1016/j.jedc.2010.06.021)

Publication date

2010

Document Version

Author accepted manuscript

Published in

Journal of Economic Dynamics & Control

[Link to publication](#)

Citation for published version (APA):

Diks, C., Panchenko, V., & van Dijk, D. (2010). Out-of-sample comparison of copula specifications in multivariate density forecasts. *Journal of Economic Dynamics & Control*, 34(9), 1596-1609. <https://doi.org/10.1016/j.jedc.2010.06.021>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.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*

June 25, 2010

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 on the competing copulas: 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-*t* copula is favored over Gaussian, Gumbel and Clayton copulas.

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

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

^{*}We thank two anonymous referees for helpful comments and suggestions. This research was partially supported under Australian Research Council's Discovery Projects funding scheme (project number DP0986718).

[†]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@ese.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), Patton (2009) and Genest *et al.* (2009) for recent surveys. The recent account of structured finance instruments by Coval *et al.* (2009) emphasizes once more that the dependence structure among random variables is crucial for pricing and risk assessment. The main attractive property of using copulas is that they allow for modeling the marginal distributions and the dependence structure of the variables of interest separately. Many parametric copula families are available, with 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 specification tests have been put forward. 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); and based on distance measures between (a function of) the empirical (nonparametric) copula (DeHeuvels, 1979) and the candidate (parametric) copula by Panchenko (2005), Genest *et al.* (2006), Genest and Rémillard (2008), and Genest *et al.* (2009). We refer the interested reader to Berg (2009) for a detailed review of these approaches and a power comparison based on extensive simulation experiments.

The related problem of selecting among alternative copula specifications from two or more parametric families has mostly been approached in an *indirect* way. Typically, one or several goodness-of-fit tests are applied to each of 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 dis-

tance from the true (but unknown) copula. Similar to the goodness-of-fit tests, these PLR tests are based on the in-sample fit of the 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, by using 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 out-of-sample KLIC scores 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, 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 in the multivariate density case considered in this paper.

Our test of equal predictive accuracy can be applied to fully parametric, semi-parametric and nonparametric copula-based multivariate density models. The test is valid under general conditions on the competing copulas, which is achieved by adopting the framework of Giacomini and White (2006). This assumes that any 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, because parameter estimation uncertainty does not play a role (it simply is part of the respective competing forecast methods). In addition, the asymptotic distribution of our test statistic in this case does not depend on

whether or not the competing copulas belong to nested families.

We examine the size and power properties of our copula predictive accuracy test via Monte Carlo simulations. Here we adopt the framework of semi-parametric copula-based multivariate dynamic (SCOMDY) models developed in Chen and Fan (2005, 2006), which combines parametric specifications for the conditional mean and conditional variance with a semi-parametric specification for the distribution of the (standardized) innovations, consisting of a parametric copula with nonparametric univariate marginal distributions. Our simulation results demonstrate that the predictive accuracy test has satisfactory size and power properties 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 - June 2008. Based on the relative predictive accuracy of one-step ahead forecasts we find that the Student- t copula is favored over Gaussian, Gumbel and Clayton copulas.

The paper is organized as follows. In Section 2 we briefly review copula-based multivariate density models. 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 Copula-based multivariate density models

Consider a vector stochastic process $\{(\mathbf{Y}'_t, \mathbf{X}'_t)\}$, $t = 1, 2, \dots$, in which $\mathbf{Y}_t = (Y_{1,t}, Y_{2,t}, \dots, Y_{d,t})'$ represents the d -dimensional vector of central interest, and $\mathbf{X}_t$ a vector of exogenous or pre-determined variables. Sklar (1959)'s theorem states that we can express the joint distribution $F(\mathbf{y})$ of $\mathbf{Y}_t$ in terms of the marginal distributions $F_1, F_2, \dots, F_d$ and a copula function C , that is

$$F(\mathbf{y}) = C(F_1(y_1), F_2(y_2), \dots, F_d(y_d)). \quad (1)$$

The decomposition in (1) clearly shows the attractiveness of the copula approach for modeling multivariate distributions. Given that the marginal distributions F_j , $j = 1, \dots, d$, only contain univariate information on the individual variables Y_j , their dependence is governed completely by the copula function C . As the choice of marginal distributions does not restrict the choice of dependence function, or vice versa, a wide range of joint distributions can be obtained by combining different marginals with different copulas.

In the context of time series variables $\mathbf{Y}_t$, instead of the unconditional distribution $F(\mathbf{y})$ it is more natural to consider the conditional distribution $F(\mathbf{y}|\mathcal{F}_{t-1})$, where $\mathcal{F}_{t-1}$ denotes the sigma-field generated by $\{\mathbf{Y}_{t-1}, \mathbf{Y}_{t-2}, \dots; \mathbf{X}_t, \mathbf{X}_{t-1}, \dots\}$. Patton (2006) provides an extension of Sklar's theorem to this case, allowing $F(\mathbf{y}_t|\mathcal{F}_{t-1})$ to be decomposed into conditional marginal distributions $F_j(y_{jt}|\mathcal{F}_{t-1})$, $j = 1, \dots, d$, and a conditional copula $C(\cdot|\mathcal{F}_{t-1})$, that is

$$F(\mathbf{y}|\mathcal{F}_{t-1}) = C(F_1(y_{1,t}|\mathcal{F}_{t-1}), F_2(y_{2,t}|\mathcal{F}_{t-1}), \dots, F_d(y_{d,t}|\mathcal{F}_{t-1})|\mathcal{F}_{t-1}). \quad (2)$$

The unconditional and conditional copula-based multivariate distributions in (1) and (2) can be specified in various ways. For instance, one may specify the (conditional) marginal distributions F_j , $j = 1, \dots, d$, and the (conditional) copula function C parametrically, up to a finite-dimensional vector of parameters. Popular parametric copulas include the Gaussian, Student- t , Clayton, and Gumbel specifications, see Joe (1997) and Nelsen (2006) for overviews. Each of these copulas has rather different dependence properties, as characterized by the presence or absence of symmetry and/or tail dependence. To summarize some of these properties for a number of copulas briefly, it is convenient to introduce the notation $\tilde{C}$ for the copula of $-\mathbf{Y}$, known as the survival copula of $\mathbf{Y}$. The lower tail dependence coefficient is defined as $\lambda_L = \lim_{q \downarrow 0} C(q, \dots, q)/q$. The upper tail dependence coefficient is defined analogously as $\lambda_U = \lim_{q \downarrow 0} \tilde{C}(q, \dots, q)/q$. The Gaussian copula is symmetric, in the sense that $C(u_1, \dots, u_d) = \tilde{C}(u_1, \dots, u_d)$ for all $0 \leq u_i \leq 1$, and has no tail dependence, that is, $\lambda_L = \lambda_U = 0$. The Student- t copula also is symmetric but with positive tail dependence. On the other hand, the Gumbel copula is asymmetric ($C(u_1, \dots, u_d) \neq \tilde{C}(u_1, \dots, u_d)$) and has upper but no lower tail dependence, while the reverse holds for (the tail dependence of) the Clayton copula.

The flexibility of the parametric approach may be further increased by allowing for time-variation in the parameters of the (conditional) copula, see Patton (2006) and Jondeau and Rockinger (2006) for instance. Alternatively, a mixture copula may be considered by taking a linear combination of several parametric copulas, possibly with time-varying weights, as in Rodriguez (2007) and Okimoto (2008), among others.

The advantage of a fully parametric copula-based multivariate density model is that maximum likelihood (ML) can be applied to obtain efficient estimates of the parameters characterizing the (conditional) marginal distributions and the (conditional) copula. The contribution to the log-likelihood function for the conditional copula specification in (2), for example, is given by

$$\sum_{j=1}^d \log f_j(y_{j,t}|\mathcal{F}_{t-1}) + \log c(F_1(y_{1,t}|\mathcal{F}_{t-1}), F_2(y_{2,t}|\mathcal{F}_{t-1}), \dots, F_d(y_{d,t}|\mathcal{F}_{t-1})|\mathcal{F}_{t-1}), \quad (3)$$

where $f_j(y_{j,t}|\mathcal{F}_{t-1})$, $j = 1, \dots, d$, are the conditional marginal densities and c is the conditional copula density, defined as

$$c(u_1, u_2, \dots, u_d|\mathcal{F}_{t-1}) = \frac{\partial^d}{\partial u_1 \partial u_2 \dots \partial u_d} C(u_1, u_2, \dots, u_d|\mathcal{F}_{t-1}),$$

which we will assume to exist throughout.

The parameters in the conditional marginal distributions and the conditional copula may be estimated simultaneously by maximizing the log-likelihood based on a sample of T observations. If the parameters of the marginal distributions can be separated from each other and from those of the copula function, a two-stage estimation procedure is an alternative. This so-called inference functions for margins (IFM) method boils down to first estimating the parameters of the conditional marginals by means of univariate ML, followed by estimating the parameters of the copula by maximizing the sum over the observations of the rightmost term in (3) conditional on the parameter estimates for the marginal distributions. Many parametric copulas have only a low-dimensional parameter vector, so that the copula estimation problem is low-dimensional if the IFM method can be used.¹

¹On the other hand, this also implies that the dependence structure of such copulas can be (too) restrictive.

The drawback of a fully parametric model obviously is that the results depend on the assumption that both the marginal distributions and the copula function are correctly specified. Misspecification of the marginal distributions, for example, can lead to severely biased estimates of the copula parameters, see Fermanian and Scaillet (2005).

This consideration in fact leads to an alternative approach for specifying a copula-based multivariate density model by combining a parametric copula specification with non- or semi-parametric marginal distributions. This avoids the adverse effects of misspecified parametric marginal distributions, which seems attractive especially in case most interest is in the dependence structure of $\mathbf{Y}_t$. The most popular nonparametric estimator in the context of copula-based models is the (rescaled) empirical CDF

$$\hat{F}_j(y) = \frac{1}{T+1} \sum_{t=1}^T I[y_{j,t} \leq y], \quad (4)$$

where $I[A]$ is an indicator function for the event A and T denotes the sample size. Estimates of the copula parameters may then be obtained by means of maximum likelihood, plugging in the estimates $\hat{F}_j(y)$ from (4). Chen *et al.* (2006) demonstrate that more efficient estimates may be obtained by using sieve estimates for the marginal distributions instead of the empirical CDF.

A compromise between a fully parametric specification of the marginal distribution and a nonparametric estimate such as (4) is offered by the class of semi-parametric copula-based multivariate dynamic (SCOMDY) models introduced by Chen and Fan (2006). We discuss this approach in some detail here, as we use it in the Monte Carlo simulations and the empirical application in subsequent sections. The SCOMDY models combine parametric specifications for the conditional mean and conditional variance of $\mathbf{Y}_t$ with a semi-parametric specification for the distribution of the (standardized) innovations, consisting of a parametric copula with nonparametric univariate marginal distributions. The general SCOMDY model is specified as

$$\mathbf{Y}_t = \boldsymbol{\mu}_t(\boldsymbol{\theta}_1) + \sqrt{H_t(\boldsymbol{\theta})} \boldsymbol{\varepsilon}_t, \quad (5)$$

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

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

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 $j = 1, \dots, d$. Applying Sklar's theorem, the joint distribution function $F(\boldsymbol{\varepsilon})$ of $\boldsymbol{\varepsilon}_t$ can be written as

$$F(\boldsymbol{\varepsilon}) = C(F_1(\varepsilon_1), \dots, F_d(\varepsilon_d); \boldsymbol{\alpha}) \equiv C(u_1, \dots, u_d; \boldsymbol{\alpha}), \quad (6)$$

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 $\boldsymbol{\alpha}$.

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. Specifically, Chen and Fan (2006) suggest the following three-stage procedure to estimate the SCOMDY model parameters. First, univariate quasi maximum likelihood under the assumptions of normality of the standardized innovations $\varepsilon_{j,t}$ is used to estimate the parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$. Second, estimates of the marginal distributions $F_j(\cdot)$ are obtained by means of the empirical CDF transformation of the residuals $\hat{\varepsilon}_{j,t} \equiv (y_{j,t} - \mu_{j,t}(\hat{\boldsymbol{\theta}}_1)) / \sqrt{h_{j,t}(\hat{\boldsymbol{\theta}})}$. Finally, the parameters of a given copula specification are estimated by maximizing the corresponding copula log-likelihood function.

In case of large uncertainty about the functional form of the copula, one may consider nonparametric estimation of the copula. The marginal distributions in this case may be modeled either parametrically or nonparametrically. DeHeuvels (1979) proposed an empirical copula estimator, which is similar to the empirical CDF estimator in (4) and is

given by

$$\hat{C}(u_1, u_2, \dots, u_d) = \frac{1}{T+1} \sum_{t=1}^T I[\hat{u}_{1,t} \leq u_1] I[\hat{u}_{2,t} \leq u_2] \dots I[\hat{u}_{d,t} \leq u_d],$$

where $\hat{u}_{j,t}$ is an estimate of the PIT $F_j(y_{j,t}|\mathcal{F}_{t-1})$. Smoother alternatives to the empirical copula include approximations with kernel estimators, as proposed by Fermanian and Scaillet (2003), with Bernstein polynomials, suggested by Sancetta and Satchell (2004), and with splines as in Shen *et al.* (2008) and Dimitrova *et al.* (2008). Nonparametric estimators of copulas generally are subject to the curse of dimensionality, such that their applicability is limited to low dimensions.

3 Equal predictive accuracy test for copulas

As mentioned in the introduction, the KLIC, measuring the divergence between the true probability density and a candidate density, has become a popular tool 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 copula-based density forecasts.

The KLIC of a one-step-ahead predictive density $\hat{f}_t$ for $\mathbf{Y}_{t+1}$ is given by

$$\text{KLIC}(\hat{f}_t) = \int_{\mathbb{R}^d} p_t(\mathbf{y}_{t+1}) \log \frac{p_t(\mathbf{y}_{t+1})}{\hat{f}_t(\mathbf{y}_{t+1})} d\mathbf{y}_{t+1} = \mathbb{E}_t \left(\log p_t(\mathbf{Y}_{t+1}) - \log \hat{f}_t(\mathbf{Y}_{t+1}) \right),$$

where p_t denotes the true but unknown conditional density of $\mathbf{Y}_{t+1}$. The KLIC is a measure of divergence between $\hat{f}_t$ and p_t ; it is non-negative, and equal to zero only if $\hat{f}_t$ corresponds with the true conditional density p_t . The smaller the KLIC or, in other words, the larger $\mathbb{E}_t(\log \hat{f}_t(\mathbf{Y}_{t+1}))$, the closer the predictive density is to the true conditional density. In practice p_t , and hence $\mathbb{E}_t(\log \hat{f}_t(\mathbf{Y}_{t+1}))$, is unknown. However, we can observe the so-called KLIC score $S_{t+1} = \log \hat{f}_t(\mathbf{Y}_{t+1})$, which has the same conditional mean: $\mathbb{E}_t(S_{t+1}) = \mathbb{E}_t(\log \hat{f}_t(\mathbf{Y}_{t+1}))$.

In case the density forecast is obtained from a copula-based model as in (1) or (2), the KLIC score can be decomposed as

$$S_{t+1} = \sum_{j=1}^d \log \hat{f}_{j,t}(Y_{j,t+1}) + \log \hat{c}_t(\hat{\mathbf{U}}_{t+1}), \quad (7)$$

where $\hat{c}_t$ is the conditional copula density associated with the density forecast, and $\hat{\mathbf{U}}_{t+1} = (\hat{F}_{1,t}(Y_{1,t+1}), \dots, \hat{F}_{d,t}(Y_{d,t+1}))'$ its multivariate conditional PIT. The terms $\log \hat{f}_{j,t}(Y_{j,t+1})$, $j = 1, \dots, d$, are KLIC scores of the conditional marginals, while the term $\log \hat{c}_t(\hat{\mathbf{U}}_{t+1})$, is a score assigned to the copula based on the obtained PIT $\hat{\mathbf{U}}_{t+1}$.

Our aim is to develop a formal test for equal predictive ability of two competing copula-based multivariate density models A and B , say, based on their relative KLIC scores S_{t+1}^A and S_{t+1}^B , respectively. We assume that the two competing multivariate density forecasts differ only in their copula specifications and have identical predictive marginal densities $\hat{f}_{j,t}$, $j = 1, \dots, d$ (but we return to this point below). The two competing copula specifications are assumed to have well-defined densities $\hat{c}_{A,t}$ and $\hat{c}_{B,t}$. The null hypothesis we consider here is that for each t the two models have equal predictive ability on average: $H_0 : \mathbb{E}(S_{t+1}^A) = \mathbb{E}(\mathbb{E}_t(S_{t+1}^A)) = \mathbb{E}(\mathbb{E}_t(S_{t+1}^B)) = \mathbb{E}(S_{t+1}^B)$, $t = 1, 2, \dots$. Since the conditional marginals are identically specified under both density forecasts, the parts of the scores in (7) that correspond to the marginals are identical and cancel out, so that an equivalent formulation of the null hypothesis is

$$H_0 : \mathbb{E}(\log \hat{c}_{A,t}(\hat{\mathbf{U}}_{t+1})) = \mathbb{E}(\log \hat{c}_{B,t}(\hat{\mathbf{U}}_{t+1})), \quad t = 1, 2, \dots \quad (8)$$

A formal test of equal predictive ability can be based on an observed sequence of score differences $d_{t+1} = S_{t+1}^A - S_{t+1}^B = \log \hat{c}_{A,t}(\hat{\mathbf{U}}_{t+1}) - \log \hat{c}_{B,t}(\hat{\mathbf{U}}_{t+1})$, for the null hypothesis $\mathbb{E}(d_{t+1}) = 0$ corresponding to (8). Assume that any parameters in the multivariate conditional density (both in the marginal distributions and in the copula function) are estimated using a moving window of fixed length R , and that P observations of d_{t+1} are available for $t+1 = R+1, \dots, T = R+P$. A test of equal KLIC scores can then be based on the test statistic

$$Q_{R,P} = \sqrt{P} \frac{\bar{d}_{R,P}}{\hat{\sigma}_{R,P}}, \quad (9)$$

where $\bar{d}_{R,P} = \frac{1}{P} \sum_{t=R}^{T-1} d_{t+1}$ is the sample mean of the KLIC score difference, and $\hat{\sigma}_{R,P}$ is a heteroskedasticity and autocovariance consistent (HAC) estimate of the asymptotic standard deviation of $\sqrt{P}\bar{d}_{R,P}$. Under the assumption of an in-sample estimation window of fixed length R and appropriate mixing conditions on $\{d_{t+1}\}$, Theorem 4 of Giacomini and White (2006) applies, which states that $Q_{R,P}$ is asymptotically standard normally

distributed. The sign of $Q_{R,P}$ obviously indicates which of the two copula specifications c_A and c_B performs better, as a higher average copula score is preferred.

The test statistic as derived above in fact is a special case of the predictive ability test of Giacomini and White (2006). Their framework is more generally applicable and, for example, allows for dynamics in the predictive distributions, and hence for time varying copulas. However, for clarity of presentation we focus on the case of two competing parametric copula families, only allowing for dynamic updating of the unknown parameters. Also, focus in this paper is on the unconditional version of Giacomini and White's test, but it is possible to extend the approach to the case of testing for equal conditional predictive accuracy. Additionally, the test statistic (9) may also be applied for comparing two copula-based multivariate density forecasts that also (or even only) have different predictive marginal densities $\hat{f}_{j,t}$, $j = 1, \dots, d$. In that case, however, the evaluation not only concerns the specification of the dependence structure, as captured by the copula function, but also the models applied for the marginal characteristics.

Note that the assumption that the two competing multivariate density forecasts differ only in their copula specifications and have identical predictive marginal densities implicitly requires that the parameters in the marginals and the copula can be separated from each other, so that they can be estimated in a multi-stage procedure such as the IFM method discussed in Section 2. Apart from this, no other restrictions are put on the marginals. In particular, they may be specified parametrically, nonparametrically, or semi-parametrically as in the SCOMDY model in (5). An advantage of using nonparametric marginal distributions is that it allows one to test for equal predictive accuracy of copula specifications without the possible side-effects of misspecification of the marginal densities.

We conclude this section with a few remarks concerning the practical implementation of the test statistic in (9). First, a standard HAC-estimator can be used for the asymptotic variance of $\sqrt{P}\bar{d}_{R,P}$, for example $\hat{\sigma}^2 = \hat{\gamma}_0 + 2 \sum_{k=1}^{K-1} a_k \hat{\gamma}_k$, 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 $K = \lfloor P^{-1/4} \rfloor$.

Second, if the conditional marginal densities are modeled parametrically, the corre-

sponding PITs $\hat{U}_{t+1} = \left(\hat{F}_{1,t}(Y_{1,t+1}), \dots, \hat{F}_{d,t}(Y_{d,t+1}) \right)'$ can be evaluated directly, or, if no closed-form expression is available for the CDFs, by numerical integration. If the marginal densities are nonparametric, as may be the case for semi-parametric models, the PITs can be estimated nonparametrically. We illustrate this point by describing how the PITs $\hat{U}_t$ are obtained for the SCOMDY model in (5). Let $\hat{\theta}_t$ denote a point estimate of the parameter vector $\theta = (\theta'_1, \theta'_2)'$ characterizing the conditional means and variances, obtained at time t using the R observations $\mathbf{Y}_{t-R+1}, \dots, \mathbf{Y}_t$. These estimates can be used to compute the sequence of in-sample residuals $\{\hat{\varepsilon}_s\}_{s=t-R+1}^t$ as

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

In addition, the one-step ahead forecast error can be obtained as $\hat{\varepsilon}_{j,t+1|t} = (Y_{j,t+1} - \mu_{j,t+1}(\hat{\theta}_{1,t})) / \sqrt{h_{j,t+1}(\hat{\theta}_t)}$. The PIT $\hat{U}_{t+1}$ associated with the forecast errors $\hat{\varepsilon}_{t+1|t}$ is then given by $\hat{U}_{t+1} = (\hat{U}_{1,t+1}, \dots, \hat{U}_{d,t+1})'$ where

$$\hat{U}_{j,t+1} = \frac{\text{rank of } \hat{\varepsilon}_{j,t+1|t} \text{ among } \{\hat{\varepsilon}_{j,t-R+1}, \dots, \hat{\varepsilon}_{j,t}, \hat{\varepsilon}_{j,t+1|t}\}}{R+2}. \quad (10)$$

The division by $R+2$ rather than by the estimation sample size $R+1$ ensures that the estimated PITs do not fall at the boundaries of the unit cube, where some copula densities diverge.

4 Monte Carlo simulations

In this section we examine the finite-sample behavior of our predictive accuracy test for comparing alternative copula specifications. We conduct our simulation experiments in the context of the SCOMDY model (5), for dimensions $d = 2, 3$ and 5 . For all data generating processes (DGPs) we use an AR(1) specification for the conditional means and a GARCH(1,1) specification for the conditional variances with coefficients that are typical for daily exchange rates, in particular

$$Y_{j,t} = 0.1Y_{j,t-1} + \sqrt{h_{j,t}}\varepsilon_{j,t} \quad (11)$$

$$h_{j,t} = 0.1 + 0.05(Y_{j,t-1} - 0.1Y_{j,t-2})^2 + 0.85h_{j,t-1}, \quad (12)$$

for $j = 1, \dots, d$. The innovations $\varepsilon_{j,t}$ are i.i.d. and drawn from the univariate Student- t distributions with 5 degrees of freedom, which are standardized to have variance equal to 1.²

In the size and power experiments reported below, we use the Gaussian copula, the Student- t copula, the Clayton copula, the survival Clayton copula, and a mixture of the latter two.

The Gaussian and Student- t copulas can be obtained using the so-called inversion method, that is

$$C^G(u_1, u_2, \dots, u_d) = F(F_1^{-1}(u_1), F_2^{-1}(u_2), \dots, F_d^{-1}(u_d)), \quad (13)$$

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 (13) 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

$$c^G(\mathbf{u}; \Sigma) = |\Sigma|^{-1/2} \exp \left(-\frac{1}{2} (\Phi^{-1}(\mathbf{u}))' (\Sigma^{-1} - \mathbf{I}_d) \Phi^{-1}(\mathbf{u}) \right), \quad (14)$$

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 $\Phi^{-1}(\cdot)$ denoting the inverse of the standard normal CDF. In the bivariate case $d = 2$, the correlation coefficient $\rho_{12} = \rho_{21}$ is the only parameter of the Gaussian copula.

The Student- t copula is obtained similarly, but using a multivariate Student- t distribution instead of the Gaussian. The corresponding copula density is given by

$$c^t(\mathbf{u}, \Sigma, \nu) = |\Sigma|^{-1/2} \frac{\Gamma([\nu + d]/2) \Gamma^{d-1}(\nu/2)}{\Gamma^d((\nu + 1)/2)} \frac{\left(1 + \frac{T_\nu^{-1}(\mathbf{u})' \Sigma^{-1} T_\nu^{-1}(\mathbf{u})}{\nu} \right)^{-(\nu+d)/2}}{\prod_{i=1}^d \left(1 + \frac{(T_\nu^{-1}(u_i))^2}{\nu} \right)^{-(\nu+1)/2}}, \quad (15)$$

²We repeat all experiments reported in this section for DGPs with constant conditional means and variances. The results of these experiments are qualitatively and quantitatively similar to results reported here for the case of time-varying 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 $T_\nu^{-1}(\cdot)$ is the inverse of the univariate Student- t CDF, Σ is the correlation matrix and ν is the number of degrees of freedom. In the bivariate case the Student- t copula has two parameters, the number of degrees of freedom ν and the correlation coefficient ρ_{12} . Note that the Student- t copula nests the Gaussian copula when $\nu = \infty$.

A major difference between the Gaussian copula and the Student- t copula is their ability to capture tail dependence, which may be important for financial applications. As mentioned in Section 2, for the Gaussian copula both tail dependence coefficients are equal to zero, while for the Student- t copula the tail dependence is symmetric and positive. Specifically, in the bivariate case $d = 2$, the tail dependence coefficients are given by

$$\lambda_L = \lambda_U = 2T_{\nu+1} \left(-\sqrt{(\nu+1)(1-\rho_{12})/(1+\rho_{12})} \right),$$

which is increasing in the correlation coefficient ρ_{12} and decreasing in the degrees of freedom ν .

The Clayton and survival Clayton copulas belong to the family of Archimedean copulas (see Nelsen (2006) for details). The d -dimensional Clayton copula is given by

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

In contrast to the Gaussian and Student- t copulas, the Clayton copula is able to capture asymmetric tail dependence. In fact, it only exhibits 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

$$c^C(\mathbf{u}, \alpha) = \left(\prod_{j=1}^d (1 + (j-1)\alpha) \right) \left(\prod_{j=1}^d u_j^{-(\alpha+1)} \right) \left(\sum_{j=1}^d u_j^{-\alpha} - d + 1 \right)^{-(\alpha^{-1}+d)}.$$

The survival Clayton copula is obtained as a mirror image of the Clayton copula, with its density function given by

$$\tilde{c}^C(\mathbf{u}, \alpha) = c^C(\mathbf{1} - \mathbf{u}, \alpha).$$

Consequently, in the bivariate case the upper tail dependence coefficient for the survival Clayton copula is $\lambda_U = 2^{-1/\alpha}$, and is increasing in the parameter α , while the lower tail dependence coefficient is zero.

Finite mixture copulas are constructed by using a convex combination of several parametric copulas. Here we use a mixture of the Clayton and survival Clayton copula, that is

$$C^{\text{mix}}(\mathbf{u}, \alpha_1, \alpha_2, w) = wC^C(\mathbf{u}, \alpha_1) + (1 - w)\tilde{C}^C(\mathbf{u}, \alpha_2), \quad (16)$$

where $0 \leq w \leq 1$ is the mixture weight for the Clayton copula. The mixture copula inherits the tail dependence properties of its basis copulas. Thus, the mixture of the Clayton and the survival Clayton copulas will exhibit both upper and lower tail dependence. The strength of the dependence is determined by the parameters of both copulas and the mixture weight, w .

We evaluate the properties of the copula predictive accuracy test for different combinations of the length of the moving in-sample window $R \in \{100, 1000\}$ to estimate the parameters in the marginal models given by (11) and (12), and the number of out-of-sample evaluations $P \in \{100, 1000, 5000\}$. The in-sample PITs $\hat{U}_s$, $s = t - R + 1, \dots, t$ are obtained from the empirical CDF transformation $\hat{U}_{j,s} = \frac{R_{j,s}}{R+1}$ 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 the competing copulas then are estimated using maximum likelihood. Finally, the PIT $\hat{U}_{t+1}$ corresponding to the one-step ahead forecast errors $\hat{\varepsilon}_{j,t+1|t}$ is obtained from (10). The number of replications in each experiment is set equal to 1000.

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 by taking the innovations $\varepsilon_{j,t}$, $j = 1, \dots, d$ in the DGP to be independent. We then test the null hypothesis of equal predictive accuracy of two Student- t copulas with the correlation parameter fixed at either $\rho_{12} = 0.3$ or $\rho_{12} = -0.3$. For the copulas with dimension $d > 2$, we set all other correlation coefficients to 0. This is done to assure equal distance from the independence (product) copula. 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.

[Table 1 about here.]

The observed rejection rates of the test for nominal sizes 0.01, 0.05 and 0.10 are reported in Table 1, where the null hypothesis is tested against the two-sided alternative that the average scores of the two copulas are not equal. Panel I in the table corresponds to case (i) where the degrees of freedom parameter ν is fixed at 5, while Panel II corresponds to case (ii) with ν being estimated. In the absence of any parameter estimation uncertainty in the copula parameters, we observe that the empirical rejection rate is close to the nominal size when the number of out-of-sample evaluations is relatively large, that is, when $P = 1000$ or 5000 . The test is somewhat liberal and rejects the correct null hypothesis too frequently when $P = 100$ and $R = 100$. Adding parameter uncertainty due to the estimation of the copula parameter ν results in larger deviations from the nominal size in the $P = 100, R = 100$ case. For other P/R combinations the results are similar. We do not observe systematic differences in size for different copula dimensions d .

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 approaching the correct specification as we vary one of the parameters of the DGP. Firstly, we consider the Gaussian copula and the Student- t copula. Hence, we focus on the question whether the $Q_{R,P}$ test statistic can distinguish between copulas with and without tail dependence.

[Figure 1 about here.]

The DGPs are specified with Student- t copulas of dimensions $d = 2, d = 3$ and $d = 5$ with all correlation coefficients set to $\rho = 0.3$ and varying values of the degrees of freedom parameter ν in the interval $[5, 50]$. We test the t copula specification (with both parameters Σ and ν estimated) against the Gaussian copula specification (with the correlation matrix

Σ estimated). The Student- t copula approaches the Gaussian copula as ν increases, and the tail dependence disappears with the coefficients λ_L and λ_U converging to zero. Intuitively, 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 1 in the form of a power plot, showing the observed rejection rates for nominal size of 0.05 as a function of the degrees of freedom parameter, ν , in the DGP. Different panels relate to different P/R combinations. The results displayed are for the null hypothesis that the Gaussian copula and Student- t copula perform equally well, against the one-sided alternative hypothesis that the correctly specified Student- t copula has a higher average score. Intuitively, one might expect the Student- t copula to perform better, since the true DGP uses the Student- t copula. Note, however, that as the number of degrees of freedom parameter ν in the copula of the DGP becomes large, the Gaussian copula might outperform the Student- t copula because the Gaussian copula is almost equivalent to the Student- t copula for large ν , but requires one parameter (ν) less to be estimated. Indeed, it can be observed that the rejection rates become smaller than the nominal size for large values of ν . This does not indicate that the test is flawed, but is merely illustrating the fact that the finite sample performance of more complex (in terms of parameters) copulas can be expected to be affected more heavily by parameter estimation uncertainty.

Figure 1 shows that the test has higher power for smaller values of ν . Increasing the dimension d of the copula improves the power for all values of ν and all P/R combinations considered. The power of the test reaches values close to one for $P = 5000$, $R = 1000$, even when the Student- t copula is relatively close to the Gaussian copula ($\nu = 20$). By comparing the panels it can be seen that the power is increasing in P as well as R . The region of ν for which the power exceeds the nominal size is smallest for $P = 100$ and $R = 100$, in which case the power exceeds the nominal size only for $\nu \leq 7$. For $P = 1000$ and $R = 100$ the situation is similar. We attribute this to the fact that a smaller in-sample size R is characterized by larger parameter estimation uncertainty. Apparently for small in-sample values and relatively large values of ν the less parsimonious Student- t copula is affected by parameter estimation uncertainty more than the parsimonious Gaussian cop-

ula, even to the extent that the latter performs better on average. The ν -regions with power exceeding the nominal size are larger for higher dimensions d , which may be explained by the fact that the ratio of the number of parameters to be estimated for the Student- t copula relative to the Gaussian copula becomes smaller.

[Figure 2 about here.]

In our final experiment we consider a DGP with varying asymmetric tail dependence by using a SCOMDY model with the mixture copula (16) based on Clayton copula with coefficient $\alpha_1 = 1.5$ and a survival Clayton copula with $\alpha_2 = 1.5$ and weight w varying within the interval $[0.5, 1]$. The conditional mean and variance are specified as in (11) and (12) with standard normal innovations ε_t . The Clayton copula specification is tested against the survival Clayton copula in a one-sided test. We examine whether the test can distinguish between the upper tail dependence and the lower tail dependence copulas. Since the survival Clayton copula density is a point reflection of the Clayton copula density, the parameter estimation uncertainty is identical for both copulas. When $w = 0.5$ both copulas are equally distant from the correct mixture copula. As w approaches unity, the Clayton copula becomes closer to the DGP. We use mixture copulas of dimensions $d = 2$, $d = 3$ and $d = 5$ in this simulation experiment. The parameter α is estimated for both copulas. Figure 2 shows power curves for this experiment for the different values of the mixture weight, w , three different dimensions d , and different combinations of P and R . When $w = 0.5$ the rejection rates are close to the nominal size of 0.05 for all P/R combinations except for $P = 100, R = 100$, in which case the test over-rejects. In all cases the power of the test increases with w . For $P = 5000$ and $R = 1000$ the power of the test quickly reaches unity, while the convergence to unity is slowest for $P = 100$ and $R = 100$. The power of the test increases with the dimension of the mixture copula for all P/R combinations considered.

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 to exchange rate returns for several major currencies. Specifically, 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 noon buying rates in New York and are obtained from the Federal Reserve Bank of 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.³

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

$$Y_{j,t} = c_j + \sum_{l=1}^5 \phi_{j,l} Y_{j,t-l} + \sqrt{h_{j,t}} \varepsilon_{j,t} \quad (17)$$

$$h_{j,t} = \kappa_j + \gamma_j \left(Y_{j,t-1} - c_j - \sum_{l=1}^5 \phi_{j,l} Y_{j,t-1-l} \right)^2 + \beta_j h_{j,t-1}, \quad (18)$$

where $\kappa_j > 0$, $\beta_j \geq 0$, $\gamma_j > 0$ and $\beta_j + \gamma_j < 1$, for $j = 1, \dots, 3$.

The joint distribution of the standardized innovations $\varepsilon_{j,t}$ is specified semi-parametrically, combining nonparametric univariate marginal distributions F_j with a parametric copula C . We consider four alternative copula specifications, which we compare in terms of their relative performance in out-of-sample density forecasting. In particular, we consider the Gaussian copula, the Student- t copula, the Clayton copula and the Gumbel copula. The Gaussian, Student- t and Clayton copula were introduced earlier in Section 4. The Gum-

³The unconditional correlations between the exchange rates of the GBP and CHF, GBP and EUR, and CHF and EUR equal 0.69, 0.72 and 0.92, respectively. The unconditional correlations between exchange rates of the CAD and JPY, CAD and EUR, and JPY and EUR equal 0.11, 0.22 and 0.54, respectively.

bel copula is another copula from the family of Archimedean copulas popular in financial time-series applications and is specified as

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

The Gumbel copula 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 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 both the Clayton copula and the Gumbel copula impose equi-dependence between all pairs of variables 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 rolling 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 the dependence between the EUR and JPY exchange rates after the introduction of the euro on January 1, 1999. We therefore partition our evaluation period at this date into two sub-periods, called the pre-euro and post-euro period, respectively, to explore whether the (relative) forecasting performance of the different copula specifications also is affected by the euro introduction. The number of observations differs slightly for the pre- and post-euro periods, being equal to $P = 2772$ and $P = 2388$, respectively.

[Table 2 about here.]

[Table 3 about here.]

Table 2 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- t copula in all cases except for the GBP-CHF-EUR group during the post-euro subperiod, where again there is insufficient evidence to reject the null of equal performance of these two copulas. For comparison, Table 3 presents the average values of the KLIC scores S_{t+1} given in (7). Thus, overall we conclude that the Student- 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.

[Figure 3 about here.]

To provide more insight into these results, in Figure 3 we show the rolling window parameter estimates of the Student- 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 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 capture these differences. This explains the inferior forecasting performance of these two copulas relative to the Gaussian and t specifications. Note that it would be too premature, however, to conclude that the exchange rates ex-

hibit symmetric rather than asymmetric tail dependence. It is possible that there is indeed asymmetric tail dependence, but that the restrictive nature of the Clayton and Gumbel copulas (which impose equi-dependence between all exchange rate pairs) leads to worse performance than the more flexible (but symmetric) Student- t copula. To truly conclude that these exchange rates have symmetric tail dependence we would need to consider a copula which is as flexible as the Student- t copula, and which allows for asymmetric tail dependence.

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 increasing correlations for CAD-EUR, CAD-JPY and JPY-EUR, while they were 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 used to maximize the likelihood, and results in somewhat unstable estimates around 2006-2008. It is also more difficult to distinguish the Gaussian copula from the Student- t copula when the correlation is close to one, as the Gaussian copula also has tail dependence in the limiting case of unit correlation.⁴ All these factors contribute to the inability of the test to reject the equal performance between the Gaussian copula and the Student- t during this subperiod.

6 Summary and discussion

In this paper we have introduced a new statistical test for comparing alternative copula specifications in the context of multivariate density forecasts. The test is based on the 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 copulas pair-wise; the copula with the smallest average KLIC is considered to be superior. The main difference with other tests for competing copula specifications is that we use an

⁴The correlation estimates obtained for the Gaussian copula specifications are very similar to the values shown in Figure 3 for the Student- t copula.

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 parameter estimation as an integral part of the forecasting method. The resulting test is valid under general conditions on the competing copulas, allowing for parameter estimation uncertainty and for comparing either nested or non-nested copula families. In addition, the conditional marginal densities, which are assumed to be the same for the two competing copula specifications, may be specified parametrically, nonparametrically, or semi-parametrically. This implies that the test may be applied to a wide variety of copula-based multivariate density models.

Monte Carlo simulation results suggest that the proposed test has satisfactory size and power properties in finite samples. Rejection rates smaller than the nominal size are observed, but this phenomenon can be understood in terms of parameter estimation uncertainty and model complexity; an incorrectly specified but more parsimonious model may perform better than a correctly specified model, when the difference in dependence structure described by these models is small.

In an empirical application to daily exchange rate returns of several major currencies against the US dollar we found that the Student- t copula is clearly favored over Gaussian, Gumbel and Clayton specifications.

In this paper, we have limited ourselves to copula specifications for which the density exists and is well-defined. This is the case, for example, when the copula function is absolutely continuous. For copulas that are not, it may also be possible to develop Kullback-Leibler type scoring rules, as these merely correspond to likelihood ratio scores. If the likelihood ratio for two competing copulas can be calculated, so can the corresponding score. We leave it for future research to investigate how the likelihood ratio scores perform for such copula specifications.

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. (2009). Copula goodness-of-fit testing: An overview and power comparison. *European Journal of Finance*, **15**, 675–701.
- 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**, 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.
- Chen, X., Fan, Y. and Tsyrennikov, V. (2006). Efficient estimation of semiparametric multivariate copula models. *Journal of the American Statistical Association*, **101**, 1228–1240.
- Cherubini, U., Luciano, E. and Vecchiato, E. (2004). *Copula Methods in Finance*. Wiley Finance Series. Wiley, Chicester.
- Coval, J.D., Jurek, J.W. and Stafford, E. (2009). Economic catastrophe bonds. *American Economic Review*, **99**, 628–666.
- 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.

- Dimitrova, D. S., Kaishev, V. K. and Penev, S. I. (2008). GeD spline estimation of multivariate Archimedean copulas. *Computational Statistics & Data Analysis*, **52**, 3570–3582.
- Fermanian, J.-D. (2005). Goodness-of-fit tests for copulas. *Journal of Multivariate Analysis*, **95**, 119–152.
- Fermanian, J.-D. and Scaillet, O. (2003). Nonparametric estimation of copulas for time series. *Journal of Risk*, **5**, 25–54.
- Fermanian, J.-D. and Scaillet, O. (2005). Some statistical pitfalls in copula modeling for financial applications. In *Capital Formation, Governance and Banking* (ed. E. Klein), pp. 57–72. Nova Science Publishers, Hauppauge NY.
- Genest, C., Gendron, M. and Bourdeau-Brien, M. (2009). The advent of copulas in finance. *European Journal of Finance*, **15**, 609–618.
- Genest, C., Quessy, J.-F. and Rémillard, 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. (2008). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Annales de l'Institut Henri Poincaré - Probabilités et Statistiques*, **44**, 1096–1127.
- Genest, C., Rémillard, B. and Beaudoin, D. (2009). Omnibus goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, **44**, 199–213.
- Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, **74**, 1545–1578.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.

- Jondeau, E. and Rockinger, M. (2006). The copula-GARCH model of conditional dependencies: An international stock market application. *Journal of International Money and Finance*, **25**, 827–853.
- 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. Springer Verlag, New York.
- Okimoto, T. (2008). New evidence of asymmetric dependence structures in international equity markets. *Journal of Financial and Quantitative Analysis*, **43**, 787–815.
- 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. (2009). 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), pp. 767–786. Springer-Verlag, Berlin.
- Rivers, D. and Vuong, Q. (2002). Model selection tests for non-linear dynamic models. *Econometrics Journal*, **5**, 1–39.
- Rodriguez, J. C. (2007). Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, **14**, 401–423.
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Annals of Mathematical Statistics*, **23**, 470–472.

- Sancetta, A. and Satchell, S.E. (2004). The Bernstein copula and its applications to modeling and approximations of multivariate distributions. *Econometric Theory*, **20**, 535–562.
- Shen, X., Zhu, Y. and Song, L. (2008). Linear B-spline copulas with applications to nonparametric estimation of copulas. *Computational Statistics & Data Analysis*, **52**, 3806–3819.
- Sklar, A. (1959). Fonction de répartitions 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.

Figure 1: Power of the test. The figure displays observed rejection rates (on the vertical axis) of a one-sided test of equal performance of the Gaussian and Student- t copula, against the alternative hypothesis that the correctly specified Student- t copula has higher average score. The horizontal axis displays the degrees of freedom parameter of the Student- t copula in the DGP. The DGP is the SCOMDY model (5) with marginal distributions specified in (17)-(18) and a Student- t copula with dimensions $d = 2$, $d = 3$ and $d = 5$, all correlation coefficients set to $\rho = 0.3$ and varying degrees of freedom ν . The test of equal predictive accuracy compares a t copula with both parameters Σ and ν estimated against a Gaussian copula with parameter Σ estimated. The four panels correspond to different P/R combinations. R denotes the number of observations in the moving in-sample window and P denotes the number of out-of-sample evaluations. The nominal size of 0.05 is indicated by the horizontal line. Results are based on 1000 replications.

Figure 2: Power of the test. The figure displays power curves, showing the rejection rate of a one-sided test against the alternative hypothesis that the correctly specified copula has higher average score (on the vertical axis) for the varying mixture copula weights (on the horizontal axis). The DGP is the SCOMDY model (5) marginal distributions specified in (17)-(18) and a mixture Clayton-survival Clayton copula of dimensions $d = 2$, $d = 3$ and $d = 5$ with coefficients $\alpha_1 = \alpha_2 = 1.5$ and the varying weight w . The test of equal predictive accuracy compares the Clayton copula against the survival Clayton copula when the parameters in both copulas are estimated. R denotes the number of observations in the moving in-sample window and P denotes the number of out-of-sample evaluations. Results are based on 1000 replications.

Figure 3: Parameter estimates of the Student- 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 first three parameters presented in the graphs are pairwise correlations between the indicated exchange rates, while the fourth is the degrees of freedom parameter. The parameter estimates are reported for every 100th day in the evaluation period. Year labels (on the horizontal axis) indicate the beginning of the corresponding year.

Table 1: Observed size of the test for equal predictive accuracy

Nominal size	P/R			
	100/100	1000/100	1000/1000	5000/1000
Panel I: ν fixed				
a) $d = 2$				
0.01	0.014	0.009	0.010	0.012
0.05	0.062	0.045	0.060	0.056
0.10	0.121	0.093	0.115	0.111
b) $d = 3$				
0.01	0.014	0.007	0.008	0.014
0.05	0.049	0.039	0.049	0.046
0.10	0.103	0.079	0.099	0.091
c) $d = 5$				
0.01	0.012	0.013	0.012	0.007
0.05	0.055	0.052	0.044	0.037
0.10	0.105	0.114	0.099	0.097
Panel II: ν estimated				
a) $d = 2$				
0.01	0.017	0.012	0.014	0.009
0.05	0.069	0.054	0.059	0.055
0.10	0.120	0.104	0.109	0.099
b) $d = 3$				
0.01	0.015	0.008	0.008	0.013
0.05	0.064	0.045	0.044	0.047
0.10	0.109	0.115	0.097	0.093
c) $d = 5$				
0.01	0.020	0.009	0.008	0.012
0.05	0.063	0.054	0.042	0.039
0.10	0.125	0.105	0.096	0.084

Note: The table presents two-sided rejection rates of the test for equal predictive accuracy of two competing copula specifications as given in (9) for nominal sizes 0.01, 0.05 and 0.10 and dimensions $d = 2$, $d = 3$ and $d = 5$. The DGP is the SCOMDY model (5) with marginal distributions following AR(1)-GARCH(1,1) process with Student- t innovations with 5 degrees of freedom. The innovations for the different series are independent. The test of equal predictive accuracy compares a Student- t copula with $\rho_{12} = 0.3$ against a t copula with $\rho_{12} = -0.3$. For dimension $d > 2$, the remaining correlation coefficients are set to 0. The results are given for two cases: (i) ν is fixed at 5 (Panel I) and (ii) ν is estimated (Panel II). R denotes the number of observations in the moving in-sample estimation window and P denotes the number of out-of-sample evaluations. Results are based on 1000 replications.

Table 2: *Pair-wise tests for out-of-sample performance of copulas*

Copula A	Copula B		
	Gumbel	Student- t	Gaussian
<u>Panel I: GBP-CHF-EUR</u>			
a) Full sample			
Clayton	−5.91***	−23.16***	−23.53***
Gumbel		−19.61***	−19.73***
Student- t			2.46**
b) Pre-euro			
Clayton	−5.34***	−17.33***	−18.26***
Gumbel		−14.50***	−14.00***
Student- t			3.17***
c) Post-euro			
Clayton	−2.71***	−13.13***	−12.69***
Gumbel		−11.69***	−12.12***
Student- t			0.52
<u>Panel II: CAD-JPY-EUR</u>			
a) Full sample			
Clayton	−1.41	−8.95***	−7.42***
Gumbel		−8.11***	−6.65***
Student- t			4.04***
b) Pre-euro			
Clayton	−1.06	−8.37***	−7.62***
Gumbel		−7.61***	−6.91***
Student- t			1.82*
c) Post-euro			
Clayton	−0.94	−4.00***	−1.90*
Gumbel		−3.35***	−1.28
Student- t			4.44***

Note: The table presents 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.

Table 3: Average out-of-sample KLIC scores for four copulas

Sample	Clayton	Gumbel	Student- t	Gaussian
<u>Panel I: GBP-CHF-EUR</u>				
Full sample	0.775	0.889	1.244	1.217
Pre-Euro	0.766	0.923	1.270	1.229
Post-Euro	0.785	0.849	1.214	1.204
<u>Panel II: CAD-JPY-EUR</u>				
Full sample	0.055	0.062	0.147	0.134
Pre-Euro	0.041	0.050	0.183	0.174
Post-Euro	0.071	0.077	0.106	0.088

Note: The table presents average values of out-of-sample KLIC scores for four considered copulas. 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.