INSTITUT DE STATISTIQUE

UNIVERSITÉ CATHOLIQUE DE LOUVAIN

DISCUSSION PAPER

0917

TIME-VARYING COPULAS: A SURVEY

MANNER H. and O. REZNIKOVA

Time-varying copulas: a survey

Hans Manner[1] Olga Reznikova[2]

September 1, 2009

Abstract

The aim of this paper is to bring together different specifications for copula models with time-varying dependence structure. Copula models are widely used now in financial econometrics and risk management. They are considered to be a competitive alternative to the Gaussian dependence structure. The dynamic structure of the dependence between the data can be modeled by allowing either the copula function or the dependence parameter to be time-varying. First, we give a brief description of eight different models, among which there are fully parametric, semiparametric and adaptive methods. The purpose of this study is to compare the applicability of each particular model in different cases. We conduct a simulation study to show the performance of model selection and goodness-of-fit measures in terms of size and power for different setups and the ability of the models to estimate the (latent) time-varying dependence parameter. Finally, we provide an illustration by applying the competing models on the same financial dataset and compare their performance by means of Value-at-Risk.

Keywords: Dynamic copula, Goodness-of-Fit test, Time-varying parameter

JEL Classification: C14, C22.

[1]Department of Quantitative Economics, Maastricht University, PoBOX 616, MD 6200, Maastricht, The Netherlands, h.manner@maastrichtuniversity.nl

[2]Institut de statistique, Université catholique de Louvain, Voie du Roman Pays 20, B-1348 Louvain-la-Neuve, Belgium, Olga.Reznikova@uclouvain.be

1 Introduction

It is well accepted that the hypothesis of (multivariate) normality is one that is usually not supported by the data for many types of variables. This has created the need to construct flexible, non-standard multivariate distributions and this task can easily be solved using a class of functions known as copulas (Sklar 1959). Any multivariate distribution function can be decomposed into the marginal distributions that describe the individual behavior of each series and the copula that fully captures the dependence between the variables. Furthermore, given a set of marginal distributions and a copula a multivariate distribution can be constructed by coupling the marginals with the copula. The flexibility of the way dependencies can be modeled independently of the marginal distributions has made copulas particularly popular for financial applications. The most important fields of applications are pricing CDO's (Li 2000), calculating the Value-at-Risk of a portfolio (Embrechts et al. 2003, Giacomini et al. 2009), the pricing of options with multiple underlying assets (van den Goorbergh et al. 2005) or portfolio construction (Patton 2004). Textbook treatments of the theory of copulas are given in Joe (1997) and Nelsen (2006). The book Cherubini et al. (2004) deals entirely with various applications of copulas in finance.

Most of the time copulas are applied to financial time series data, but often they are treated to be constant over time. However, it has become a stylized fact that correlations between asset returns are not constant through time, a finding that has been documented by, among many others, Erb et al. (1994), Longin and Solnik (1995) or Engle (2002). Some notable parametric models to model these time-varying correlations in multivariate volatility models are the DCC GARCH model, simultaneously proposed by Engle (2002) and Tse and Tsui (2002), a stochastic volatility model with stochastic correlations by Yu and Meyer (2006) and the regime switching model for dynamic correlations by Pelletier (2006). Hafner et al. (2006) propose a semiparametric model for correlation dynamics.

Even though copulas allow for more general dependence structures than simple linear correlation it seems unrealistic to treat dependence as constant, given that correlations have been found to be time-varying. To our knowledge the first papers allowing copulas to be time-varying were Patton (2006), who extended Sklar's theorem for conditional distributions and proposed a parametric model to describe the evolution of the copula parameter, and Dias and Embrechts (2004) who proposed a test for structural breaks in the copula parameter. Subsequently, a large number of studies has dealt with the application of time-varying copulas and the development of new models and tests to appropriately model time-varying dependencies. Some contributions to this fast growing field of research are van den Goorbergh et al. (2005), Jondeau and Rockinger (2006), Giacomini et al. (2009), Guégan and Zhang (2009), Chollete et al. (2008), Creal et al. (2008), Hafner and Manner (2008) and Hafner and Reznikova (2008).

In this paper we want to offer a survey over the existing models for time-varying copulas by focusing on the specification, estimation and properties of a number of models. Furthermore, we discuss how the best fitting time-varying copula can be chosen among a number of competing ones and how the goodness-of-fit of a candidate model can be tested. A Monte Carlo study compares the performance of the model selection and goodness-of-fit criteria for competing specifications of dynamics of the copula parameter, and shows how well the competing time-varying copula models are able to estimate the (latent) dependence process. In an empirical application alternative models are estimated for two financial data sets and in addition to statistical model selection the ability of the models to correctly estimate the Value-at-Risk is tested.

The rest of the paper is organized as follows. In Section 2 copulas and their estimation are reviewed. Section 3 provides a survey over existing time-varying copula models followed by a simulation study in Section 4. An empirical application is provided in Section 5 and, finally, Section 6 provides conclusions and an outlook to future developments.

2 Copulas

In this section we shortly discuss the basic theory of copulas and some ways to estimate their parameters. For a complete introduction to copulas see Joe (1997).

Lets consider the bivariate stochastic process $\{X_t\}_{t=1}^T$ with $X_t = (X_{1t}, X_{2t})'$. Let $F(X_{1t}, X_{2t})$ be the joint distribution, whereas F_i and f_i will denote the marginal *cdf* and *pdf* respectively for $i = 1, 2$. Then by Sklar's theorem there exists a copula function $C(\cdot, \cdot) : [0, 1]^2 \to [0, 1]$ mapping the marginal distributions of X_{1t} and X_{2t} to their joint distribution through

$$F(X_{1t}, X_{2t}) = C\left(F_1(X_{1t}), F_2(X_{2t})\right). \tag{1}$$

We assume that the marginals can be modeled parametrically, thus the probability transform is given by $U_{it} = F_i(X_{it}; \phi_i)$, where ϕ_i is the vector of parameters completely describing the individual behavior of the series. $F_i(X_{it}; \phi_i)$ can be a conditional distribution and in financial econometrics X_{it} is usually modeled by an ARMA-GARCH type model, whose residuals are treated as *iid* random variables. We also assume that the copula belongs to a parametric family $C_\theta, \theta \in \Theta \subset \mathbb{R}^K$. Some examples of parametric copulas are given in the appendix.

Given that the copula function and the marginals are absolutely continuous, the following expression for the joint *pdf* holds

$$f(X_{1t}, X_{2t}) = c(U_{1t}, U_{2t}; \theta) \prod_{i=1}^{2} f_i(X_{it}; \phi_i), \tag{2}$$

where $c(\cdot,\cdot)$ denotes the copula density. Assume a sample $X_t,\ t = 1, \ldots, T$. The log-likelihood function is given by

$$\begin{aligned} L(\theta, \phi) &= \sum_{t=1}^{T} \{\log c(U_{1t}, U_{2t}; \theta) + \log f_1(X_{1t}; \phi_1) + \log f_2(X_{2t}; \phi_2)\} && (3) \\ &= L_C(\theta, \phi) + L_{X_1}(\phi_1) + L_{X_2}(\phi_2), && (4) \end{aligned}$$

where $\phi = (\phi_1', \phi_2')'$. Thus, the full log-likelihood function $L(\theta, \phi)$ can be split into two parts, copula likelihood $L_C(\theta, \phi)$ and likelihood of the marginals $L_{X_1}(\phi_1)$ and $L_{X_2}(\phi_2)$.

There are several ways to estimate θ and ϕ. One possible method is to estimate the parameters simultaneously by *full maximum likelihood*

$$(\widehat{\theta}, \widehat{\phi}) = \arg\max_{\theta,\phi} L(\theta, \phi). \tag{5}$$

This estimation method is conceptually straightforward. However, in some situations it may be computationally rather burdensome.

Another approach is to use a two stage estimator. At the first stage only the parameters from the marginal distributions are estimated

$$\widehat{\phi}_i = \arg\max_{\phi} L_{X_i}(\phi_i),\ i = 1, 2. \tag{6}$$

At the second stage the dependence parameter is estimated from the copula likelihood

$$\widehat{\theta} = \arg\max_{\theta} L_C(\theta, \widehat{\phi}). \tag{7}$$

However, the estimation of the parameters in two steps leads to a loss in efficiency and standard errors cannot be obtained as the inverse of the Fisher Information Matrix anymore. By applying one step of the Newton-Rhapson algorithm to the full likelihood function using the two step estimators, statistical efficiency can be achieved (see van der Vaart (1998), Ch.5).

Alternatively when the marginal model is unknown Genest et al. (1995) suggest modeling the marginal distribution with the empirical *cdf* and estimating the copula on the ranks of the data. Again, the problem of loss of efficiency occurs and calculation of the standard errors of the estimated copula parameter is quite tedious. On the other hand, this method is robust to the misspecification of the marginals, which can cause biased estimates of the copula parameter.

3 Survey

In this section we will give an overview of the time-varying copula models that have been proposed in the literature. We focus our attention on the specification of the dynamics of the copula parameter and the estimation of the models. For the sake of brevity a complete description of the properties and many details of the procedures involved must be omitted. The interested reader is referred to the original papers.

Note that the following paragraphs describe only the specification and estimation of the copula, whereas the marginals are assumed to be appropriately modeled and the data is assumed to be transformed into the $U(0,1)$ variables U_{1t} and U_{2t}. In general the time-varying dependence parameter of the copula will be called θ_t, and for the correlation coefficient of the Gaussian copula ρ_t is reserved.

3.1 Observation driven models

Patton (2006) and Creal et al. (2008) propose similar observation driven copula models for which the time-varying dependence parameter of a copula is a parametric function of transformations of the lagged data and an autoregressive term.

The model of Patton for the dynamics of the correlation for Gaussian or Student copula has the following form,

$$\rho_t = \Lambda_1\left(\omega + \beta\Lambda_1^{-1}(\rho_{t-1}) + \alpha\frac{1}{m}\sum_{i=1}^{m}\Phi^{-1}(U_{1,t-i})\Phi^{-1}(U_{2,t-i})\right), \tag{8}$$

$$\Lambda_1(x) = \frac{1-exp(-x)}{1+exp(-x)}, \tag{9}$$

where $\Lambda_1(\cdot)$ is a transformation function which holds the correlation parameter ρ_t in the interval $(-1,1)$, $\Phi(\cdot)$ is the standard normal *cdf* and m is an arbitrary window length. If the data is positively dependent, the inverse of marginal transforms of both variables will have the same sign. Thus, in case of positive dependence the parameter α should be positive.

For the non-Gaussian case Patton suggests modeling the tail dependence parameters (λ^U and λ^L) of the Symmetrized Joe-Clayton (SJC) copula, where λ^U and λ^L are stand-alone monotonic transformations to copula parameters[1]. In general, the model for the evolution of a dependence parameter (or tail dependence) of a copula is

$$\theta_t = \Lambda_2 \left(\omega + \beta \Lambda_2^{-1}(\theta_{t-1}) + \alpha \frac{1}{m} \sum_{j=0}^{m-1} |U_{1,t-j} - U_{2,t-j}| \right), \tag{10}$$

where $\Lambda_2(x)$ is an appropriate transformation function to ensure the parameter always remains in its domain: $(1 + exp(-x))^{-1}$ for tail dependence, $exp(x)$ for Clayton copula and $(exp(x) + 1)$ for Gumbel copula. In case of perfect positive dependence the forcing variable $|U_{1,t} - U_{2,t}|$ is close to zero, therefore the parameter α is expected to be negative.

Creal et al. (2008) developed a unifying framework named Generalized Autoregressive Score (GAS) for time series processes with time varying parameters. A scaled score vector is used as an updating mechanism for the observation driven part of a model. In general, the model GAS(p,q) for a time-varying parameter f_t looks as follows

$$f_t = \omega + \sum_{j=1}^{q} \beta_j f_{t-j} + \sum_{i=0}^{p-1} \alpha_i s_{t-i}, \tag{11}$$

where $s_t = S_{t-1} \cdot \nabla_t$ is the scaled score of the log-likelihood function of the model of interest. ∇_t is the first derivative of the log-likelihood with respect to the parameter, whereas S_{t-1} is the scaling matrix, which is approximated by the inverse of Fisher information matrix. The GAS(1,1) model for the correlation coefficient of the Gaussian copula is

$$f_t = \omega + \beta f_{t-1} + \alpha \frac{2(y_t - \rho_{t-1} - \rho_{t-1}(1 + \rho_{t-1}^2)^{-1}(z_t - 2))}{(1 - \rho_{t-1}^2)}, \tag{12}$$

$$\rho_t = \Lambda_1(f_t), \tag{13}$$

where $y_t = \Phi^{-1}(U_{1t}) \cdot \Phi^{-1}(U_{2t})$ and $z_t = \Phi^{-1}(U_{1t})^2 + \Phi^{-1}(U_{2t})^2$.

Such a specification is more sensitive to the off-diagonal observations than the Patton model and the correlation parameter more rapidly adjusts to the decrease in dependence as illustrated nicely by Creal et al. (2008). The GAS model is also shown to be more

[1] The Joe-Clayton copula is such a transformation of the Clayton copula that possesses upper and lower dependence and it is characterized by two parameters; the SJC allows for the special case of the symmetry in the dependence.

sensitive to observations in the lower and upper tail.

This approach is also applicable to Archimedean copulas and, unlike Patton's model, it can be used for multivariate data. However, the problems with computing $s_t = \nabla_t \mathcal{I}_{t-1}^{-1}$ term might occur. A numerical approximation is suggested for obtaining the Fisher information matrix $\mathcal{I}_{t-1} = E_{t-1}[(\nabla_t)^2]$. The conducted simulation study shows that GAS model provides an estimator, which is closer to the true parameter but has a higher variation.

A further paper dealing with dynamic copulas is Jondeau and Rockinger (2006), who model time-varying correlations for Gaussian and Student copulas in three different ways. Two of them, DCC correlations and regime-switching correlations, will be described in Sections 3.2 and 3.7. The third way can be seen as a discrete variation of the forcing equation by Patton (2006). For this the unit square is split into a number of subsets $\mathcal{A}_j$, $j = 1, ..., 16$. The choice of the subintervals can be chosen by the modeler and the authors suggest using 16 equally sized sub-squares over the grid 0, 0.25, 0.5, 0.75, 1. Correlation then is given by

$$\rho_t = \sum_{j=1}^{16} d_j \mathbf{I}[(U_{1t-1}, U_{2t-1}) \in \mathcal{A}_j], \tag{14}$$

with $d_j \in [-1, 1]$ and $\mathbf{I}$ the indicator function. Thus correlation at time t is driven by the concordance of the observations at $t-1$.

Estimation of the observation driven models is based on the maximization of the copula log-likelihood as in (7), having the vector of parameters as an argument and treating the evolution function of θ_t as a constraint.

3.2 DCC copulas

Engle (2002) proposed a multivariate GARCH model with dynamic conditional correlations (DCC), where the correlations are driven by the cross product of the lagged standardized residuals and an autoregressive term. Estimation is done, similarly as for copula models, by first estimating the GARCH parameters for the individual series and then estimating the parameters driving the correlation dynamics. This specification can easily be adapted to model the dynamics of copula parameters. Let $Y_{it} = \Phi^{-1}(U_{it})$, where Φ denotes the *cdf* of the standard normal distribution and $Y_t = (Y_{1t}, Y_{2t})'$. Then the DCC model specifies the correlation matrix R_t as

$$R_t = \text{diag}\{Q_t\}^{-1/2} Q_t \text{diag}\{Q_t\}^{-1/2}, \tag{15}$$

where Q_t follows

$$Q_t = \Omega(1 - \alpha - \beta) + \alpha Y_{t-1} Y'_{t-1} + \beta Q_{t-1} \tag{16}$$

and Ω is the unconditional covariance matrix of Y_{t-1}. This specification ensures positive definiteness of the correlation matrix and that the correlation coefficient ρ_t, which is the off-diagonal element of R_t, lies in $[-1, 1]$ at all times. Heinen and Valdesogo (2008) suggest how this approach can be extended to some non-elliptical copulas. They propose transforming the correlations into Kendall's tau through

$$\tau_t^K = \frac{2}{\pi} \arcsin(\rho_t).$$

Some copulas have a one-to-one relation between Kendall's tau and the dependence parameter θ and using this relationship the τ_t^K is mapped into θ_t. As some copulas only allow for positive dependence, Heinen and Valdesogo (2008) overcome this potential problem by replacing the off-diagonal elements of Q_t by $max(0, q_t)$ to ensure that the copula parameter always remains in its domain. Thus, the negative dependence is treated by setting the corresponding copula to the independence copula. This can be seen as a potential drawback, but as the authors mention when the conditional correlation is below zero a large fraction of the time, models only allowing for positive dependence are likely to have bad fit and will not be considered to be appropriate very often. Another disadvantage of the DCC copula specification is that it is not obvious how to generalize it to copulas that have more than one parameter.

Estimation can be done by treating the copula parameter θ_t as an observable function of α, β and $\mathcal{F}_{t-1}$, the information at time $t-1$. The copula likelihood (7) is then maximized over the parameters α and β that drive the dependence parameter.

3.3 Stochastic autoregressive copulas (SCAR)

Hafner and Manner (2008) suggest a time-varying copula model where dynamics of the copula parameter are not driven by the observations as in the DCC or the Patton model, but where the copula parameter is driven by an independent stochastic process. Formally, $\theta_t = \Lambda(\lambda_t)$, where $\Lambda : \mathbb{R} \to \Theta$ is an appropriate transformation to ensure that the copula parameter remains in its domain and whose functional form depends on the choice of copula. The underlying process $\{\lambda_t\}_{t=1}^T$, which is latent, is assumed to follow a Gaussian autoregressive process of order one,

$$\lambda_t = \omega + \beta \lambda_{t-1} + \sigma_\eta \eta_t, \tag{17}$$

where η_t is an i.i.d. $N(0,1)$ innovation and $|\beta| < 1$ to ensure stationarity of λ_t. For the Frank and the Plackett copulas the transformation Λ is simply $\Lambda(x) = x$, implying normality of the copula parameter, for the Clayton copula it is $\Lambda(x) = \exp(x)$, and for the Gumbel copula $\Lambda(x) = \exp(x) + 1$, implying log-normality of θ_t for these two families. For the Gaussian and the Student copulas the inverse Fisher transform $\Lambda(x) = (\exp(2x) - 1)/(\exp(2x) + 1)$ is the most natural choice, since the Fisher transform is the variance stabilizing transformation for the correlation coefficient (van der Vaart 1998).
Estimation of the parameter vector $(\omega, \beta, \sigma_\eta)$ is not straightforward since the process $\{\lambda_t\}_{t=1}^T$ is unobservable. Hafner and Manner (2008) propose to integrate it out of the likelihood function of the copula. Denote $U_1 = \{U_{1t}\}_{t=1}^T$, $U_2 = \{U_{2t}\}_{t=1}^T$, $\lambda = \{\lambda_t\}_{t=1}^T$ and let $f(U_1, U_2, \lambda; \omega, \beta, \sigma_\eta)$ be the joint density of the observable variables (U_1, U_2) and the latent process $\{\lambda_t\}_{t=1}^T$. Then the likelihood function is given by

$$\mathcal{L}(\omega, \beta, \sigma_\eta; U_1, U_2) = \int f(U_1, U_2, \lambda; \omega, \beta, \sigma_\eta) d\lambda. \tag{18}$$

Hafner and Manner (2008) discuss how the efficient importance sampler (EIS) by Liesenfeld and Richard (2003) and Richard and Zhang (2007) can be adapted to evaluate this T-dimensional integral by simulation. The simulated likelihood function can then be maximized over the parameter vector $(\omega, \beta, \sigma_\eta)$. As a byproduct one obtains a smoothed estimate $\hat{\lambda}_t$ of the underlying latent process and thus also a smoothed estimate $\hat{\theta}_t$ of the time-varying copula parameter.

3.4 Semiparametric dynamic copula (SDC)

Hafner and Reznikova (2008) propose a semiparamteric approach to model the time-varying behavior of the dependence parameter of a copula treating θ as a smooth function of time. On the second stage of the estimation the log-likelihood function from (7) is locally weighted around location τ

$$L(\theta; h, \tau) = \sum_{t=1}^{T} \log c(U_{1t}, U_{2t}; \theta) \cdot K_h(t/T - \tau), \tag{19}$$

where $K(\cdot)$ is a kernel function, $K_h(\cdot) = (1/h)K(\cdot)$, $h > 0$ is a bandwidth and $\tau \in [0,1]$ is an appropriate grid. Then the locally estimated dependence parameter takes the form:

$$\widehat{\theta}(\tau) = \arg\max_{\theta} L(\theta; h, \tau). \tag{20}$$

In the case when $K(\cdot)$ is a symmetric function, the estimator can possess a considerable bias at the boundaries, which is a well known problem of kernel estimation techniques (see Simonoff (1996), Ch.3). A possible solution is to approximate θ by a higher order polynomial, e.g. by simply taking the local linear function

$$\theta(t/T) \approx \theta(\tau) + \theta'(\tau)\left(\frac{t}{T} - \tau\right). \tag{21}$$

The important step prior to estimation of θ is the bandwidth selection. The MSE-optimal bandwidth is

$$\widehat{h} = \arg\min_h \left\{ \int \widehat{MSE}(x; h) w(x) dx \right\}, \tag{22}$$

where $\widehat{MSE}(\tau; h) = \widehat{bias}^2(\tau; h) + \widehat{var}(\tau; h)$ and $w(x)$ is any weight function. To obtain the estimators of the bias and variance one needs first to select the pilot bandwidth h^*, which is the minimum of the integrated Extended Residual Square Criterion (ERSC) of Fan et al. (1998)

$$ERSC(\tau; h) = J_T^{-2}(\tau) s_T(\tau) \left\{ 1 + \frac{||K||^2}{nh} \right\}, \tag{23}$$

where $J_T(\tau) = \ell''_{[\tau T]}(\widehat{\theta}(\tau))$, $s_T(\tau) = \frac{\sum_{t=1}^T (\ell'_{\tau T}(\theta^*(t/T)))^2 K_h(t/T-\tau)}{\sum_{t=1}^T K_h(t/T-\tau)}$ with $\ell_t(\theta) = \log c(U_{1t}, U_{2t}; \theta)$ and $\theta^*(t/T)$ is estimated for the local quadratic function.
If T is not equal to the number of grid subintervals, then the estimated $\theta(\tau)$ is extrapolated on $[1, T]$. Hafner and Reznikova (2008) also provide the asymptotic theory for the θ estimator.

3.5 Structural breaks

Another possibility to allow for changing dependence over time is to test for a structural break in the copula parameter at a given point in time t^* as suggested by Dias and Embrechts (2004). Let the distribution of $U_t = (U_{1t}, U_{2t})'$ be $C(U_{1t}, U_{2t}, \theta_t)$, where $t = 1, \ldots, T$.
Formally, the null hypothesis of no structural break in the copula parameter becomes

$$H_0 : \theta_t = \theta, \tag{24}$$

whereas the alternative hypothesis of the presence of a single structural break is formulated as:

$$H_1 : \theta_t = \begin{cases} \theta_1 & 1 \le t \le t^* \\ \theta_2 & t^* < t \le T. \end{cases} \tag{25}$$

In the case of a known break-point t^*, the test statistics can be derived as a generalized likelihood ratio test. Let $L_1(\theta)$, $L_2(\theta)$ and $L(\theta)$ be the log-likelihood functions of the copula using the first t^* observations, the observations from t^*+1 to T and all observations, respectively. Then the likelihood ratio statistic can be written as

$$LR_{t^*} = 2[L_1(\hat{\theta}_1) + L_2(\hat{\theta}_2) - L(\hat{\theta})], \tag{26}$$

where a hat denotes the maximizer of the corresponding likelihood function. Note that $\hat{\theta}_1$ and $\hat{\theta}_2$ denote the estimates of θ before and after the break, whereas $\hat{\theta}$ is the estimate of θ using the full sample. For t^* fixed this statistic follows a χ^2 distribution with the number of degrees of freedom equal to the dimension of θ. In the case of an unknown break date t^*, a procedure similar to the one proposed in Andrews (1993) can be applied. The test statistic proposed by Dias and Embrechts (2004) is the supremum of the sequence of statistics for known t^*

$$Z_T = \max_{1 \leq t^* < T} LR_{t^*} \tag{27}$$

and the asymptotic critical values of Andrews (1993) can be used.

Candelon and Manner (2007) have extended the procedure to additionally allow for a breakpoint in the parameters of the marginal distribution at a (possibly) different point in time and they propose a bootstrap procedure to obtain critical values of the test statistic.

3.6 Adaptive estimation method (LCP)

Giacomini et al. (2009) propose to estimate the time-varying parameters of the copula adaptively by means of local parametric fitting. The main idea is that the varying copula parameter θ_t can be well approximated by a constant θ on an interval of homogeneity I_t. The crucial point is how to estimate the length of each interval $\forall t$. This distinguishes the model from the simple case of moving window estimator, as for this method the length of the window is determined by a data driven procedure.

The Local Change Point (LCP) method developed by Mercurio and Spokoiny (2004) determines the largest interval where the dependency parameter is invariant. The method tests the hypothesis of homogeneity for the interval $I_t = [t - m_t, t)$ with the right end-point t. As soon as the length of the interval m_t is estimated, the parameter θ_t is approximated by a constant estimator $\widehat{\theta}_{\widehat{I}_t}$. The method is carried out in the counter direction for $t = T, \ldots, 1$.

The length of the interval of homogeneity I_t is estimated as follows. First, a family of nested intervals is defined as $\mathcal{I} = \{I_k = [t - m_k, t), k = 1, 2, \ldots\}$, such that $m_{k+1} > m_k$. Then, within an interval I_k a set of internal points $\mathcal{T}_k \subset I_k$ is selected. This set of points $\mathcal{T}_k$ is suspected to contain a break-point t^*. The procedure works as follows:

1. Test the hypothesis of homogeneity on $\mathcal{T}_k \subset I_k$. The null and the alternative hypothesis are similar to (24) and (25). As for the likelihood ratio test in (26), here the point $t^* \in \mathcal{T}_k$ divides the testing interval I_k in two disjoint intervals I_1 and I_2. Thus, the likelihoods are calculated for I_k, I_1 and I_2 with the ML estimators $\widetilde{\theta}_k, \widetilde{\theta}_1$ and $\widetilde{\theta}_2$. The corresponding Z_{I_k} statistics from (27) is then compared to the critical value. The hypothesis of homogeneity of θ is rejected when Z_{I_k} exceeds the critical value.

2. If H_0 for k is not rejected, then the next interval I_{k+1} is tested for homogeneity.

3. If H_0 is rejected on I_k, then the interval of homogeneity is the last accepted interval $\widehat{I}_t = I_{k-1}$.

If a large window is selected the estimate of dependence is not sensitive and reacts to changes in dependence with high delay. On the contrary, if a window is small, the estimate is quite unstable with high perturbation. This is also the case for the first observations, for which the window is forced to be small. The size of the window depends on the choice of the critical values and other parameters, described in Giacomini et al. (2009) and Mercurio and Spokoiny (2004).

3.7 Regime switching copulas (RSC)

A further way to specify a copula model in which both the degree and the type of dependence change over time is to allow for a number of states, each being characterized by a different copula. These copulas can be from the same family but allowing for different parameters. They may, however, also change their functional forms implying different states having entirely different dependence structures, a possibility we do not consider here, but that allows for interesting modeling of financial data. One may think of a model distinguishing tranquil and crisis time, the former being characterized by a Gaussian copula, whereas during the latter data is being generated by a copula allowing for lower tail dependence. To our knowledge the first to allow for regime switching in correlations is Pelletier (2006). Garcia and Tsafack (2008) and Chollete et al. (2008) have explicitly modeled copulas in a regime switching framework. Let k_t be a latent random variable that takes on the value $k = 1, ..., K$ when regime k is the current state. Then

$$(U_{1t}, U_{2t} | k_t = k) \sim C(U_{1t}, U_{2t}; \theta_k) \tag{28}$$

and k_t is assumed to follow a Markov chain of order one with π_{ij} the probability of moving to regime j in period t conditional on being in state i at time $t-1$. Usually the number of states K is taken to be equal to two or three. $K = 2$ is the more common choice which we focus on in this study. Estimation can be done using the Expectation Maximization (EM) algorithm as outlined in Hamilton (1994) Ch. 22. Define the matrix of transition probabilities

$$P = \begin{pmatrix} \pi_{11} & 1-\pi_{11} \\ 1-\pi_{22} & \pi_{22} \end{pmatrix},$$

and let $\hat{\xi}_{t|t}$ be a (2×1) vector containing the estimated probabilities of being in each state at time t given the information at time t. Further $\hat{\xi}_{t|t-1}$ are the same estimated probabilities only using information until time $t-1$. Then the system is described by

$$\hat{\xi}_{t|t} = \frac{\hat{\xi}_{t|t-1} \odot \eta_t}{\mathbf{1}'(\hat{\xi}_{t|t-1} \odot \eta_t)}, \tag{29}$$

$$\hat{\xi}_{t+1|t} = P'\hat{\xi}_{t|t}, \tag{30}$$

$$\eta_t = \begin{pmatrix} c_1(U_{1t}, U_{2t}; \theta_1) \\ c_2(U_{1t}, U_{2t}; \theta_2) \end{pmatrix}, \tag{31}$$

with $\mathbf{1}$ a vector of ones and $\odot$ the Hadamard product[2]. For a given starting value $\hat{\xi}_{1|0}$ and copula parameters θ_1, θ_2 and transition probabilities π_{11} and π_{22} one can iterate over (29) and (30) to obtain the log-likelihood function of the copula

$$LL_C(\theta_1, \theta_2, \pi_{11}, \pi_{22}; U_{1t}, U_{2t}) = \sum_{t=1}^{T} \log(\mathbf{1}'(\hat{\xi}_{t|t-1} \odot \eta_t)). \tag{32}$$

Formulas to estimate the smoothed probabilities of being in each state at time t, $\hat{\xi}_{t|T}$, can be found in Hamilton (1994).

3.8 Other approaches

In this section we shortly review additional approaches for testing for and modeling time-varying copulas that have been proposed in the literature. However, we will skip most of

[2]The Hadamard product denotes element by element multiplication of two equally sized matrices.

the details for the sake of brevity.

van den Goorbergh et al. (2005):
In this paper time-varying copulas are used to price options with multiple underlying assets and it is found that the option prices implied by time-varying copulas are quite different from those using static copulas. The relation between the parameter of some one-parameter copulas and Kendall's tau is exploited to estimate the copula parameter by a moment type estimator. A time-varying measure of Kendall's tau then implies a time-varying copula parameter. It is assumed that dependence is driven by the volatility of the assets, which is reasonable as it is implied by factor models for asset pricing and this relation has been confirmed in a number of studies. Let h_{it} be the conditional variance of asset i (e.g. the GARCH variance). Then τ_t is assumed to follow

$$\tau_t = \gamma_0 + \gamma_1 \log\{\max(h_{1t}, h_{2t})\}. \tag{33}$$

The parameters γ_0 and γ_1 are estimated by regressing a rolling window estimate of τ_t on a constant and the maximum of the logarithm of the maximum of the GARCH variances. The window size is chosen to be equal to about 40 days, although it is found that the results are robust to the choice of the window size.
The main difference of this approach to the ones presented so far is that the copula parameter is assumed to depend on the marginal distribution through the conditional variance, whereas all the other approaches assume that the copula parameter behaves independently of the parameters of the marginal distributions.

Guégan and Zhang (2009):
The difference between this approach to the majority of the competing approaches is that the authors do not only test for a change in the relationship between the variables of interest, but also whether the copula remains the same and only the degree of dependence changes, or whether additionally also the type of copula changes at a given point in time[3]. The main idea is to compare a parametric copula to a nonparametric estimate of the copula density at m distinct points in time using the goodness-of-fit tests by Fermanian (2005). By applying the test to a conditional copula one can check whether the copula family changes. When the copula family changes the authors suggest using a binary segmentation procedure to detect the change points and the type of copula on

[3]One exception is the regime switching copula presented in Section 3.7.

each sub-interval, otherwise they suggest using the structural break test by Dias and Embrechts (2004) to detect the change points of the copula parameters. For the details of the procedure and the test statistics we refer the interested reader to the original paper.

Harvey (2008):
A further technique worth mentioning is that of Harvey (2008), who treats the problem of changing copulas by noting that it is related to estimating time-varying quantiles. The method is non-parametric and very different to the other techniques described here. Busetti and Harvey (2008) build on the same methods to construct a formal test for changing dependence. A description of the approaches is beyond the scope of this paper.

4 Model selection and simulations

In this section we study how to measure the goodness-of-fit for time-varying copulas, how to select the best fitting model and how well the competing specifications for time-varying dependence presented in the previous section are able to estimate the underlying dependence process.

4.1 Specification testing

Assume for a given time series of observations $(U_{1t}, U_{2t}),\ t = 1, \ldots, T$ copula model C_i has been estimated, where i denotes a candidate parametric copula, and an estimate for the sequence of dependence parameters $\hat{\theta}_{it},\ t = 1, \ldots, T$ has been obtained. The first thing we are interested in is which of the competing models C_i fits the data at hand best. Even though the models are usually non-nested and standard likelihood ratio test cannot be performed a very simple and (as we shall see) reliable way to select the best fitting model is to compare the value of the log-likelihood function LL_i. The model with the highest likelihood is considered to be the best fitting one[4].
The model maximizing the LL statistics, however, must not necessarily provide a satisfac-

[4]It is theoretically more sound to use the Akaike Information Criterion (AIC) to compare the fit of non-nested models, but since we only compare the fit within each specification for the time-variation, the number of parameters is always the same and hence it is equivalent to looking at the value of the log-likelihood function.

tory fit for the data being analyzed. Thus, for an estimate $\hat{\theta}_{it},\ t = 1, \ldots, T$ the hypothesis of interest is whether the data has actually been generated by C_i. Let $C_0(U_{1t}, U_{2t}, \theta_t^0)$ be the true copula where θ_t^0 denotes the true parameter at time t. Then formally the null hypothesis is

$$H_0 : C_i(U_{1t}, U_{2t}, \hat{\theta}_{it}) = C_0(U_{1t}, U_{2t}, \theta_t^0). \tag{34}$$

Note that this means that we are testing both the copula specification C_i and the estimate of the latent dependence parameter of model i, $\hat{\theta}_{it}$, and rejecting H_0 does not necessarily mean that the data was not generated by C_i. We test the hypothesis in (34) by testing whether the copula of U_1 given U_2 is uniformly distributed, which is an application of the Rosenblatt probability integral transformation. In our case this means

$$\hat{z}_t = C_i(U_{1t}|U_{2t}, \hat{\theta}_{it}) = \frac{\partial C_i(U_{1t}, U_{2t}, \hat{\theta}_{it})}{\partial U_{2t}} \sim U(0, 1). \tag{35}$$

We test this hypothesis by applying the Anderson-Darling (Anderson and Darling (1952)) test, which is given by

$$T_{AD} = \sup_x \frac{\sqrt{T}|\hat{\mathbb{F}}(x) - F(x)|}{\sqrt{F(x)(1 - F(x))}}, \tag{36}$$

where $\hat{\mathbb{F}}(\cdot)$ denotes the empirical *cdf* of $\hat{z}_t$ and $F(x)$ is the $U(0, 1)$ *cdf*. For this statistics tabulated critical values must be used. Contrary to applying the test in the static copula setting for the time-varying case we are actually not only testing the functional form, but, as mentioned above, also the quality of the estimate $\hat{\theta}_{it},\ t = 1, \ldots, T$, which may cause size distortions and influence the power of the tests.

4.2 Monte Carlo study

The simulation setup is as follows. We randomly draw a sample $(U_{1t}, U_{2t})_{t=1}^T$ from a Gaussian copula with time-varying correlation coefficient. The correlations follow three alternative processes, two of which are deterministic and one is stochastic:

1. Step: $\rho_t = 0.2 + 0.6 I_{t>500}$
2. Sine: $\rho_t = 0.5 + 0.4 \cos(2\pi t/400)$
3. AR(1): $\rho_t = (\exp(2\lambda_t) - 1)/(\exp(2\lambda_t) + 1)$ with $\lambda_t = 0.02 + 0.97\lambda_{t-1} + 0.1\varepsilon_t$,

Table 1: MSE for estimating the underlying correlation

MSE	Const	DCC	PATT	SDC	LCP	SCAR	RSC
Step	0.092	0.016	0.053	0.007	0.017	0.008	0.004
	(0.002)	(0.004)	(0.005)	(0.003)	(0.005)	(0.003)	(0.003)
Sine	0.082	0.021	0.048	0.006	0.047	0.010	0.020
	(0.002)	(0.004)	(0.005)	(0.003)	(0.007)	(0.003)	(0.005)
AR(1)	0.076	0.040	0.052	0.035	0.063	0.025	0.036
	(0.022)	(0.009)	(0.011)	(0.008)	(0.014)	(0.006)	(0.010)

Note: Table 1 reports the MSE for estimating the underlying correlation process for data that has been generated by Gaussian copulas with correlation following a Step, Sine and AR(1) processes. Monte Carlo standard errors are given in parenthesis. The sample size is 1000 and the number of Monte Carlo replications is equal to 1000 for Const, DCC, PATT and RSC, 250 for SDC and SCAR, and 100 for LCP.

where $\varepsilon_t \sim N(0,1)$. Note that the average correlation is 0.5 for each of the data generating processes. We decided to leave out the case of data generated by a model with constant correlation, but we note that the models seem to be able to deal well with the case of constant dependence. Some simulation results for this situation can be found in Hafner and Manner (2008) and Hafner and Reznikova (2008). For each artificial data set we estimate the Gaussian, Frank, Gumbel and Clayton copulas with the following method to allow for time variation: Constant, DCC (§3.2), PATT (§3.1), SDC (§3.4), LCP (§3.6), SCAR (§3.3), and RSC (§3.7). For each estimation technique and each model $\hat{\theta}_{it}$, $t = 1, \ldots, T$ and LL_i is obtained[5]. The sample size is equal to $T = 1000$, corresponding to 4 years of daily data, and the number of Monte Carlo replications is 1000 in general, although due to the extremely high computational complexity it was only 250 for the SCAR and SDC models, and 100 for the LCP specification.

In order to get an idea of how well the competing time-varying copula models introduced above are able to estimate the underlying dependence parameter θ_t at each point in time we compute the mean square distance between the true dependence parameter and its estimate

$$MSE = \frac{1}{K}\sum_{k=1}^{K}\frac{1}{T}\sum_{t=1}^{T}(\hat{\theta}_t^k - \theta_t^{0k})^2, \tag{37}$$

[5]For the regime switching copula $\hat{\theta}_{it}$ is computed as the smoothed probabilities of being in each of the two states times the parameter in that state.

Table 2: Model selection by the log-likelihood statistic

	Const	DCC	PATT	SDC	LCP	SCAR	RSC
				Step			
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.212	0.997	0.011	0.996	0.720	0.968	0.990
Clayton	0.008	0.001	0.001	0.000	0.000	0.020	0.001
Frank	0.697	0.000	0.824	0.000	0.160	0.008	0.000
Gumbel	0.083	0.002	0.164	0.004	0.120	0.004	0.009
				Sine			
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.212	0.981	0.007	1.000	0.350	1.000	0.999
Clayton	0.008	0.002	0.002	0.000	0.010	0.000	0.000
Frank	0.697	0.006	0.488	0.000	0.260	0.000	0.000
Gumbel	0.083	0.011	0.503	0.000	0.380	0.000	0.001
				AR(1)			
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.318	0.925	0.327	0.956	0.190	0.962	0.991
Clayton	0.004	0.001	0.002	0.000	0.060	0.034	0.000
Frank	0.608	0.043	0.312	0.024	0.360	0.004	0.005
Gumbel	0.070	0.031	0.359	0.02	0.390	0.000	0.004

Note: Table 2 reports the fraction of times each estimated copula has the highest log-likelihood statistics for data that has been generated by Gaussian copulas with correlation following a Step, Sine and AR(1) processes. Monte Carlo standard errors are given in parenthesis. The sample size is 1000 and the number of Monte Carlo replications is equal to 1000 for Const, DCC, PATT and RSC, 250 for SDC and SCAR, and 100 for LCP.

Table 3: Size and power of the Anderson-Darling test based on the probability integral transform

	Step						
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.352	0.058	0.254	0.056	0.040	0.056	0.048
	(0.015)	(0.007)	(0.014)	(0.007)	(0.006)	(0.007)	(0.007)
Clayton	0.643	0.864	0.619	0.600	0.480	0.737	0.838
	(0.015)	(0.011)	(0.015)	(0.015)	(0.016)	(0.014)	(0.012)
Frank	0.051	0.207	0.254	0.268	0.190	0.267	0.062
	(0.007)	(0.013)	(0.014)	(0.014)	(0.012)	(0.014)	(0.008)
Gumbel	0.539	0.621	0.584	0.564	0.420	0.649	0.571
	(0.016)	(0.015)	(0.016)	(0.016)	(0.016)	(0.015)	(0.016)
	Sine						
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.352	0.129	0.324	0.068	0.260	0.060	0.041
	(0.015)	(0.011)	(0.015)	(0.008)	(0.014)	(0.007)	(0.006)
Clayton	0.643	0.898	0.635	0.640	0.770	0.790	0.762
	(0.015)	(0.010)	(0.015)	(0.015)	(0.013)	(0.013)	(0.013)
Frank	0.051	0.142	0.134	0.212	0.110	0.329	0.130
	(0.007)	(0.011)	(0.011)	(0.013)	(0.010)	(0.015)	(0.011)
Gumbel	0.539	0.625	0.561	0.552	0.520	0.671	0.595
	(0.016)	(0.015)	(0.016)	(0.016)	(0.016)	(0.015)	(0.016)
	AR(1)						
	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.291	0.160	0.249	0.100	0.370	0.076	0.054
	(0.014)	(0.012)	(0.014)	(0.019)	(0.048)	(0.017)	(0.007)
Clayton	0.656	0.897	0.661	0.708	0.730	0.810	0.858
	(0.015)	(0.010)	(0.015)	(0.014)	(0.014)	(0.012)	(0.011)
Frank	0.078	0.118	0.126	0.216	0.080	0.397	0.134
	(0.008)	(0.010)	(0.010)	(0.013)	(0.009)	(0.015)	(0.011)
Gumbel	0.555	0.652	0.605	0.636	0.590	0.741	0.668
	(0.016)	(0.015)	(0.015)	(0.015)	(0.016)	(0.014)	(0.015)

Note: Table 3 reports the rejection frequency of the null hypothesis of correct copula specification using the Anderson-Darling test based on the probability integral transform at a 5% nominal level. Data has been generated by Gaussian copulas with correlation following a Step, Sine and AR(1) processes. Monte Carlo standard errors are given in parenthesis. The sample size is 1000 and the number of Monte Carlo replications is equal to 1000 for Const, DCC, PATT and RSC, 250 for SDC and SCAR, and 100 for LCP.

where K is the number of Monte Carlo replications, and θ_t^{0k} and $\hat{\theta}_t^k$ denote the true and estimated dependence paths at replication k, respectively. Table 1 reports the average MSE between the true and the estimated correlation processes for the Gaussian copula for the static and each of the time-varying copula specifications and for the different correlation dynamics. As expected, all models lead to substantial improvements over the constant copula model. However, the RSC, SCAR and SDC models are superior to the competing ones in all cases, the RSC being better for the Step correlation, SCAR for the AR(1) correlation and the SDC for the Sine correlation, as to be expected. The DCC performs worse than all three, but better than both PATT and LCP. The latter specification naturally does not do too well for the Sine and AR(1) correlations as the assumption of intervals of homogeneity is violated. Surprisingly, although the performance for the Step correlation is acceptable the MSEs are still higher than those of the DCC, SDC and SCAR models. This is puzzling insofar as this DGP should strongly favor the nature of the LCP procedure.

The fraction of times each copula is the selected as the best fitting one in terms of the highest LL statistics can be found in Table 2. One has to keep in mind that the comparison of different copulas using the LL statistic is only possible within the same specification for the time-variation, but cannot generally be used to compare different models for the dynamics in dependence. When ignoring the time-variation of dependence the Frank copula is chosen quite often. This suggests that the unconditional copula corresponding to a time-varying Gaussian copula is closer to the static Frank than to the Gaussian copula. When using the RSC, DCC, SDC and the SCAR models the LL statistics turns out to be a very reliable model selection criterion. It does, however, become quite unreliable for both PATT and LCP, although for the latter the results for the step correlation are still acceptable.

The size and the power for the AD test at a 5% nominal level are reported in Table 3. Monte Carlo standard errors are included, since for different models a different number of Monte Carlo replication was chosen. In terms of size the RSC model performs best closely followed by the SCAR model, which is slightly oversized for all cases. The SDC has higher size distortions, but still does quite well and the other models are all severely oversized with the exception of the DCC and the LCP models for the Step correlation. The power of the tests, which is not corrected for the size distortions, is best for the RSC, DCC and the SCAR models, although the SDC also has good power properties. All models have problems rejecting the Frank copula and in some cases the power against the Frank copula is even below the size. This is probably due to the fact that the Frank copula is quite similar to the Gaussian copula having no tail dependence and a symmetric dependence structure. The power against the asymmetric copulas looks better for all models.

Overall, we can conclude the RSC, SDC and SCAR specifications for time-varying copulas are superior to the competing specifications. These models do not only perform very well for the DGPs that clearly favor the models, namely Step for RSC, AR(1) for the SCAR and Sine for SDC, but also for the other DGP's. This shows the flexibility of these approaches. For the SDC it is due to the non-parametric nature of the parameter changes and the local estimation of the model. The SCAR model most likely performs well due to the high flexibility allowed for by including a random error term in the dependence process and the fact that the importance sampler exploited for its estimation makes efficient use of the information contained in the data. The usefulness and flexibility of the regime switching approach has already been shown for many other models and it seems to work equally well for copulas. Still, the DCC model also shows a rather good performance having the big advantage that it is easy to implement and that it does not require heavy computations, which in fact is also the case for the RSC.

Which model to use depends on the assumptions one is willing to make on the time-evolution of the dependence parameter, SDC being more suitable for smoothly changing processes, whereas the DCC and SCAR models are more appropriate for autoregressive correlations and regimes switching naturally applying when one believes in different states of the world. The simulation result showed that even for misspecified correlation dynamics these models perform well. Still, formal techniques to decide which method provides the best fit on a given data set need to be developed. From a practical point of view the choice of the model is also a matter of taste and software availability.

Note that although we only considered the Gaussian copula as the data generating process unreported simulations suggest that our findings continue to hold when the data is generated by different copulas.

5 Empirical illustration

For the empirical example we consider two data sets. The first data set are daily returns of the exchange rates of Yen-USD and Euro-USD. It contains 1564 observations from 31 December 1999 till 30 December 2005. The second data set are weekly returns of Morgan Stanley Capital International (MSCI) indexes of Korea and Singapore (in US Dollars) with 1039 observations from 10 May 1989 till 29 April 2009. With these examples we want to check the ability of the copula models to describe both data in tranquil and crisis times, and also to find out how much information is hidden in volatility vs. dependence. The log-returns of Yen and Euro do not show any unusual behavior due to the selected observation period, with the skewness (-0.06 and -0.08) and kurtosis (3.61 and 4.27), respectively. The log-returns of Korea and Singapore MSCI indexes, on the other hand, show vivid evidence of the clusters of volatility (Dec'97 and Nov'08). The descriptive statistics also suggest that the observations should be filtered: both series posses negative skewness (-0.48 and -0.40) and large kurtosis (9.64 and 5.87) respectively. The Jarque-Bera test for normality clearly rejects the null hypothesis for all series.

At the first stage of estimation of the models we model the marginal distributions of the data. We use AR(p)-GARCH(1,1) models with Student-t error terms to correct the log-returns for the presence of autocorrelation and conditional heteroscedasticity. The number of lags p of the AR(p) model is selected by Bayesian information criterion (BIC). Thus, the model for the log-returns X_{it} looks as follows:

$$X_{it} = \alpha_{i0} + \sum_{j=1}^{p} \alpha_{ip} X_{i,t-j} + \varepsilon_{it} \tag{38}$$

$$\varepsilon_{it} = \sqrt{h_{it}} z_{it} \tag{39}$$

$$h_{it} = \omega_i + \alpha_i \varepsilon_{i,t-1}^2 + \beta_i h_{i,t-1}, \tag{40}$$

where i is the index of the analyzed data series and z_{it} are standard-t distributed with ν_i degrease of freedom. The first stage estimators are given in Table 4. The adequacy of the estimated models is tested by applying the Ljung-Box test on the estimated residuals. Thus, we estimate $\widehat{z}_{it} = \varepsilon_{it}/\sqrt{\widehat{h}_{it}}$, where $\widehat{z}_{it}\sqrt{\frac{\nu_i}{\nu_i-2}}$ follows a Student-t distribution with ν_i degrees of freedom.

5.1 Copula model for exchange rates of Euro-USD and Yen-USD

Next we estimate the dependence structure between Euro-USD and Yen-USD with six types of copulas: two symmetric with no tail dependency (Gaussian, Frank), two with upper tail dependency (Gumbel, rotated Clayton), and two with lower tail dependency (Clayton, rotated Gumbel). These copulas and their properties are reviewed in the appendix. The models for the time evolution of the parameter are Constant, DCC, PATT,

Table 4: First stage estimators: AR(p)-GARCH(1,1) model

	AR(p)	**GARCH(1,1)**			**d.o.f.**
	$\alpha_0, \alpha_1, \ldots \alpha_p$	ω	α	β	ν
Euro	$-9.7E{-}05$, -0.06	$3.5E{-}07$	0.02	0.97	28.83
	$(1.7E{-}04)$ (0.03)	$(1.3E{-}07)$	(0.01)	(0.01)	(12.03)
Yen	$9.8E{-}05$, -0.04	$5.3E{-}07$	0.02	0.96	7.11
	$(1.5E{-}04)$ (0.03)	$(1.5E{-}07)$	(0.01)	(0.01)	(1.15)
Singapore	$5.3E{-}04$, 0.06	$1.6E{-}05$	0.11	0.88	7.48
	$(9.9E{-}04)$ (0.03)	$(7.9E{-}06)$	(0.03)	(0.03)	(1.58)
Korea	$-4.8E{-}02$, 0.03, 0.13	$5.9E{-}05$	0.12	0.86	10.60
	$(3.1E{-}02)$ (0.03) (0.03)	$(2.3E{-}05)$	(0.03)	(0.03)	(3.26)

Note: Table 4 reports the estimated parameters and standard errors of the AR(p)-GARCH(1,1)-t model for the log-returns of exchange rates Euro-USD and Yen-USD (daily observations, Dec'99 - Dec'05) and MSCI indexes of Singapore and Korea (weekly observations, May'89 - Apr'09).

SCDM, LCP, SCAR and RSC discussed in previous sections. Table 5(a) reports the log-likelihoods of the estimated models. For each model the best fitting type of copula in terms of the likelihood is marked out in bold. As it is seen from the table the likelihoods of Constant, PATT and LCP models favor Frank copula, whereas DCC, SDC, SCAR and RSC models point to Gaussian copula. However, the log-likelihoods for the Frank copula are virtually identical in latter cases. Taking into account the finding of the Monte Carlo study that the DCC, SDC, SCAR and RSC models are more reliable when selecting the best fitting copula using the log-likelihood either the Frank or the Gaussian copula could be selected. Recall that in general it is not possible to compare the fit across different specification by looking at the log-likelihood, as not all models have the same number of parameters. However, the fit of the DCC, PATT and SCAR models may in fact be compared, because they do have the same number of parameters.

The goodness-of-fit of the estimated models is then checked with Anderson-Darling (AD) test of correct copula specification, described in Section 4. The p-values of the test are presented in Table 5(b). The Frank copula passed the test for all estimated specifications, whereas all the other copulas are rejected. Taken as a whole, these findings strongly favor the Frank copula as the best fitting copula.

Figure 1 presents the dependence paths, estimated from Frank copula and transformed to Kendall's tau for the sake of comparison. The estimated paths of SDC and SCAR models are very close. Dependence estimated with Patton, DCC and RSC models show similar

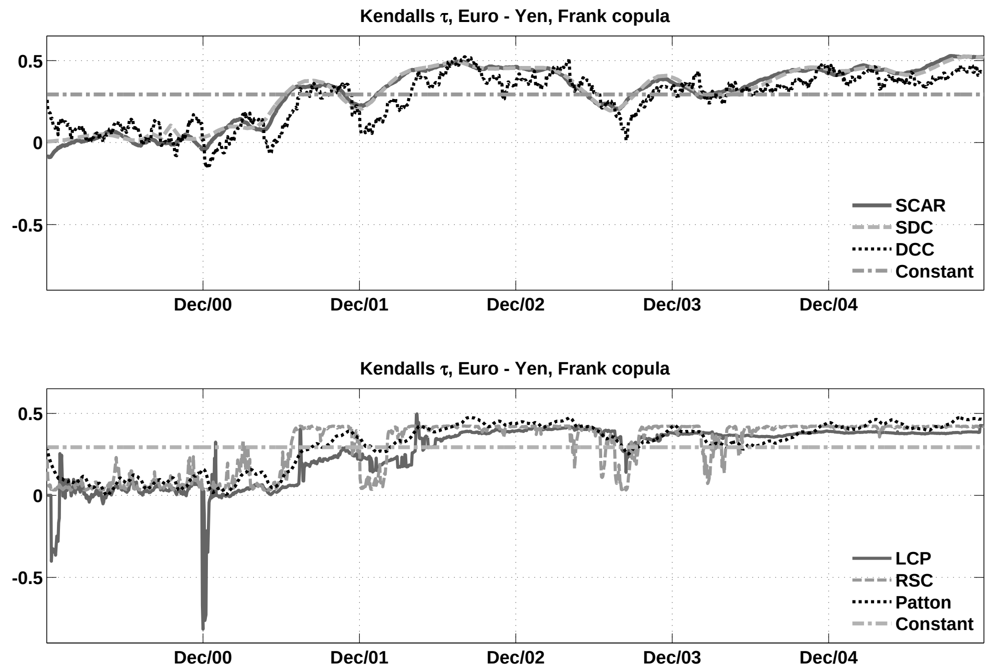

Figure 1: Estimated dependence for the pair of exchange rates Euro-USD and Yen-USD. Frank copula. Dependence paths are transformed to Kendall's tau. Daily observations, Dec'99 - Dec'05.

behavior as SDC and SCAR, but shifted to the right. This can be explained by the fact that the SDC and SCAR take into account the information of the full sample to estimate dependence at time t, whereas the other specifications only rely on past information. Finally, the erratic behavior of the dependence path estimated for LCP in 2001 suggests the presence of a sudden change. Indeed, the data seem to be independent until January 2001 and then the dependence grows considerably. This corresponds to the expectations of the introduction of the Euro in January 2002.

Finally, the last measure that we use to test the adequacy of the estimated models is the Value-at-Risk (VaR) of an equally weighted portfolio. $VaR_t(\alpha)$ is the α-quantile of the conditional distribution of portfolio returns at time t, which can be obtained by simulation. Table 5(c) reports the results of the Dynamic Quantile (DQ) test of Engle and Manganelli (2004). The null hypothesis of the DQ test states that the model is correctly specified and that VaR is not under or over-estimated. The test is based on F statistics and tests $H_0 : \delta_0 = \delta_1 = \ldots = \delta_6 = 0$ for the regression:

$$hit_t^\alpha - \alpha = \delta_0 + \delta_1 hit_{t-1}^\alpha + \ldots + \delta_5 hit_{t-5}^\alpha + \delta_6 VaR_t(\alpha) + \nu_t, \tag{41}$$

where $hit_t^{\alpha} = \mathbb{I}(X_t \leq VaR_t(\alpha))$ and X_t is the return of the portfolio. The results of the DQ test are shown in Table 5(c). For Gaussian and Frank copulas the estimated VaR has no autocorrelation in the hits for five and four out of seven models, respectively. For other types of copulas the P-values are in general close to zero. Thus, we can conclude that for this data example not only the properly estimated volatilities of the marginals matters, but also the dynamics of the joint dependence structure of the assets. However, as it will be shown in the next section, DQ test for VaR as a models' goodness-of-fit criterion should be used with care.

5.2 Copula model for MSCI indexes of Korea and Singapore

In the second application we consider weekly observations of the MSCI indexes of Korea and Singapore. As in the example above we estimate time-varying copula specifications for the same types of copulas. The results of the evaluated log-likelihoods can be found in Table 6(a). The log-likelihoods unambiguously point to the Gaussian copula as the best fitting copula type. As for the second best choice, for all seven models it is a rotated Gumbel copula. This provides some evidence of lower tail dependence, which is not a surprise given a financial crisis occurred in 1997 and stock market returns tend to have more dependence for losses than for gains.

The AD test results are reported in Table 6(b). The test rejects only Gumbel and rotated Clayton copulas, but approves all the other types. Given that it produces the highest log-likelihood statistic and that it is not rejected by the AD test the Gaussian copula seems to be the best fitting model, although one may argue in favor of the rotated Gumbel copula. The transforms to Kendall's tau of the dependence paths based on Gaussian copula are shown in figure 2. The estimated paths of dependence for the SDC and SCAR models are very close and look quite smooth. The correlation estimated from DCC model is also very close to SDC and SCAR with some deviations. The dependence estimated from Patton's model is this time noisier than of DCC model and compared to the SDC/SCAR models lies closer to the unconditional dependence parameter throughout the sample. The RSC estimator vividly shows the periods of constancy of the dependence. However, the main shift in the dependence for this model falls on the year 2000. The dependence path estimated from the LCP model deviates a lot from the other models when the dependence increases due to the Asian crisis in 1997. Note that this increase in dependence provides evidence for financial contagion as studied using copulas in Rodriguez (2007) and Candelon and Manner (2007).

Table 5: Model selection for Euro-USD and Yen-USD data

(a) Log-likelihood

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	132.6	**194.3**	170.3	**228.9**	151.9	**202.2**	**207.63**
Gumbel	123.7	176.5	161.0	200.6	169.9	173.7	178.53
Clayton	113.4	145.2	142.9	161.9	135.3	149.5	151.86
Frank	**146.5**	194.2	**194.9**	226.8	**183.1**	201.8	205.32
rot Gumbel	134.4	182.9	169.5	198.3	169.3	177.6	169.04
rot Clayton	95.3	131.1	128.4	161.2	140.7	110.6	144.10

(b) AD test

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.00	0.00	0.00	0.00	0.00	0.03	0.03
Gumbel	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Clayton	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Frank	0.14	0.16	0.51	0.48	0.17	0.32	0.25
rot Gumbel	0.00	0.00	0.00	0.00	0.00	0.00	0.00
rot Clayton	0.00	0.00	0.00	0.00	0.00	0.00	0.00

(c) EM test

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.04	0.07	0.17	0.07	0.04	0.64	0.07
Gumbel	0.03	0.03	0.00	0.29	0.02	0.10	0.03
Clayton	0.19	0.44	0.00	0.23	0.00	0.02	0.26
Frank	0.27	0.04	0.05	0.17	0.03	0.16	0.03
rot Gumbel	0.08	0.00	0.02	0.00	0.03	0.00	0.02
rot Clayton	0.02	0.01	0.05	0.04	0.03	0.11	0.07

Note: Table 5 reports the log-likelihood (a), the p-values of the Anderson-Darling test for correct copula specification (b) and the p-values of the Engle-Manganelli test for the correct specification of the Value-at-Risk (c). The data are log-returns of the exchange rates Euro-USD and Yen-USD (daily observations, Dec'99 - Dec'05).

Table 6: Model selection for Singapore-Korea MSCI indexes

(a) Log-likelihood

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	**98.5**	**134.6**	**120.4**	**149.1**	**117.7**	**133.9**	**127.39**
Gumbel	88.9	121.9	109.7	133.5	106.6	118.6	117.65
Clayton	81.7	109.4	105.6	128.2	95.9	114.1	109.43
Frank	87.9	117.0	106.1	130.6	101.8	116.2	106.31
rot Gumbel	93.4	125.6	115.8	139.7	112.2	122.2	122.88
rot Clayton	71.1	94.4	89.9	110.4	83.6	92.7	90.94

(b) AD test

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.17	0.30	0.30	0.41	0.27	0.38	0.19
Gumbel	0.02	0.00	0.01	0.01	0.01	0.00	0.00
Clayton	0.75	0.04	0.49	0.15	0.60	0.14	0.15
Frank	0.21	0.05	0.10	0.05	0.08	0.06	0.11
rot Gumbel	0.34	0.13	0.59	0.14	0.53	0.17	0.51
rot Clayton	0.03	0.00	0.01	0.00	0.02	0.00	0.00

(c) EM test

	Const	**DCC**	**PATT**	**SDC**	**LCP**	**SCAR**	**RSC**
Gaussian	0.76	0.45	0.14	0.14	0.14	0.79	0.86
Gumbel	0.45	0.24	0.66	0.66	0.66	0.83	0.24
Clayton	0.32	0.55	0.39	0.39	0.39	0.34	0.44
Frank	0.01	0.11	0.17	0.17	0.17	0.49	0.91
rot Gumbel	0.26	0.66	0.71	0.71	0.71	0.65	0.47
rot Clayton	0.03	0.90	0.09	0.09	0.09	0.84	0.47

Note: Table 5 reports the log-likelihood (a), the p-values of the Anderson-Darling test for correct copula specification (b) and the p-values of the Engle-Manganelli test for the correct specification of the Value-at-Risk (c). The data are log-returns of the MSCI indexes of Singapore and Korea (weekly observations, May'89 - Apr'09).

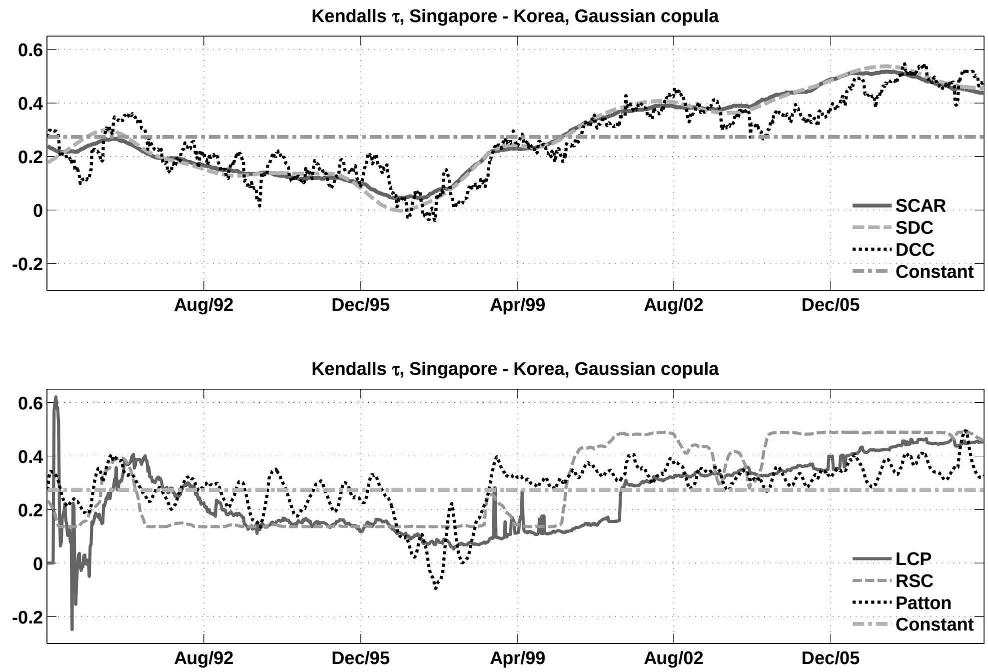

Figure 2: Estimated dependence for the pair of MSCI indexes of Singapore and Korea. Gaussian copula. Dependence paths are transformed to Kendall's tau. Weekly observations, May'89 - Apr'09.

Finally, Table 6(c) provides the DQ test results. For this data example the DQ test approved the estimated VaR for almost all types of models and copulas. Such a result shows us that most of the risk information is hidden in the volatilities of the individual data series and less in the joint dependence structure. Thus, though it is demonstrated that the associated countries tend to be more dependent during the crisis period and even after, the risk hidden in the dependence structure is not always relevant. Hence, DQ test is not a bona fide goodness-of-fit measure, but just an auxiliary method.

6 Conclusions

In this paper we have provided a survey over existing copula models allowing for time-varying dependencies that have been proposed in recent years. Correctly modeling the dependence between financial assets plays a crucial role for measuring risks and pricing derivatives and since there is strong evidence that dependencies change over time, appropriately modeling and measuring these changes is not only interesting for its own sake,

Table 7: Comparison of the presented models

	DCC	PATT	SDCM	LCP	SCAR	RSC
estimating θ_t	+/0	-	+	-	+	+
GoF testing	+/0	-	+/0	-	+	+
computations	+	+	-	-	-	+
flexibility	0	0	+	-	+	+

but also has important economic implications.

The different time-varying copula models we reviewed rely on different assumptions about the way dependence may change over time ranging from structural breaks in dependence, the existence of different dependence regimes, smooth changes or copula parameters behaving like an independent stochastic process. Since one cannot directly observe the dependence parameter and hence no *a priori* type of dynamics can be favored a natural question is how robust the competing models are to a misspecification of these dynamics. Our simulation results suggest that the RSC relying on a regime switching framework, the SDCM assuming smoothly changing dependence parameter and SCAR model assuming autoregressive stochastic dependence seem to work better than those competing techniques that have been studied in more detail, also in situations when they are clearly misspecified. However, the DCC-copula model also performs quite well and given that its estimation is easy its use can be recommended in many situations. Table 7 gives an overview of the properties of the techniques under different criteria. Overall, if we had to recommend a single model it would be the RSC since in addition to good performance in the simulations it is easy to program and does not require heavy computations.

For assessing the goodness-of-fit we recommend comparing the log-likelihood statistics in addition to performing the Anderson-Darling test on the data transformed by the Rosenblatt probability integral transform, which has acceptable size and power properties for a number of models. However, ignoring the time-variation of the dependence when deciding which copula best fits the data is not recommendable as it will most likely lead to false conclusions.

In our empirical application we found that when allowing for time-varying dependence parameters symmetric copulas that do not allow for tail dependence offer the best fit, which is in contrast to what has been found in the literature for the static case, where

usually copulas that feature tail dependence and asymmetry seem appropriate. Thus it appears that part of the asymmetry may be generated by time-varying parameters. The lack of tail dependence may partially be offset by the possibility of large overall dependence, which would explain why the Gaussian and Frank copulas fit the data so well. Finally, the models we studied seem to be reliable when estimating the Value-at-Risk.
The most important challenge for future research is to develop time-varying copula models in dimensions larger than two. This is crucial in order to make these models applicable for practical purposes. For Gaussian and Student copulas techniques from multivariate volatility modeling such as the DCC model and the model by Asai and McAleer (2009) look promising. Nevertheless, for non-elliptical dependence structures extensions are far from obvious and more research needs to be done. Further, methods to obtain multi-step forecasts of the dependence parameter have not been studied thoroughly in the literature, with the exception of the DCC-GARCH and the SCAR models, for which known results on autoregressive models can be used. Finally, goodness-of-fit techniques that help deciding which specification for the time-variation to chose need to be developed to avoid making to strong assumptions on the way dependence changes over time.

Acknowledgements

We thank Christian Hafner, Franz Palm and Jean-Pierre Urbain for helpful discussions and suggestions. O.Reznikova gratefully acknowledges support from the contract "Projet d'Actions de Recherche Concertées" nr. 07/12/002 of the "Communauté française de Belgique", granted by the "Académie universitaire Louvain" as well as from the IAP research network grant P 6/03 of the Belgian government (Belgian Science Policy). H.Manner gratefully acknowledges financial support from METEOR.

Appendix A: Examples of Copulas

In this appendix we introduce the most important families of copulas used and describe some of their properties. For more details we refer to Nelsen (2006).

Kendall's tau is a widely used rank correlation coefficient which can be directly represented by copulas. In general, for the jointly distributed but independent from each other variables $(U_{1i}, U_{2i}), i = 1 \ldots n$, the empirical Kendall's tau is given by

$$\tau^K = \frac{C_n - D_n}{0.5n(n-1)},$$

where C_n and D_n are the numbers of concordant and discordant pairs respectively. For a copula Kendall's tau can be shown to be:

$$\tau^K = 4 \int_0^1 \int_0^1 C(u_1, u_2) dC(u_1, u_2) - 1.$$

Upper and lower tail dependence coefficients can be interpreted as follows: for a pair of random variables U_1 and U_2 upper tail dependence means that for high values of U_1 we expect also high values of U_2. More precisely, for $(U_1, U_2) \in [0, 1]^2$ upper and lower tail dependence coefficients are defined as

$$\begin{aligned} \lambda^U &= \lim_{u_1 \to 1^-} P(U_1 > u_1 | U_2 > u_2) = \lim_{u_1 \to 1^-} \frac{1 - 2u_1 + C(u_1, u_1)}{1 - u_1} \\ \lambda^L &= \lim_{u_1 \to 0^+} P(U_1 \le u_1 | U_1 \le u_2) = \lim_{u_1 \to 0^+} \frac{C(u_1, u_1)}{u_1}, \end{aligned}$$

provided that the limit exists and $\lambda^U, \lambda^L \in [0, 1]$. If $\lambda^U = 0$ ($\lambda^L = 0$), then U_1 and U_2 are asymptotically independent in the upper (lower) tail.

Elliptical copulas are simply the copulas of elliptical distributions. They share a number of properties of the multivariate normal distribution. The most common example is the Gaussian copula, which can easily be derived from the bivariate normal distribution and has the following distribution function

$$C_{Gaussian}(u_1, u_2) = \int_{-\infty}^{\phi^{-1}(u_1)} \int_{-\infty}^{\phi^{-1}(u_2)} \frac{1}{2\pi\sqrt{1 - \rho^2}} \exp\left\{ -\frac{s^2 - 2\rho st + t^2}{2(1 - \rho^2)} \right\} ds\ dt,$$

where ρ is the linear correlation coefficient of the corresponding bivariate normal distribution. Note that it can be shown that the Gaussian copula does not have tail dependence. The expression for Kendall's tau is given by $\tau^K = \frac{2}{\pi} \arcsin(\rho)$.

Archimedean copulas form a large family of copulas with a number of convenient properties and they allow for a large number of dependence structures. Most have closed form expressions, which turns out to be very useful for estimation. Some of these copulas allow for both lower and upper tail dependence, others for only one of them or none. For transformations of Archimedean copulas, for which upper and lower tail dependence can have special forms we refer to Joe (1997). Archimedean copulas are, unlike many other copulas, not constructed from multivariate distributions using Sklar's theorem. Here we report the three most commonly used ones.

Clayton copula for $\theta > 0$ allows for lower tail dependence. The coefficient of lower tail dependence is given by $\lambda^L = 2^{-1/\theta}$, whereas $\lambda^U = 0$. The expression of Kendall's tau can be shown to be $\tau^K = \frac{\theta}{\theta+2}$. Its distribution function is

$$C_\theta^{Clayton}(u_1, u_2) = \max\left[(u_1^{-\theta} + u_2^{-\theta} - 1)^{\frac{-1}{\theta}}, 0\right].$$

Gumbel copula requires $\theta > 1$ and generates upper tail dependence with the coefficient $\lambda^U = 2 - 2^{1/\theta}$ and no lower tail dependence $\lambda^L = 0$. The Kendall's tau for the Gumbel copula is $\tau^K = 1 - \frac{1}{\theta}$. The distribution function is

$$C_\theta^{Gumbel}(u_1, u_2) = \exp\left(-[(-\log(u_1))^\theta + (-\log(u_2))^\theta]^{1/\theta}\right).$$

Frank copula displays the property of radial symmetry and does not have any tail dependence. The Kendall's tau coefficient is $\tau^K = 1 - \frac{4(1-D_1)(\theta)}{\theta}$, where D is the Debye function $D_k(x) = \frac{k}{x^k}\int_0^x \frac{t^k}{e^t-1}dt$. Its distribution function is

$$C_\theta^{Frank}(u_1, u_2) = -\frac{1}{\theta}\log\left(1 + \frac{(e^{-\theta u_1} - 1)(e^{-\theta u_2} - 1)}{e^{-\theta} - 1}\right).$$

The survival (rotated) copula. For a given copula $C(u_1, u_2)$ its survival copula $\widehat{C}(u_1, u_2)$ is defined as

$$\widehat{C}(u_1, u_2) = C(1 - u_1, 1 - u_2) + u_1 + u_2 - 1.$$

Its density is given by $\overline{c}(1 - u_1, 1 - u_2) = c(1 - u_1, 1 - u_2)$, so basically it is the original copula rotated by 180^o. A copula is called rotationally symmetric if it is equal to its survival copula.

References

T. W. Anderson and D. A. Darling. Asymptotic theory of certain goodness-of-fit criteria based on stochastic processes. *Annals of Mathematical Statistics*, 23:193–212, 1952.

D. W. K. Andrews. Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61:821–856, 1993.

M. Asai and M. McAleer. The structure of dynamic correlations in multivariate stochastic volatility models. *Journal of Econometrics*, 150:182–192, 2009.

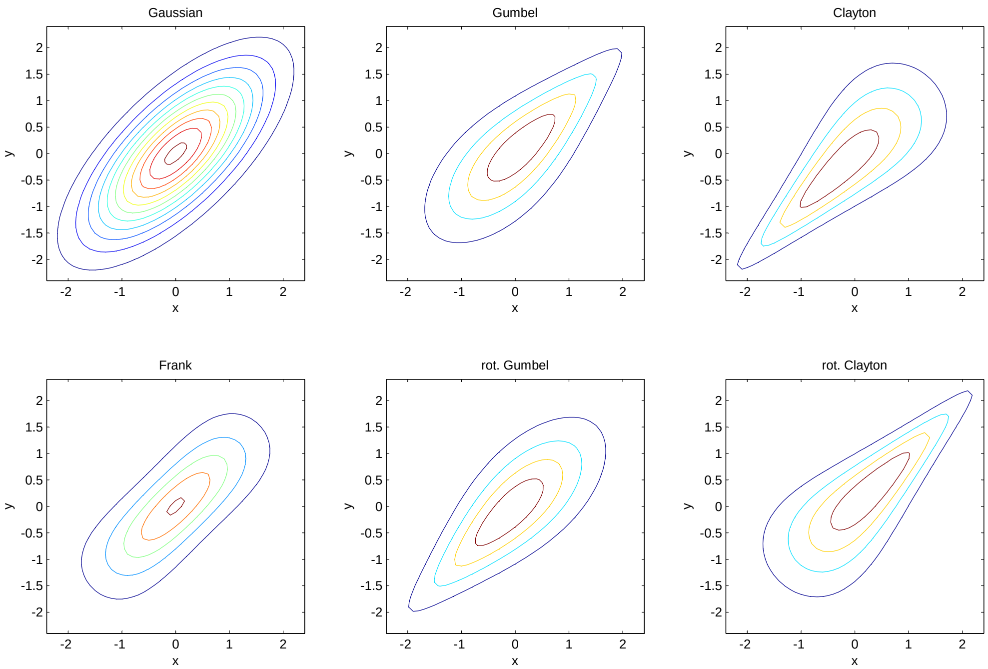

Figure 3: Contours plots of some copulas for the Kendall's tau equal 0.5

F. Busetti and A. Harvey. When is a copula constant? a test for changing relationships. Cambridge Working Papers in Economics 0841, 2008.

B. Candelon and H. Manner. Testing for asset market linkages: A new approach based on time-varying copulas. METEOR Research Memorandum RM/07/052, Maastricht University, 2007.

G. Cherubini, E. Luciano, and W. Vecchiato. *Copula Methods in Finance*. John Wiley and Sons Ltd, 2004.

L. Chollete, A. Heinen, and A. Valdesogo. Modeling international financial returns with a multivariate regime switching copula. CORE Discussion Paper 2008/13, 2008.

D. Creal, S. J. Koopman, and S. J. Lucas. A general framework for observation driven time-varying parameter models. Tinbergen Institute Discussion Paper 08-108/4, 2008.

A. Dias and P. Embrechts. Change-point analysis for dependence structures in finance and insurance. In G. Szegoe, editor, *Risk Measures for the 21st Century*, pages 321–335. Wiley Finance Series, 2004.

P. Embrechts, A. Höing, and A. Juri. Using copulae to bound value-at-risk for functions of dependent risks. *Finance and Stochastics*, 7:145–167, 2003.

R. F. Engle. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics*, 20:339–350, 2002.

R. F. Engle and S. Manganelli. CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics*, 22:367–381, 2004.

C. B. Erb, C. R. Harvey, and T. E. Viskante. Forecasting international equity correlations. *Financial Analysts Journal*, 50:32–45, 1994.

J. Fan, M. Farmen, and I. Gijbels. Local maximum likelihood estimation and inference. *Journal of the Royal Statistical Society B*, 60:591–608, 1998.

J.-D. Fermanian. Goodness of fit tests for copulas. *Journal of Multivariate Analysis*, 95: 119–152, 2005.

R. Garcia and G. Tsafack. Dependence structure and extreme comovements in international equity and bond markets with portfolio diversification effects. Working paper, EDHEC Risk Asset Management Research Centre, 2008.

C. Genest, K. Ghoudi, and L.-P. Rivest. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82:543–552, 1995.

E. Giacomini, W. Härdle, and V. Spokoiny. Inhomogeneous dependency modelling with time varying copulae. *Journal of Business and Economic Statistics*, 27:224234, 2009.

D. Guégan and J. Zhang. Change analysis of dynamic copula for measuring dependence in multivariate financial data. *Quantitative Finance*, forthcoming, 2009.

C. M. Hafner and H. Manner. Dynamic stochastic copula models: Estimation, inference and applications. METEOR Research Memorandum RM/08/043, Maastricht University, 2008.

C. M. Hafner and O. Reznikova. Efficient estimation of a semiparametric dynamic copula model. Manuscript, Institute of Statistics, UCL, 2008.

C. M. Hafner, D. van Dijk, and P. H. Franses. Semiparametric modelling of correlation dynamics. In T. Fomby and C. Hill, editors, *Advances in Econometrics*, volume 20, pages 59–103, Department of Economics, Louisiana State University, 2006.

J. Hamilton. *Time Series Analysis.* Princeton University Press, 1994.

A. Harvey. Dynamic distributions and changing copulas. Cambridge Working Papers in Economics 0839, 2008.

A. Heinen and A. Valdesogo. Asymmetric capm dependence for large dimensions: The canonical vine autoregressive copula model. available at ssrn: http://ssrn.com/abstract=1297506, CORE, 2008.

H. Joe. *Multivariate Models and Dependence Concepts.* Chapman & Hall/CRC, 1997.

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

D. X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, 9:43–54, 2000.

R. Liesenfeld and J. F. Richard. Univariate and multivariate volatility models: Estimation and diagnostics. *Journal of Empirical Finance*, 10:505–531, 2003.

F. Longin and B. Solnik. Is the correlation in international equity returns constant: 1960-1990. *Journal of International Money and Finance*, 14:3–26, 1995.

D. Mercurio and V. Spokoiny. Statistical inference for time-inhomogeneous volatility models. *Annals of Statistics*, 32:577–602, 2004.

R. B. Nelsen. *An Introduction to Copulas.* Springer, 2006.

A. Patton. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *Journal of Financial Econometrics*, 2:130–168, 2004.

A. Patton. Modelling asymmetric exchange rate dependence. *International Economic Review*, 47:527–556, 2006.

D. Pelletier. Regime switching for dynamic correlations. *Journal of Econometrics*, 131: 445–473, 2006.

J. F. Richard and W. Zhang. Efficient high-dimensional importance sampling. *Journal of Econometrics*, 141:1385–1411, 2007.

J. C. Rodriguez. Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, 14:401–423, 2007.

J. S. Simonoff. *Smoothing methods in statistics*. Springer series in statistics, 1996.

A. Sklar. Fonctions de repartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de L'Université de Paris*, 8:229–231, 1959.

Y. K. Tse and A. K. C. Tsui. A multivariate garch model with time-varying correlations. *Journal of Business and Economic Statistics*, 20:351–362, 2002.

R. W. J. van den Goorbergh, C. Genest, and B. J. M. Werker. Bivariate option pricing using dynamic copula models. *Insurance: Mathematics and Economics*, 37:101–114, 2005.

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, New York, 1998.

J. Yu and R. Meyer. Multivariate stochastic volatility models: Bayesian estimation and model comparison. *Econometric Reviews*, 25:361–384, 2006.