

TI 2011-125/4
Tinbergen Institute Discussion Paper



Forecasting Volatility with Copula-Based Time Series Models

Oleg Sokolinskiy
*Dick van Dijk**

Erasmus University Rotterdam, Tinbergen Institute.

** Econometric Institute*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

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

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

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

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Forecasting Volatility with Copula-Based Time Series Models*

Oleg Sokolinskiy[†]

Tinbergen Institute
Erasmus University Rotterdam

Dick van Dijk[‡]

Econometric Institute
Erasmus University Rotterdam

September 2, 2011

Abstract

This paper develops a novel approach to modeling and forecasting realized volatility (RV) measures based on copula functions. Copula-based time series models can capture relevant characteristics of volatility such as nonlinear dynamics and long-memory type behavior in a flexible yet parsimonious way. In an empirical application to daily volatility for S&P500 index futures, we find that the copula-based RV (C-RV) model outperforms conventional forecasting approaches for one-day ahead volatility forecasts in terms of accuracy and efficiency. Among the copula specifications considered, the Gumbel C-RV model achieves the best forecast performance, which highlights the importance of asymmetry and upper tail dependence for modeling volatility dynamics. Although we find substantial variation in the copula parameter estimates over time, conditional copulas do not improve the accuracy of volatility forecasts.

Keywords: Nonlinear dependence, long memory, copulas, volatility forecasting
JEL classification: C22, C53, C58, G17.

*We thank participants at the 3rd International Conference on Computational and Financial Econometrics (Cyprus, October 2009) for useful comments and suggestions.

[†]Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands. E-mail: oleg.sokolinskiy@gmail.com

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

1 Introduction

Forecasts of asset price volatility are of crucial importance in portfolio management, risk management, and derivatives pricing. In recent years, the increasing availability of high-frequency asset price data has led to the development of various different measures of daily volatility based on intraday prices, see McAleer and Medeiros (2008b) and Andersen et al. (2009) for recent surveys. Considerable research effort has also been spent on designing suitable time series models for forecasting these so-called realized volatility measures. In several empirical applications, the heterogeneous autoregressive (HAR) model of Corsi (2009) has been found most successful for this purpose, see Andersen et al. (2007), Corsi (2009), and Busch et al. (2011), among others. The HAR model's appeal is due to its parsimony and ability to capture the stylized fact of long-memory in (realized) volatility. In addition, the model can be extended in straightforward ways to incorporate other typical features such as leverage effects, jumps, seasonality, as well as impacts of macro-economic announcements, see Martens et al. (2009).

A crucial feature of the HAR model is its linearity, in the sense that the dependence between realized volatility on consecutive days is assumed to be constant over time and independent of the level of volatility. In principle, these assumptions can be relaxed by allowing for structural breaks or regime-switching behavior within the HAR framework, see McAleer and Medeiros (2008a), but in that case the model quickly becomes heavily parameterized and, thus, loses its attractive property of parsimony.

In this paper, we propose an alternative approach to modeling possibly nonlinear dynamics in realized volatility based on copula functions. The key idea underlying this copula-based model for realized volatility (C-RV model) is that we can decompose the joint distribution of current volatility and its first lag into their marginal distributions and a copula function, with the latter characterizing the temporal dependence. As the marginal distributions and the dependence structure can be modeled separately, the copula-based approach allows for a great deal of flexibility in the construction of an appropriate multivariate distribution.

Not surprisingly, copulas have quickly gained popularity in economics and particularly fi-

nance, see Cherubini et al. (2004), Patton (2009) and Genest et al. (2009) for recent surveys. Although in finance copulas have been used mostly to describe the contemporaneous dependence between returns on different assets, they can also be used to model nonlinear time series dependence of a single variable. In fact, by combining different marginal distributions with different copula functions, a wide variety of marginal characteristics (including skewness and excess kurtosis) can be modeled, in addition to dependence characteristics such as clustering, asymmetry and tail dependence. Recently, Ibragimov and Lentzas (2008) demonstrate that copula-based time series models can also display long memory properties, see also Chen et al. (2009) and Beare (2010). Together with the ability to capture nonlinear dependence in a flexible and parsimonious way, this makes copula-based time series models a possible contender to conventional approaches for modeling realized volatility, such as the HAR model. In this paper we examine whether this indeed is the case, in particular from a forecasting perspective.¹

We evaluate the forecasting performance of the C-RV model in an empirical application to daily volatility of the S&P500 futures, over the period from January 1995 to December 2006. We employ the realized range developed by Martens and van Dijk (2007) and Christensen and Podolskij (2007) to measure daily volatility. As we focus on the ability of the C-RV model to capture long-memory and nonlinearity in the time series dependence of realized volatility, we adopt the semi-parametric approach advocated by Chen and Fan (2006). This combines non-parametric estimation of the marginal distributions with a parametric copula specification. The adverse effects of misspecified marginals (Fermanian and Scaillet, 2005) are thus avoided, while retaining consistency of estimates of important characteristics of the multivariate distribution, such as moments and quantiles.

In the empirical analysis, we address the following two key issues in the specification of C-RV models. First, different copula functions imply different types of time series dependence in volatility, in terms of (a)symmetry and tail (in)dependence. This makes the choice of a copula specification an important issue in practice. We consider a variety of copula functions and

¹Independent, contemporaneous research by Ning et al. (2010) also suggests copula-based time series models for describing the dynamics of realized volatility measures, but does not consider out-of-sample forecasting.

examine which specifications yield the most accurate out-of-sample forecasts of volatility.

Second, although the C-RV model allows for possibly nonlinear time series dependence in volatility, the dependence is assumed to be ‘stable’, that is, constant over time. We allow for the possibility of changes in the dependence in the C-RV model by using conditional copulas with time-varying parameters, as in Patton (2006).

Our empirical results can be summarized as follows. We find that the C-RV model outperforms the HAR for one-day ahead volatility forecasts in terms of accuracy and in terms of efficiency. Among the copula specifications considered, the Gumbel C-RV model achieves the best forecast performance, which highlights the importance of asymmetry and upper tail dependence for modeling volatility. Although we find substantial variation in the copula parameter estimates over time, conditional copulas do not improve the accuracy of volatility forecasts.

The rest of the paper is organized as follows. In Section 2, we introduce the C-RV model and briefly describe the semi-parametric estimation procedure of Chen and Fan (2006). In Section 3 we compare the out-of-sample forecasting performance of the C-RV and HAR models for the volatility of S&P500 index futures. We conclude in Section 4. An Appendix provides details on the parametric copula specifications used in the empirical analysis.

2 The copula realized volatility model

Let S_t denote the price of a financial asset and assume that it is determined by the stochastic differential equation

$$d \log S_t = \mu dt + \sigma_t dW_t, \tag{1}$$

where μ is the constant drift, σ_t is the possibly stochastic spot volatility, and W_t is a standard Brownian motion that is independent of σ_t .

We focus on daily volatilities. Hence, for convenience we normalize time such that the unit

interval corresponds to a day. We may then define the daily integrated variance (volatility),

$$IV_t = \int_{t-1}^t \sigma_s^2 ds. \quad (2)$$

The integrated volatility IV_t is not directly observable, but can be estimated consistently using a ‘realized’ volatility measure based on high-frequency intraday prices, as discussed in Section 3.1. In the remainder of this section we describe the copula-based time series model for volatility in terms of IV_t , with the implicit understanding that in empirical applications it should be replaced by an estimate.

The essence of the C-RV model is the dependence structure between consecutive observations of the integrated volatility process. Following the convention in the stochastic volatility literature, we assume that IV_t is a first order Markov process. Hence, its statistical properties are completely characterized by the joint distribution function of IV_{t-1} and IV_t , denoted as $F(IV_{t-1}, IV_t)$. Sklar’s (1959) theorem states that we can represent any such bivariate distribution by means of a copula function C and the marginal distribution of IV_t , that is,

$$F(IV_{t-1}, IV_t) = C(F(IV_{t-1}), F(IV_t)), \quad (3)$$

where F is the marginal distribution of integrated volatility, which is assumed to be time-invariant.

It is more common to express stochastic volatility models in terms of the conditional distribution of IV_t given IV_{t-1} . This may also be done within the framework of copula functions. Note that it follows from (3) that the joint density of IV_{t-1} and IV_t is given by

$$f(IV_{t-1}, IV_t) = c(F(IV_{t-1}), F(IV_t))f(IV_{t-1})f(IV_t),$$

where $f(\cdot)$ is the marginal density of IV_t and $c(\cdot, \cdot) = \partial^2 C(\cdot, \cdot) / \partial u_{t-1} \partial u_t$ is the density of the copula function C with $u_t = F(IV_t)$ denoting the probability integral transform of integrated

volatility. Hence, the conditional density of IV_t given IV_{t-1} can be expressed as

$$f(IV_t|IV_{t-1}) = f(IV_{t-1}, IV_t)/f(IV_{t-1}) = c(F(IV_{t-1}), F(IV_t))f(IV_t). \quad (4)$$

The representation in (4) shows the attractiveness of the copula approach for modeling nonlinear time series dependence in a flexible way. As the conditional density of IV_t given IV_{t-1} can be written as the product of the copula density and the marginal density, we can separate the temporal dependence structure from the marginal behavior. As the choice of the marginal distribution does not restrict the choice of the dependence function, or vice versa, a wide range of conditional distributions can be obtained by combining different marginals $F(\cdot)$ with different copulas $C(\cdot, \cdot)$. As we aim to focus on the usefulness of copulas to model the nonlinear temporal dependence in realized volatility, we concentrate on possible specifications of the copula function C . As discussed in more detail below, the marginal distribution F is estimated nonparametrically.

A wide range of parametric copula specifications is available, with different implications for the dependence structure of volatility, see Joe (1997) and Nelsen (2006) for overviews. Different copula functions may usefully be compared in terms of their so-called ‘quantile dependence’ and the limiting case of tail dependence. A particular measure of quantile dependence is the exceedance probability. This is defined as the conditional probability that IV_t exceeds a given quantile q of its marginal distribution given that IV_{t-1} exceeds that quantile. Specifically, the exceedance probabilities are given by

$$\tau(q) = \begin{cases} P(u_t < q | u_{t-1} < q) = C(q; q)/q & \text{for } q \leq 0.5, \\ P(u_t > q | u_{t-1} > q) = (1 - 2q + C(q; q))/(1 - q) & \text{for } q > 0.5. \end{cases} \quad (5)$$

where $u_t = F(IV_t)$ is the probability integral transform of integrated volatility. The lower and upper tail dependence coefficients are then defined as the limits of this quantile dependence measure, that is, $\tau_L = \lim_{q \downarrow 0} \tau(q)$ and $\tau_U = \lim_{q \uparrow 1} \tau(q)$.

Different copula specifications have different quantile and tail dependence characteristics. Some copulas (such as the Gaussian, Student's t and Frank copulas) are symmetric in the sense that $\tau(q) = \tau(1 - q)$ for all $0 \leq q \leq 0.5$, while others (such as the Gumbel and Clayton copulas) are asymmetric. Some copulas have no tail dependence, i.e. $\tau_L = \tau_U = 0$, while others have positive tail dependence. In the latter case, the dependence can be positive in either the lower tail or in the upper tail, or in both. Which type of dependence structure is most appropriate for realized volatility is not clear *a priori*. For that reason, we examine a broad range of copula specifications with different quantile and tail dependence properties, and evaluate their relative out-of-sample forecasting performance. In the Appendix we describe the specifications of the copula functions that we consider, together with their quantile and tail dependence properties.

2.1 Estimation and forecasting

As discussed before, in the copula approach the temporal dependence structure, represented by the copula function itself, is separated from the marginal distributions of volatility. However, in a sense the choice of marginal distribution does affect the copula as $F(IV_{t-1})$ and $F(IV_t)$ are the arguments of the copula specification, see (3). As demonstrated by Fermanian and Scaillet (2005), among others, misspecification of the marginal distribution may lead to spurious results for the copula specification, for example in the form of biased parameter estimates. In order to avoid this issue, we adopt the semi-parametric approach of Chen and Fan (2006) and estimate the marginal distribution $F(IV_t)$ nonparametrically. Specifically, given a time series $\{IV_t\}_{t=1}^T$ we use the empirical distribution function (EDF) defined as

$$\hat{F}(x) = \frac{1}{T+1} \sum_{t=1}^T I[IV_t \leq x],$$

where $I[A]$ is an indicator function for the event A and T denotes the sample size.² The copula parameters are then estimated by maximum likelihood using the copula part of the log likelihood

²We follow the convention in the copula literature to divide by $T + 1$ rather than T , to accommodate the fact that some copula densities $c(\cdot, \cdot)$ are not defined when one of the arguments takes the value 0 or 1.

function corresponding to (4), that is,

$$\hat{\theta} = \operatorname{argmax}_{\theta} \sum_{t=2}^T \log c(\hat{F}(IV_{t-1}), \hat{F}(IV_t); \theta). \quad (6)$$

We obtain one step ahead forecasts of volatility in period $T + 1$ from the C-RV model as follows. First, given the parameter estimates $\hat{\theta}$, we simulate B draws of the adjusted ranks $\hat{F}(IV_{T+1})$ from the fitted copula $C(\hat{F}(\cdot), \hat{F}(\cdot); \hat{\theta})$, conditional on the known value of IV_T . Second, we use the inverse empirical distribution to transform each of the simulated ranks into values of realized volatility. The mean across the B simulations is taken as the realized volatility forecast, denoted $\widehat{IV}_{T+1|T}$.

Finally, it is useful to note that in our semiparametric set-up the C-RV model specification is invariant to strictly monotonic transformations of IV_t . Applying the model to the integrated variance as defined in (2) or its square root or its logarithm will render identical results. This is in contrast to conventional approaches such as the HAR model, which may be quite sensitive to which transformation of IV_t is used.

3 Application to S&P 500 index futures

In this section we apply the C-RV methodology to forecasting daily volatility of the S&P 500 index futures. In Section 3.1, we start with a brief description of the data and the estimation of the integrated volatility (2) based on high-frequency intra-day prices, highlighting characteristics that suggest the possibility of nonlinear dependence in these volatility measures. We discuss the empirical results in Section 3.2, comparing the out-of-sample forecasting performance of the C-RV model with the HAR model of Corsi (2009). Here we also address the issues which type of copula specification renders most accurate volatility forecasts, and whether conditional copulas with time-varying parameters may improve forecast performance.

3.1 Data and integrated volatility measures

We obtain intraday transaction prices for futures contracts on the S&P500 (traded on the Chicago Mercantile Exchange with trading hours from 8:30am - 3:15pm) for the period from January 3, 1995 to December 29, 2006, for a total of 3031 daily observations.³ The S&P500 futures contract has maturities in March, June, September and December. At any given day we use prices of the most liquid contract. Typically, this is the nearby contract until approximately one week before maturity, when the trading volume in the second nearby contract becomes larger. We make sure that when changing from one contract to the next, we never compute returns based on prices from two different contracts.

The most straightforward estimator of the integrated volatility IV_t in (2) is the realized variance, defined as

$$RV_t^M = \sum_{m=1}^M r_{t,m}^2 + r_{t,o}^2, \quad (7)$$

where M is the number of intra-day intervals used, $r_{t,m}$ is the return during the m -th interval on day t , and $r_{t,o}$ is the overnight return between the closing price on day t and the opening price on day $t + 1$. The latter is incorporated as Martens (2002) documents that the overnight volatility represents an important fraction of total daily volatility, see also Fleming et al. (2003) and Hansen and Lunde (2005) for discussion.

An alternative measure is the realized range, introduced by Martens and van Dijk (2007) and Christensen and Podolskij (2007), defined as

$$RR_t^M = \sum_{m=1}^M \frac{1}{4 \log 2} (\log H_{t,m} - \log L_{t,m})^2 + r_{t,o}^2, \quad (8)$$

where the high price $H_{t,m}$ and the low price $L_{t,m}$ are defined as the maximum and minimum of all transaction prices observed during the m -th interval on day t . The realized range exploits the complete price path in the intra-day intervals, while the realized variance only uses the first and last price observations. Consequently, the realized range is a more efficient estimator than

³The data is obtained from Tick Data, <http://www.tickdata.com/>.

the realized variance based on the same sampling frequency.⁴ For this reason, we focus on the realized range.⁵ In the implementation of (8) we adopt the popular five-minute frequency, corresponding to $M = 81$.

Market microstructure effects hamper the use of high-frequency data for estimating daily variances. In particular, the standard realized variance in (7) suffers from an upward bias due to the presence of bid-ask bounce.⁶ The same effect occurs for the realized range in (8), as the observed high price in a given interval is an ask and the observed low price is a bid with probability close to 1 when trading is continuous. The realized range therefore overestimates the true daily variance by an amount equal to the squared bid-ask spread times the number of intraday intervals M . Additionally, however, infrequent trading (which does not affect the realized variance) leads to a downward bias in the realized range. When the continuous underlying price process is only observed at discrete points in time, the observed high price during a given intraday interval underestimates the true maximum. Similarly, the observed low price overestimates the true minimum in that case. Martens and van Dijk (2007) suggest to deal with the “net” bias due to the combined effects of bid-ask bounce and infrequent trading on the realized range by applying a multiplicative bias-correction. Specifically, the scaled realized range is defined by

$$RR_{S,t}^M = \left(\frac{\sum_{q=1}^Q RR_{t-q}}{\sum_{q=1}^Q RR_{t-q}^M} \right) RR_t^M, \quad (9)$$


where $RR_t \equiv RR_t^1$ is the daily range. Hence, the multiplicative correction factor is the ratio of the average daily range estimator and the average of the realized range over the past Q days. Adopting this bias-correction with $Q = 66$ (corresponding to approximately three months of trading days), we find similar forecasting results as for the ‘unscaled’ realized range RR_t^M . To save space, we therefore only report detailed results for the latter measure.

⁴Specifically, assuming that the asset price S_t follows a geometric Brownian motion with constant instantaneous volatility σ , the variance of the realized range estimator is equal to $0.407\sigma^4/M^2$, compared to $2\sigma^4/M^2$ for the realized variance.

⁵Results using the realized variance RV_t are qualitatively similar and are available upon request.

⁶This has resulted in a variety of approaches for bias-correcting realized volatility estimators on the one hand, and for determining the ‘optimal’ sampling frequency (based on the trade-off between accuracy and bias) for the standard realized variance estimator on the other hand, see the overview in McAleer and Medeiros (2008b).

Figure 1: Realized range for S&P500 index futures




Note: The graph shows the realized range RR_t^M (converted to annualized volatility in percentage points) obtained from (8) based on 5-minute intraday intervals ($M = 81$). The sample period runs from January 3, 1995 to December 29, 2006.

Figure 1 shows the time series of the realized range RR_t^M obtained from (8), converted to annualized volatility in percentage points. While the level of volatility is rather low at approximately 4 percent at the start of the sample period, it gradually increases during the second half of the 1990s, with short outbursts during the Asian crisis in 1997 and the Russian default in 1998. Similar relatively short periods of high volatility occur during the collapse of the dot-com bubble in 2000, and following September 11, 2001. During the final three years of the sample period, volatility declines again to moderate levels below 10 percent. Due to the short periods of high volatility, the realized range shows large positive skewness and excess kurtosis. The asymmetry and fat-tailedness of the distribution of RR_t are removed almost completely by taking logs. This is confirmed by Figure 2, which shows that the log of the realized range is close to being normally distributed.

Figure 3 shows a scatterplot of the normalized EDF of the log realized range on the vertical axis against its first lag on the horizontal axis. The strong positive dependence is obvious from this graph. The empirical first-order autocorrelation of $\log RR_t^M$ is equal to 0.81. We examine the presence of nonlinear and asymmetric dependence by computing exceedance probabilities and exceedance correlations for quantiles $q = 0.05, 0.075, \dots, 0.95$. Figure 4 shows these for


Figure 2: Histogram of log realized range for S&P500 index futures



Note: The graph shows the histogram of the logarithm of the realized range RR_t^M obtained from (8) based on 5-minute intraday intervals ($M = 81$). The sample period runs from January 3, 1995 to December 29, 2006. The dashed line is a normal density with the same mean and variance.

the log scaled realized range, together with the corresponding values for a bivariate normal distribution with correlation equal to the empirical first-order autocorrelation of $\log RR_t^M$, which is equal to 0.81. The latter are obtained by simulating time series with length $T = 3031$ from this bivariate normal distribution. The graph shows the mean as well as the 5th and 95th percentiles of the distribution of exceedance probability and exceedance correlation estimates in this simulation, using 10,000 replications. The exceedance correlations display a pronounced asymmetric pattern. For quantiles above the median, the empirical exceedance correlations correspond reasonably well with the values obtained for the normal distribution. By contrast, for values below the median the empirical exceedance correlations decay much more rapidly than for the normal distributions as q becomes smaller. The empirical exceedance probabilities are more symmetric, in the sense that the values for quantiles q and $1 - q$ are approximately equal for all $q \in (0, 0.5)$. At the same time, the empirical probabilities are larger than those for the normal distribution. Taken together, these features suggest that a bivariate normal distribution is not adequate to characterize the empirical temporal dependence in the log realized range. In particular, the dependence is (possibly) asymmetric and stronger than implied by the normal distribution, which makes modeling the dependence by means of copulas an attractive alternative

Figure 3: Scatter of log realized range for S&P500 index futures




Note: The graph shows a scatterplot of the normalized EDF of the realized range RR_t^M on the vertical axis against its first lag on the horizontal axis. RR_t^M is obtained from (9) based on 5-minute intraday intervals ($M = 81$). The sample period runs from January 3, 1995 to December 29, 2006.

approach.


We assess the stability of these dependence characteristics by computing the exceedance probabilities for a moving window of 500 observations. Figure 5 shows the resulting exceedance probabilities at $q = 0.10$ and 0.90 for the log realized range, where the date displayed on the horizontal axis refers to the mid-point of the moving window. We observe substantial time-variation in both exceedance probabilities. For both quantiles, the exceedance probabilities fluctuate between 0.3 and 0.7. The changes in the dependence in the left tail and the right tail of the joint distribution of RR_{t-1} and RR_t do not seem to be closely related. Sometimes the exceedance probabilities move up or down together (e.g. during 1997 and 2003-4), but sometimes the change in the opposite direction (e.g. during 1998-9 and 2001). The observed instability in the exceedance probabilities motivates us to consider *conditional* copulas with time-varying parameters.

Figure 4: Exceedance correlations and probabilities of log realized range for S&P500 index futures



Note: The graph shows the exceedance correlations of and exceedance probabilities for the log realized range and its first lag (line marked with solid circles). The sample period runs from January 3, 1995 to December 29, 2006. The line marked with open circles shows the average exceedance correlations (probabilities) across 10,000 simulations of series with length $T = 3031$ from a bivariate normal distribution with correlation equal to the empirical first-order autocorrelation of $\log RR_t^M$ (equal to 0.81). The dashed lines are 5th and 95th percentiles of the distribution of exceedance correlation (probability) estimates in this simulation.

Figure 5: Exceedance probabilities of log realized range for S&P500 index futures



Note: The graph shows exceedance probabilities at $q = 0.10$ (solid line) and $q = 0.90$ (dashed line) for the log realized range and its first lag, for rolling windows of 500 observations. RR_t^M is obtained from (9) based on 5-minute intraday intervals ($M = 81$). The sample period runs from January 3, 1995 to December 29, 2006 ($T = 3031$ observations).

3.2 Out-of-sample forecasting performance

We use the period from January 2, 1997 until December 29, 2006 ($P = 2530$ observations) to assess the potential usefulness of the C-RV modeling approach from an out-of-sample forecasting perspective. We re-estimate all models for each trading day using a rolling window of 500 observations and construct a one-step ahead forecast of the realized range, using the algorithm described in Section 2.1. We have three specific purposes with this forecasting experiment. First, we intend to gain some insight into which type of copula function captures the time series dependence in realized volatility most adequately. For that reason, we compare the forecasting performance of the C-RV model for a variety of copula specifications, namely Gaussian, Student's t , Frank, Gumbel, Clayton, and Clayton-survival, as well as mixtures of the Clayton or Gumbel copula with its survival counterpart.

Second, we examine the forecasting performance of C-RV models relative to conventional modeling approaches for realized measures. As a benchmark, we construct forecasts from the logarithmic HAR model of Corsi (2009), given by

$$\log RR_t = \beta_0 + \beta_1 \log RR_{t-1,1} + \beta_2 \log RR_{t-1,5} + \beta_3 \log RR_{t-1,22} + \varepsilon_t, \quad (10)$$

where $\log RR_{t-1,L} \equiv \frac{1}{L} \sum_{l=1}^L \log RR_{t-l}$ is the average logarithmic realized range between days $t - L$ and $t - 1$ (Corsi et al., 2008). The HAR model has been found to be most successful for forecasting empirical measures of integrated volatility in several empirical applications (see Andersen et al. (2007), Corsi et al. (2008), Corsi (2009), and Busch et al. (2011), among others).

Third, the instability in the exceedance probabilities documented in Section 3.1 motivates us to consider *conditional* copulas with time-varying parameters. Following Patton (2006), we specify the dynamics of the copula parameters as a measurable function of past observations, such that estimation and inference for these conditional copulas does not become more complicated than for standard copulas with constant parameters.⁷ To save space, we only report results

⁷Alternative approaches to conditional copulas allow the parameters in a given copula function to vary over time in the form of an autoregressive or Markov switching process (see Jondeau and Rockinger, 2006; Bartram

for conditional Gumbel copula with time-varying parameter θ_t given by

$$\theta_t = \theta_0 + \theta_1 \hat{u}_{t-1} \hat{u}_{t-2},$$

where \hat{u}_t is the normalized EDF of the log realized range for day t .

We assess the out-of-sample forecasting results by means of several performance measures to evaluate and compare the various copula specifications and the HAR benchmark model. First, we consider the mean prediction error (ME), that is

$$ME = \frac{1}{P} \sum_{t=R+1}^{R+P} (y_t - \hat{y}_{t|t-1}),$$

where $\hat{y}_{t|t-1}$ denotes the one-step ahead forecast of the realized measure y for day t . The quality of individual forecasts is further assessed by the Mincer-Zarnowitz type regression

$$y_t = b_0 + b_1 \hat{y}_{t|t-1} + \eta_t. \quad (11)$$

In (11), b_0 and b_1 should be equal to 0 and 1, respectively, for the forecast to be considered efficient. We follow the suggestion of Patton and Sheppard (2009) to estimate the Mincer-Zarnowitz regression using Generalized Least Squares (GLS). In this case, this effectively boils down to estimating b_0 and b_1 with OLS in the regression specification

$$\frac{y_t}{\hat{y}_{t|t-1}} = b_0 \frac{1}{\hat{y}_{t|t-1}} + b_1 + \eta_t.$$

The accuracy of the volatility forecasts is evaluated using the Mean Squared Prediction Error (MSPE), that is

$$MSPE = \frac{1}{P} \sum_{t=R+1}^{R+P} (y_t - \hat{y}_{t|t-1})^2.$$

et al., 2007; Hafner and Manner, 2011, among others). Instead of the copula parameters the copula function itself may also be allowed to vary over time (as in Okimoto, 2008; Chollete et al., 2009; Garcia and Tsafack, 2011). Manner and Reznikova (2011) provide a recent survey.

We directly compare the copula-based forecasts with the benchmark HAR specifications by testing the null hypothesis of equal predictive accuracy with the Diebold and Mariano (1995) statistic. Specifically, let $\hat{y}_{\text{C-RV},t|t-1}$ and $\hat{y}_{\text{HAR},t|t-1}$ denote the two competing one-step ahead forecasts of y_t , and define the loss differential $d_t = e_{\text{HAR},t|t-1}^2 - e_{\text{C-RV},t|t-1}^2$, where $e_{\cdot,t|t-1} = y_t - \hat{y}_{\cdot,t|t-1}$ is the forecast error of the HAR and C-RV models. We then test the null hypothesis of equal predictive accuracy, which corresponds to $E[d_t] = 0$, by means of the t -statistic

$$\text{DM} = \frac{\bar{d}}{\sqrt{\widehat{V}(d_t)/P}}, \quad (12)$$

where \bar{d} is the sample mean of the loss differential d_t and $\widehat{V}(d_t)$ is an estimate of the variance of d_t . Finally, we estimate the forecast encompassing regression

$$y_t = b_0 + b_1 \hat{y}_{\text{HAR},t|t-1} + b_2 \hat{y}_{\text{C-RV},t|t-1} + \eta_t. \quad (13)$$

In case $b_1 = 0$ and $b_2 \neq 0$ the copula-based forecast encompasses the HAR forecast, and vice versa in case $b_1 \neq 0$ and $b_2 = 0$. Again we adopt the suggestion of Patton and Sheppard (2009) to use GLS to estimate the encompassing regression.

Table 1 shows the results based on the complete forecasting period from January 2, 1997 until December 29, 2006. The ME and MSPE are based on forecasts of the square root of the realized range expressed in terms of annualized percentage points. These results lead to fairly clear-cut conclusions concerning the three goals of the forecasting experiment mentioned before. First, we find that the Gumbel copula performs best among the C-RV specifications with constant parameters. Its forecasts are (i) unbiased in the sense that the ME is not significantly different zero, (ii) efficient as we cannot reject the null hypothesis $b_0 = 0$ and $b_1 = 1$ in (11), and (iii) most accurate in the sense that they minimize the MSPE at 12.41. The R^2 of the MZ regression is equal to 0.658, indicating that the Gumbel forecasts explain a fairly large fraction of the variation in the realized range estimate of volatility. In terms of forecast accuracy, the closest competitor is the Clayton survival copula, which achieves an MSPE of 12.55. The

forecasts based on the Clayton survival copula have the undesirable feature of being biased though (as shown by the significantly positive ME, and the null $b_0 = 0$ in (11) being rejected). The superior performance of the Gumbel and Clayton survival copulas suggests that allowing for upper tail dependence is crucial for obtaining accurate volatility forecasts. Interestingly, extending the Gumbel copula to a mixture with the Gumbel survival copula does not improve forecast performance. While the forecasts remain unbiased and efficient, the MSPE increases to 12.85. Allowing for lower tail dependence only, as with the Clayton copula, clearly is not appropriate, leading to an MSPE of 17.48, which is about 40% larger than the MSPE achieved by the Gumbel copula. The elliptical Gaussian and Student's t copulas are not appropriate for describing the temporal dependence of the realized range, in the sense that they produce forecasts that are severely biased ($ME \approx -1.6$) and are rather inaccurate with an MSPE that is twice as large as the MSPE of the Gumbel copula. For the Gaussian copula this may not come as a surprise given that it does not allow for tail dependence, but the Student's t copula may have been expected to perform better.

Second, the better performing copula specifications significantly outperform the benchmark HAR model in terms of forecast accuracy. The HAR forecasts have an MSPE of 13.83, which is significantly larger than the MSPE of the Gumbel copula as well as the mixtures of the Clayton or Gumbel copulas according to the DM test statistic (based on a one-sided 5% significance level). Interestingly, in the encompassing regression (13) for these three copula specifications, we find estimates of b_1 and b_2 that are significantly different from 0 and 1, rejecting the hypothesis that the C-RV model encompasses the HAR model or vice versa. This suggests that it may be worthwhile to combine these forecasts to further improve accuracy. Note that the R^2 of (13) exceeds the R^2 of the MZ regression (11) by only 0.03, such that the possible gains in terms of MSPE seem limited though. We leave the issue of forecast combination for further research.

Third, inspecting the moving window estimates of the copula parameters reveals quite substantial variation over time.⁸ Incorporating this instability in the dependence explicitly in the

⁸Results are not shown here to save space, but are available upon request.

Table 1: Out-of-sample forecast evaluation, January 1997-December 2006.

| | MZ regression | | | | | Encompassing regression | | | | |
|---------|-------------------|-------------------|------------------|-------|-------|-------------------------|-------------------|------------------|------------------|-------|
| | ME | b_0 | b_1 | R^2 | MSPE | DM | b_0 | b_1 | b_2 | R^2 |
| HAR | 0.291 (0.074) | 0.788 (0.221) | 0.958 (0.021) | 0.621 | 13.83 | | | | | |
| Ga | -1.600 (0.094) | 0.741 (0.281) | 0.832 (0.022) | 0.385 | 24.66 | -14.69 | 0.096 (0.225) | 0.099 (0.025) | 0.901 (0.029) | 0.625 |
| St- t | -1.625 (0.094) | 0.763 (0.277) | 0.829 (0.022) | 0.385 | 24.77 | -14.84 | 0.104 (0.224) | 0.097 (0.025) | 0.902 (0.029) | 0.625 |
| Fr | 0.214 (0.073) | -0.157 (0.186) | 1.028 (0.018) | 0.628 | 13.58 | 0.47 | -0.416 (0.178) | 0.605 (0.035) | 0.447 (0.035) | 0.666 |
| Cl | 0.329 (0.083) | 1.021 (0.286) | 0.937 (0.026) | 0.516 | 17.48 | -4.98 | -0.149 (0.196) | 0.335 (0.026) | 0.699 (0.034) | 0.637 |
| Cl-s | 0.098 (0.072) | -0.039 (0.209) | 1.011 (0.019) | 0.638 | 13.13 | 1.17 | -0.518 (0.191) | 0.539 (0.034) | 0.518 (0.030) | 0.680 |
| Cl/Cl-s | 0.181 (0.070) | 0.456 (0.205) | 0.973 (0.019) | 0.653 | 12.55 | 2.59 | -0.177 (0.178) | 0.609 (0.034) | 0.420 (0.032) | 0.681 |
| Gu | 0.137 (0.070) | 0.044 (0.189) | 1.008 (0.018) | 0.658 | 12.41 | 2.44 | -0.367 (0.178) | 0.624 (0.034) | 0.421 (0.031) | 0.688 |
| Gu/Gu-s | 0.121 (0.071) | 0.130 (0.193) | 0.997 (0.018) | 0.646 | 12.85 | 1.94 | -0.317 (0.183) | 0.559 (0.037) | 0.481 (0.035) | 0.677 |
| Gu-c | 0.364 (0.070) | -0.070 (0.189) | 1.034 (0.018) | 0.658 | 12.58 | 2.42 | -0.438 (0.180) | 0.640 (0.035) | 0.421 (0.031) | 0.686 |

Note: The table reports the mean prediction error (ME), GLS estimation results for the Mincer-Zarnowitz (MZ) type regression in (11), the Mean Squared Prediction Error (MSPE), the Diebold-Mariano (DM) statistic defined in (12), and GLS estimation results for the encompassing regression in (13), for one-day ahead forecasts of the square root of the realized range RR_t^M , obtained from (9) based on 5-minute intraday intervals ($M = 81$). The forecasting period runs from January 2, 1997 to December 29, 2006 ($P = 2530$ observations). Acronyms used for the copula specifications are: Ga - Gaussian; St- t - Student's t ; Fr - Frank; Cl - Clayton; Cl-s - Clayton survival; Cl/Cl-s - Clayton-Clayton survival mixture; Gu - Gumbel; Gu-s - Gumbel survival; Gu/Gu-s - Gumbel-Gumbel survival mixture; Gu-c - conditional Gu copula with the time-varying parameter.


copula specification does not improve forecast accuracy. As shown in Table 1, the conditional Gumbel specifications achieves an MSPE of 12.58, slightly larger than the MSPE of the Gumbel copula with constant parameters. Note that the forecasts of the conditional Gumbel copula are biased downward, with a significant ME of 0.358.

Finally, we examine the (in)stability of the (relative) forecasting performance of the C-RV model with its different copula specifications and the log-HAR model. For this purpose, panel (a) in Figure 6 shows the ratio of the MSPE of the volatility forecasts obtained from the C-RV model with the Gumbel and conditional Gumbel copula specifications over the MSPE of the forecasts obtained from the log-HAR model, computed for moving windows of 500 observations. The relative MSPEs are below one except for a short period with moving windows ending in 2004-5, indicating that these C-RV specifications outperform the log-HAR model quite consistently. To put this result into perspective, panel (b) of Figure 6 shows the MSPE of the forecasts obtained from the log-HAR model together with the mean and variance of the square root of the realized range, again computed for moving windows of 500 observations. This graph shows that during 2004-5, both the mean and variance of volatility were rapidly declining. The MSPE of the log-HAR model declines correspondingly. Apparently, the C-RV specifications adjust to this change in market conditions more slowly, such that they are temporarily outperformed by the log-HAR model. On the upside, the copula-based volatility forecasts are considerably more accurate when they matter most, that is, during turbulent times with high and volatile volatility.


4 Conclusions

This paper has introduced a novel approach to modeling and forecasting measures of daily volatility based on copula functions. Copula-based time series models can capture relevant characteristics of volatility such as nonlinear dynamics and long-memory type behavior in a flexible yet parsimonious way. In an empirical application to the daily realized range for S&P500 index futures, we find that the copula-based model outperforms the popular HAR approach for one-day ahead volatility forecasts in terms of accuracy and efficiency. Among the specifications

Figure 6: Forecasting performance over time



(a) Relative MSPEs over moving windows of 500 observations



(b) Mean and variance of RR with MSPE of HAR forecasts

Note: Panel (a) shows the ratio of the MSPE of the volatility forecasts obtained from the C-RV model with the Gumbel and conditional Gumbel copula specifications over the MSPE of the forecasts obtained from the log-HAR model, computed for moving windows of 500 observations. The date on the horizontal axis indicates the end-point of the moving window.

considered, the Gumbel copula achieves the best forecast performance, which highlights the importance of asymmetry and upper tail dependence for modeling volatility dynamics. Although we find substantial variation in the copula parameter estimates over time, conditional copulas do not improve the accuracy of volatility forecasts.

Appendix: Copula specifications

In this appendix we describe the specifications of the bivariate copula functions that we consider in the empirical analysis, together with their quantile and tail dependence properties.

Gaussian copula The Gaussian copula can be obtained using the so-called inversion method, that is

$$C^{\text{Ga}}(u_{t-1}, u_t) = F_b(F(u_{t-1}), F^{-1}(u_t)),$$

where $F^{-1}(u) = \min\{x | u \leq F(x)\}$ is the (quasi)-inverse of the marginal CDF F . The Gaussian copula is obtained by taking F_b to be the bivariate normal distribution with mean zero, unit variances, correlation ρ , and a standard normal marginal F . The corresponding copula density is given by

$$c^{\text{Ga}}(u_{t-1}, u_t) = \frac{1}{\sqrt{1-\rho^2}} \exp\left(-\frac{\Phi^{-1}(u_{t-1})^2 + \Phi^{-1}(u_t)^2 - 2\rho\Phi^{-1}(u_{t-1})\Phi^{-1}(u_t)}{2(1-\rho^2)} - \frac{\Phi^{-1}(u_{t-1})^2 + \Phi^{-1}(u_{t-1})^2}{2}\right), \quad (\text{A.1})$$

The Gaussian copula is symmetric, with $\tau(q) = \tau(1-q)$ for all $0 \leq q \leq 0.5$, but has neither lower nor upper tail independence.

Student's t copula The Student's t copula is also obtained by the inversion method, but using a bivariate Student's t distribution with ν degrees of freedom and correlation ρ instead of the Gaussian. The Student's t copula is symmetric, with $\tau(q) = \tau(1-q)$ for all $0 \leq q \leq 0.5$, and has equal lower and upper tail dependence with $\tau_L = \tau_U = 2T_{\nu+1}(-\sqrt{(\nu+1)(1-\rho)/(1+\rho)})$.

Frank copula The Frank (Fr) copula is given by

$$C^{\text{Fr}}(u_{t-1}, u_t) = -\frac{1}{\theta} \log \left(1 + \frac{(\exp(-\theta u_{t-1}) - 1)(\exp(-\theta u_t) - 1)}{\exp(-\theta) - 1} \right), \quad \text{with } \theta \in \mathbb{R}. \quad (\text{A.2})$$

The Frank copula is symmetric, with $\tau(q) = \tau(1 - q)$ for all $0 \leq q \leq 0.5$, and has lower and upper tail independence.

Clayton copula The Clayton (Cl) copula is defined by

$$C^{\text{Cl}}(u_{t-1}, u_t) = (u_{t-1}^{-\theta} + u_t^{-\theta} - 1)^{-1/\theta}, \quad \text{with } 0 < \theta < \infty. \quad (\text{A.3})$$

The Clayton copula is asymmetric, with $\tau(q) > \tau(1 - q)$ for all $0 \leq q \leq 0.5$; it has zero upper tail dependence and positive lower tail dependence with $\tau_L = 2^{-1/\theta}$.

Clayton survival copula The Clayton survival (Cl-s) copula is the mirror image of the Clayton copula:

$$C^{\text{Cl-S}}(u_{t-1}, u_t) = u_{t-1} + u_t - 1 + C^{\text{Cl}}(1 - u_{t-1}, 1 - u_t). \quad (\text{A.4})$$

It follows from the properties of the Clayton copula that the Clayton survival copula is asymmetric with $\tau(q) < \tau(1 - q)$ for all $0 \leq q \leq 0.5$. Cl-s has zero lower tail dependence and positive upper tail dependence with $\tau_U = 2^{-1/\theta}$.

Clayton-Clayton survival mixture copula The Clayton-Clayton survival mixture (Cl-mix) copula is a convex sum of Cl and Cl-s copulas:

$$C^{\text{Cl-mix}}(u_{t-1}, u_t) = (1 - \lambda)C^{\text{Cl}}(u_{t-1}, u_t) + \lambda C^{\text{Cl-S}}(u_{t-1}, u_t), \quad (\text{A.5})$$

where $0 \leq \lambda \leq 1$ denotes the mixing parameter, and the parameters θ_{Cl} and $\theta_{\text{Cl-S}}$ in the Clayton and Clayton survival copulas are allowed to be different. The quantile dependence properties of the Clayton-Clayton survival mixture copula depend on the mixture weight λ . It can accommodate both lower and upper tail dependence with $\tau_L = (1 - \lambda)2^{-1/\theta_{\text{Cl}}}$ and $\tau_U = \lambda 2^{-1/\theta_{\text{Cl-S}}}$.

Gumbel copula The Gumbel (Gu) copula is given by

$$C^{\text{Gu}}(u_{t-1}, u_t) = \exp\left(-\left[(-\log u_{t-1})^\theta + (-\log u_t)^\theta\right]^{\frac{1}{\theta}}\right), \quad \text{with } 1 \leq \theta < \infty. \quad (\text{A.6})$$

The Gumbel copula is asymmetric, with $\tau(q) < \tau(1 - q)$ for all $0 \leq q \leq 0.5$, and has zero lower tail dependence and positive upper tail dependence with $\tau_U = 2 - 2^{1/\theta}$.

The Gumbel survival and Gumbel-Gumbel survival mixture copulas are defined analogously to the case of the Clayton copula.

References

- Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics*, 89:701–720.
- Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2009). Parametric and nonparametric measurement of volatility. In Aït-Sahalia, Y. and Hansen, L. P., editors, *Handbook of Financial Econometrics*, pages 67–138. Elsevier: Amsterdam.
- Bartram, S. M., Taylor, S. J., and Wang, Y.-H. (2007). The euro and European financial market dependence. *Journal of Banking and Finance*, 31:1461–1481.
- Beare, B. K. (2010). Copulas and temporal dependence. *Econometrica*, 78:395–410.
- Busch, T., Christensen, B. J., and Nielsen, M. Ø. (2011). The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets. *Journal of Econometrics*, 160:48–57.
- Chen, X. and Fan, Y. (2006). Estimation of copula-based semiparametric time series models. *Journal of Econometrics*, 130:307–335.
- Chen, X., Wu, W. B., and Yi, Y. (2009). Efficient estimation of copula-based semiparametric markov models. *Annals of Statistics*, 37:4214–4253.
- Cherubini, U., Luciano, E., and Vecchiato, W. (2004). *Copula Methods in Finance*. Chicester: John Wiley.
- Chollete, L., Heinen, A., and Valdesogo, A. (2009). Modelling international financial returns with a multivariate regime switching copula. *Journal of Financial Econometrics*, 7:437–480.
- Christensen, K. and Podolskij, M. (2007). Realized range-based estimation of integrated variance. *Journal of Econometrics*, 141:323–349.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7:174–196.
- Corsi, F., Mittnik, S., and Pigorsch, U. (2008). The volatility of realized volatility. *Econometric Reviews*, 27:46–78.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13:253–263.
- Fermanian, J.-D. and Scaillet, O. (2005). Some statistical pitfalls in copula modeling for financial applications. In Klein, E., editor, *Capital Formation, Governance and Banking*, pages 57–72. Nova Science Publishing.
- Fleming, J., Kirby, C., and Ostdiek, B. (2003). The economic value of volatility timing using realized volatility. *Journal of Financial Economics*, 67:473–509.

- Garcia, R. and Tsafack, G. (2011). Dependence structure and extreme comovements in international equity and bond markets. *Journal of Banking and Finance*, 35:1954–1970.
- Genest, C., Gendron, M., and Bourdeau-Brien, M. (2009). The advent of copulas in finance. *European Journal of Finance*, 15:609–618.
- Hafner, C. M. and Manner, H. (2011). Dynamic stochastic copula models: Estimation, inference and applications. *Journal of Applied Econometrics*, page forthcoming.
- Hansen, P. R. and Lunde, A. (2005). A realized variance for the whole day based on intermittent high-frequency data. *Journal of Financial Econometrics*, 3:525–554.
- Ibragimov, R. and Lentzas, G. (2008). Copulas and long memory. Harvard Institute of Economic Research Discussion Paper No. 2160.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Monographs on Statistics and Applied Probability 73, London: Chapman & Hall.
- Jondeau, E. and Rockinger, M. (2006). The copula-GARCH model of conditional dependencies: An international stock market application. *Journal of International Money and Finance*, 25:827–853.
- Manner, H. and Reznikova, O. (2011). A survey on time-varying copulas: Specification, simulations, and estimation. *Econometric Reviews*, page forthcoming.
- Martens, M. (2002). Measuring and forecasting s&p500 index-futures volatility using high-frequency data. *Journal of Futures Markets*, 22:497–518.
- Martens, M. and van Dijk, D. (2007). Measuring volatility with the realized range. *Journal of Econometrics*, 138:181–207.
- Martens, M., van Dijk, D., and de Pooter, M. (2009). Forecasting S&P 500 volatility: Long memory, level shifts, leverage effects, day-of-the-week seasonality, and macro-economic announcements. *International Journal of Forecasting*, 25:282–303.
- McAleer, M. and Medeiros, M. C. (2008a). A multiple regime smooth transition heterogeneous autoregressive model for long memory and asymmetries. *Journal of Econometrics*, 147:104–119.
- McAleer, M. and Medeiros, M. C. (2008b). Realized volatility: A review. *Econometric Reviews*, 27:10–45.
- Nelsen, R. B. (2006). *An Introduction to Copulas*. New York: Springer-Verlag, second edition.
- Ning, C., Xu, D., and Wirjanto, T. (2010). Modeling asymmetric volatility clusters using copulas and high frequency data. Ryerson Working Paper No. 006.
- Okimoto, T. (2008). New evidence of asymmetric dependence structures in international equity markets. *Journal of Financial and Quantitative Analysis*, 43:787–815.

- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47:527–556.
- Patton, A. J. (2009). Copula-based models for financial time series. In Andersen, T., Davis, R., Kreiss, J.-P., and Mikosch, T., editors, *Handbook of Financial Time Series*, pages 767–786. Berlin: Springer-Verlag.
- Patton, A. J. and Sheppard, K. (2009). Evaluating volatility and correlation forecasts. In Andersen, T., Davis, R., Kreiss, J.-P., and Mikosch, T., editors, *Handbook of Financial Time Series*, pages 801–838. Berlin: Springer-Verlag.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231.