




Contents lists available at SciVerse ScienceDirect

Journal of Multivariate Analysis

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

A review of copula models for economic time series

Andrew J. Patton

Department of Economics, Duke University, 213 Social Sciences Building, Box 90097, Durham NC 27708–0097, USA

ARTICLE INFO

Article history:

Available online xxxx

AMS subject classifications:

62-02
62H20
62P20
91B84
62E20

Keywords:

Correlation
Inference
Multivariate models
Semiparametric estimation
Time series

ABSTRACT

This survey reviews the large and growing literature on copula-based models for economic and financial time series. Copula-based multivariate models allow the researcher to specify the models for the marginal distributions separately from the dependence structure that links these distributions to form a joint distribution. This allows for a much greater degree of flexibility in specifying and estimating the model, freeing the researcher from considering only existing multivariate distributions. The author surveys estimation and inference methods and goodness-of-fit tests for such models, as well as empirical applications of these copulas for economic and financial time series.

© 2012 Elsevier Inc. All rights reserved.

1. Introduction

This survey reviews some of the methods, models, and results for copula-based time series models, with an emphasis on economic and financial applications. Recall from Sklar [136] that an n -dimensional joint distribution can be decomposed into its n univariate marginal distributions and an n -dimensional copula. To be specific, let $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ be a random vector with cumulative distribution function F and, for $i \in \{1, \dots, n\}$, let F_i denote the marginal distribution of Y_i . Then there exists a copula $C : [0, 1]^n \rightarrow [0, 1]$ such that, for all $\mathbf{y} = (y_1, \dots, y_n) \in \mathbb{R}^n$,

$$F(\mathbf{y}) = C\{F_1(y_1), \dots, F_n(y_n)\}. \quad (1)$$

A copula C of the random vector \mathbf{Y} is thus a function that maps the univariate marginal distributions F_1, \dots, F_n to the joint distribution F , and we write $\mathbf{Y} \sim F = C(F_1, \dots, F_n)$.

For continuous random variables, the copula is unique. It is the joint distribution of the “probability integral transforms” (PITs) of the original variables, which are defined, for $i \in \{1, \dots, n\}$, as $U_i = F_i(Y_i)$. Then $\mathbf{U} = (U_1, \dots, U_n)^\top \sim C$. It is well known (see, e.g., [58]) that if Y_i is continuous, then $U_i \sim \mathcal{U}(0, 1)$, and in this case a copula is any joint distribution function with $\mathcal{U}(0, 1)$ margins. In this paper, we will assume that the variables of interest are continuous, simplifying some of the descriptions; see [67] for issues that arise when considering copulas for discrete random variables.

What makes the representation in (1) particularly useful for applied researchers is the converse of Sklar’s theorem: given any collection of univariate distributions F_1, \dots, F_n and any copula C , the function F defined by (1) above defines a valid joint distribution with marginal distributions F_1, \dots, F_n . For example, one might combine a Normally distributed variable with an Exponentially distributed variable via a Clayton copula, and obtain an unusual but valid bivariate distribution. The ability to combine marginal distributions with a copula model allows the researcher to directly draw on the large body of research on modeling univariate distributions, leaving only the task of modeling the dependence structure.

E-mail address: andrew.patton@duke.edu.

One particularly useful feature of this decomposition is that different estimation methods can be used for the different components of the model. The leading example of this is the use of a nonparametric estimator of the marginal distributions (usually the empirical distribution function) and a maximum likelihood estimator of the unknown parameter of a model for the copula. This semiparametric approach is considered in Genest et al. [66] and Shih and Louis [135] for i.i.d. data, and by Chen and Fan [25], Chen and Fan [26] and Rémillard [125] for time series data. Below we will review both fully parametric and semiparametric copula-based models.

Copula-based models have been considered for both univariate time series processes and for multivariate time series processes. In the former case, the structure in (1) generates a stationary Markov chain of order n , under certain conditions. See [39,26,29,95,10,139] for this and related results.

In multivariate time series applications, we use a version of Sklar's theorem for conditional joint distributions presented in Patton [120], where we consider some information set \mathcal{F}_{t-1} (for example, that generated by $\{\mathbf{Y}_{t-j}, j = 1, 2, \dots\}$), and decompose the conditional distribution of \mathbf{Y}_t given \mathcal{F}_{t-1} into its conditional marginal distributions and the conditional copula. For $t \in \{1, \dots, T\}$, let

$$\mathbf{Y}_t | \mathcal{F}_{t-1} \sim F(\cdot | \mathcal{F}_{t-1}), \quad Y_{it} | \mathcal{F}_{t-1} \sim F_i(\cdot | \mathcal{F}_{t-1}).$$

Then

$$F(\mathbf{y} | \mathcal{F}_{t-1}) = C\{F_1(y_1 | \mathcal{F}_{t-1}), \dots, F_n(y_n | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}\}. \quad (2)$$

Note that the same information set must be used in each of the marginals *and* for the copula in order for the resulting function to be a multivariate conditional joint distribution. When different information sets are used, the resulting function $F(\cdot | \cdot)$ is *not* generally a joint distribution with the specified conditional marginal distributions; see Fermanian and Wegkamp [57].

Several other surveys of copula theory and applications have appeared in the literature to date. Nelsen [112] and Joe [98] are two key textbooks on copula theory, providing clear and detailed introductions to copulas and dependence modeling, with an emphasis on statistical foundations. McNeil et al. [110] contains an overview of copula methods in the context of risk management, as does Denuit et al. [43]. Cherubini et al. [32] present an introduction to copulas using methods from mathematical finance, updated in [30] along with new results for convolution-based copulas. Choros et al. [34] provide a concise survey of estimation methods, both parametric and nonparametric, for copulas for both i.i.d. and time series data. Genest and Favre [65] present a description of semiparametric inference methods for i.i.d. data with a detailed empirical illustration, and Patton [122] reviews copula-based methods for economic forecasting, and includes detailed empirical examples illustrating some commonly-used methods.

In the sections that follow, we first consider estimation and inference for copula models for time series, covering both fully parametric and semiparametric models. We then review goodness-of-fit tests for copula-based models, and finally we present a brief survey of some of the numerous applications of copula-based models that have appeared in the economics and finance literature.

2. Estimation and inference

This section describes estimation and inference methods for copula-based models that have been proposed in the literature. We will specialize to the bivariate case for ease of exposition, and note any instances where the generalization to arbitrary n is not immediate. When considering copula-based models for multivariate time series, we will assume that the marginal distributions are of the form

$$Y_{it} = \mu_i(\mathbf{Z}_{t-1}; \boldsymbol{\phi}_i) + \sigma_i(\mathbf{Z}_{t-1}; \boldsymbol{\phi}_i) \varepsilon_{it},$$

where, for $i = 1, 2$,

$$\begin{aligned} \mathbf{Z}_{t-1} &\in \mathcal{F}_{t-1}, \quad \varepsilon_{it} | \mathcal{F}_{t-1} \sim F_{it}, \\ \boldsymbol{\varepsilon}_t &\equiv (\varepsilon_{1t}, \varepsilon_{2t})^\top | \mathcal{F}_{t-1} \sim F_{\varepsilon t} = C_t(F_{1t}, F_{2t}). \end{aligned}$$

That is, we will allow each series to have potentially time-varying conditional mean and variance, each parametrically modeled, and we will assume that the standardized residual, ε_{it} , has a conditional distribution F_{it} (with mean zero and variance one, for identification). When the distribution of the standardized residuals, F_{it} , is modeled parametrically, it may be modeled as time-varying; however when it is estimated nonparametrically it is assumed constant, and so $F_{it} = F_i$ for all t . Following Chen and Fan [26,125], we assume that the marginal distribution parameters, $\boldsymbol{\phi}_i$, are \sqrt{T} -consistently estimable. This is a mild assumption and allows for a wide variety of models for the conditional mean and variance. The estimated standardized residuals are then obtained, for $i = 1, 2$, as

$$\hat{\varepsilon}_{it} \equiv \frac{Y_{it} - \mu_i(\mathbf{Z}_{t-1}; \hat{\boldsymbol{\phi}}_i)}{\sigma_i(\mathbf{Z}_{t-1}; \hat{\boldsymbol{\phi}}_i)},$$

where $\hat{\boldsymbol{\phi}}_i$ is the vector of estimated parameters for the models of the conditional mean and conditional variance. We will consider both parametric and nonparametric models for F_{it} , and in the parametric case these parameters will become part of the vector $\boldsymbol{\phi}_i$. Many choices are possible for the parametric model for F_{it} , including the Normal, the standardized Student's t

(as in Bollerslev [17]), the skewed t (as in Patton [118]), and others. For the nonparametric estimate of F_{it} , we will assume that $F_{it} = F_i$ for all t , and use the empirical distribution function (EDF), and obtain the estimated probability integral transform variable, \hat{U}_{it} , viz.

$$\hat{F}_i(\varepsilon) \equiv \frac{1}{T+1} \sum_{t=1}^T \mathbf{1}(\hat{\varepsilon}_{it} \leq \varepsilon), \quad \hat{U}_{it} = \hat{F}_i(\hat{\varepsilon}_{it}). \quad (3)$$

2.1. Dependence summary statistics

Prior to the specification and estimation of a copula model, it is often of interest to compute some simple measures of dependence. Inference on estimated dependence measures can be conducted quite easily using a bootstrap. In the case of i.i.d. data, the validity of an i.i.d. bootstrap approach is presented in Fermanian et al. [55]. In multivariate time series applications when the conditional copula is constant, Rémillard [125] shows that the estimated parameters from the conditional mean and variance do not affect the asymptotic distribution of estimated dependence measures such as Spearman's rank correlation and Kendall's tau. This is a surprising result, and, under the conditions stated in [125], implies that we can ignore the error resulting from the estimation of the marginal distribution parameters. It further implies that we can use an i.i.d. bootstrap, applied to the estimated standardized residuals, $\hat{\varepsilon}_{it}$, for inference in time series applications. When the conditional copula is time-varying, the parameter estimation error from the models for the conditional mean and variance *cannot* generally be ignored (see [125]) and so an i.i.d. bootstrap cannot be used. To our knowledge, this case has not (yet) been considered in the literature.

For univariate copula-based time series models, one may be interested in such measures as rank autocorrelation. Ferguson et al. [54] and Genest and Rémillard [70] consider rank-based tests of independence for serially dependent data. Under stationarity, Gaißer et al. [62] suggest a block bootstrap approach to conduct inference on dependence measures for serially dependent data, and Ruppert [130] proposes a block multiplier technique for inference on the empirical copula process of time series data.

2.1.1. Time-varying dependence

The econometrics literature contains a preponderance of evidence that the conditional volatility of economic time series changes through time; see Andersen et al. [5] for example. This motivates us to consider whether the conditional dependence structure also varies through time. Before specifying a functional form for a time-varying conditional copula model (examples of which are reviewed in Section 4.4), it is informative to test for the presence of time-varying dependence.

As usual, tests for time-varying dependence will maintain a constant conditional copula under the null, and thus the asymptotic distribution theory from Rémillard [125] may be used to obtain the limiting distribution of the test statistics. Rémillard [125] considers a test for a one-time change (or “structural break”) in the copula at some time in the sample period, and his test statistic is similar to a Kolmogorov–Smirnov statistic comparing the empirical copula before and after a conjectured break date. The break date can be known or unknown in this framework. Gaißer et al. [62] consider testing for a change in the dependence structure by looking for a change in Hoeffding's Φ^2 dependence measure. One could naturally also consider testing for a change using alternative dependence measures, such as rank correlation.

2.1.2. Empirical illustration

Consider the daily returns on two equity indices: the S&P 100 index of the largest US firms and the S&P 600 index of small firms, over the period 17 August 1995–30 May 2011. An AR(2) model was found to be adequate for the conditional mean of the S&P 100 index, while simply a constant (AR(0)) was used for the S&P 600 index. For both indices, the GJR-GARCH model of Glosten et al. [78] was used to estimate the conditional volatility, and obtain the estimated standardized residuals. See Patton [122] for a more detailed description of this data and for more discussion of the steps involved in building a copula-based model for the conditional joint distribution.

The rank correlation between these two series of standardized residuals is estimated at 0.781, and a 90% i.i.d. bootstrap confidence interval is [0.769, 0.793], indicating that the dependence between these two series is positive and relatively strong. The upper panel of Fig. 1 presents the estimated “quantile dependence” plot, for $q \in [0.025, 0.975]$, along with 90% (pointwise) i.i.d. bootstrap confidence intervals. Lower quantile dependence is given by $\lambda_L^q = \Pr(U_1 \leq q, U_2 \leq q) / q$, for $q \in (0, 0.5)$ and the upper quantile dependence is given by $\lambda_U^q = \Pr(U_1 > q, U_2 > q) / (1 - q)$, for $q \in [0.5, 1)$. The lower panel of Fig. 1 presents the difference between the upper and lower portions of this plot, along with a pointwise confidence interval for this difference. As expected, the confidence intervals are narrower in the middle of the distribution (values of q close to 1/2) and wider near the tails (values of q near 0 or 1). This figure shows that observations in the lower tail are somewhat more dependent than observations in the upper tail, with the difference between corresponding quantile dependence probabilities being as high as 0.1. The confidence intervals show that these differences are borderline significant at the 0.10 level, with the upper bound of the confidence interval on the difference lying around zero for most values of q .

In univariate copula-based time series models, one may be interested in quantities such as rank autocorrelations, to measure the strength and sign of serial dependence of the series. As an illustration, Table 1 presents rank autocorrelations for the returns and squared returns on these two equity indices. The stationary block bootstrap of Politis and Romano [124],

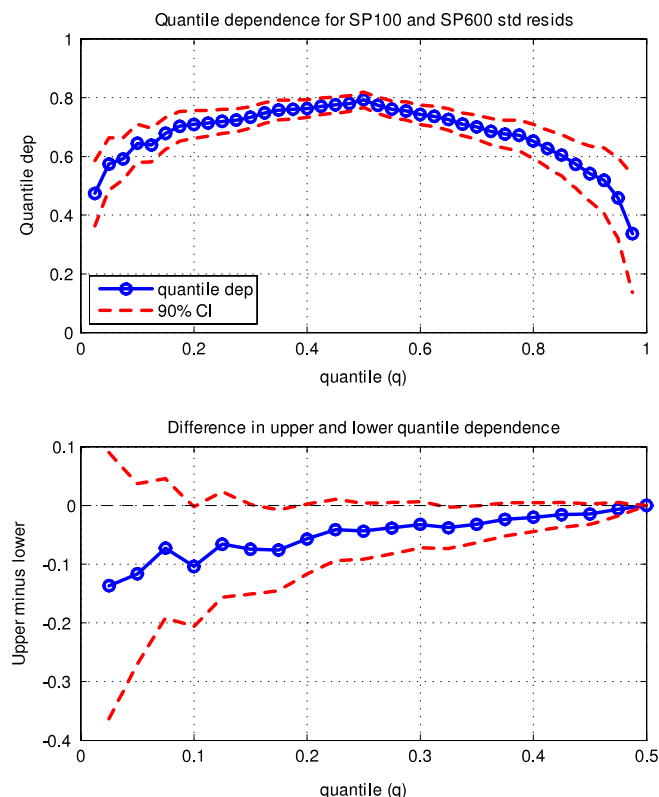


Fig. 1. The upper panel shows the estimated quantile dependence between the standardized residuals for the S&P 100 index and the S&P 600 index along with 90% bootstrap confidence intervals. The lower panel presents the difference between corresponding upper and lower quantile and tail dependence estimates, along with a 90% bootstrap confidence interval for this difference.

Table 1

Rank autocorrelation in returns and squared returns.

	S&P 100		S&P 600	
	Levels	Squares	Levels	Squares
1	−0.0506 (0.0196)	0.1313 (0.0449)	0.0312 (0.0253)	0.1473 (0.0539)
2	−0.0345 (0.0285)	0.2003 (0.0835)	−0.0327 (0.0203)	0.1878 (0.0670)
3	−0.0153 (0.0317)	0.1730 (0.0282)	0.0227 (0.0192)	0.1850 (0.0348)
4	0.0056 (0.0258)	0.1885 (0.0566)	0.0005 (0.0285)	0.1877 (0.0672)
5	−0.0224 (0.0249)	0.2160 (0.0531)	−0.0256 (0.0249)	0.1741 (0.0534)
6	−0.0055 (0.0247)	0.1864 (0.0686)	−0.0345 (0.0277)	0.1838 (0.0468)
7	−0.0230 (0.0262)	0.1923 (0.0515)	−0.0110 (0.0224)	0.1743 (0.0603)
8	−0.0125 (0.0361)	0.1810 (0.0356)	−0.0011 (0.0269)	0.1647 (0.0277)
9	0.0037 (0.0205)	0.1671 (0.0665)	−0.0111 (0.0184)	0.1555 (0.0512)
10	0.0317 (0.0243)	0.1985 (0.0581)	−0.0106 (0.0226)	0.1633 (0.0507)

Notes: This table presents sample rank autocorrelations, for lags 1–10, for daily returns and squared daily returns on the S&P 100 and S&P 600 equity indices. A block bootstrap with block length of 30 observations is used to obtain the standard errors reported in parentheses below the estimates.

with an average block length of 30 observations, is used to obtain the standard errors on each of these estimates. We find significant negative rank autocorrelation at the first lag for the SP100 index, but no significant rank autocorrelations for the SP600 index. Consistent with a large literature on volatility clustering in asset returns, we find strongly significant positive rank autocorrelations in squared returns, for all lags between 1 and 10.

A time series plot of rolling 60-day rank correlation, along with pointwise bootstrap standard errors (correct only under the null that this correlation is not changing), is presented in Fig. 2. This figure shows that the rank correlation hovered around 0.6–0.7 in the early part of the sample, rising to around 0.9 during the financial crisis of 2008–09. Having no *a priori*




Fig. 2. This figure shows the rank correlation between the standardized residuals for the S&P 100 index and the S&P 600 index over a 60-day moving window, along with 90% bootstrap confidence intervals.

Table 2
Testing for time-varying dependence.

	Break date (t/T)			Anywhere
	0.15	0.50	0.85	
p -val	0.667	0.373	0.045	0.269

Notes: This table presents bootstrap p -values from tests for a change in the rank correlation between the standardized residuals for the S&P 100 index and the S&P 600 index.

dates to consider for the timing of a break, for illustration we consider three tests for a break at three different points in the sample, namely at $t^*/T \in \{0.15, 0.50, 0.85\}$, which correspond to the dates 23-Dec-1997, 7-July-2003, and 8-Jan-2009. Table 2 presents p -values from these tests, based on a bootstrap method proposed by Rémillard [125]. Only for the latter date is evidence of a break in rank correlation found, with a p -value of 0.045. Thus it appears that the rank correlation towards the end of the sample is different from that during the earlier part of the sample. However, given a lack of a reason for choosing a break date of 8-Jan-2009, a more appropriate test is one where the break date is treated as unknown and *estimated*, and using that test the p -value is 0.269, indicating no evidence against a constant rank correlation in the direction of a one-time break.

2.2. Estimating copula-based multivariate time series models

The majority of applications of copula models for multivariate time series build the model in stages, with the models for the marginal distributions (means, variances, and distribution of the standardized residual, ε_{it}) estimated separately from the copula model. The conditional distribution of ε_{it} is treated in one of two ways, either parametrically or nonparametrically. In the former case, this distribution may vary through time as a (parametric) function of \mathcal{F}_{t-1} -measurable variables (e.g., the time-varying skewed t distribution of Hansen [86]), or may be constant. In the nonparametric case, the majority of the literature assumes that the conditional distribution is constant and estimable via the empirical distribution function (EDF). The choice of a parametric or nonparametric model for the distribution of the standardized residuals leads to different inference procedures for the copula parameters.

2.2.1. Fully parametric

When all components of the multivariate model are parametric, the most efficient estimation method is maximum likelihood. Under regularity conditions, see White [145] for example, standard results for parametric time series models can be used to show that the MLE is \sqrt{T} -consistent and asymptotically Normal, and a consistent estimator of the asymptotic covariance matrix can also be obtained using standard methods. The drawback of this approach is that even for relatively simple bivariate models, the number of parameters to be estimated simultaneously can be large, creating a computational burden. This burden is of course even greater in higher dimensions.

A more common approach is to estimate the model in stages. This requires that the parameter vector can be separated in parameters for the first margin, second margin, and the copula, which is often satisfied for models used in practice. In this case one can estimate the parameters of the marginal distributions separately, and then estimate the copula parameters conditioning on the estimated marginal distribution parameters, greatly simplifying the estimation problem. This estimation method is sometimes called “inference functions for margins” in this literature, see Joe [98] and Joe and Xu [99], though

more generally this is known as multi-stage maximum likelihood (MSML) estimation. Clearly, MSMLE is asymptotically less efficient than one-stage (full) MLE (except in the special case that the variables are independent). Simulation studies in Joe [97] and Patton [119], however, indicate that this loss is not great in many cases. As for one-stage MLE, under regularity conditions, see White [145] or Patton [119], the MSMLE is asymptotically normal but the asymptotic covariance matrix now takes a non-standard form. See Patton [119,122] for details on how to estimate this matrix. Newey and McFadden [113] discuss a one-step adjustment of the MSMLE that achieves full efficiency, and Song et al. [140] present an iterative multi-stage estimation procedure that achieves full efficiency.

2.2.2. Semiparametric

An attractive feature of the copula decomposition of a joint distribution is that it allows the marginal distributions and copula to be estimated separately, potentially via different methods. Semiparametric copula-based models exploit this feature and employ a nonparametric model for the marginal distributions and a parametric model for the copula. In such cases, the estimation of the copula parameter is usually conducted via maximum likelihood and in this case this estimator is sometimes called “canonical maximum likelihood” in this literature. It has also been called “pseudo maximum likelihood”, see [66,102], though their use of this phrase is different from its use in the econometrics literature; see [80]. The asymptotic distribution of this estimator was studied by Genest et al. [66] and Shih and Louis [135] for i.i.d. data and by Chen and Fan [25] and Chen and Fan [26] for time series data.

The difficulty relative to the parametric case is that the copula likelihood depends on the infinite-dimensional parameters F_1, \dots, F_n , as well as the marginal distribution parameters α . Standard maximum likelihood methods cannot be applied. Chen and Fan [25] provided conditions under which an asymptotic normal distribution is obtained, and provide a method for estimating the asymptotic covariance matrix. See Chen et al. [23] for a detailed proof.

A surprising feature of the result of Chen and Fan [25] is that the asymptotic variance of the MLE of the copula parameter depends on the estimation error in the EDF but *not* on the estimated parameters in the marginal distributions. Thus in this case the researcher can estimate the models for the conditional means and variances, compute the standardized residuals, and then *ignore*, for the purposes of copula estimation and inference, the estimation error from the mean and variance models.

Two important caveats are worth noting here: Firstly, this only applies for *constant* conditional copula models; if the conditional copula is time-varying, then Rémillard [125] shows that the estimation error from the models for the conditional mean and variance will affect the asymptotic distribution of the copula parameter estimate. Second, this only applies when the marginal distributions of the standardized residuals are estimated nonparametrically; with parametric marginal distribution models the estimation error from the models for the conditional mean and variance will, in general, affect the distribution of the copula parameter estimate, and methods from Section 2.2.1 should be used. Chen et al. [27] show that efficient estimation of this semiparametric copula-based model can be obtained by using the method of sieves for the marginal distributions and estimating these along with the copula parameter in a single estimation step.

2.2.3. Nonparametric

Fully nonparametric estimation of the copula in the i.i.d. case was studied by Genest and Rivest [73] and Genest et al. [69] for Archimedean copulas, and by Genest and Segers [74] for extreme value copulas. Nonparametric copula estimation using time series data is studied by Fermanian and Scaillet [56], Fermanian et al. [55], Sancetta and Satchell [132] and Ibragimov [95].

2.2.4. Other estimation methods

While maximum likelihood estimation is the most prevalent in the literature, other methods have been considered. Method of moments-type estimators, where the parameter of a given family of copulas has a known, invertible, mapping to a dependence measure (such as rank correlation or Kendall's tau) are considered in [64,75,125,68], among others. Generalized method of moments estimation, where the number of dependence measures may be greater than the number of unknown parameters, and an analogous simulation-based method, are considered in Oh and Patton [115]. Minimum distance estimation is considered by Tsukahara [143]. Bayesian estimation of copula models is considered in [111,123,139,138]; see Smith [137] for a review.

2.2.5. Empirical illustration, continued

To illustrate some of the different estimation methods described above, we consider now the problem of estimating a parametric model for the copula of the standardized residuals on the SP100 and SP600 indices, described in Section 2.1.2. To illustrate a fully parametric model we combine the AR-GARCH models described above with the skewed t distribution of Hansen [86] for the marginal distributions. For the semiparametric model we will use the EDF to estimate the marginal distributions.

A variety of different standard errors are computed for the estimated parameters, implementation details of which are presented in Patton [122]. For both the parametric and semiparametric cases, we will consider (i) naïve standard errors, where the estimation error from the earlier stages of estimation (AR, GARCH and marginal distributions) is ignored; (ii) multi-stage MLE or multi-stage semiparametric MLE (MSML) standard errors, using the asymptotic distribution theory

Table 3
Standard errors on estimated copula parameters—multivariate model.

		Parametric				Semi-parametric		
		Naïve	MSML	Boot	Sim.	Naïve	MSML	Boot
Normal	$\hat{\rho}$			0.7959			0.7943	
	s.e.	0.0046	0.0106	0.0099	0.0062	0.0046	0.0061	0.0065
	$\ln \mathcal{L}$			1991.8			1978.3	
Clayton	$\hat{\kappa}$			2.0279			2.0316	
	s.e.	0.0451	0.0951	0.0862	0.0664	0.0449	0.0545	0.0580
	$\ln \mathcal{L}$			1720.5			1723.1	
Rotated Gumbel	$\hat{\kappa}$			2.3715			2.3673	
	s.e.	0.0310	0.0603	0.0595	0.0386	0.0309	0.0421	0.0344
	$\ln \mathcal{L}$			2013.6			2008.4	
Student's t	$\hat{\rho}$			0.8019			0.8005	
	s.e.	0.0053	0.0100	0.0096	0.0070	0.0053	0.0055	0.0054
	$\hat{\nu}^{-1}$			0.1455			0.1428	
	$\ln \mathcal{L}$	0.0172	0.0202	0.0222	0.0186	0.0172	0.0182	0.0169
				2057.4			2041.9	

Note: This table presents the estimated parameters of four different copula models for the standardized residuals for the S&P 100 index and the S&P 600 index, when the marginal distributions are estimated parametrically (left panel) or nonparametrically (right panel). For the parametric model four different estimators of the standard error on the estimated parameter are presented, and for the semiparametric model three different standard errors are presented. For all models the log-likelihood at the estimated parameter is also presented.

for these estimators in Patton [119] or Chen and Fan [25], respectively; (iii) bootstrap standard errors, based on either a stationary block bootstrap Politis and Romano [124] of the original returns and estimation of all stages on the bootstrap sample (parametric case), based on Gonçalves and White [79] (Joe [98], suggests a jackknife method for this), or an i.i.d. bootstrap of the standardized residuals and estimation only of the EDF and the copula (semiparametric case), based on Chen and Fan [25] and Rémillard [125].

For the parametric case we can consider one further type of standard error, namely a simulation-based standard error, where the model is simulated many times using the estimated parameters, and on each of the simulated samples the parameters are re-estimated. With a sufficient number of simulations, this latter approach yields correct *finite-sample* standard errors, unlike the other approaches which are all based on asymptotic arguments. For the bootstrap and the simulation-based standard errors 1000 replications are used. For the parametric model, the computation times for these (using MATLAB R2011a, on a 3 GHz machine) were 1.5 and 9.4 h respectively. The semiparametric model bootstrap standard errors had computation time of 11 min. The results are presented in Table 3.

Table 3 shows that the naïve standard errors are too small relative to the correct MSML standard errors, a predictable outcome given that naïve standard errors ignore the additional estimation error arising from the estimation of marginal distribution parameters. In the parametric case the naïve standard errors are on average about half as large as the MSML standard errors (average ratio is 0.54), while for the semiparametric case the ratio is 0.84. The relatively better performance in the semiparametric case is possibly attributable to the fact that the MSML standard errors in that case can, correctly, ignore the estimation error coming from the AR–GARCH models for the conditional mean and variance; only the estimation error from the EDF needs to be accounted for. In the fully parametric case, the marginal distribution shape parameters and the parameters of the AR–GARCH models must be accounted for.

In both the parametric and the semiparametric cases, the bootstrap standard errors are very close to the MSML standard errors, with the ratio of the former to the latter being 0.98 and 0.97 respectively. This is what we expect asymptotically, and confirms that the researcher may use either “analytical” MSML standard errors or more computationally-intensive bootstrap standard errors for inference on the estimated copula parameters.

In the parametric case, where simulation-based standard errors can be computed, we see that these are smaller than the MSML and bootstrap standard errors, with the average ratio being around 0.7. Asymptotically we expect this ratio to go to 1, but in finite samples this value of ratio will depend on the particular model being used. While Table 3 shows the ratio to be less than 1 for all four models considered, this is *not* a general result and need not hold in other applications.

2.3. Estimating copula-based univariate time series models

Consider the following model for a stationary first-order Markov process:

$$(Y_t, Y_{t-1})^\top \sim F = C(F, F).$$

Fully parametric models of this sort are considered in Chapter 8 of Joe [98], and under regularity conditions on the copula the estimation of such models is straightforward using maximum likelihood methods. Chen and Fan [26] consider the estimation of semi-parametric copula-based Markov models, where the marginal distributions are estimated nonparametrically while the copula is assumed to belong to some parametric family and estimated in a second stage via ML. With an estimate

Table 4

Standard errors on estimated copula parameters—univariate model.

		Naïve	MSML	Sim.
Normal	$\hat{\rho}$		−0.0566	
	s.e.	0.0159	0.0159	0.0153
	$\ln \mathcal{L}$		6.3758	
Half Rotated Clayton	$\hat{\kappa}$		0.1292	
	s.e.	0.0189	0.0174	0.0197
	$\ln \mathcal{L}$		30.6174	
Half Rotated Gumbel	$\hat{\kappa}$		1.0665	
	s.e.	0.0100	0.0089	0.0106
	$\ln \mathcal{L}$		38.8145	
Student's t	$\hat{\rho}$		−0.0534	
	s.e.	0.0180	0.0160	0.0175
	$\hat{\nu}^{-1}$		0.2070	
	s.e.	0.0194	0.0244	0.0249
	$\ln \mathcal{L}$		76.8829	

Note: This table presents the estimated parameters of four different copula models for a semiparametric first-order Markov model of daily returns on the S&P 100 index. Three different standard errors are presented. For all models the log-likelihood at the estimated parameter is also presented.

of the copula parameter and the marginal CDFs, it is then possible to compute (and conduct inference on) functionals of these, such as conditional moments and conditional quantiles. Chen and Fan [26] present conditions on the copula under which the above Markov process is β -mixing, and, with further regularity conditions, establish the asymptotic Normality of the estimated copula parameter. Other work on the dependence properties of copula-based Markov processes includes Beare [10], Beare [11], Beare and Seo [12], Bouyé and Salmon [20], Chen et al. [28], Chen et al. [29], Gagliardini and Gouriéroux [61], Ibragimov [95], and Ibragimov and Lentzas [96].

Papers that consider both the temporal *and* the cross-sectional dependence via copulas include Abegaz and Naik-Nimbalkar [2], Yi and Liao [146] and Rémillard et al. [126]. Papers that consider copula-based models for longitudinal data include Frees and Wang [59], Sun et al. [141], and Smith et al. [138]. Bonhomme [18] use copulas to model the dynamics in an earnings panel data set.

2.3.1. Empirical illustration, continued

Consider now the estimation of a semiparametric first-order Markov copula model for the S&P 100 index returns described in Section 2.1.2. As in Chen and Fan [26], we use the EDF to estimate the marginal distribution of these returns. Given the rank autocorrelation estimates presented in Table 1, it is clear that we should allow for *negative* dependence between consecutive returns. This is accommodated easily by some copula models (such as the Normal and Student's t , where the correlation parameter can be either positive or negative) but requires more work for other copula models, such as the commonly-used Clayton and Gumbel–Hougaard copulas. To accommodate negative dependence with a Clayton copula, one may consider a “half rotation” of the copula (where we “flip” just the first variable) so that

$$(1 - U_t, U_{t-1})^\top \sim C_{\text{Clayton}}(\kappa).$$

Larger values of κ in this model imply stronger negative dependence, and with this rotation it implies tail dependence in the second quadrant (i.e., as $u_{t-1} \rightarrow 0$ and $u_t \rightarrow 1$) rather than the third quadrant as for the usual Clayton copula. The other “half rotation” of this copula also implies negative dependence, but with tail dependence in the fourth quadrant (as $u_{t-1} \rightarrow 1$ and $u_t \rightarrow 0$). Estimating both of these on the data indicated strongly that the first rotation is preferred (log-likelihood of 30.6 vs 2.2), and both of these are preferred to the original Clayton copula. A similar rotation improves the fit of the Gumbel–Hougaard copula on this data, though given that the original Gumbel–Hougaard copula implies tail dependence in the joint upper tail (rather than the joint lower tail for the Clayton copula) a rotation of the u_{t-1} variable provides the better fit. That is,

$$(U_t, 1 - U_{t-1})^\top \sim C_{\text{Gumbel}}(\kappa).$$

These two copulas, along with the Normal and Student's t copulas, are estimated on the S&P 100 returns and the results are presented in Table 4. The values of the log-likelihood indicate that the Student's t copula is preferred, with a log-likelihood of 76.9, compared with the next best model, the “half-rotated” Gumbel–Hougaard with a log-likelihood of 38.8.

Table 4 also presents three different estimates of the standard errors for the copula parameters. The first is a “naïve” estimate based solely on the inverse Hessian of the copula likelihood, where the estimation error from the use of the EDF in the first stage is ignored. The second estimate is that of [26], which does account for the estimation error due to the EDF. The third estimate is based on 1000 simulations of this copula model, and which also incorporates the impact of estimation error from the EDF. (The computation time for these standard errors was 2.3 h.) Unlike the naïve standard errors in the multivariate application considered in Section 2.2.5, which were found to be too small relative to the correct standard errors, in this univariate application the naïve standard errors are not very different from the correct standard errors.

Of course, it should be noted that this need not be true in other applications, and the current application is perhaps special in that the overall dependence between consecutive returns is relatively weak (first order rank autocorrelation of just -0.05).

3. Goodness-of-fit tests

In this section we consider the problem of goodness-of-fit (GoF) testing. A GoF test looks for evidence that the copula is misspecified, i.e., different in some way from the unknown true copula. As for parameter estimation, inference for GoF tests differ depending on whether the model under analysis is parametric or semiparametric, and we will consider these two cases separately. We will focus on in-sample (full sample) tests of GoF; see Chen [24] for analysis of out-of-sample GoF tests.

Two tests that are widely used for GoF tests of copula models are the Kolmogorov–Smirnov (KS) and the Cramér–von Mises (CvM) tests, see Rémillard [125], both of which are based on comparing the fitted copula CDF to the empirical copula

$$\hat{C}_T(\mathbf{u}) \equiv \frac{1}{T} \sum_{t=1}^T \prod_{i=1}^n \mathbf{1}(\hat{U}_{it} \leq u_i),$$

where \hat{U}_{it} is defined in Eq. (3). An alternative, related, GoF test is based on Rosenblatt’s transform, which is a multivariate “probability integral transformation”; see [45,125]. In this approach the data are first transformed so that, if the model is correct, the data are independent $\mathcal{U}(0, 1)$ random variables, and then KS and CvM tests are applied to the transformed data. GoF tests that use the empirical copula of the data rely on the assumption that the true conditional copula (defined in Eq. (2)) is constant, and so are inappropriate for time-varying copula models. GoF tests based on the Rosenblatt transform can be used to test both constant and time-varying copula models, as well as copula-based Markov models.

Genest et al. [72] provide a comprehensive review of the many copula GoF tests available in the literature, and compare these tests to tests of their own via a simulation study. Across a range of data generating processes, they conclude that a Cramér–von Mises test (applied to the empirical copula or to the Rosenblatt transform of the original data) is the most powerful, a finding that is supported by Berg [14] who considers some further tests. Berg and Quesy [15] study the local power of copula GoF tests.

3.1. Fully parametric

For fully parametric copula-based models, GoF testing is a relatively standard problem, as these models are simply non-linear time series models; see the review article on evaluating predictive densities by Corradi and Swanson [37], and Bontemps et al. [19] and Chen [24] on GoF tests for multivariate distributions via moment conditions.

A difficulty in obtaining critical values for GoF test statistics, such as the KS and CvM test statistics, is that they depend not only on the estimated copula parameter, but also on the estimated marginal distribution parameters. As discussed in the context of obtaining standard errors on estimated copula parameters, the parameter estimation error coming from the marginal distributions cannot in general be ignored. GoF tests can be implemented in various ways, but for fully parametric models a simple (computationally intensive) procedure is always available, based on simulating from the model, estimating the parameters of the marginal distributions and the copula on the simulated data, and then computing the GoF test statistics. Repeating this many times provides the distribution of the test statistic under the null that the model is correct, and from this distribution a p -value can be obtained; see Genest and Rémillard [71] for example.

3.2. Semi-parametric

Rémillard [125] considers GoF tests for semi-parametric copula-based multivariate models for time series, and shows the surprising and useful result that the asymptotic distributions of GoF copula tests are unaffected by the estimation of marginal distribution parameters (as was the case for the asymptotic distribution of the estimated copula parameters, shown by Chen and Fan [26]). The estimation error coming from the use of the empirical distribution functions *does* matter, and he proposes a simple simulation-based method that is similar, but simpler, than the simulation for the fully parametric case described above. Chen and Fan [26] describe a similar simulation-based method that can be used to obtain critical values for GoF tests of semi-parametric copula-based Markov models.

3.3. Empirical illustration, continued

In Table 5 we present the results of GoF tests applied to the four copula models estimated on the standardized residuals of the S&P 100 and S&P 600 stock index returns, described in Section 2.2.5. The subscripts “C” and “R” on the column headings indicate whether the test was applied using the empirical copula directly or the empirical copula of the Rosenblatt transforms. In all cases a simulation-based approach was used, as described above, with 100 simulations. For the fully parametric models we see that the Clayton copula is rejected by all four tests. The Normal copula is rejected by both tests using a Rosenblatt transform, while the rotated Gumbel–Hougaard and Student’s t copulas are each rejected by one out of the four tests. For the semiparametric models all four copulas are strongly rejected by all four tests. Thus we do not find empirical support for any of these copulas.

Table 5

Goodness-of-fit tests for multivariate copula models.

	KS_C	CvM_C	KS_R	CvM_R
Parametric				
Normal	0.10	0.09	0.00	0.00
Clayton	0.00	0.00	0.00	0.01
Rot. Gumbel	0.09	0.02	0.09	0.06
Student's t	0.00	0.13	0.04	0.07
Semi-parametric				
Normal	0.00	0.00	0.00	0.00
Clayton	0.00	0.00	0.00	0.01
Rot. Gumbel	0.00	0.00	0.02	0.00
Student's t	0.00	0.00	0.02	0.00

Note: This table presents the p -values from various tests of goodness-of-fit for four different copula models for the standardized residuals for the S&P 100 index and the S&P 600 index, when the marginal distributions are estimated parametrically (top panel) or nonparametrically (lower panel). KS and CvM refer to the Kolmogorov–Smirnov and Cramér–von Mises tests, respectively. The subscripts C and R refer to whether the test was applied to the empirical copula of the standardized residuals, or to the empirical copula of the Rosenblatt transform of these residuals. The p -values are based on 100 simulations.




Fig. 3. This figure shows the sample quantile dependence between the standardized residuals for the S&P 100 index and the S&P 600 index and 90% bootstrap confidence intervals (shaded), as well as the quantile dependence implied by four copula models. The vertical axis measures the probability of one variable lying below (for $q < 0.5$), or above (for $q > 0.5$), its q quantile, given that the other variable also lies below/above its q quantile.

It is at this point that a drawback of “blanket” tests, such as KS and CvM tests, becomes clear: these tests are consistent (in that they have power against all possible deviations from the null of correct specification), but upon rejecting a model they provide no guidance on the direction in which the model might be improved. An alternative approach to GoF testing involves looking for misspecification in particular directions, such as tests based on particular moments or measures of dependence (see Bontemps et al. [19] for example), which may be more informative but are not consistent.

To try to see why these models are all rejected, one might look at the quantile dependence plot implied by each of these models, compared with the sample quantile dependence plot. This is presented in Fig. 3. From this figure we can see that the Clayton copula is too asymmetric compared with the data, and this is also the case for the rotated Gumbel–Hougaard copula. The Normal and Student's t copulas appear to fit the sample upper quantile dependence estimates, but are (slightly) too low in the lower tail. Thus it appears that a more flexible copula, allowing for some asymmetry but not imposing too much, may be required for this data.

In Table 6 we report GoF tests for the copula models of semiparametric first-order Markov copula model for the S&P 100 index returns described in Section 2.3.1. In this case we focus only the KS and CvM tests based on the Rosenblatt transform, given the time series nature of the model. We again use a simulation-based method to obtain critical values, based on 100 simulations. Table 6 shows that the Normal and “half rotated” Clayton copulas are both rejected, while neither the “half rotated” Gumbel–Hougaard nor the Student's t copula models can be rejected using either test.

Table 6

Goodness-of-fit tests for univariate copula models.

	KS_R	CvM_R
Normal	0.02	0.07
HRot. Clayton	0.01	0.03
HRot. Gumbel	0.13	0.11
Student's t	0.22	0.42

Note: This table presents the p -values from various tests of goodness-of-fit for four different semiparametric first-order Markov copula models for daily returns on the S&P 100 index. KS and CvM refer to the Kolmogorov–Smirnov and Cramér–von Mises tests, respectively. The p -values are based on 100 simulations.

4. Applications of copula methods for economic time series

The past decade has witnessed an ever-growing array of applications of copula methods in empirical economic research, driven by wide-ranging evidence against the assumption of a Normal copula (a benchmark model) for many economic variables, particularly financial asset returns. For example, without actually drawing on copula theory, Erb et al. [53], Longin and Solnik [106], Ang and Chen [6] and Bae et al. [8] all document evidence against the Normal copula. In this section we briefly review some of the applications of copulas for economic time series, divided into broad groups according to the nature of the application, and then review two of the frontiers of research on copulas for economic time series, namely models for time-varying conditional copulas, and models for high-dimensional copulas.

4.1. Risk management

One of the first areas of application of copulas in economics and finance was risk management. The focus of risk managers on Value-at-Risk (VaR), and other measures designed to estimate the probability of large losses, leads to a demand for flexible models of the dependence between sources of risk. See Komunjer [103] for a recent review of VaR methods. Hull and White [94], Cherubini and Luciano [31], Embrecht et al. [50], Embrechts et al. [49] and Embrechts and Hoing [48] study the VaR of portfolios. Rosenberg and Schuermann [129] use copulas to consider “integrated” risk management problems, where market, credit and operational risks must be considered jointly. McNeil et al. [110] and Alexander [4] provide clear and detailed textbook treatments of copulas and risk management.

4.2. Derivative contracts

Another early application of copulas was to the pricing of credit derivatives (credit default swaps and collateralized debt obligations, for example), as these contracts routinely involve multiple underlying sources of risk. Li [105] was first to use copulas in a credit risk application (and is behind the headline of “the formula that killed Wall Street”). See also [60,133,77,92] for applications to default risk. Duffie [47] argues that copulas are too restrictive for certain credit risk applications. Applications of copulas in other derivatives markets include Cherubini et al. [32], Rosenberg [128], Bennett and Kennedy [13], van den Goorbergh et al. [144], Salmon and Schleicher [131], Grégoire et al. [82] and Taylor and Wang [142].

4.3. Portfolio decision problems

Considering a portfolio decision problem in its most general form involves finding portfolio weights that maximize the investor's expected utility, and thus requires a predictive multivariate distribution for the assets being considered. Applications of copulas in portfolio problems include Patton [118], who considers a bivariate equity portfolio problem using time-varying copulas; Hong et al. [93] consider an investment decision involving eleven equity portfolios under “disappointment aversion” preferences; Christoffersen and Langlois [36] consider portfolio decisions involving four common equity market factors; and Garcia and Tsafack [63] consider portfolio decisions involving stocks and bonds in two countries.

4.4. Time-varying copula models

As noted earlier in this survey, the econometrics literature contains a wealth of evidence that the conditional volatility of economic time series changes through time, motivating the consideration of models that also allow the conditional copula to vary through time. Various models have been proposed in the literature to date. Patton [117,118,120], Jondeau and Rockinger [100], Ausin and Lopes [7], Christoffersen et al. [35] and Creal et al. [38] consider models of time-varying copulas where the copula functional form is fixed and its parameter is allowed to vary through time as a function of lagged information, similar to the famous ARCH model for volatility, see Engle [51] and Bollerslev [16]. “Stochastic copula” models, analogous to stochastic volatility models, see Shephard [134], were proposed by Hafner and Manner [84] and further studied

in Manner and Segers [108]. “Locally constant” copula models are considered by Giacomini et al. [76], Guégan and Zhang [83], Dias and Embrechts [44], Harvey [87], Rémillard [125] as well as in Busetti and Harvey [22]. Regime switching models, as in Hamilton [85], for the conditional copula allow the functional form of the copula to vary through time and are considered by Rodriguez [127], Okimoto [116], Chollete et al. [33], Markwat et al. [109], Garcia and Tsafack [63]. See Manner and Reznikova [107] for a survey specifically focused on time-varying copula models.

4.5. High-dimension copula applications

Early applications of copulas in economics were almost all bivariate in nature, however much recent work has focused on how to handle higher dimension applications of copulas. While bivariate and low dimension ($d < 10$) applications are still prevalent, many authors now consider dimensions greater than this, up to around one hundred variables. For example, Daul et al. [40] proposed a “grouped t ” copula and show that this copula can be used in applications of up to 100 variables. Hofert and Scherer [92] and Hering et al. [90] consider nested Archimedean copulas for modeling credit default swaps on 125 companies; see also Hofert et al. [91] for a study of numerical issues associated with estimation in this context. Smith et al. [138] use a skew t copula to model groups of up to 15 variables, and Christoffersen et al. [35] combine the skew t copula of Demarta and McNeil [42] with DCC Engle [52] dynamics for correlations in their study of 33 developed and emerging equity market indices. Multivariate “vine” copulas (or “pair copula constructions”) are constructed by sequentially applying bivariate copulas to build up a higher dimension copula, see [1,89,111] for example; see Acar et al. [3] for an important critique of vine copulas. Oh and Patton [114] apply a new class of “factor copulas” to a collection of 100 daily stock returns.

4.6. Other applications

There are several noteworthy economic applications of copulas that do not neatly fit into one of the above categorizations. Breyman et al. [21] and Dias and Embrechts [44] study the copulas of financial assets using intra-daily data sampled at different frequencies; Granger et al. [81] use copulas to provide a definition of a “common factor in distribution” for macroeconomic time series; Bartram et al. [9] use a time-varying conditional copula model to study financial market integration between seventeen European stock market indices; Heinen and Rengifo [88] use copulas to model multivariate time series of counts; Rodriguez [127] uses copulas to study financial contagion; Dearden et al. [41] and Bonhomme and Robin [18] use copulas to model dynamics in a panel of earnings data; Lee and Long [104] use copulas to model the (uncorrelated) residuals of a multivariate GARCH model; Patton [121], Dudley and Nimalendran [46] and Kang et al. [101] apply copulas to study dependence between hedge funds and other assets; and Zimmer [147] studies the role of copulas in the recent US housing crisis.

5. Conclusion

This survey reviews the growing literature on copula-based models for economic and financial time series. Models for multivariate time series, where copulas are used to model cross-sectional dependence, and univariate time series, where copulas model the serial dependence, are discussed. Estimation of these models is commonly conducted either parametrically, with the marginal distributions and the copula specified as belonging to parametric families, or semiparametrically, where the marginal distributions are estimated nonparametrically. Inference methods differ according to this choice, and both cases are reviewed. Goodness-of-fit tests for copula-based models are reviewed, and a brief survey of the many applications of copulas in the economics and finance literature is provided. Two simple empirical examples illustrate some of the methods presented.

Acknowledgments

The author thanks the Editor, Christian Genest, an Associate Editor and two anonymous referees for comments and suggestions. The author is also grateful to Yanqin Fan, Dominique Guégan, Bruno Rémillard, seminar participants at the Federal Reserve Bank of Saint Louis, and the audience at the Workshop on Copula Modeling and Dependence (Montréal, Canada, June 2011) for their helpful remarks. Dong Hwan Oh provided outstanding research assistance.

References

- [1] K. Aas, C. Czado, A. Frigessi, H. Bakken, Pair-copula constructions of multiple dependence, *Insurance Math. Econom.* 44 (2009) 182–198.
- [2] F. Abegaz, U.V. Naik-Nimbalkar, Modeling statistical dependence of Markov chains via copula models, *J. Statist. Plann. Inference* 138 (2008) 1131–1146.
- [3] E.F. Acar, C. Genest, J. Nešlehová, Beyond simplified pair-copula constructions, *J. Multivariate Anal.* (2012) in press (<http://dx.doi.org/10.1016/j.jmva.2012.02.001>).
- [4] C. Alexander, *Market Risk Analysis*, Volume III, Wiley, London, 2008.
- [5] T.G. Andersen, T. Bollerslev, P. Christoffersen, F.X. Diebold, Volatility and correlation forecasting, in: *Handbook of Economic Forecasting*, Vol. 1, Elsevier, Oxford, 2006.
- [6] A. Ang, J. Chen, Asymmetric correlations of equity portfolios, *J. Finan. Econ.* 63 (2012) 443–494.

- [7] M.C. Ausin, H.F. Lopes, Time-varying joint distribution through copulas, *Comput. Statist. Data Anal.* 54 (2010) 2383–2399.
- [8] K.-H. Bae, G.A. Karolyi, R.M. Stulz, A new approach to measuring financial contagion, *Rev. Finan. Stud.* 16 (2003) 717–764.
- [9] S.M. Bartram, S.J. Taylor, Y.-H. Wang, The Euro and European financial market dependence, *J. Bank. Financ.* 51 (2007) 1461–1481.
- [10] B.K. Beare, Copulas and temporal dependence, *Econometrica* 78 (2010) 395–410.
- [11] B.K. Beare, Archimedean copulas and temporal dependence, *Econom. Theory* (in press).
- [12] B.K. Beare, J. Seo, Time irreversible copula-based Markov models, Working Paper, University of California, San Diego, 2012.
- [13] M.N. Bennett, J.E. Kennedy, Quanto pricing with copulas, *J. Derivatives* 12 (2004) 26–45.
- [14] D. Berg, Copula goodness-of-fit testing: an overview and power comparison, *European J. Finance* 15 (2009) 675–701.
- [15] D. Berg, J.-F. Quessy, Local power analyses of goodness-of-fit tests for copulas, *Scandinavian J. Statist.* 36 (2009) 389–412.
- [16] T. Bollerslev, Generalized autoregressive conditional heteroskedasticity, *J. Econometrics* 31 (1986) 307–327.
- [17] T. Bollerslev, A conditionally heteroskedastic time series model for speculative prices and rates of return, *Rev. Econom. Statist.* 69 (1987) 542–547.
- [18] S. Bonhomme, J.-M. Robin, Assessing the equalizing force of mobility using short panels: France, 1990–2000, *Rev. Econom. Stud.* 76 (2009) 63–92.
- [19] C. Bontemps, B. Feunou, N. Meddahi, Testing distributional assumptions: the multivariate case, Working Paper, École d'économie de Toulouse, Toulouse, France, 2011.
- [20] É. Bouyé, M. Salmon, Dynamic copula quantile regressions and tail area dynamic dependence in forex markets, *European J. Finance* 15 (2009) 721–750.
- [21] W. Breymann, A. Dias, P. Embrechts, Dependence structures for multivariate high-frequency data in finance, *Quant. Finance* 3 (2003) 1–16.
- [22] F. Busetti, A. Harvey, When is a copula constant? a test for changing relationships, *J. Finan. Econ.* 9 (2011) 106–131.
- [23] N.-H. Chan, J. Chen, X. Chen, Y. Fan, L. Peng, Statistical inference for multivariate residual copula of GARCH models, *Statist. Sinica* 19 (2009) 53–70.
- [24] Y.-T. Chen, Moment tests for density forecast evaluation in the presence of parameter estimation uncertainty, *J. Forecast.* 30 (2011) 409–450.
- [25] X. Chen, Y. Fan, Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification, *J. Econometrics* 135 (2006) 125–154.
- [26] X. Chen, Y. Fan, Estimation of copula-based semiparametric time series models, *J. Econometrics* 130 (2006) 307–335.
- [27] X. Chen, Y. Fan, V. Tsyrennikov, Efficient estimation of semiparametric multivariate copula models, *J. Amer. Statist. Assoc.* 101 (2006) 1228–1240.
- [28] X. Chen, R. Koenker, Z. Xiao, Copula-based nonlinear quantile autoregression, *Econom. J.* 12 (2009) 50–67.
- [29] X. Chen, W.B. Wu, Y. Yi, Efficient estimation of copula-based semiparametric Markov models, *Ann. Statist.* 37 (2009) 4214–4253.
- [30] U. Cherubini, F. Gobbi, S. Mulinacci, S. Romagnoli, Dynamic Copula Methods in Finance, Wiley, England, 2012.
- [31] U. Cherubini, E. Luciano, Value at risk trade-off and capital allocation with copulas, *Econom. Notes* 30 (2011) 235–256.
- [32] U. Cherubini, E. Luciano, W. Vecchiato, Copula Methods in Finance, Wiley, England, 2004.
- [33] L. Chollote, A. Heinen, A. Valdesogo, Modeling international financial returns with a multivariate regime-switching copula, *J. Finan. Econ.* 7 (2009) 437–480.
- [34] B. Choros, R. Ibragimov, E. Permiakova, Copula estimation, in: Workshop on Copula Theory and its Applications, Springer, Berlin, 2010.
- [35] P. Christoffersen, V. Errunza, K. Jacobs, H. Langlois, Is the potential for international diversification disappearing?, Working Paper, Rotman School of Management, University of Toronto, Canada, 2011.
- [36] P. Christoffersen, H. Langlois, The joint dynamics of equity market factors, Working Paper, Rotman School of Management, University of Toronto, Canada, 2011.
- [37] V. Corradi, N.R. Swanson, Predictive density evaluation, in: Handbook of Economic Forecasting, Vol. 1, North Holland, Amsterdam, 2006.
- [38] D. Creal, S.J. Koopman, A. Lucas, Generalized autoregressive score models with applications, *J. Appl. Econometrics* (in press).
- [39] W.F. Darsow, B. Nguyen, Copulas and Markov processes, *Illinois J. Math.* 36 (1992) 600–642.
- [40] S. Daul, E.D. Giorgi, F. Lindskog, A.J. McNeil, The grouped t -copula with an application to credit risk, *J. Risk* 16 (2003) 73–76.
- [41] L. Dearden, E. Fitzsimons, A. Goodman, G. Kaplan, Higher education funding reforms in England: the distributional effects and the shifting balance of costs, *Econom. J.* 118 (2008) 100–125.
- [42] S. Demarta, A.J. McNeil, The t copula and related copulas, *Internat. Statist. Rev.* 73 (2005) 111–129.
- [43] M. Denuit, J. Dhaene, M. Goovaerts, R. Kaas, Actuarial Theory for Dependent Risks: Measures, Orders and Models, Wiley, England, 2005.
- [44] A. Dias, P. Embrechts, Modeling exchange rate dependence dynamics at different time horizons, *J. Internat. Money Finance* 29 (2010) 1687–1705.
- [45] F.X. Diebold, J. Hahn, A.S. Tay, Multivariate density forecast evaluation and calibration in financial risk management: high frequency returns on foreign exchange, *Rev. Econom. Statist.* 81 (1999) 661–673.
- [46] E. Dudley, M. Nimalendran, Margins and hedge fund contagion, *J. Finan. Quant. Anal.* 46 (2011) 1227–1257.
- [47] D. Duffie, Measuring Corporate Default Risk, Oxford University Press, Oxford, 2011.
- [48] P. Embrechts, A. Höing, Extreme VaR scenarios in higher dimensions, *Extremes* 9 (2006) 177–192.
- [49] P. Embrechts, A. Höing, A. Juri, Using copulae to bound the value-at-risk for functions of dependent risks, *Finance Stoch.* 7 (2003) 145–167.
- [50] P. Embrechts, A.J. McNeil, D. Straumann, Correlation and dependence properties in risk management: properties and pitfalls, in: Risk Management: Value at Risk and Beyond, Cambridge University Press, Cambridge, 2002.
- [51] R.F. Engle, Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation, *Econometrica* 50 (1982) 987–1007.
- [52] R.F. Engle, Dynamic conditional correlation—a simple class of multivariate GARCH models, *J. Bus. Econom. Statist.* 20 (2002) 339–350.
- [53] C.B. Erb, C.R. Harvey, T.E. Viskanta, Forecasting international equity correlations, *Financ. Anal. J.* 50 (1994) 32–45.
- [54] T.S. Ferguson, C. Genest, M. Hallin, Kendall's tau for serial dependence, *Canad. J. Statist.* 28 (2000) 587–604.
- [55] J.-D. Fermanian, D. Radulović, M.H. Wegkamp, Weak convergence of empirical copula processes, *Bernoulli* 10 (2004) 847–860.
- [56] J.-D. Fermanian, O. Scaillet, Nonparametric estimation of copulas for time series, *J. Risk* 5 (2003) 25–54.
- [57] J.-D. Fermanian, M.H. Wegkamp, Time-dependent copulas, *J. Multivariate Anal.* (2012) this issue (<http://dx.doi.org/10.1016/j.jmva.2012.02.018>).
- [58] R.A. Fisher, Statistical Methods for Research Workers, Oliver & Boyd, Edinburgh, 1932.
- [59] E.W. Frees, P. Wang, Copula credibility for aggregate loss models, *Insurance Math. Econom.* 38 (2006) 360–373.
- [60] R. Frey, A.J. McNeil, Modelling dependent defaults, <http://e-collection.ethbib.ethz.ch/show?type=bericht&nr=273>, 2001.
- [61] P. Gagliardini, C. Gouriéroux, Duration time-series models with proportional hazard, *J. Time Ser. Anal.* 29 (2007) 74–124.
- [62] S. Gaißer, M. Ruppert, F. Schmid, A multivariate version of Hoeffding's phi-square, *J. Multivariate Anal.* 101 (2010) 2571–2586.
- [63] R. Garcia, G. Tsafack, Dependence structure and extreme comovements in international equity and bond markets, *J. Bank. Financ.* 35 (2011) 1954–1970.
- [64] C. Genest, Frank's family of bivariate distributions, *Biometrika* 74 (1987) 549–555.
- [65] C. Genest, A.-C. Favre, Everything you always wanted to know about copula modeling but were afraid to ask, *J. Hydraul. Eng.* 12 (2007) 347–368.
- [66] C. Genest, K. Ghoudi, L.-P. Rivest, A semiparametric estimation procedure of dependence parameters in multivariate families of distributions, *Biometrika* 82 (1995) 543–552.
- [67] C. Genest, J. Nešlehová, A primer on copulas for count data, *ASTIN Bull.* 37 (2007) 475–515.
- [68] C. Genest, J. Nešlehová, N. Ben Ghorbal, Estimators based on Kendall's tau in multivariate copula models, *Aust. N. Z. J. Stat.* 53 (2011) 157–177.
- [69] C. Genest, J. Nešlehová, J. Ziegel, Inference in multivariate Archimedean copula models (with discussion), *TEST* 20 (2011) 223–292.
- [70] C. Genest, B. Remillard, Tests of independence and randomness based on the empirical copula process, *TEST* 13 (2004) 335–369.
- [71] C. Genest, B. Remillard, Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models, *Ann. Inst. H. Poincaré Probab. Statist.* 44 (2008) 1096–1127.
- [72] C. Genest, B. Remillard, D. Beaudoin, Goodness-of-fit tests for copulas: a review and a power study, *Insurance Math. Econom.* 44 (2009) 199–213.
- [73] C. Genest, L.-P. Rivest, Statistical inference procedures for bivariate Archimedean copulas, *J. Amer. Statist. Assoc.* 88 (1993) 1034–1043.
- [74] C. Genest, J. Segers, Rank-based inference for bivariate extreme-value copulas, *Ann. Statist.* 37 (2009) 2990–3022.

- [75] K. Ghoudi, B. Rémillard, Empirical processes based on pseudo-observations II: the multivariate case, in: *Asymptotic Methods in Stochastics: Festschrift for Miklós Csörgő*, Amer. Math. Soc., Providence, RI, 2004.
- [76] E. Giacomini, W. Härdle, V. Spokoiny, Inhomogeneous dependence modeling with time-varying copulae, *J. Bus. Econom. Statist.* 27 (2009) 224–234.
- [77] K. Giesecke, Correlated default with incomplete information, *J. Bank. Financ.* 28 (2004) 1521–1545.
- [78] L.R. Glosten, R. Jagannathan, D.E. Runkle, On the relation between the expected value and the volatility of the nominal excess return on stocks, *J. Financ.* 48 (1993) 1779–1801.
- [79] S. Gonçalves, H. White, Maximum likelihood and the bootstrap for nonlinear dynamic models, *J. Econometrics* 119 (2004) 199–220.
- [80] C. Gouriéroux, A. Monfort, A. Trognon, Pseudo-maximum likelihood methods: theory, *Econometrica* 52 (1984) 681–700.
- [81] C.W.J. Granger, T. Teräsvirta, A.J. Patton, Common factors in conditional distributions for bivariate time series, *J. Econometrics* 132 (2006) 43–57.
- [82] V. Grégoire, C. Genest, M. Gendron, Using copulas to model price dependence in energy markets, *Energy Risk* 5 (2008) 58–64.
- [83] D. Guégan, J. Zhang, Change analysis of dynamic copula for measuring dependence in multivariate financial data, *Quant. Finance* 10 (2009) 421–430.
- [84] C.M. Hafner, H. Manner, Dynamic stochastic copula models: estimation, inference and applications, *J. Appl. Econometrics* 27 (2012) 269–295.
- [85] J.D. Hamilton, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57 (1989) 357–384.
- [86] B.E. Hansen, Autoregressive conditional density estimation, *Internat. Econom. Rev.* 35 (1994) 705–730.
- [87] A. Harvey, Tracking a changing copula, *J. Emp. Financ.* 17 (2010) 485–500.
- [88] A. Heinen, E. Rengifo, Multivariate autoregressive modeling of time series count data using copulas, *J. Emp. Financ.* 14 (2007) 564–583.
- [89] A. Heinen, A. Valdesogo, Asymmetric CAPM dependence for large dimensions: the canonical vine autoregressive model, Discussion Paper 2009–069, Center for Operations Research and Econometrics, Université catholique de Louvain, Louvain-la-Neuve, Belgium, 2009.
- [90] C. Hering, M. Hofert, J.-F. Mai, M. Scherer, Constructing hierarchical Archimedean copulas with Lévy subordinators, *J. Multivariate Anal.* 101 (2010) 1428–1433.
- [91] M. Hofert, M. Mächler, A.J. McNeil, Likelihood inference for Archimedean copulas in high dimensions under known margins, *J. Multivariate Anal.* (2012) this issue (<http://dx.doi.org/10.1016/j.jmva.2012.02.019>).
- [92] M. Hofert, M. Scherer, CDO pricing with nested Archimedean copulas, *Quant. Finance* 11 (2011) 775–787.
- [93] Y. Hong, J. Tu, G. Zhou, Asymmetries in stock returns: statistical tests and economic evaluation, *Rev. Finan. Stud.* 20 (2007) 1547–1581.
- [94] J. Hull, A. White, Value at risk when daily changes in market variables are not normally distributed, *J. Derivatives* 5 (1998) 9–19.
- [95] R. Ibragimov, Copula-based characterizations for higher-order Markov processes, *Econometric Theory* 25 (2009) 819–846.
- [96] R. Ibragimov, G. Lentzas, Copulas and long memory, Harvard Institute of Economic Research, Discussion Paper No. 2160, 2009.
- [97] H. Joe, Asymptotic efficiency of the two-stage estimation method for copula-based models, *J. Multivariate Anal.* 94 (1995) 401–419.
- [98] H. Joe, *Multivariate Models and Dependence Concepts*, Chapman & Hall, London, 1997.
- [99] H. Joe, J.J. Xu, The estimation method of inference functions for margins for multivariate models, Working Paper, Department of Statistics, University of British Columbia, Vancouver, Canada, 1996.
- [100] É. Jondeau, M. Rockinger, The copula-GARCH model of conditional dependencies: an international stock market application, *J. Internat. Money Finance* 25 (2006) 827–853.
- [101] B.U. Kang, F. In, G. Kim, Y.S. Kim, A longer look at the asymmetric dependence between hedge funds and the equity market, *J. Finan. Quant. Anal.* 45 (2010) 763–789.
- [102] C.A.J. Klaassen, J.A. Wellner, Efficient estimation in the bivariate normal copula model: normal margins are least favourable, *Bernoulli* 3 (1997) 55–77.
- [103] I. Komunjer, Quantile prediction, in: *Handbook of Economic Forecasting*, Vol. 2, Elsevier, Oxford, 2011.
- [104] T.-H. Lee, X. Long, Copula-based multivariate GARCH model with uncorrelated dependent standardized returns, *J. Econometrics* 150 (2009) 207–218.
- [105] D.X. Li, On default correlation: a copula function approach, *J. Fixed Income* 9 (2000) 43–54.
- [106] F. Longin, B. Solnik, Extreme correlation of international equity markets, *J. Financ.* 56 (2001) 649–676.
- [107] H. Manner, O. Reznikova, A survey on time-varying copulas: specification, simulations and estimation, *Econometric Rev.* (2012) (in press).
- [108] H. Manner, J. Segers, Tails of correlation mixtures of elliptical copulas, *Insurance Math. Econom.* 48 (2011) 153–160.
- [109] T.D. Markwat, E. Kole, D.J.C. van Dijk, Time variation in asset return dependence: strength or structure?, Working Paper, Econometric Institute, Erasmus University Rotterdam, The Netherlands, 2009.
- [110] A.J. McNeil, R. Frey, P. Embrechts, *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton University Press, Princeton, NJ, 2005.
- [111] A. Min, C. Czado, Bayesian inference for multivariate copulas using pair-copula constructions, *J. Financ. Econ.* 8 (2010) 450–480.
- [112] R.B. Nelsen, *An Introduction to Copulas*, second ed., Springer, 2006.
- [113] W.K. Newey, D. McFadden, Large sample estimation and hypothesis testing, in: *Handbook of Econometrics*, North-Holland, Amsterdam, 1994.
- [114] D.-H. Oh, A.J. Patton, Modelling dependence in high dimensions with factor copulas, Working Paper, Duke University, Durham, NC, 2011.
- [115] D.-H. Oh, A.J. Patton, Simulated method of moments estimation for copula-based multivariate models, Working Paper, Duke University, Durham, NC, 2011.
- [116] T. Okimoto, New evidence of asymmetric dependence structure in international equity markets, *J. Finan. Quant. Anal.* 43 (2008) 787–815.
- [117] A.J. Patton, Applications of copula theory in financial econometrics, Unpublished Ph.D. Dissertation, University of California, San Diego, CA, 2002.
- [118] A.J. Patton, On the out-of-sample importance of skewness and asymmetric dependence for asset allocation, *J. Financ. Econ.* 2 (2004) 130–168.
- [119] A.J. Patton, Estimation of multivariate models for time series of possibly different lengths, *J. Appl. Econometrics* 21 (2006) 147–173.
- [120] A.J. Patton, Modelling asymmetric exchange rate dependence, *Internat. Econom. Rev.* 47 (2006) 527–556.
- [121] A.J. Patton, Are “market neutral” hedge funds really market neutral? *Rev. Finan. Stud.* 22 (2009) 2495–2530.
- [122] A.J. Patton, Copula methods for forecasting multivariate time series, in: *Handbook of Economic Forecasting*, Vol. 2, Elsevier, Oxford, 2011.
- [123] M. Pitt, D. Chan, R.J. Kohn, Efficient Bayesian inference for Gaussian copula regression models, *Biometrika* 93 (2006) 537–554.
- [124] D.N. Politis, J.P. Romano, The stationary bootstrap, *J. Amer. Statist. Assoc.* 89 (1994) 1303–1313.
- [125] B. Rémillard, Goodness-of-fit tests for copulas of multivariate time series, Working Paper, HEC Montréal, Canada, 2010.
- [126] B. Rémillard, N. Papageorgiou, F. Soustra, Copula-based semiparametric models for multivariate time series, *J. Multivariate Anal.* (2012) this issue (<http://dx.doi.org/10.1016/j.jmva.2012.03.001>).
- [127] J.C. Rodriguez, Measuring financial contagion: a copula approach, *J. Empirical Finance* 14 (2007) 401–423.
- [128] J.V. Rosenberg, Nonparametric pricing of multivariate contingent claims, *J. Derivatives* 10 (2003) 9–26.
- [129] J.V. Rosenberg, T. Schuermann, A general approach to integrated risk management with skewed, fat-tailed risks, *J. Financ. Econ.* 79 (2006) 569–614.
- [130] M. Ruppert, Consistent testing for a constant copula under strong mixing based on the tapered block multiplier technique, Working Paper, Wirtschafts- und Sozialwissenschaftliche Fakultät, Universität zu Köln, Germany, 2011.
- [131] M. Salmon, C. Schleicher, Pricing multivariate currency options with copulas, in: *Copulas: From Theory to Application in Finance*, Risk Books, London, 2006.
- [132] A. Sancetta, S. Satchell, The Bernstein copula and its applications to modeling and approximations of multivariate distributions, *Econom. Theory* 20 (2004) 535–562.
- [133] P. Schönbucher, D. Schubert, Copula dependent default risk in intensity models, Working Paper, Universität Bonn, Germany, 2001.
- [134] N. Shephard, *Stochastic Volatility: Selected Readings*, Oxford University Press, Oxford, 2005.
- [135] J.H. Shih, T.A. Louis, Inferences on the association parameter in copula models for bivariate survival data, *Biometrics* 51 (1995) 1384–1399.
- [136] A. Sklar, Fonctions de répartition à n dimensions et leurs marges, *Publ. Inst. Statist. Univ. Paris* 8 (1959) 229–231.
- [137] M.S. Smith, Bayesian approaches to copula modelling, in: *Hierarchical Models and MCMC: A Tribute to Adrian Smith*, Oxford University Press (in press).
- [138] M.S. Smith, Q. Gan, R.J. Kohn, Modeling dependence using skew t copulas: Bayesian inference and applications, *J. Appl. Econometrics* 27 (2012) 500–522.

- [139] M.S. Smith, A. Min, C. Almeida, C. Czado, Modeling longitudinal data using a pair-copula decomposition of serial dependence, *J. Amer. Statist. Assoc.* 105 (2010) 1467–1479.
- [140] P. Song, Y. Fan, J.D. Kalbfleisch, Maximization by parts in likelihood inference, *J. Amer. Statist. Assoc.* 100 (2005) 1145–1158.
- [141] J. Sun, E.W. Frees, M. Rosenberg, Heavy-tailed longitudinal data modeling using copulas, *Insurance Math. Econom.* 42 (2008) 817–830.
- [142] S.J. Taylor, Y. Wang, Option prices and risk-neutral densities for currency cross-rates, *J. Futures Markets* 30 (2010) 324–360.
- [143] H. Tsukahara, Semiparametric estimation in copula models, *Canad. J. Statist.* 33 (2005) 357–375.
- [144] R. van den Goorbergh, C. Genest, B.J.M. Werker, Multivariate option pricing using dynamic copula models, *Insurance Math. Econom.* 37 (2005) 101–114.
- [145] H. White, *Estimation, Inference and Specification Analysis*, Cambridge University Press, Cambridge, 1994.
- [146] W. Yi, S.S. Liao, Statistical properties of parametric estimators for Markov chain vectors on copula models, *J. Statist. Plann. Inference* 140 (2010) 1465–1480.
- [147] D. Zimmer, The role of copulas in the housing crisis, *Rev. Econom. Statist.* (2012) (in press).