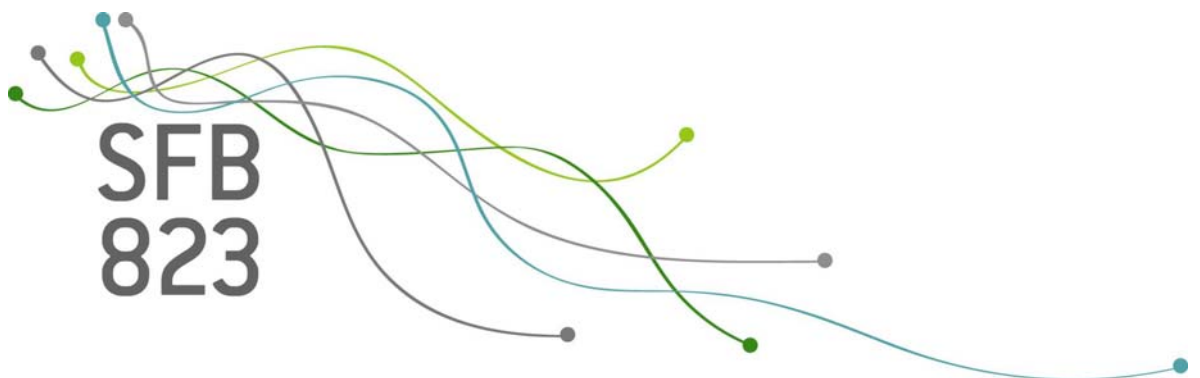# Estimation of risk measures in energy portfolios using modern copula techniques

Stefan Jäschke

SFB
823

Estimation of risk measures in energy portfolios using modern copula techniques

Stefan Jäschke

Email: stefan.jaeschke@rwe.com

Phone: + 49 162 2596816

Department of Statistics

University of Dortmund

Vogelpothsweg 78

44227 Dortmund

Germany

RWE Supply & Trading GmbH

Risk Valuation

Altenessener Str. 27

45141 Essen

Germany

October 7, 2012

Abstract

The present paper analyses the dependence structure between WTI and Brent crude oil spot log-returns using modern copula techniques. In a first step we apply several single equation models to the marginals to account for autocorrelation and volatility clustering. Second, to select both copulas and tail copulas characterising the joint dynamics between the time series we implement and evaluate newly introduced bootstrap-based goodness-of-fit tests. Based on each approach, a comprehensive backtesting is performed by simulating and comparing the risk measures Value-at-Risk and Expected Shortfall with historical values.

Keywords: Crude oil; Joint dynamics; Tail copulas; Goodness-of-fit tests; Backtesting; Risk measures

JEL classification: C22; C46; C51; C63; G11; G32

1 Introduction

The modelling of stochastic dependence in energy and commodity markets via copulas has become increasingly common in applications. Grégoire et al. (2008) analyse the dependence structure of log-returns of futures on crude oil and natural gas, Accioly and Aiube (2008) study the co-movement of crude oil and gasoline prices, while Reboredo (2011) focuses on the dependence structure between crude oil benchmark prices. Using weekly data the article examines whether crude oil markets are rather globalised or regionalised. Having estimated each log-return time series individually, the author accounts for the dependence between the different crude oil grades by fitting various copula families to the error terms. The unknown parameters of the copula functions are obtained via maximum likelihood whereas the decision which model performs best relies on an adjusted information criterion and a pseudo-likelihood ratio test. Additionally, and directly linked to the estimated copula parameters, coefficients for upper and lower tail dependence are provided.

According to Jäschke et al. (2012) a general goodness-of-fit test for copulas does not necessarily provide a good model of tail dependence as most of these procedures take the whole support of the distribution into account and therefore, adopting copula inference techniques for modelling joint extreme events can be very misleading for risk management purposes. Using tail copulas Jäschke et al. (2012) account for all possible scenarios of joint extreme outcomes and thus decouples the decision to select an appropriate model describing the overall dependence from the analysis of the joint dynamics in the tail of the underlying distribution.

The present paper follows this route and extends the current state of knowledge within a couple of areas. First and foremost, we apply the partial derivatives multiplier bootstrap goodness-of-fit test for tail copulas (Bücher and Dette, 2012) to the log-returns of two crude oil grades. To better compare the findings to a traditional copula fit we present both the best model characterising the overall dependence structure using a bootstrap-based goodness-of-fit test for copulas and the newly introduced Copula Information Criterion (Grønneberg and Hjort, 2012). Moreover, from a risk management perspective it is indispensable to capture the joint behaviour of certain assets within energy portfolios. More precisely, following Jäschke et al. (2012), a thorough understanding of energy portfolio risk requires an adequate assessment of the probability that large negative co-movements occur together, i. e. lower tail dependence. Therefore, the present paper introduces a wide and comprehensive backtesting framework for two of the most

commonly applied risk measures, i. e. Value-at-Risk and Expected Shortfall. To the best of our knowledge, the problem of modelling and assessing risk measures in energy portfolios using the described modern copula techniques has not yet been addressed before. The paper's claim to cover a topic of broad applicability and high practical relevance is backed by extensive guidelines from a practitioner's point of view as well as detailed comments concerning the empirical implementation.

The present analysis uses a large data set of daily quotes of WTI Cushing Crude Oil Spot and the Bloomberg European Dated Brent. Quotes of WTI are commonly used as a reference spot price for U.S. crude oil whereas the price for Brent serves as a benchmark for European crude oil. Both oil grades belong to the class of light and sweet crude oils, i. e. they are characterised by their low density and their low sulphur content. Although Brent is not as light or as sweet as WTI, it is still a high-grade crude for subsequent processing and thus, both crude oils are not completely, but to a great extent, generally treated as interchangeable.

Due to these differences in quality, WTI futures were usually traded at a small premium to Brent futures. From early on in 2011 the development of the spread between the two futures has attracted the energy and finance markets' attention. In fact, Brent did not only trade over WTI, but also the spread of the futures widened notably, which is reason enough to analyse the dependence structure between the two crude oil grades from a probabilistic point of a view.

The paper is organised as follows: Section 2 briefly introduces the well-known copula framework and reviews the closely related concept of tail copulas with their corresponding non-parametric estimates. Section 3 models the individual time series of both crude oil grades, accounting for serial dependence in the data. Section 4 describes the applied rank-based goodness-of-fit tests for copulas and tail copulas, respectively. Section 5 provides an instruction how to put the described theory into practise. Section 6 backtests the findings using the example of risk measures. Finally, Section 7 summarises the results and the used methods.

2   Preliminaries

2.1   Copulas and tail dependence

The theory of copulas investigates the dependence structure of multivariate distribution functions. As this article focuses on co-movements between the WTI and Brent crude oil futures, we state all further def-

4

initions and results for the bivariate case only. From a probabilistic perspective, a copula is a joint distribution function with uniformly distributed margins on the interval $[0, 1]$. To begin with, we consider a random vector (X, Y) with continuous marginal distribution functions $F(x) := \mathbb{P}[X \leq x]$ and $G(y) := \mathbb{P}[Y \leq y]$, $x, y \in \mathbb{R}$, respectively. The theoretical foundation for the application of copulas is provided by Sklar's theorem (Sklar, 1959), according to which there exists a unique copula C, called the copula of X and Y, such that

$$\mathbb{P}[X \leq x, Y \leq y] = C(F(x), G(y)), \tag{1}$$

for all $x, y \in \mathbb{R}$. Conversely, if C is a copula and F and G are distribution functions, then the function defined by Equation (1) is a bivariate distribution function with margins F and G, respectively. It can be shown that

$$C(u, v) = \mathbb{P}[U \leq u, V \leq v] = \mathbb{P}[X \leq F^{(-1)}(u), Y \leq G^{(-1)}(v)], \tag{2}$$

for all $u, v \in [0, 1]$, where $F^{(-1)}$ and $G^{(-1)}$ denote the quasi-inverses of F and G, respectively, i. e.

$$F^{(-1)}(u) = \inf\{x \in \mathbb{R} \mid F(x) \geq u\}, \tag{3}$$

for all $u \in [0, 1]$ (analogously for G). Summing up, copulas allow for a step-wise modelling of the joint distribution whose dependence structure is independent of the respective marginal distributions. Further details can be found in Nelsen (2006).

Having set the copula framework, we are now able to introduce the concept of tail copulas which provides a generalised approach to model dependencies of extreme events. Following Schmidt and Stadtmüller (2006), the lower tail copula Λ_L associated with X and Y is a function of their copula C and is defined by

$$\Lambda_L(x, y) := \lim_{t \to 0^+} \frac{C(tx, ty)}{t}, \tag{4}$$

if the above limit exists for all $(x, y) \in \mathbb{R}_+^2 := [0, \infty]^2 \setminus \{(\infty, \infty)\}$. Considered analytically, Λ_L describes the directional derivative of the copula C along the vector $(x, y) \in \mathbb{R}_+^2$ at the point $(0, 0)$. According to Schmidt and Stadtmüller (2006), the lower tail copula admits the homogeneous structure among its components, i. e.

$$\Lambda_L(sx, sy) = s\Lambda_L(x, y), \quad s \geq 0, \tag{5}$$

for all $(x, y) \in \mathbb{R}_+^2$. Regarding this property, we conclude that evaluating the lower tail copula on the unit circle in the first quadrant contains

all information about the directional derivatives of the copula C at the origin. Once these values are calculated the missing ones can be obtained by linear continuation according to Equation (5). Moreover, in the following we assume that the lower tail copula is non-zero in some point $(x_0, y_0) \in \mathbb{R}_+^2$ and thus non-zero everywhere on $\mathbb{R}_+^2$ (see Schmidt and Stadtmüller, 2006, Theorem 1).

The next part allows for embedding the well-known concept of tail dependence (see among others Caillault and Guéégan, 2005; Frahm et al., 2005; Chui and Wu, 2009) into the framework of tail copulas. More precisely, two random variables X and Y are called lower tail dependent if the limit

$$\lambda_L := \Lambda_L(1,1) = \lim_{u \to 0^+} \mathbb{P}[X \leq F^{(-1)}(u) \mid Y \leq G^{(-1)}(u)] \qquad (6)$$

exists and lies in the interval $(0,1]$. In case $\lambda_L = 0$, X and Y exhibit no lower tail dependence. Regarding the scope of this article, tail dependence can be viewed as the limiting likelihood of an commodity return falling below its Value at Risk at a certain level, given that another commodity return has fallen below its Value at Risk at the same level. As pointed out in Jäschke et al. (2012), the so-called lower tail dependence coefficient λ_L is not able to thoroughly capture the joint behaviour in the lower tail of the distribution. Nevertheless, it is widely used as a simple and intuitive scalar measure for dependence between extreme losses of commodity returns.

2.2 Non-parametric estimates

In preparation for the rank-based goodness-of-fit tests in Section 4 we introduce non-parametric estimators for the copula C and the lower tail copula Λ_L, respectively. For this purpose we consider a sample $(X_1, Y_1), \ldots, (X_T, Y_T)$ from a pair (X, Y) of continuous random variables. Then the normalised rank-based representation (u_t, v_t) of the sample data (X_t, Y_t), $t = 1, \ldots, T$, is defined by

$$u_t = \frac{\mathrm{rank}(X_t)}{T+1} \quad \text{and} \quad v_t = \frac{\mathrm{rank}(Y_t)}{T+1}. \qquad (7)$$

Assuming a time-independent dependence structure, the empirical copula $\widehat{C}_T$ (Deheuvels, 1979) is given by

$$\widehat{C}_T(u,v) = \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}(u_t \leq u, v_t \leq v), \qquad (8)$$

where $\mathbb{1}(A)$ denotes the indicator function of a set A. As argued in Genest and Favre (2007), the empirical Copula $\widehat{C}_T$ is the best sample-

based non-parametric representation of the copula C, which itself characterises the dependence structure of the random vector (X, Y). Moreover, for any given pair (u, v), $\widehat{C}_T(u, v)$ serves as a consistent rank-based estimator of $C(u, v)$ that is asymptotically normally distributed with mean $C(u, v)$ (see Segers, 2012a).

To be able to adequately assess the risk of joint extreme losses, we now introduce the lower empirical tail copula $\hat{\Lambda}_L$ according to Schmidt and Stadtmüller (2006). This non-parametric estimator is defined by

$$
\begin{aligned}
\hat{\Lambda}_L(x, y) &:= \frac{T}{k} \widehat{C}_T \left(\frac{kx}{T}, \frac{ky}{T} \right) \\
&\approx \frac{1}{k} \sum_{t=1}^{T} \mathbb{1} \left(u_t \leq \frac{kx}{T+1}, v_t \leq \frac{ky}{T+1} \right)
\end{aligned}
\tag{9}
$$

with some parameter $k \in \{1, \dots, T\}$ to be chosen by the statistician. As proved in Schmidt and Stadtmüller (2006), the estimator $\hat{\Lambda}_L$ exhibits weak convergence and strong consistency provided that $k = k(T) \to \infty$ and $k/T \to 0$ for $T \to \infty$ (and other regularity conditions, quod vide Huang, 1992). From an analytical perspective the lower empirical tail copula $\hat{\Lambda}_L$ can be viewed as the slope of the secant of the empirical copula $\widehat{C}_T$ containing the points $(0, 0)$ and $(kx/T, ky/T)$ for $x^2 + y^2 = 1$, $x, y > 0$.

At the same time, the lower empirical tail copula yields a convenient method for the non-trivial task of estimating the lower tail dependence coefficient λ_L (see Schmidt and Stadtmüller, 2006). To be precise,

$$
\hat{\lambda}_L := \hat{\Lambda}_L(1, 1),
\tag{10}
$$

which again illustrates that the lower tail copula is an intuitive generalisation of the lower tail dependence coefficient.

3 Modelling marginal time series

Since the copula is a function of the marginal distributions, an adequate modelling of the individual time series is crucial for estimating the dependence structure between two commodities. Recently, several empirical studies have analysed the modelling of univariate time series for crude oil spot and futures prices considering different data frequencies (see among others Kang et al., 2009; Mohammadi and Su, 2010; Chang et al., 2010). For the purpose of the present study it is preferable to use daily returns over weekly returns to investigate joint extreme events, as weekly returns tend to smoothen returns especially in the tails of the distribution whose structure we are trying to

capture here. Consequently, our data set covers daily closing quotes of the WTI Cushing Crude Oil Spot and the Bloomberg European Dated Brent from October 2, 2006 to October 1, 2010, collected from Bloomberg's Financial Information Services.

First, to test the time series for weak stationarity we perform unit root tests on the logarithmic spot prices. Following the strategy proposed in Perron (1988), the realised Augmented Dickey-Fuller tests indicate that the null hypothesis of a unit root cannot be rejected for both WTI Crude Oil and European Brent. In a second step we apply the test to the first differences of the log time series. Here, the null hypothesis of nonstationarity can clearly be rejected at the 0.1% level, i. e. the log time series of both crude oil commodities are integrated of order one. Accordingly, the following investigation focuses on the log-returns $r_t = \log(P_t) - \log(P_{t-1})$. Figure 1 (see over) shows the log-returns for both series.

To check for temporal dependence within the individual time series, we apply standard Ljung-Box tests to the observed and squared observed log-returns for three different lags (lag 1, lag 5 and lag 10). While for both series, in the case of squared observations, the null hypothesis that none of the autocorrelation coefficients up to the specified lag are different from zero can clearly be rejected at the 0.1% level of significance, the p-values for the non-squared observations are given by $(0.239, 0.017, 0.004)$ for the WTI log-returns and $(0.167, 0.121, 0.052)$ for the Brent log-returns, respectively. Furthermore, applying the Lagrange multiplier test proposed by Engle (1982) rejects the null hypothesis of no ARCH effects for all lags at the 0.1% level.

Summing up, the findings show that the assumption of an i.i.d. sample is unrealistic. Thus, to account for autocorrelation and volatility clustering in the marginal series we employ ARMA(a,b)-processes, $a, b \in \{0, 1\}$, for modelling the conditional mean equation in combination with GARCH(p,q), APARCH(p,q), GJRGARCH(p,q) and EGARCH(p,q)-models where $p, q \in \{1, 2, 3\}$. Additionally, following Liu (2011), we introduce exogenous variables into the conditional mean equation to explain the log-return time series. As energy prices are highly affected by imbalances of supply and demand, a representative indicator for this balance is given by the change in the energy inventory. Consequently, the applied model uses weekly U.S. stock data of commercial crude oil which is released by the Department of Energy. While Liu (2011) fits weekly published figures of the energy inventory data are fitted to daily granularity using piecewise linear interpolation, we assume that the market is merely influenced on a single trading day, i. e. the weekly publication date. For details see Kemp (2010) and Fattouh (2010). Besides, the modelling allows various distributions
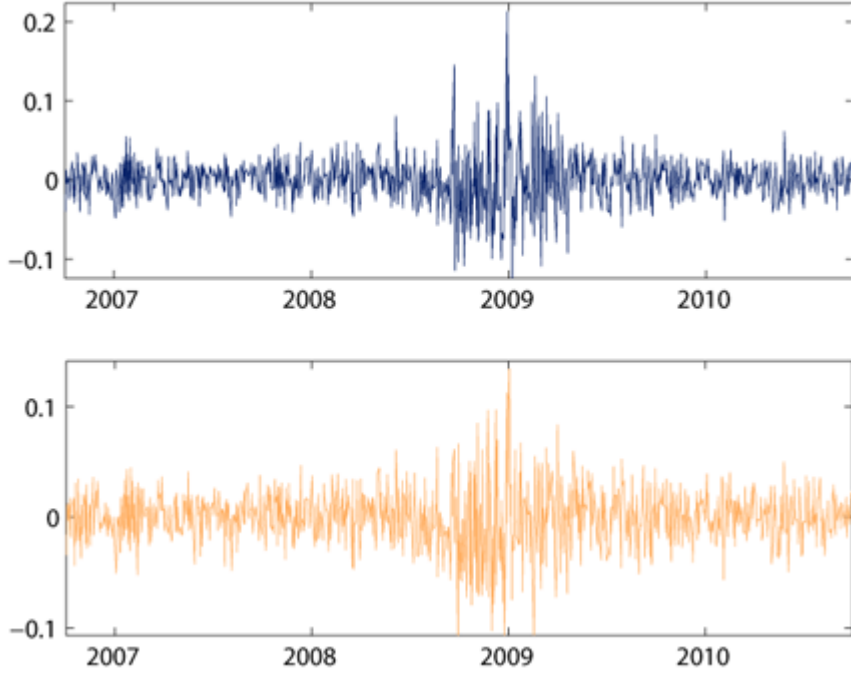
Figure 1: Log-returns of WTI Cushing Crude Oil Spot (top) and Bloomberg European Dated Brent (bottom).

for the error term, specifically the normal, the generalized error, the Student's t-distribution and their skewed extensions.

For illustration, we next consider the ARMA(a, b)-EGARCH(p, q)-model from Nelsen (1991) with underlying fundamental factors:

$$r_t = \mu + \delta\mu_t + \sum_{i=1}^{a} \theta_i r_{t-i} + \varepsilon_t + \sum_{j=1}^{b} \psi_j \varepsilon_{t-j} \tag{11}$$

where

$$\mu, \delta \in \mathbb{R}, \quad a, b \in \mathbb{N},$$
$$\theta_i \in \mathbb{R}, \quad i = 1, \ldots, a, \quad \psi_j \in \mathbb{R}, \quad j = 1, \ldots, b$$

and μ_t denotes the log-returns of the U.S. crude oil stock. The conditional variance process of the innovations $\varepsilon_t = \sigma_t z_t$ is given by

$$
\begin{aligned}
\log\left(\sigma_t^2\right) = {} & \omega + \sum_{i=1}^{p} \alpha_i z_{t-i} + \gamma_i(|z_{t-i}| - E(|z_{t-i}|)) \\
& + \sum_{j=1}^{q} \beta_j \log\left(\sigma_{t-j}^2\right)
\end{aligned}
\tag{12}
$$

9

where

$$p, q \in \mathbb{N}, \quad \omega > 0,$$
$$\alpha_i \geq 0, \quad i = 1, \ldots, p, \quad \beta_j \geq 0, \quad j = 1, \ldots, q,$$
$$\gamma_i \in \mathbb{R}, \quad i = 1, \ldots, p,$$

and $E(|z_t|)$ denotes the unconditional expected value of the absolute standardised innovations $|z_t|$.

The selection of an appropriate model requires criteria which are a priori specified. To be precise, all estimated parameters should be significant at the 1% level. Furthermore, standard Box-Pierce and Ljung-Box tests on the standardised and squared standardised residuals for three different lags (lag 1, lag 5 and lag 10) should indicate that the null hypothesis, that none of the autocorrelation coefficients up to the certain lag are different from zero, cannot be rejected at the 10% level. The same threshold should be valid when applying Engle's Lagrange multiplier test for all lags up to and including lag 10. In addition we consider the information criteria AIC, BIC, SIC and HQIC as well as the corresponding QQ-plot of the standardised residuals. The modelling is carried out using the functionality of Ghalanos (2011).

Taking these previously mentioned criteria into account we conclude that the ARMA(0,0)-EGARCH(2,3)-model including the explanatory variable and the skewed generalized error distribution adequately describes the data generating process $W_t, t = 1, \ldots, T$ of the daily WTI spot log-returns:

$$W_t = 2.596 \times 10^{-4} - 3.216 \times 10^{-1} \mu_t$$
$$+ \sigma_t X_t,$$

$$\log \left(\sigma_t^2 \right) = -3.438 \times 10^{-1} - 1.260 \times 10^{-1} X_{t-1}$$
$$- 1.267 \times 10^{-1} X_{t-2}$$
$$+ 2.458 \times 10^{-1} (|X_{t-1}| - 7.961 \times 10^{-1})$$
$$+ 2.243 \times 10^{-1} (|X_{t-2}| - 7.961 \times 10^{-1})$$
$$- 6.701 \times 10^{-1} \log \left(\sigma_{t-1}^2 \right) + 9.745 \times 10^{-1} \log \left(\sigma_{t-2}^2 \right)$$
$$+ 6.494 \times 10^{-1} \log \left(\sigma_{t-3}^2 \right). \tag{13}$$

Consequently, the standardised residuals $X_1, \ldots, X_T$ can be viewed as a random sample from a skewed generalized error distribution with skewness parameter $\xi = 0.958$ and shape parameter $\nu = 1.960$; for details on the skewed generalized error distribution see Würtz et al. (2006).

At the same time, the ARMA(1,1)-EGARCH(2,3)-model including the explanatory variable and the skewed generalized error distribution provides the best fit to the daily Brent spot log-returns
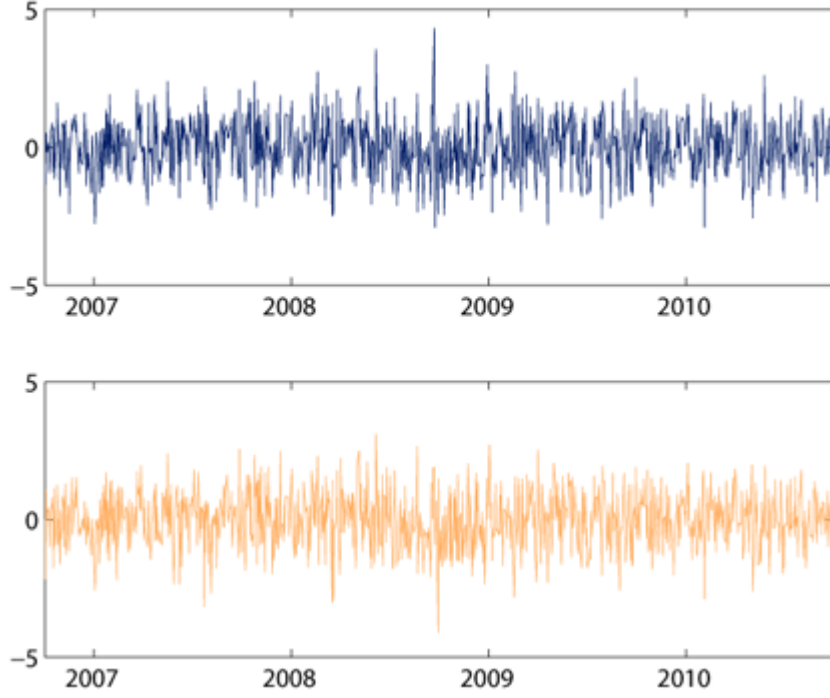
Figure 2: Standardised residuals of both WTI Cushing Crude Oil Spot (top) and Bloomberg European Dated Brent (bottom) log-return time series.

$B_t, t = 1, \ldots, T$:

$$
\begin{aligned}
B_t = {}& 3.064 \times 10^{-4} - 2.315 \times 10^{-1}\mu_t \\
& - 7.893 \times 10^{-1}B_{t-1} \\
& + \sigma_t Y_t + 7.641 \times 10^{-1}\sigma_{t-1}Y_{t-1},
\end{aligned}
$$

$$
\begin{aligned}
\log\left(\sigma_t^2\right) = {}& -2.803 \times 10^{-1} - 1.026 \times 10^{-1}Y_{t-1} \\
& - 1.076 \times 10^{-1}Y_{t-2} \\
& + 1.343 \times 10^{-1}(|Y_{t-1}| - 7.920 \times 10^{-1}) \\
& + 8.270 \times 10^{-2}(|Y_{t-2}| - 7.920 \times 10^{-1}) \\
& - 7.586 \times 10^{-1}\log\left(\sigma_{t-1}^2\right) + 9.098 \times 10^{-1}\log\left(\sigma_{t-2}^2\right) \\
& + 8.117 \times 10^{-1}\log\left(\sigma_{t-3}^2\right).
\end{aligned}
\tag{14}
$$

Therefore, we conclude that the standardised residuals $Y_1, \ldots, Y_T$ are randomly drawn from a skewed generalized error distribution with skewness parameter $\xi = 0.921$ and shape parameter $\nu = 1.873$. Figure 2 exhibits the standardised residuals which provide the basis for

11

the following investigations for both WTI and Brent log-return time series. The corresponding QQ-plots are shown in Figure 3 (see over). Although a few outliers do not fit the straight line, it seems safe to conclude that the respective residuals follow the specified distributions.

4 Rank-based goodness-of-fit tests

Having modelled the individual time series, we are now able to investigate the dependence structure between the WTI and Brent crude oil log-returns. For this, we first apply a one-level bootstrap-based goodness-of-fit test for copulas (Genest and Rémillard, 2008) which takes the whole support of the joint distribution into account. In a second step, we directly estimate the lower tail copula with a multiplier bootstrap along the lines of Bücher and Dette (2012). For an introduction to bootstrap-based goodness-of-fit tests see Stute et al. (1993). The next subsections show how to implement these procedures.

4.1 Bootstrap procedures for copulas

There are several rank-based tests of the appropriateness of copula families when modelling the dependence structure between two random variables. Following Genest and Rémillard (2008) we consider an open set $\mathcal{O} \subseteq \mathbb{R}$ and the parametric copula class

$$\mathcal{C} = \{C_\theta \mid \theta \in \mathcal{O}\}.$$

Hence, the test for the null hypothesis that C_θ belongs to a certain parametric class is given by

$$\mathcal{H}_0 \colon C_\theta \in \mathcal{C}, \quad \mathcal{H}_1 \colon C_\theta \notin \mathcal{C}. \tag{15}$$

The test statistic sums up the squared deviations between the empirical copula $\widehat{C}_T$ and $C_{\hat{\theta}_T}$, vis.

$$S_T^* = \sum_{t=1}^{T} \left(\widehat{C}_T(u_t, v_t) - C_{\hat{\theta}_T}(u_t, v_t) \right)^2, \tag{16}$$

where (u_t, v_t) denotes the rank-based representation of the standardised residuals (X_t, Y_t), see Equation (7).

To compute the optimal parameter estimate $\hat{\theta}_T$ for a given copula family we use a monotone relation between the parameter θ and Kendall's tau, one of the commonly applied rank-based correlation
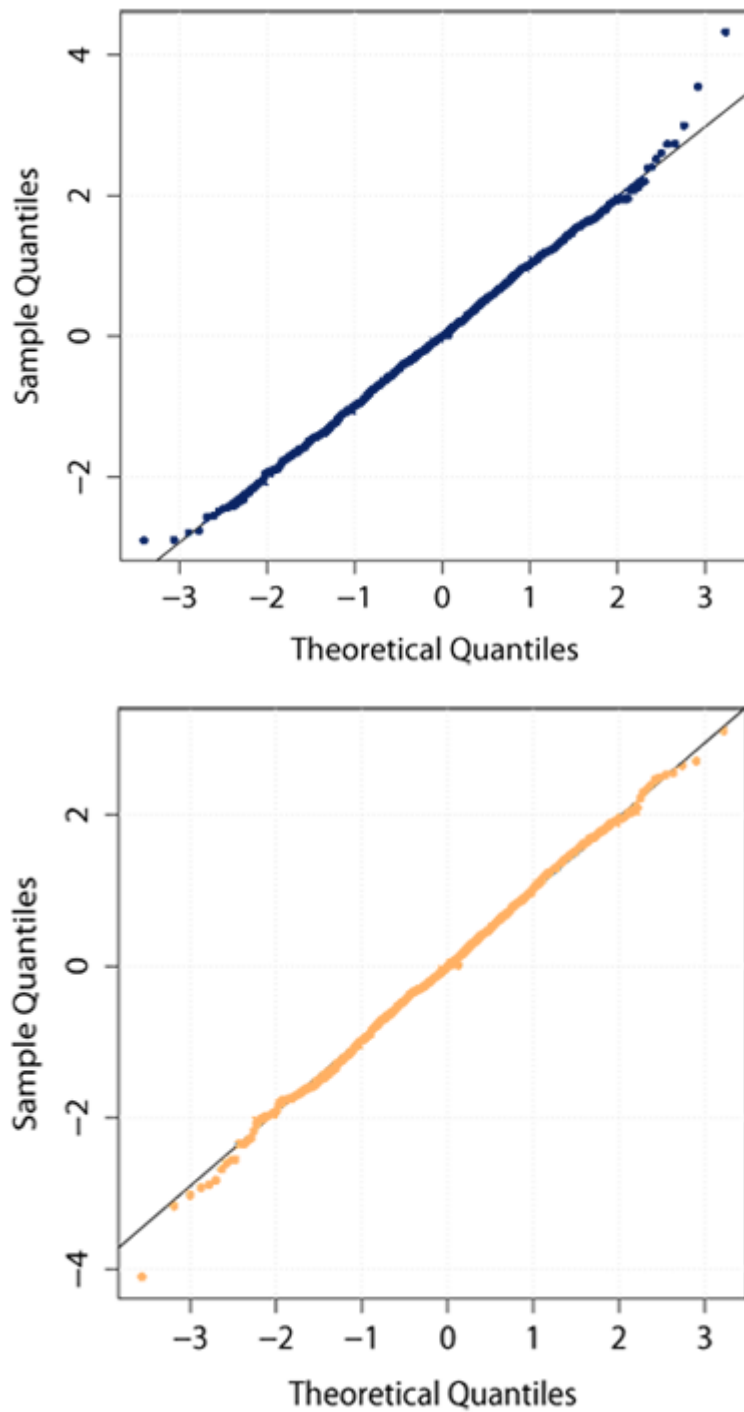
Figure 3: QQ-plots of standardised residuals for WTI Cushing Crude Oil Spot (top) and Bloomberg European Dated Brent (bottom).

13

measures (Nelsen (2006), chap. 5). For example, in case of all non-degenerated elliptical copulas – in particular the Gaussian and the t-copula – the correlation parameter θ can be calculated by

$$\theta = \sin\left(\tau\pi/2\right). \tag{17}$$

For the t-copula we estimate the degrees of freedom ν by maximum likelihood, keeping the correlation parameter θ constant, i. e.

$$\hat{\nu}_T = \arg\max_{\nu\in(2,\infty)} \sum_{t=1}^{T} \log\left(c_{\theta,\nu}^t(u_t, v_t)\right), \tag{18}$$

where $c_{\theta,\nu}^t$ denotes the density of a t-copula with correlation parameter θ and degrees of freedom ν. This method was chosen for its consistency with the other estimators as described above and also because it was found to give very similar estimates to the full maximum likelihood procedure (Demarta and McNeil, 2005). Evaluating the empirical equivalent of Kendall's tau yields $\hat{\tau}_T = 0.705$, and thus $\hat{\theta}_T = 0.895$ and $\hat{\nu}_T = 0.531$, respectively. For comparison only, we find Spearman's rho, another rank-based correlation measure, $\hat{\rho}_T = 0.878$, and Pearson's product-moment correlation coefficient $\hat{r}_T = 0.875$. See Genest and Favre (2007) for a detailed description of these quantities.

Figure 4 (see over) shows a scatter plot of the 1,001 pairs (u_t, v_t) of standardised residuals. It reveals a strong tendency of u_t and v_t (and thus of X_t and Y_t) to vary together, without regarding their marginal distributions. This is not surprising as the log-returns originate from similar types of crude oil. More importantly, on the other hand, we notice that the dependence structure between the two commodities differs clearly from the Fréchet-Hoeffding upper bound, namely $C(u, v) = \min(u, v)$, which would indicate that Y is almost surely an increasing function of X. The latter fact suggests a small opportunity for diversification in the current framework, i. e. reducing the unsystematic risk by investing in a variety of assets.

The following one-level parametric bootstrap-based goodness-of-fit test yields approximated p-values assuming that the copula belongs to the selected parametric copula class. The algorithm (Genest and Rémillard, 2008) proceeds as follows:

1. transform the standardised residuals $(X_t, Y_t), t = 1, \ldots, T$ into their rank-based representation vectors (u_t, v_t), $t = 1, \ldots, T$;

2. evaluate the empirical copula $\widehat{C}_T$ for all rank vectors (u_t, v_t), $t = 1, \ldots, T$:

$$\widehat{C}_T(u_t, v_t) = \frac{1}{T} \sum_{i=1}^{T} \mathbb{1}(u_i \le u_t, v_i \le v_t); \tag{19}$$
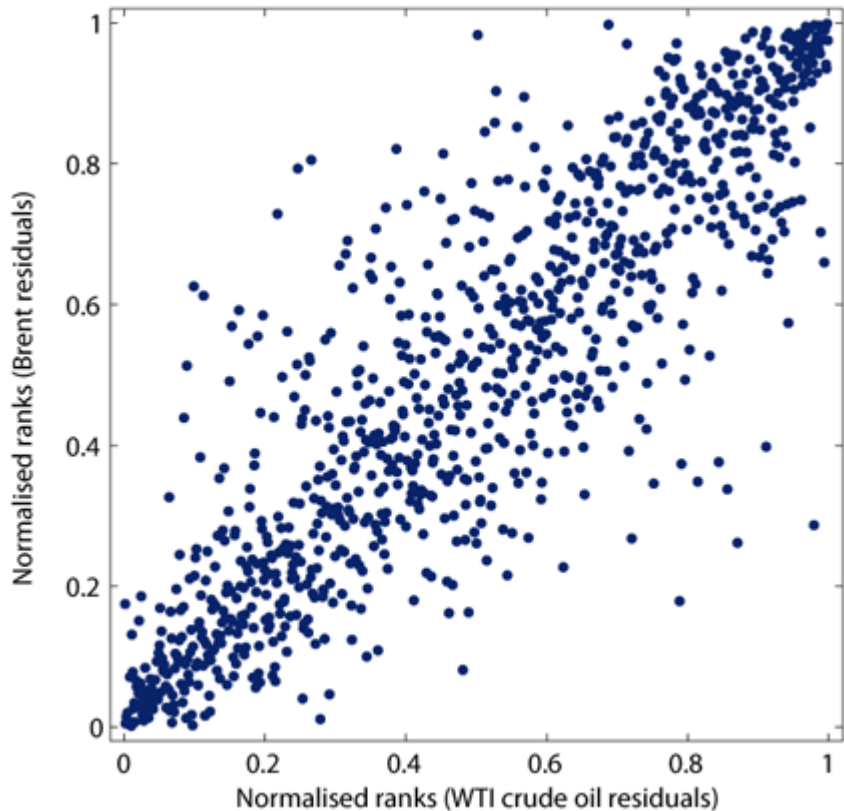
Figure 4: Scatter plot of the rank-based representation of the standardised residuals.

3. compute the optimal parameter estimate $\hat{\theta}_T$ by inversion of the empirical equivalent of Kendall's tau $\hat{\tau}_T$;

4. calculate S_T^* according to (16);

5. choose $N \in \mathbb{N}$ sufficiently large and repeat the following steps for each $k \in \{1, \ldots, N\}$:

   (a) generate a random sample $(X_{k,t}, Y_{k,t}), t = 1, \ldots, T$, of the copula $C_{\hat{\theta}_T}$ and compute the corresponding rank vectors $(u_{k,t}, v_{k,t})$ for $t = 1, \ldots, T$;

   (b) evaluate the empirical copula for all simulated rank vectors $(u_{k,t}, v_{k,t}), t = 1, \ldots, T$:

$$\widehat{C}_{k,T}(u_{k,t}, v_{k,t}) = \frac{1}{T} \sum_{i=1}^{T} \mathbb{1}(u_{k,i} \leq u_{k,t}, v_{k,i} \leq v_{k,t}); \qquad (20)$$

15

(c) use the same method as in Step 3 to construct a rank-based estimator $\hat{\theta}_{k,T}$ for θ;

(d) compute

$$S_{k,T} = \sum_{t=1}^{T} \left(\widehat{C}_{k,T}(u_{k,t}, v_{k,t}) - C_{\hat{\theta}_{k,T}}(u_{k,t}, v_{k,t}) \right)^2 ; \quad (21)$$

6. calculate

$$\frac{1}{N} \sum_{k=1}^{N} \mathbb{1} \left(S_{k,T} > S_T^* \right) \quad (22)$$

to obtain an approximated p-value of the goodness-of-fit test.

For an overview of goodness-of-fit tests for copulas see Genest et al. (2009).

4.2 Multiplier bootstrap

Initially, let $\hat{\Lambda}_L$ be the empirical lower tail copula of our dataset of standardised residuals (X_t, Y_t), $t = 1, \ldots, T$, and Λ_L be an arbitrary lower tail copula. Following Bücher and Dette (2012) we introduce a distance ϱ between $\hat{\Lambda}_L$ and Λ_L by

$$\varrho(\hat{\Lambda}_L, \Lambda_L) := \int_0^{\pi/2} \left(\hat{\Lambda}_L(\cos\varphi, \sin\varphi) - \Lambda_L(\cos\varphi, \sin\varphi) \right)^2 d\varphi$$

$$= \int_0^{\pi/2} \left(\hat{\Lambda}_L^{\angle}(\varphi) - \Lambda_L^{\angle}(\varphi) \right)^2 d\varphi, \quad (23)$$

where

$$\hat{\Lambda}_L^{\angle}(\varphi) = \hat{\Lambda}_L(\cos\varphi, \sin\varphi), \quad \Lambda_L^{\angle}(\varphi) = \Lambda_L(\cos\varphi, \sin\varphi).$$

Analytically speaking, ϱ describes the squared distance between the empirical lower tail copula $\hat{\Lambda}_L$ and the arbitrary lower tail copula Λ_L, corresponding to the metric induced by the L^2-norm evaluated on the unit circle in the first quadrant. In the following this set is denoted by $K_+ = \{(\cos\varphi, \sin\varphi) \colon \varphi \in [0, \pi/2]\}$.

For the purpose of estimating the lower tail copula of the time series (X_t, Y_t), $t = 1, \ldots, T$, we consider an open set $\mathcal{O} \subseteq \mathbb{R}$ and the one-parameter class $\mathcal{L} = \{\Lambda_L(\cdot; \theta) \mid \theta \in \mathcal{O}\}$. Then, the test for the hypothesis that Λ_L belongs to a certain parametric class is given by

$$\mathcal{H}_0 \colon \Lambda_L \in \mathcal{L}, \quad \mathcal{H}_1 \colon \Lambda_L \notin \mathcal{L}. \quad (24)$$

We estimate the optimal parameter θ by applying the minimum distance estimator

$$\hat{\theta}_T = \arg\min_{\theta \in \mathcal{O}} \varrho\left(\hat{\Lambda}_L, \Lambda_L(\cdot; \theta)\right), \tag{25}$$

where ϱ denotes the distance defined in (23) and $\hat{\Lambda}_L$ the empirical lower tail copula. The reference statistic of the test is given by

$$
\begin{aligned}
S_T^{*pdm} &:= k\varrho\left(\hat{\Lambda}_L, \Lambda_L(\cdot; \hat{\theta}_T)\right) \\
&= k\int_0^{\pi/2}\left(\hat{\Lambda}_L^{\angle}(\varphi) - \Lambda_L^{\angle}(\varphi; \hat{\theta}_T)\right)^2 d\varphi.
\end{aligned} \tag{26}
$$

To construct multiplier bootstrap critical values, we introduce i.i.d. nonnegative random variables ξ_t, independent of (X_t, Y_t), $t = 1, \ldots, T$, with mean μ in $(0, \infty)$ and finite variance σ^2 which satisfy the condition $\int_0^\infty \sqrt{\mathcal{P}(|\xi_t| > x)}dx < \infty$. A multiplier bootstrap analogue of Equation (9) can then be defined by

$$\hat{\Lambda}_L^\xi(x, y) := \frac{1}{k}\sum_{t=1}^T \frac{\xi_t}{\bar{\xi}_T}\mathbb{1}\left(u_t \leq \frac{kx}{T+1}, v_t \leq \frac{ky}{T+1}\right), \tag{27}$$

where $\bar{\xi}_T = T^{-1}\sum_{t=1}^T \xi_t$ denotes the mean of $\xi_1, \ldots, \xi_T$. As proposed in Bücher and Dette (2012) we will use $Laplacian(0, 2)$ multipliers, more precisely, the discrete random variables ξ_t are i.i.d. with probability density function $\mathbb{P}(0) = \mathbb{P}(2) = 1/2$, which obviously fulfil the mentioned conditions.

Having determined the multipliers, we are now able to formulate the process

$$
\begin{aligned}
\beta_T^{pdm}(x, y) &= \frac{\mu}{\sigma}\sqrt{k}\left(\hat{\Lambda}_L^\xi(x, y) - \hat{\Lambda}_L(x, y)\right) \\
&= \frac{\mu}{\sigma}\frac{1}{\sqrt{k}}\sum_{t=1}^T\left(\frac{\xi_t}{\bar{\xi}_T} - 1\right)\mathbb{1}\left(u_t \leq \frac{kx}{T+1}, v_t \leq \frac{ky}{T+1}\right),
\end{aligned} \tag{28}
$$

which will be part of the partial derivatives multiplier bootstrap (*pdm*-bootstrap) process. Furthermore, following Rémillard and Scaillet (2009), we compute consistent estimates for the partial derivatives of the lower tail copula as follows:

$$\widehat{\partial_x\Lambda_L}(x, y) := \begin{cases} \dfrac{\hat{\Lambda}_L(x+h, y) - \hat{\Lambda}_L(x-h, y)}{2h}, & h \leq x < \infty, \\[2ex] \dfrac{\hat{\Lambda}_L(x+2h, y) - \hat{\Lambda}_L(0, y)}{2h}, & x < h, \end{cases} \tag{29}$$

where $h \sim k^{-1/2}$ tends to 0 with increasing sample size. In case $x = \infty$, we set $\widehat{\partial_x \Lambda_L}(\infty, y) = 0$. In the same vain, we define $\widehat{\partial_y \Lambda_L}(x, y)$. Summing up, this yields the process

$$
\begin{aligned}
\alpha_T^{pdm}(x, y) = \beta_T^{pdm}(x, y) & \\
& - \widehat{\partial_x \Lambda_L}(x, y) \beta_T(x, \infty) \\
& - \widehat{\partial_y \Lambda_L}(x, y) \beta_T(\infty, y),
\end{aligned}
\tag{30}
$$

which only depends on the standardised residuals of the log-return time series and the multipliers $\xi_1, \ldots, \xi_T$.

Finally, let $\delta_\theta(x, y) = \partial_\theta \Lambda_L(x, y; \theta)$ denote the partial derivative of the lower tail copula with respect to $\theta \in \mathcal{O}$. Given some regularity conditions it can be easily shown that this partial derivative function is also homogeneous. The pdm-bootstrap statistic is obtained as

$$
S_T^{pdm} = \int_0^{\pi/2} \left(\alpha_T^{pdm\angle}(\varphi) - \delta_{\hat{\theta}_T}^{\angle}(\varphi) \int_0^{\pi/2} \gamma_{\hat{\theta}_T}(\varphi) \alpha_T^{pdm\angle}(\varphi) d\varphi \right)^2 d\varphi,
\tag{31}
$$

where

$$
\gamma_{\hat{\theta}_T}(\varphi) = \hat{A}_{\hat{\theta}_T}^{-1} \delta_{\hat{\theta}_T}^{\angle}(\varphi), \quad \alpha_T^{pdm\angle}(\varphi) = \alpha_T^{pdm}(\cos \varphi, \sin \varphi)
$$

and

$$
\hat{A}_{\hat{\theta}_T} := \int_0^{\pi/2} \left(\delta_{\hat{\theta}_T}^{\angle}(\varphi) \right)^2 d\varphi, \quad \delta_{\hat{\theta}_T}^{\angle}(\varphi) = \delta_{\hat{\theta}_T}(\cos \varphi, \sin \varphi).
\tag{32}
$$

Note that we only consider one-parameter families of tail copulas and therefore Equation (32) is also one-dimensional. For a detailed description of the required regularity conditions and a more theoretical view see Bücher and Dette (2012).

The following goodness-of-fit test based on the partial derivatives multiplier bootstrap (Bücher and Dette, 2012) yields approximated p-values assuming that the respective lower tail copula has a specific parametric form. The detailed algorithm proceeds as follows:

1. transform the standardised residuals $(X_t, Y_t), t = 1, \ldots, T$ into their rank-based representation vectors $(u_t, v_t), t = 1, \ldots, T$;

2. discretise the unit circle in the first quadrant K_+ and choose the parameter k in relation to the sample size T for calculating the empirical tail copula $\hat{\Lambda}_L^{\angle}(\varphi)$ according to (9);

3. for the purpose of evaluating the reference statistic (26) compute the optimal parameter estimate $\hat{\theta}_T$ by applying the minimum distance estimator (25);

18

4. to prepare for the next steps estimate the partial derivatives of the selected lower tail copula via (29) and approximate (32) using central difference quotients, if applicable;

5. choose $N \in \mathbb{N}$ sufficiently large and repeat the following steps for each $j \in \{1, \ldots, N\}$:

   (a) generate a random sample $\xi_{j,1}, \ldots, \xi_{j,T}$ of the random variables $\xi_1, \ldots, \xi_T$ as described above;

   (b) compute the statistics $\alpha_{j,T}^{pdm} = \alpha_T^{pdm}(\xi_{j,1}, \ldots, \xi_{j,T})$ for every $(x, y) \in K_+$ using the updated multiplier bootstrap analogue of the empirical lower tail copula (27), the evaluated process β_T^{pdm} from Equation (28) and the calculated partial derivatives from Step 4;

   (c) evaluate

   $$S_{j,T}^{pdm} = \int \left(\alpha_{j,T}^{pdm\angle}(\varphi) - \delta_{\hat{\theta}_T}^{\angle}(\varphi) \int \gamma_{\hat{\theta}_T}(\varphi) \alpha_{j,T}^{pdm\angle}(\varphi) d\varphi \right)^2 d\varphi$$

   according to (31);

6. calculate

   $$\frac{1}{N} \sum_{j=1}^{N} \mathbb{1} \left(S_{j,T}^{pdm} > S_T^{*pdm} \right) \tag{33}$$

   to obtain an approximated p-value of the goodness-of-fit test.

5 Empirical results

The following subsections provide guide how to implement the methodology above. First, we apply the parametric bootstrap-based goodness-of-fit test to suitable classes of copulas. Having rejected certain copula models according to a predefined level of significance, we use a newly introduced information criterion for copulas to find the best fit to the crude oil log-returns. In a second step a closed-form expression of the lower tail copula is derived and several possibly suitable one-parametric models are presented. The results of the goodness-of-fit test based on the partial derivatives multiplier bootstrap finalise the tail copula fit.

5.1 Copula selection

The goodness-of-fit test from Section 4.1 requires a preselection of possible one-parametric copula classes. The structure of the scatter plot in Figure 4 suggests to limit the range to symmetric copulas, i. e. $C(u, v) = C(v, u)$ for all (u, v) in $[0, 1]^2$. In addition to two elliptical

copula	p-**value**	$\hat{\theta}_T$-**value**
Frank	0.000	11.649
(4.1.14)	0.001	2.891
Plackett	0.012	47.150
Gaussian	0.052	0.895
(4.1.12)	0.065	2.261
t ($\hat{\nu}_T = 5.313$)	0.074	0.895

Table 1: Test results ($N = 2,000$ bootstrap iterations) for the goodness-of-fit of different copula models in ascending order (p-value).

copulas, Gaussian and t, we fit the widely used Plackett copula and three Archimedean copulas. Within the latter class we choose the Frank family and two relatively unknown strict copulas from (4.1.12) and (4.1.14) in Nelsen (2006), Table 4.1, with corresponding functions given by

$$C_\theta(u, v) = \left(1 + \left[(u^{-1} - 1)^\theta + (v^{-1} - 1)^\theta \right]^{1/\theta} \right)^{-1} \qquad (34)$$

and

$$C_\theta(u, v) = \left(1 + \left[(u^{-1/\theta} - 1)^\theta + (v^{-1/\theta} - 1)^\theta \right]^{1/\theta} \right)^{-\theta}, \qquad (35)$$

respectively, $\theta \in [1, \infty)$. Both copulas share the same limit case $C_\infty = \min(u, v)$, which is not Archimedean anymore.

Table 1 summarises the results (p-values) of the goodness-of-fit tests for the above mentioned copula classes together with their estimated optimal parameter $\hat{\theta}_T$. For the Plackett copula, where no explicit expression for the parameter θ is available, we apply penalised splines to approximate the function $\theta(\tau)$ using an appropriately chosen grid of θ values (Kojadinovic and Yan, 2010). Figure 5 (see over) presents scatter plots of $n = 1,001$ sampling points from simulations for all copula families under consideration, $\hat{\theta}_T$ as listed in Table 1. From these scatter plots alone, it is nearly impossible to draw definite conclusion about the goodness-of-fit.

The results of the parametric bootstrap procedure indicate that neither the Frank family nor the copula (4.1.14), and to a great extent, the Plackett copula seem to be appropriate to model the dependence structure between the residuals of the crude oil log-returns. On the contrary, at a 5% level of significance, the null hypothesis cannot be rejected for any of the remaining three candidates. At this point we
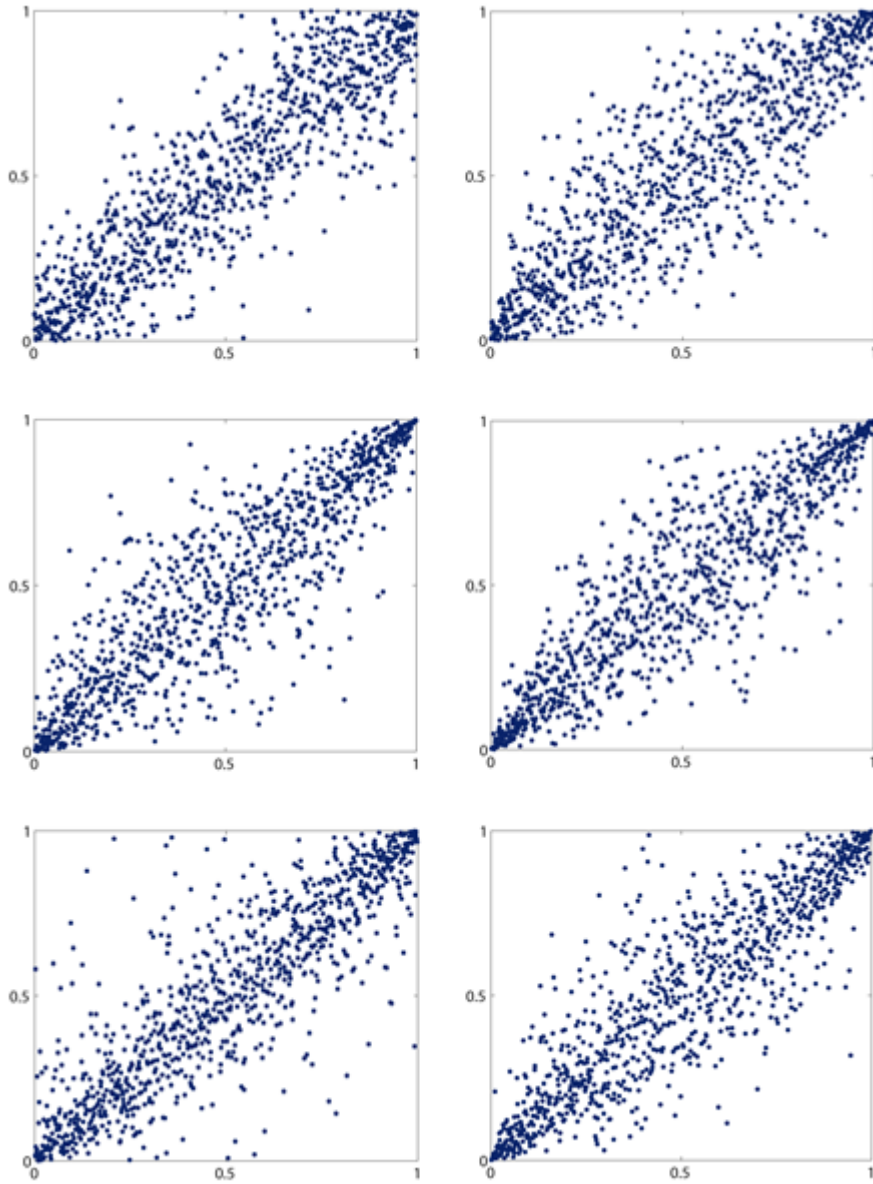
Figure 5: Simulation of Frank copula (top left), copula (4.1.14) (center left), Plackett copula (bottom left), Gaussian copula (top right), copula (4.1.12) (center right) and t-copula (bottom right).

copula	CIC-value	$\hat{\theta}$-value	$\ell_T(\hat{\theta})$-value
Gaussian	1,454	0.877	729.75
(4.1.12)	1,471	2.088	743.61
t $(\hat{\nu} = 5.101)$	1,522	0.880	764.36

Table 2: Cross-validation Copula Information Criterion values for the remaining copula models in ascending order (CIC-value).

emphasise, as the actual investigation focuses on the behaviour of joint extreme events, it is fairly unlikely to achieve p-values higher than 10%, while, at the same time, not trimming or winsorising the original data in order to mitigate the effects of (marginal) extreme outliers.

However, as the p-values are 'only' inverse measures of the strength of evidence against the null hypothesis $\mathcal{H}_0$, we need to select the most appropriate model from the remaining candidates. Following Grønneberg and Hjort (2012), we implement the generally applicable cross-validation Copula Information Criterion for the left three copula families. In contrast to the common used formula AIC $= 2\ell_{T,\max} - 2$, where $\ell_{T,\max}$ is defined as the maximiser of the pseudo likelihood

$$\ell_T(\theta) = \sum_{t=1}^{T} \log\left(c_\theta(u_t, v_t)\right),$$

the cross-validation CIC corrects for finite-sample bias introduced by the rank-based non-parametric modelling of the marginals. For a more detailed view on the different model selection criteria see Claeskens and Hjort (2008).

Table 2 shows the maximum pseudo likelihood estimator together with its corresponding likelihood and the cross-validation Copula Information Criterion value. Using the maximal CIC-value as a criterion to select the model with the best fit to the data, we conclude that the t-copula, with correlation parameter $\hat{\theta}_T = 0.895$ and degrees of freedom $\hat{\nu}_T = 5.313$, best describes the dependence between the WTI and Brent crude oil log-returns. Note that in this case evaluating the frequently applied AIC formula would lead to the same decision.

5.2 Tail copula fit

Implementing the goodness-of-fit test for tail copulas described in Section 4.2 requires some background. In order to derive a closed-form expression of the lower tail copula (in contrast to the limit definition

in Equation (4)), let (U, V) be a random vector with uniformly distributed margins on the interval $[0, 1]$ and joint distribution

$$C_A(u, v) = \exp\left(\log(uv) A\left(\frac{\log(v)}{\log(uv)}\right)\right) \qquad (36)$$

where $A\colon [0, 1] \to [1/2, 1]$ is called Pickands dependence function. C_A can be identified with the class of extreme-value copulas which arise naturally as limits of copulas of componentwise maxima in independent random samples and provide a convenient framework for modelling the dependence structure between extreme events (see Gudendorf and Segers, 2010). Following Pickands (1981) the extreme-value copula $C_A(u, v)$ is a copula if and only if the function $A\colon [0, 1] \to [1/2, 1]$ is a convex function which satisfies $\max(t, 1 - t) \le A(t) \le 1$. Next, it can be shown that the joint distribution of the random vector $(1 - U, 1 - V)$ is given by the survival copula

$$\widehat{C}_A(u, v) = u + v - 1 + C_A(1 - u, 1 - v). \qquad (37)$$

From the homogeneous structure of the lower tail copula in Equation (5) it is sufficient to consider the restriction of these functions to the unit simplex $\Delta_1 = \{(u, v) \in [0, 1]^2 : u + v = 1\}$ (see Segers, 2012b). By substituting Equation (37) in Equation (4) and applying l'Hôpital's rule, a tedious but straightforward calculation yields

$$\Lambda_L(1 - x, x) = \lim_{t \to 0^+} \frac{\widehat{C}_A(t(1 - x), tx)}{t} = 1 - A(x) \qquad (38)$$

for the lower tail copula of the random vector $(1 - U, 1 - V)$. Finally, by renaming variables and the homogeneity property the entire lower tail copula can be obtained as $\Lambda_L(1 - t, t) = 1 - A(t)$ with Pickands dependence function $A(t)$. The latter implies the existence of an one-to-one and onto relationship between lower tail copulas and Pickands dependence functions.

Consequently, the next part focuses on different classes of Pickands dependence functions. The structure of the empirical estimate $\hat{\Lambda}_L$, see Equation (9), evaluated on the unit circle in the first quadrant, in Figure 6 (see Page 25) suggests to limit the range to symmetric lower tail copulas. Below, we seek to cover a spectrum as wide as possible via

(i) the logistic or Gumbel model (Gumbel, 1960), defined by

$$\Lambda_L(1 - t, t) = 1 - \left((1 - t)^\theta + t^\theta\right)^{1/\theta}, \quad \theta \in [1, \infty),$$

which belongs, amongst other, to the survival copula of the well-known Gumbel-Hougaard copula,

(ii) the negative logistic or Galambos model (Galambos, 1975), given by

$$\Lambda_L(1-t,t) = \left((1-t)^{-\theta} + t^{-\theta}\right)^{-1/\theta}, \quad \theta \in (0,\infty),$$

that is the lower tail copula of, amongst other, the Clayton copula, the survival copula of the Galambos copula, and the copula family (4.1.12),

(iii) the mixed model (Tawn, 1988), defined by

$$\Lambda_L(1-t,t) = \theta(1-t)t, \quad \theta \in [0,1],$$

which constitutes the complete class of quadratic functions for differentiable parametric models; in the case $\theta = 1$, Λ_L belongs, amongst other, to the copula family (4.1.14),

(iv) the Hüsler-Reiss model (Hüsler and Reiss, 1989), vis.

$$\Lambda_L(1-t,t) = 1 - (1-t)\Phi\left(\theta + \frac{1}{2\theta}\log\left(\frac{1-t}{t}\right)\right)$$
$$- t\Phi\left(\theta + \frac{1}{2\theta}\log\left(\frac{t}{1-t}\right)\right), \quad \theta \in (0,\infty),$$

with Φ denoting the standard normal cumulative distribution function; Λ_L is, amongst other, the lower tail copula of the survival copula of the Hüsler-Reiss copula.

Note that the list does not contain the well-known class of t-extreme-value copulas C_A^t, (see Demarta and McNeil, 2005). Following their remarks the correspondent Pickands dependence function is not particularly convenient for practical application. Demarta and McNeil (2005) show empirically, that for each fixed (θ, ν) in the t-model, there exists a parameter θ from either the Gumbel or the Galambos model such that the resulting curves are indistinguishable.

Referring back to the applied copula models in Section 5.1 and digressing, the relationship between the lower tail copula of the t-copula and the t-extreme-value copula can be obtained as:

$$\Lambda_L^t(1-s,s) = \log\left(C_A^t\left(e^{-(1-s)}, e^{-s}\right)\right)$$
$$= 1 - A(s) \quad s \in [0,1]. \tag{39}$$

The latter can be proven (also in a general setting) using the natural definition of extreme-value copulas (Gudendorf and Segers, 2010), applying the Taylor expansion and finally, considering the identity $\widehat{C} = C$ for radially symmetric copulas (Nelsen, 2006). Equation (39) shows
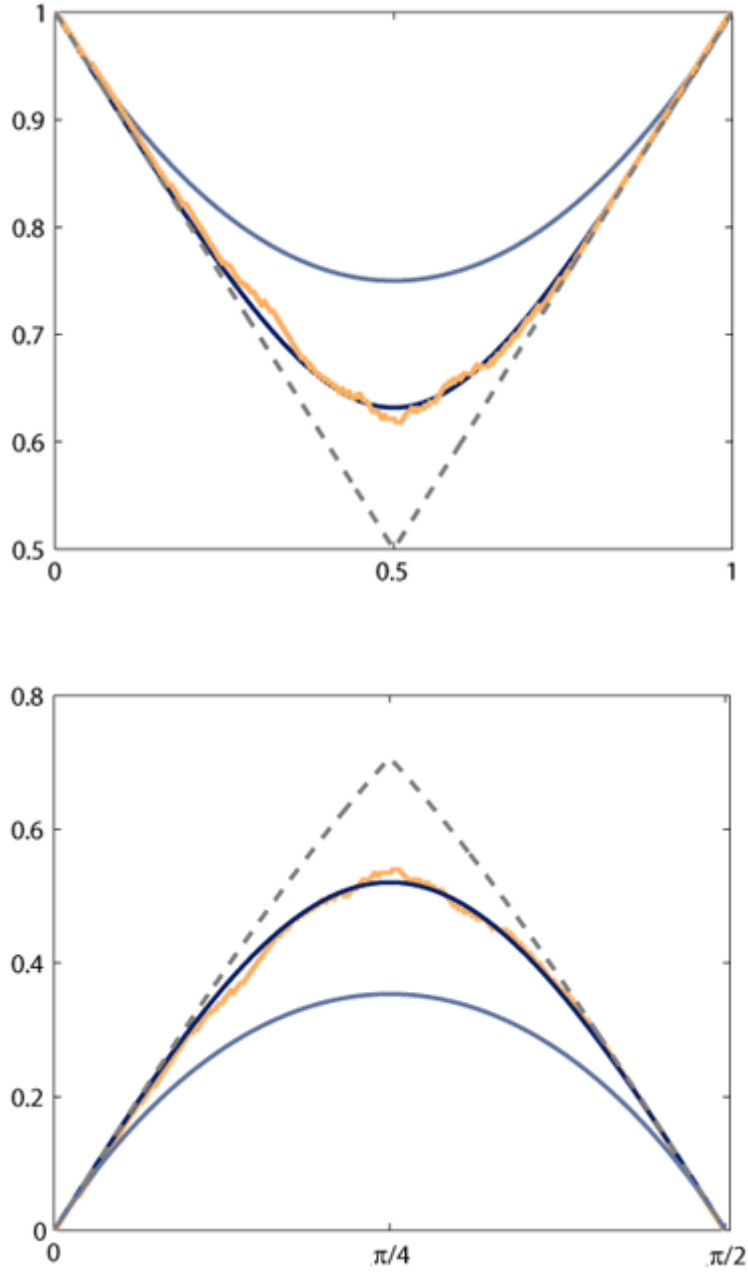
Figure 6: Plot of the empirical lower tail copula $\hat{\Lambda}_L$ (amber), Gumbel model (blue line), mixed model (light blue) and the overall range (grey dashed line) for both the Pickands framework (top) and the tail copula approach (bottom).

25

model	p-value	$\hat{\theta}_T$-value	$\hat{\lambda}_L(\hat{\theta}_T)$-value
Mixed	0.000	1.000	0.500
Hüsler-Reiss	0.146	0.343	0.732
Galambos	0.242	2.251	0.735
Gumbel	0.290	2.961	0.736

Table 3: Test results ($N = 10,000$ bootstrap iterations) for the derivatives multiplier bootstrap of different lower tail copulas in ascending order (p-value).

that, in order to avoid the complexity of the univariate t-distribution, the lower tail copula of the t-copula can again be accurately approximated by the simpler Gumbel or Galambos models. Furthermore, neither the Frank copula, nor the Plackett copula, nor the Gaussian copula allow for lower tail dependence and thus $\Lambda_L = 0$ for all $(x, y) \in \mathbb{R}_+^2$.

Care should be taken when implementing the derivatives multiplier bootstrap with respect to the different applied norms. Whereas the Pickands dependence functions are defined within the L^1-norm framework, the minimum distance estimator in Equation (25) is formulated using the L^2-norm. Furthermore, the transformation of the radian coordinate $\varphi \in [0, \pi/2]$ to $t \in [0, 1]$ can be obtained by $t = \tan(\varphi)/(1 + \tan(\varphi))$. According to Bücher and Dette (2012) we choose the parameter k equal to 200. The results of this computationally efficient goodness-of-fit test for tail copulas based on the partial derivatives multiplier bootstrap are summarised in Table 3. In addition to the p-value, the table lists the minimum distance estimator $\hat{\theta}_T$ and the estimated lower tail dependence coefficient $\hat{\lambda}_L(\hat{\theta}_T)$, which is obtained by evaluating the lower tail copula at the point $t = (1/2, 1/2)$ multiplied with the factor $\sqrt{2}$.

While the mixed model is definitely not appropriate for the joint behaviour in the lower tails, the null hypothesis, at a 5% level of significance, cannot be rejected for any of the remaining three candidates. We decide to use the smallest distance between the empirical lower tail dependence coefficient $\hat{\lambda}_L = 0.757$ and the estimated lower tail dependence coefficient $\hat{\lambda}_L(\hat{\theta}_T)$ as criterion to select the model with the best fit to the data in the lower tails, which turns out to be the logistic or Gumbel model with $\hat{\theta}_T = 2.961$. Note that both the Hüsler-Reiss and the Galambos models also seem appropriate to model the structure of the lower tail dependence.

Although the lower tail copula of the t-copula can be accurately approximated with the Gumbel model, it is a far cry from saying that

the joint dynamics in the lower tail of the t-copula ($\hat{\theta}_T = 0.895$, $\hat{\nu}_T = 5.313$) from Section 5.1 are identical to the corresponding behaviour of the Gumbel model ($\hat{\theta}_T = 2.961$). This can be explained by the completely different ways of estimating the model parameters in both subsections. Figure 6 (see Page 25) shows the plot of the empirical lower tail copula $\hat{\Lambda}_L$, the selected Gumbel model, the unsuitable mixed model and the overall range for both the Pickands framework and the tail copula setting.

# 6 Risk measures and backtesting

Returning to the initial question, and having selected both the copula and tail copula model, we are now interested in the risk measurement of a portfolio consisting of an arbitrary convex combination of the WTI and Brent crude oil. For this purpose, we simulate the common risk parameters Value-at-Risk and Expected Shortfall. According to Jorion (2006) the Value-at-Risk summarises the worst loss over a target horizon that will not be exceeded with a given level of confidence, or more formally, it describes the quantile of the projected distribution of gains or losses over the target horizon. If α is the selected level of confidence, the Value-at-Risk corresponds to the $1 - \alpha$ lower tail level. Extending this concept, the Expected Shortfall (also known as Conditional Value-at-Risk) for a given level of confidence α is defined as the expected loss, given the loss is larger or equal to the Value-at-Risk (Deutsch, 2009).

Following Rank and Siegl (2002) we first generate n pairs (u, v) of observations of $U(0, 1)$ distributed random variables U and V whose joint distribution function is either given by the t-copula ($\hat{\theta}_T = 0.895$, $\hat{\nu}_T = 5.313$) from Section 5.1 or the survival copula of the Gumbel copula ($\hat{\theta}_T = 2.961$) from Section 5.2. The simulation is carried out using the functionality of Kojadinovic (2010) and Yan (2007). To obtain the required marginal distributions we apply the inverse distribution functions, estimated in Section 3, to the sample (u, v). Accordingly, the n obtained pairs (x_i, y_i), $i = 1, \ldots, n$, form scenarios of possible logarithmic changes of the two crude oil grades, while at the same time, having preserved the desired dependence structure.

For the purpose of simulating the Value-at-Risk we consider a simple linear portfolio position with weighting factors $\omega_1, \omega_2 \in [0, 1]$ and $\omega_2 = 1 - \omega_1$. The logarithmic change a_t of the (simulated) portfolio at time t is given by

$$a_t = \omega_1 W_t + \omega_2 B_t, \tag{40}$$

where the univariate WTI crude oil model W_t is described in Equa-
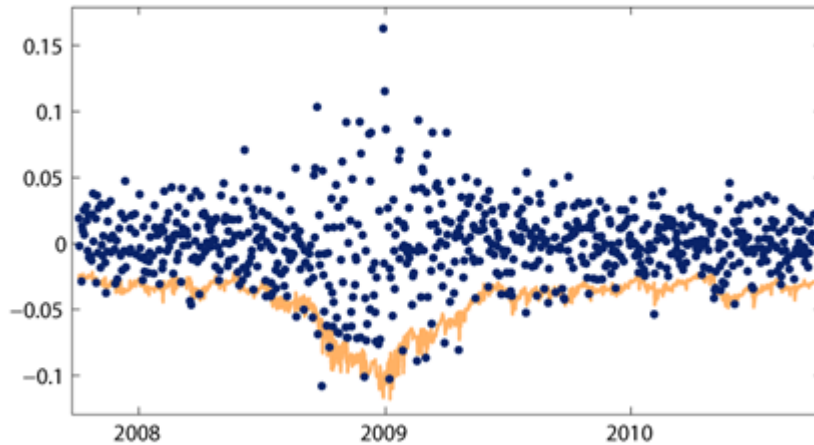
Figure 7: Out of sample one day-ahead simulated portfolio ($\omega_1 = \omega_2 = 0.5$) Value-at-Risk for a given level of confidence $\alpha = 95\%$ for the Gumbel model (amber) and the realised (observed) returns (blue).

tion (13) and the Brent analogue B_t in Equation (14), respectively. To sum up, the applied Monte Carlo method generates n scenarios for each day t, and evaluates the present value change of a portfolio under each scenario. Finally, calculating the $1 - \alpha$ quantile yields the daily Value-at-Risk with level of confidence α. Note that within a non-rolling target horizon framework, and having once modelled the univariate time series, the only remaining task is to simulate the standardised residuals (x_i, y_i), $i = 1, \ldots, n$, and substitute these values into the corresponding model equation.

Figure 7 shows the out of sample one day-ahead simulated portfolio ($\omega_1 = \omega_2 = 0.5$) Value-at-Risk for a given level of confidence $\alpha = 95\%$ for the estimated Gumbel model in combination with the realised (observed) returns (quod vide Palaro and Hotta, 2006). To be able to judge the effectiveness of the applied Value-at-Risk model, we compare the simulated (predicted) and empirical number of outliers, where the actual loss exceeds the Value-at-Risk (Rank and Siegl, 2002). Consequently, from a risk management point of view, this backtesting approach is easily comprehensible and allows for the required transparency. To benchmark the simulated Value-at-Risk with a fully deterministic approach we apply historical simulation. The freedom from model assumptions is one of the primary advantages of this method (Deutsch, 2009). Based on the given time series and the target horizon of $T = 250$ business days we compute 751 daily Value-at-Risk estimates for the confidence level vector α_i ($\alpha_1 = 95\%, \alpha_2 = 97.5\%, \alpha_3 = 99\%$).

28

$1-\alpha$	ω_1	ω_2	t-**copula**	**Gumbel**	**historical**
0.050	0.50	0.50	0.059 (44)	0.056 (42)	0.067 (50)
0.025	0.50	0.50	0.021 (16)	0.021 (16)	0.052 (39)
0.010	0.50	0.50	0.008 (06)	0.008 (06)	0.021 (16)
0.050	0.25	0.75	0.056 (42)	0.055 (41)	0.063 (47)
0.025	0.25	0.75	0.017 (13)	0.019 (14)	0.048 (36)
0.010	0.25	0.75	0.011 (08)	0.009 (07)	0.025 (19)
0.050	0.75	0.25	0.056 (42)	0.055 (41)	0.069 (52)
0.025	0.75	0.25	0.027 (20)	0.024 (18)	0.045 (34)
0.010	0.75	0.25	0.011 (08)	0.011 (08)	0.016 (12)
0.050	**average**		0.007	0.005	0.016
0.025	**average**		0.005	0.004	0.023
0.010	**average**		0.001	0.001	0.011
average	**average**		0.003	0.002	0.015
ranking			2	1	3

Table 4: Relative number (absolute number) of backtest outliers $1-\hat{\alpha}$ for the different applied Value-at-Risk simulations with level of confidence $1-\alpha$, weighted average error $|\hat{\alpha}-\alpha|$ and error ranking.

Note that the latter approach takes a rolling data set of 250 (observed) values (for each one day-ahead estimated Value-at-Risk) into account whereas the Value-at-Risk simulations via copulas depend on the entire considered time horizon. Consequently, a comparison between the applied models is only of limited use.

Table 4 summarises the results of the backtesting for the set of scenarios $w_{1,i}$ ($\omega_{1,1} = 0.50$, $\omega_{1,2} = 0.25$, $\omega_{1,3} = 0.75$) with a given level of confidence α. Following Rank and Siegl (2002) we define the prediction error as the absolute difference between the relative number of outliers $1-\hat{\alpha}$ and the predicted relative number $1-\alpha$. While the average over the portfolios uses equal weights, the average over the level of confidence α emphasises the tails by a weighting scheme β_i ($\beta_1 = 1, \beta_2 = 2, \beta_3 = 5$). The Value-at-Risk simulation via copulas is carried out generating $n = 100,000$ samples of pairs (u, v). As expected the tail copula fit from Section 5.2 approach yields the best result whereas both the t-copula and the survival copula of the Gumbel copula obviously outperform the historical simulation. The latter should be regarded with caution as outlined before. Furthermore, although the Gumbel model yields a slightly better result than the t-copula both models seem to pass the backtesting stage.

To further verify the validity of the t-copula and the Gumbel model we apply the traffic light approach used by the supervising authorities when auditing banks' internal risk management models (Deutsch, 2009). For this, a confidence interval for the observed outliers has to be constructed. It is quite obvious to assume that the number of outliers follow a binomially distributed random variable with probability $p = 1 - \alpha$ and number of trials n. Within this framework the probability of k or more outliers is

$$1 - \sum_{i=0}^{k} B_{n,p}(i) = \sum_{i=k+1}^{n} \binom{n}{i} p^i (1-p)^{n-i}. \tag{41}$$

This is equal to the probability of making a type-I-error (the rejection of a correct model) when the hypothesis is rejected if k or more outliers are observed. Based on this probability the supervising authorities establish the boundaries for three different zones, i. e. a lower bound of the red zone (equals the upper bound of the yellow zone) for a type-I-error of less than 0.01% and a lower bound of the yellow zone (equals the upper bound of the green zone) for a type-I-error less than 5%. For $p = 1 - 0.95 = 0.05$ and $n = 751$ this translates into the following intervals as a function of realised outliers k

$$f(k) = \begin{cases} \text{green zone,} & k \leq 47, \\ \text{yellow zone,} & 48 \leq k \leq 61, \\ \text{red zone,} & k \geq 62. \end{cases} \tag{42}$$

Table 4 clearly shows that both models (for each type of portfolio weighting) are in the green zone and thus no adjustments are necessary. Completely analogous calculations yield the same results for the cases $\alpha = 0.025$ and $\alpha = 0.01$.

Having assessed the Value-at-Risk for a given level of confidence α it is straightforward to estimate the corresponding Expected Shortfall, i. e. averaging the $(1 - \alpha)$-quantile. Table 5 (see over) lists the average relative errors between the observed outliers and the Expected Shortfall based on the full backtesting period (751 days). Note that here the absolute values of the relative errors are taken into account. Finally, the concluding averages use the same weights as in the Value-at-Risk approach.

As it can be seen from Table 5, the results are consistent with the previous Value-at-Risk measures. Here as well, the Gumbel model yields the overall best result having in mind that the t-copula results are very close. For the sake of completeness, the historical simulation shows the worst fit.

$1-\alpha$	ω_1	ω_2	t-copula	Gumbel	historical
0.050	0.50	0.50	0.141	0.144	0.235
0.025	0.50	0.50	0.129	0.131	0.284
0.010	0.50	0.50	0.056	0.061	0.374
0.050	0.25	0.75	0.168	0.170	0.261
0.025	0.25	0.75	0.138	0.135	0.298
0.010	0.25	0.75	0.097	0.087	0.548
0.050	0.75	0.25	0.140	0.142	0.353
0.025	0.75	0.25	0.111	0.108	0.415
0.010	0.75	0.25	0.056	0.057	0.282
0.050	**average**		0.150	0.152	0.283
0.025	**average**		0.126	0.125	0.332
0.010	**average**		0.070	0.068	0.401
average	**average**		0.094	0.093	0.369
ranking			2	1	3

Table 5: Average relative error between the observed outliers and the Expected Shortfall for the different applied Expected Shortfall simulations with level of confidence $1-\alpha$, weighted averages and error ranking.

7 Concluding remarks

The present paper studies the dependence structure of log-returns of two crude oil grades, WTI Cushing Crude Oil Spot and Bloomberg European Dated Brent, from a probabilistic point of view. The investigation focuses on lower tail dependence, i. e. on assessing the probability that large losses occur together. As a rough but not exhaustive indicator (see Jäschke et al., 2012), the lower tail dependence coefficient is estimated as the value of the empirical lower tail copula at the point $(1, 1)$ (Schmidt and Stadtmüller, 2006). The resulting estimate $\hat{\lambda}_L = 0.757$ reveals a relatively high level of lower tail dependence.

The empirical analysis can be summarised as follows: the overall dependence between the log-returns sample of the two crude oil grades can be adequately described by a t-copula. When directly modelling the joint dynamics in the lower tail via tail copulas the Gumbel model performs best. Here, for the first time in the literature, the partial derivatives multiplier goodness-of-fit test for tail copulas (Bücher and Dette, 2012) is applied to energy portfolios.

From a risk management perspective it is crucial to backtest the findings using commonly applied risk measures. Therefore, the present

paper simulates the Value-at-Risk and Expected Shortfall for both models and compares the results for different portfolio weights and certain levels of confidence (Rank and Siegl, 2002). As suspected and confirmed by the backtesting the Gumbel model slightly outperforms the t-copula. Finally, the traffic light approach verifies the validity of both the Gumbel model and the t-copula.

Acknowledgements

References

Accioly, R. M. S. and F. A. L. Aiube (2008). Analysis of crude oil and gasoline prices through copulas. *Cadernos do IME – Série Estatística 24*, 15–28.

Bücher, A. and H. Dette (2012). Multiplier bootstrap of tail copulas - with applications. *Bernoulli*. (in press).

Caillault, C. and D. Guéégan (2005). Empirical estimation of tail dependence using copulas: application to asian markets. *Quantitative Finance 5*, 489–501.

Chang, C.-L., M. McAleer, and R. Tansuchat (2010). Analyzing and forecasting volatility spillovers, asymmetries and hedging in major oil markets. *Energy Economics 32*, 1445–1455.

Chui, C. and X. Wu (2009). Exponential series estimation of empirical copulas with application to financial returns. In Q. Li and J. S. Racine (Eds.), *Nonparametric econometric methodes (Advances in econmetrics, Volume 25)*, pp. 263–290. Emerald Group Publishing Limited.

Claeskens, G. and N. L. Hjort (2008). *Model selection and model averaging*. Cambridge: Cambridge University Press.

Deheuvels, P. (1979). La fonction de dépendance empirique et ses propriétés. *Acad. R. Belg., Bull. Cl. Sci. 65*, 274–292.

Demarta, S. and A. J. McNeil (2005). The *t*-copula and related copulas. *International Statistical Review 73*(1), 111–129.

Deutsch, H.-P. (2009). *Derivatives and internal models* (4th ed.). Palgrawe Macmillan.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of variance of united kingdom inflation. *Econometrica 50*, 987–1008.

Fattouh, B. (2010). Oil market dynamics through the lens of the 2002–2009 price cycle. *Oxford Institute for Energy Studies*.

Frahm, G., M. Junker, and R. Schmidt (2005). Estimating the tail-dependence coefficient: Properties and pitfalls. *Insurance: Mathematics and Economics 37*, 80–100.

Galambos, J. (1975). Order statistics of samples from multivariate distributions. *J. Amer. Statist. Assoc. 70*(351), 674–680.

Genest, C. and A. C. Favre (2007, July). Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering 12*(4), 347–368.

Genest, C. and B. Rémillard (2008). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Annales de l'Institut Henri Poincaré - Probabilités et Statistique 44*(6), 1096–1127.

Genest, C., B. Rémillard, and D. Beaudoin (2009, April). Goodness-of-fit tests for copulas: A review and a power study. *Insurance Math. Econom. 44*(2), 199–213.

Ghalanos, A. (2011). *rgarch: Flexible GARCH modelling in R.* R package version 1.91.

Grégoire, V., C. Genest, and M. Gendron (2008). Using copulas to model price dependencies in energy markets. *Energy Risk 5*(5), 58–64.

Grønneberg, S. and N. L. Hjort (2012). The copula information criterion. *Scandinavian Journal of Statistics*. (in press).

Gudendorf, G. and J. Segers (2010). Extreme-value copulas. In F. Durante, W. Härdle, P. Jaworski, and T. Rychlik (Eds.), *Workshop on Copula Theory and its Applications*. Springer.

Gumbel, E. J. (1960). Bivariate logistic distributions. *J. Amer. Statist. Assoc. 56*, 335–349.

Hüsler, J. and R. Reiss (1989). Maxima of normal random vectors: Between independence and complete dependence. *Statistics und Probability Letters 7*, 283–286.

Huang, X. (1992). *Statistics of bivariate extreme values*. Ph. D. thesis, Thesis Publishers and Tinbergen Institute.

Jorion, P. (2006). *Value at Risk: The new benchmark for managing financial risk* (3rd ed.). McGraw-Hill.

Jäschke, S., K. F. Siburg, and P. A. Stoimenov (2012). Modelling dependence of extreme events in energy markets using tail copulas. *Journal of Energy Markets*. (in press).

Kang, A. H., S.-M. Kang, and S.-M. Yoon (2009). Forecasting volatitlity of crude oil markets. *Energy Economics 31*, 119–125.

Kemp, J. (2010). Oil inventories lose their influence on prices. *Reuters News*. March 24, 2010.

Kojadinovic, I. (2010). Modeling multivariate distributions with continuous margins using the copula R package. *Journal of Statistical Software 34*(9), 1–20.

Kojadinovic, I. and J. Yan (2010). Comparison of three semiparametric methods for estimating dependence parameters in copula models. *Insurance: Mathematics and Economics 47*, 52–63.

Liu, J. (2011). *Extreme value theory and copula theory: A risk management application with energy futures*. Ph. D. thesis, University of Victoria.

Mohammadi, H. and L. Su (2010). International evidence on crude oil price dynamics: Applications of arima-garch models. *Energy Economics 32*, 1001–1008.

Nelsen, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–370.

Nelsen, R. B. (2006). *An introduction to copulas* (2nd ed.). New York: Springer.

Palaro, H. P. and L. K. Hotta (2006). Using conditional copula to estimate value at risk. *Journal of Data Science 4*, 95–113.

Perron, P. (1988). Trends and random walks in macroeconomic time series: Further evidence from a new approach. *Journal of Economic Dynamics and Control 12*, 297–332.

Pickands, J. (1981). Multivariate extreme value distributions (with a discussion). In *Proceedings of the 43rd Session of the International Statistical Institute. Bull. Inst. Internat. Stat. 49*, pp. 859–878.

Rank, J. and T. Siegl (2002). Applications of copulas for the calculation of value-at-risk. In W. Härdle, T. Kleinow, and G. Stahl (Eds.), *Applied Quantitative Finance*, pp. 35–47. Berlin–Heidelberg: Springer.

Reboredo, J. C. (2011). How do crude oil prices co-move?: A copula approach. *Energy Economics 33*(5), 95–113.

Rémillard, B. and O. Scaillet (2009). Testing for equality between two copulas. *Journal of Multivariate Analysis 100*, 377–386.

Schmidt, R. and U. Stadtmüller (2006). Non-parametric estimation of tail dependence. *Scand. J. Statist. 33*(2), 307–335.

Segers, J. (2012a). Asymptotics of empirical copula processes under nonrestrictive smoothness assumptions. *Bernoulli 18*(3), 764–782.

Segers, J. (2012b). Max-stable models for multivariate extremes. *Rev-Stat - Statistical Journal 10*(1), 61–82.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris 8*, 229–231.

Stute, W., W. G. Manteiga, and M. P. Quindimil (1993). Bootstrap based goodness-of-fit tests. *Metrika 40*, 243–256.

Tawn, J. A. (1988). Bivariate extreme value theory: Models and estimation. *Biometrika 75*(3), 397–415.

Würtz, D., Y. Chalabi, and L. Luksan (2006). Parameter estimation of ARMA models with GARCH/APARCH errors – an R and SPlus software implementation. (unpublished).

Yan, J. (2007). Enjoy the joy of copulas: With a package copula. *Journal of Statistical Software 21*(4), 1–21.