TIME-VARYING COPULA MODELS FOR FINANCIAL TIME SERIES

BY RÜDIGER KIESEL, MAGDA MROZ AND ULRICH STADTMÜLLER

Abstract

We perform an analysis of the potential time inhomogeneity in the dependence between multiple financial time series. To this end, we use the framework of copula theory and tackle the question of whether dependencies in such a case can be assumed constant throughout time or rather have to be modeled in a time-inhomogeneous way. We focus on parametric copula models and suitable inference techniques in the context of a special copula-based multivariate time series model. A recent result due to Chan *et al.* (2009) is used to derive the joint limiting distribution of local maximum-likelihood estimators on overlapping samples. By restricting the overlap to be fixed, we establish the limiting law of the maximum of the estimator series. Based on the limiting distributions, we develop statistical homogeneity tests, and investigate their local power properties. A Monte Carlo simulation study demonstrates that bootstrapped variance estimates are needed in finite samples. Empirical analyses on real-world financial data finally confirm that time-varying parameters are an exception rather than the rule.

Keywords: Multivariate time series; copula; binomial test; extremal test

2010 Mathematics Subject Classification: Primary 62M10
Secondary 62G32; 91B84

1. Introduction

Our epigraph 'copulas form a most useful concept for a lot of applied modeling [...]' is part of a citation in [16, p. 13] in support of a historical perspective when applying copulas. Recent textbooks such as [9], [27], and [28] contain theoretical background and a multitude of examples for applications supporting this view. Also, [32] contains a recent overview with an exhaustive list of references. We are interested in the dependency of multivariate time series and in particular in possible changes in the dependency structure. In the spirit of [16, p. 13] we develop our methodology based on some well-known and well-studied mathematical concepts, while focusing on suitable inference techniques in the context of a special copula-based multivariate time series model.

Owing to its special structure, the semiparametric copula-based multivariate dynamic (SCOMDY) model due to [6] directly retains many findings presented in the univariate case. Although there are more general approaches for multivariate heteroscedasticity, we decided in favor of the SCOMDY model due to its parsimonious parametrization and its statistical tractability. In particular, we were able to derive the tail behavior of the finite-dimensional distribution of the SCOMDY model by embedding it into a class of multivariate models for which the behavior was recently shown in [17]. A possible time-varying dependence structure is captured by the parameter of the underlying copula family in the SCOMDY model. Time-dependent copula models have been used in [8] for analysing equity markets. For multivariate time series, Bücher *et al.* [4] developed tests on local nonparametric copula estimates, but

doi:10.1017/apr.2016.48

strong mixing was assumed. For an overview on goodness-of-fit tests for copulas and further references, we refer the reader to [21].

In Section 2 we start by defining a series of local estimators on overlapping subsamples of length $b = \mathrm{o}(T)$, and in Section 2.2 we derive the basic joint distributional results as the overall observation horizon T tends to ∞. Section 3 contains our original contribution in pursuit of the objective to detect inhomogeneities in the dependence in multivariate time series. In Section 3.1 a binomial test based on independent subsamples is constructed. We discuss the notion of statistical size and local power in detail, and show that the binomial test has local power of order $\sqrt{b}$. The extremal test presented in Section 3.2 is based on subsamples that overlap by a certain fixed amount of observations. The considerations concerning its local power of order $\sqrt{b \log(T/b)}$ entail an extreme-value limit result for asymptotically normal random variables. A third test that uses the whole of the moving-window estimators and has local power of order $\sqrt{T}$ is presented in Section 4.4 of [30]. A summary of the advantages and drawbacks of the different homogeneity tests concludes Section 3.

In Section 4 we apply the binomial and extremal tests to real-world financial data. We find that inhomogeneities in the dependence structure arise with unusual market events and can be remedied by careful univariate modeling. This is in accordance with the general perception that changes in the dependency are an exception rather than the rule, and stem from extreme economic events. In most cases a copula that is constant throughout time is sufficient to capture the dependence among the components of a multivariate time series model, and there is no gain in choosing a more complicated and statistically unhandy model with time-varying copula parameters.

2. The SCOMDY model and its estimators

2.1. Definition of the SCOMDY model

We now define a multivariate time series model which is linked to copula theory. As we cannot provide every technical detail here, we have to refer on several occasions to [30]. The model was first introduced in [7], and subsequently used for further statistical developments, such as value-at-risk calculations in financial portfolios in [25] or local change-point tests of the copula parameter in [22], among others.

Definition 1. Let $\{\boldsymbol{Y}_t\}_{t\geq 1}$ be a vector-valued process such that $\boldsymbol{Y}_t = \boldsymbol{\Sigma}_t^{1/2}\boldsymbol{\varepsilon}_t$ with vector-valued residuals $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \dots, \varepsilon_{dt})$ and conditional covariance matrices $\boldsymbol{\Sigma}_t = (\sigma^2_{ijt})_{i,j=1,\dots,d}$.

(i) Let $\boldsymbol{\Sigma}_t = \operatorname{diag}(\sigma^2_{1t}, \dots, \sigma^2_{dt})$ be such that, for every $j = 1, \dots, d$, the following univariate GARCH(p, q) specification holds:

$$\sigma^2_{jt} = \alpha_{j0} + \sum_{i=1}^{p} \alpha_{ji} Y^2_{j,t-i} + \sum_{k=1}^{q} \beta_{jk}\sigma^2_{j,t-k}.$$

Here $\gamma_j = (\alpha_{j0}, \alpha_{j1}, \dots, \alpha_{jp}, \beta_{j1}, \dots, \beta_{jq}) \in (0, \infty) \times [0, \infty)^{p+q}$. Furthermore, let $\boldsymbol{\varepsilon}_t$, $t \geq 1$, be independent and identically distributed (i.i.d.) according to a d-variate distribution with distribution function $H = C(F_1, \dots, F_d; \theta)$, where $F_1, \dots, F_d$ are arbitrary marginal distributions and $C(\cdot; \theta)$ belongs to a parametric copula family with parameter $\theta \in \Theta \subseteq \mathbb{R}$. Let $\mathbb{E}(\boldsymbol{\varepsilon}_t) = \mathbf{0}$ and $\operatorname{cov}\boldsymbol{\varepsilon}_t = R_t = (r_{ijt})_{i,j=1,\dots,d}$ such that the variables are normalized, i.e. for all $i = 1, \dots, d$,

$$r_{iit} = \operatorname{var}\varepsilon_{it} = 1. \tag{1}$$

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

(ii) Let $\tilde{\gamma}_j$ be the quasi-maximum-likelihood estimator defined using finite sums (see [30, Section 3.1]), and let $\tilde{\varepsilon}_{jt} = Y_{jt}/\tilde{\sigma}_{jt}$ be the empirical residuals in every component $j = 1, \ldots, d$. With $\nu = \nu(T)$ being an offset, set

$$\tilde{U}_{jt} = \frac{1}{T - \nu + 1} \sum_{s=\nu}^{T} \mathbf{1}_{\{\tilde{\varepsilon}_{js} \leq \tilde{\varepsilon}_{jt}\}}, \tag{2}$$

and consider $\tilde{\boldsymbol{U}}_t = (\tilde{U}_{1t}, \ldots, \tilde{U}_{dt})$ for $t = \nu, \ldots, T$ as a pseudo-sample from the copula. As an estimator for the copula parameter θ, consider

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} \frac{1}{T - \nu + 1} \sum_{t=\nu}^{T} \log c(\tilde{U}_{1t}, \ldots, \tilde{U}_{dt}; \theta), \tag{3}$$

a canonical maximum-likelihood estimator offset by ν.

The process $\{\boldsymbol{Y}_t\}_{t \geq 1}$ with the time series specification in (i) and the estimation procedure described in (ii) is called a SCOMDY model.

Remark 1. In Definition 1(i) the parameter of the copula family is univariate so that the testing procedures can be applied immediately—in case of a multidimensional parameter space, i.e. $\boldsymbol{\theta} \in \boldsymbol{\Theta} \subseteq \mathbb{R}^d$ with $d \geq 1$, joint confidence areas have to be constructed. As usual, the sizes of the marginal tests have to be adjusted in such a way that the overall size of the multiple tests meets the theoretical size.

Remark 2. Note that the residuals $\boldsymbol{\varepsilon}_t$ are no longer white noise with identity covariance matrix, but i.i.d. random vectors with mean zero and covariance matrix R which corresponds to the correlation matrix, due to (1). Apart from a different covariance matrix, all white noise properties are preserved, so that we denote the residual process by $\boldsymbol{\varepsilon}_t \sim \mathrm{WN}(\mathbf{0}, R)$. Condition (1) is required for identifiability of the parameters in the univariate GARCH(p, q) models and can be implemented into the SCOMDY model by re-scaling the components by the respective variances. For this, consider $\boldsymbol{\varepsilon}_t^* = (\varepsilon_{1t}^*, \ldots, \varepsilon_{dt}^*) \sim H^* = C^*(F_1^*, \ldots, F_d^*)$ with $\operatorname{var} \varepsilon_{jt}^* = v_j$, and set $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}^*/v_1, \ldots, \varepsilon_{dt}^*/v_d)$. Then the joint distribution for all $(x_1, \ldots, x_d) \in \mathbb{R}^d$ is

$$H(x_1, \ldots, x_d) = H^*(F_1^*(v_1 x_1), \ldots, F_d^*(v_d x_d)),$$

and the marginal distributions and their (generalized) inverses for all $x_j, y_j \in \mathbb{R},\ j = 1, \ldots, d$, are

$$F(x_j) = F_j^*(v_j x_j) \quad \text{and} \quad F_j^-(y_j) = \frac{F_j^{*-}(y_j)}{v_j}.$$

Owing to the invariance principle (scaling invariance), the copula of the distribution of $\boldsymbol{\varepsilon}_t$ is the same as that of $\boldsymbol{\varepsilon}_t^*$, namely C. This implies that restricting the matrix R to have a unit diagonal is merely a condition on the marginal distributions (which can easily be accomplished by re-scaling), and does not impinge on the freedom of choice as to the desired copula.

Remark 3. There are various dependencies in the data (2). First, $\tilde{U}_{jt}$ and $\tilde{U}_{j,t+h}$ are dependent for $0 < h < \nu$, but we take care of that below. That we use pseudo data causes dependencies and distortion as described in [15]. So the size of the Rosenblatt transformation test does not correspond to the theoretical size. We are estimating parameters and this is there reflected by a larger asymptotic variance than with copula data (see [20]), but we still have asymptotic normality and that is what we need.

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

Consistency and an asymptotic normality result due to [6] can be shown. We re-state the respective theorems in Theorem 1 below; the precise prerequisites and notation can be found in the appendix of [30], which we silently assume to hold in the sequel. Furthermore, for $\boldsymbol{x} = (x_1, \dots, x_d)$ and $\boldsymbol{y} = (y_1, \dots, y_d)$, let the real-valued Gaussian process $U(\boldsymbol{x})$ with $\mathbb{E}U(\boldsymbol{x}) = 0$ and $\mathbb{E}(U(\boldsymbol{x})U(\boldsymbol{y})) = \prod\{x_i \wedge y_i\} - \prod x_i y_i$ be given.

Theorem 1. *Let* $\{\boldsymbol{Y}_t\}_{t\geq 1}$ *and* $\hat{\theta}$ *be given as in Definition 1.*

(i) *Suppose that* $\nu(T)/T \to 0$ *as* $T \to \infty$. *Then, as* $T \to \infty$, *we have weak consistency of* $\hat{\theta}$, *i.e.*

$$\hat{\theta} \xrightarrow{\mathrm{P}} \theta.$$

(ii) *If, as* $T \to \infty$, *we have* $\nu(T)/\log T \to \infty$ *and* $\nu(T)/T \to 0$, *then* $\sqrt{T}(\hat{\theta} - \theta)$ *converges in distribution to Z distributed as*

$$\begin{aligned} -\Sigma^{-1}(\theta)\bigg\{ \int \delta(T_d^{-}(x_1, \dots, x_d); \theta)\, \mathrm{d}U(x_1, \dots, x_d) \\ + \sum_{j=1}^{d} \int \delta_j(x_1, \dots, x_d; \theta) U((1, \dots, 1, x_j, 1, \dots, 1)^{\top}) \\ \times c(x_1, \dots, x_d; \theta)\, \mathrm{d}x_1 \cdots \mathrm{d}x_d \bigg\}, \end{aligned}$$

where $\Sigma(\theta)$, δ_j, *and* T_d *are given by*

$$\Sigma(\theta) = \bigg(\mathbb{E}\bigg(\frac{\partial^2}{\partial\theta_i \partial\theta_j} \log c(F_1(\varepsilon_1), \dots, F_d(\varepsilon_d); \theta)\bigg)\bigg)_{i,j=1,\dots,m},$$

$$\delta_j(u_1, \dots, u_d; \theta) = \frac{\partial}{\partial u_j} c(u_1, \dots, u_d; \theta), \tag{4}$$

$$T_j(u_1, \dots, u_j) = (C_1(u_1), C_2(u_2 \mid u_1), \dots, C(u_j \mid u_1, \dots, u_{j-1})). \tag{5}$$

Despite the necessity to fall back on bootstrapping as an auxiliary variance estimation, the SCOMDY model is our model of choice and we assume for all further considerations that we work on the residual vectors $\tilde{\varepsilon}_t$. The model is numerically tractable in the sense that quasi-maximum likelihood for the GARCH parameters and canonical maximum likelihood for the copula parameters are well established and available in many statistical software implementations. As the evolution in the single components depends only on the respective component itself, the stationarity conditions for univariate GARCH(p, q) models transfer directly to the marginal conditional variance models in SCOMDY. The same is true for the asymptotic results with respect to the quasi-maximum-likelihood estimators for the respective parameter sets. As to the finite-dimensional distributions, the SCOMDY model embeds itself in the class of constant-conditional-correlation GARCH(p, q) models, and therefore the heavy-tailedness result of [17] transfers across. There the authors showed in Theorem 5 that stationarity and heavy tailedness of multivariate constant-conditional-correlation GARCH(p, q) models are guaranteed by analogous assumptions as in the univariate case. A proof can be found in [30, Theorem 3.1.5] which develops the arguments along the lines of [29, Theorem 2.1].

Assume that $\boldsymbol{Y}_1, \dots, \boldsymbol{Y}_T$, $T \geq 1$, is a finite sample coming from a SCOMDY process $\{\boldsymbol{Y}_t\}_{t\geq 1}$ as given in Definition 1(i). According to Definition 1(ii), we obtain the empirical residuals $\tilde{\boldsymbol{\varepsilon}}_1, \dots, \tilde{\boldsymbol{\varepsilon}}_T$ as pseudo-realizations of the residual distribution H, and pseudo-realizations

$\tilde{\boldsymbol{U}}_1, \ldots, \tilde{\boldsymbol{U}}_T$ of the copula C_θ through a probability integral transform in every component of the vectors. As we aim at testing for homogeneity of the copula parameter throughout time, consider, for the moment,

$$\tilde{\boldsymbol{U}}_t \sim C_{\theta_t},$$

where θ_t is the present parameter for every point in time $t = 1, \ldots, T$. For an arbitrary $\theta_0 \in \Theta$, the pair of hypotheses we want to test is then

$$\mathbb{H}_0^*\colon \{\theta_t \equiv \theta_0 \text{ for all } t\} \quad \text{versus} \quad \mathbb{H}_A^*\colon \{\text{there exists } t\colon \theta_t \neq \theta_0\}. \tag{6}$$

We address this task nonparametrically through a purely data-driven approach in Section 2.2. Based on the resulting estimators and their asymptotic distribution, three different tests will be constructed and analyzed with respect to their ability to detect actual deviations from the homogeneity hypotheses. In Section 3.1 we present a binomial multiple test based on independent subsamples of the given data, and in Section 3.2 we examine whether the extremes of a parameter series based on overlapping subsamples exceed certain confidence levels.

2.2. Moving window estimators

In order to assess θ_t for $1 \leq t \leq T$, we define maximum-likelihood estimators on certain subsamples of a bandwidth b in the following.

Definition 2. Let $\tilde{\boldsymbol{U}}_1, \ldots, \tilde{\boldsymbol{U}}_T$, $T \geq 1$, be pseudo-realizations of a parametric copula C_θ according to Definition 1(ii). With a bandwidth $1 \leq b = b(T) \leq T$, the *local* or *moving-window maximum-likelihood estimator* $\hat{\theta}_t$ of θ for every $t = b, \ldots, T$ is defined by

$$\hat{\theta}_t = \operatorname*{argmax}_{\theta \in \Theta} \sum_{s=t-b+1}^{t} \log c(\tilde{U}_{1s}, \ldots, \tilde{U}_{ds}; \theta). \tag{7}$$

In the following proposition the consistency and asymptotical normality result for the copula estimators in the SCOMDY model are transferred to the local estimators $\hat{\theta}_t$ in terms of the bandwidth $b = b(T)$.

Proposition 1. *Under the conditions of Theorem 1, and under the assumption that $b = b(T)$ is chosen such that $b \to \infty$ and $b/T \to 0$ as $T \to \infty$, in shorthand notation denoted by $b = \mathrm{o}(T)$, the following statements hold as $T \to \infty$.*

(i) $\hat{\theta}_t \xrightarrow{\mathrm{P}} \theta$.

(ii) $\sqrt{b}(\hat{\theta}_t - \theta) \xrightarrow{\mathrm{D}} \mathcal{N}(0, \Sigma)$, *where Σ is the asymptotic variance in canonical maximum-likelihood estimation as in [20].*

Proof. Statements (i) and (ii) correspond to Theorem 1(i) and (ii), respectively. By imposing the condition $b = \mathrm{o}(T)$ $(T \to \infty)$ on the bandwidth b, the respective offset conditions on ν are fulfilled. □

Now we consider the joint limiting distribution of (two) moving window estimators and show that the covariance structure is nondegenerate if the overlap of the subsamples, which the respective estimators are based on, is asymptotically constant. For points in time $1 \leq s < t \leq T$, consider $\hat{\theta}_t$ and $\hat{\theta}_s$ based on subsamples of bandwidth b according to (7). Their *lag* is defined as $l := |s - t|$, whereas the *overlap* of the subjacent subsamples is $b - l$.

Proposition 2. *Let $1 \leq s, t \leq T$, and let $\hat{\theta}_t$ and $\hat{\theta}_s$ be the respective moving-window estimators. Let $1 \leq l := |s - t| < b(T)$ be such that $l/b(T) \to \mu_l \in (0,1)$ as $T \to \infty$. Under the conditions of Theorem 1, the following joint limiting law holds:*

$$\sqrt{b}((\hat{\theta}_s, \hat{\theta}_t)^\top - (\theta, \theta)^\top) \xrightarrow{\mathrm{D}} \mathcal{N}(\mathbf{0}, \Sigma \bar{M})$$

with matrix

$$\bar{M} = \begin{pmatrix} 1 & 1 - \mu_l \\ 1 - \mu_l & 1 \end{pmatrix}$$

and Σ the scalar variance as in [20].

Proof. The proof can be done by first-order Taylor expansion of the log-likelihood and the classical central limit theorem (for details, see Proposition 4.1.3 of [30]). □

The moving-window estimators $\hat{\theta}_t$ are defined for all $t = b, \dots, T$ and a given realization $\boldsymbol{Y}_1, \dots, \boldsymbol{Y}_T$ of a SCOMDY model. For the subsequent sections, we will regard the estimators as an estimator time series, and proceed in testing the null hypothesis (6) depending on the index set for the times t. To this end, consider the following definitions.

Definition 3. Let a SCOMDY sample of size $T \geq 1$ and the respective moving-window estimators $\hat{\theta}_t$ with bandwidth $b = b(T)$ and for all $t = b, \dots, T$ be given. Consider the *thinning constant* $c \in [1/b, 1]$ and the *series length* $n^{(c)} = \lfloor (T - b)/(cb) \rfloor$, where $\lfloor x \rfloor := \max\{k \in \mathbb{Z} \colon k \leq x\}$ for all $x \in \mathbb{R}$. Set the *index set* as $I^{(c)} := \{t_k = \lfloor T - (n^{(c)} - k)cb \rfloor \colon k = 1, \dots, n^{(c)}\}$ and consider the *estimator series* $\{\hat{\theta}_t\}_{t \in I^{(c)}}$.

(i) For $c = 1$, $\{\hat{\theta}_t\}_{t \in I^{(c)}}$ is called the *independent* estimator series. In particular, we set $n := n^{(1)} = \lfloor (T - b)/b \rfloor$ and $\hat{\theta}_k := \hat{\theta}_{t_k}$ for all $t_k = T - (n - k)b$, $k = 1, \dots, n$.

(ii) For fixed c with $1/b < c < 1$, the estimator series $\{\hat{\theta}_t\}_{t \in I^{(c)}}$ is called *thinned*.

(iii) For $c = c(b) = 1/b$, we refer to $\{\hat{\theta}_t\}_{t \in I^{(c)}}$ as the *full* estimator series. In this case, we stress the dependence on the bandwidth by writing $n(b) := n^{(1/b)} = T - b$ and $t_k(b) = t_k^{(1/b)} = k + b$.

Along with this three-step definition, we proceed in Section 3.1 by presenting a suitable approach to test for parameter homogeneity in the independent estimator series.

3. Homogeneity tests

3.1. Binomial test

Let $\hat{\theta}_1, \dots, \hat{\theta}_n$ be an independent estimator series as in Definition 3(i). The general null hypothesis is replaced by a first approximation in order to take advantage of the independence of the individual estimators $\hat{\theta}_k$. We aim at testing

$$\mathbb{H}_0 \colon \{\theta_{t_1} = \cdots = \theta_{t_n} = \theta_0 \colon t_k = T - (n-k)b\} \quad \text{versus} \quad \mathbb{H}_A \colon \{\text{there exists } t_k \colon \theta_{t_k} \neq \theta_0\}. \quad (8)$$

We clarify the notation and the statistical concepts for one individual estimator in Section 3.1.1. In Section 3.1.2 we perform a multiple simultaneous test of the partial null hypotheses to arrive at the hypothesis given in (8).

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

3.1.1. *Partial tests.* Consider an arbitrary $k = 1, \dots, n$, and set $\hat{\theta} = \hat{\theta}_k$ as well as $\theta = \theta_{t_k}$ in the following. Furthermore, consider Σ as the asymptotic variance in canonical maximum-likelihood estimation given in [20]. Note that the copula parameter estimator $\hat{\theta}$ is determined locally on subsamples preceding the point in time t_k. The asymptotically normal test statistic

$$T_b := \frac{\sqrt{b}(\hat{\theta} - \theta_0)}{\sqrt{\Sigma}} \xrightarrow{\text{D}} \mathcal{N}(0,1) \tag{9}$$

for $b(T) \to \infty$ $(T \to \infty)$ can be constructed to test the partial null hypothesis $\text{H}_0\colon \{\theta = \theta_0\}$. In applications θ_0 is replaced by the global estimate $\hat{\theta}$ (see (3)) without disturbing the asymptotic results below since $\sqrt{b}(\hat{\theta} - \theta_0) \xrightarrow{\text{P}} 0$ as $T \to \infty$ under the null hypotheses.

As noted in [26, p. 433], local power considerations for a given test have to be conducted in every special case of test statistic, asymptotic distribution, and alternative hypothesis. Subsequently, we elaborate the definitions in terms of the test statistic T_b given in (9) and the partial null hypothesis $\text{H}_0\colon \{\theta = \theta_0\}$. For the power calculations, we consider a fixed and a local alternative hypothesis. On the one hand, we have $\text{H}_A(\eta)\colon \{\theta = \theta(\eta) = \theta_0 + \eta\}$ as the fixed alternative; on the other hand, we have $\bar{\text{H}}_A(\eta)\colon \{\theta = \theta_b(\eta) = \theta_0 + \eta/\sqrt{b}\}$ as the local alternative. Both can be regarded as *one sided* if $\eta > 0$ or $\eta < 0$, respectively, and as *two sided* if $\eta \neq 0$.

For the one-sided alternatives, H_0 is rejected to the significance level $\alpha \in (0,1)$ if $T_b > q_\alpha$ or $T_b < q_{-\alpha}$, whereas, for the two-sided alternative, rejection occurs if $|T_b| > q_{\alpha/2}$. Here we denote by q_α the $(1-\alpha)$-quantile of the standard normal distribution, i.e. $\mathbb{P}(Z \leq q_\alpha) = 1 - \alpha$ for $Z \sim \mathcal{N}(0,1)$. For the critical region K, this implies that

$$K(\eta) = \begin{cases} (q_\alpha, \infty) & \text{for } \eta > 0, \\ (-\infty, -q_\alpha) & \text{for } \eta < 0, \\ (-\infty, -q_{\alpha/2}) \cup (q_{\alpha/2}, \infty) & \text{for } \eta \neq 0. \end{cases}$$

Depending on whether the alternative is local or fixed, we derive the power of both the one- and the two-sided tests in the following lemma. Without loss of generality, we assume that $\theta_0 = 0$ and $\Sigma = 1$ for notational ease. The results remain unchanged under general $\theta_0 \in \Theta$ and $\Sigma > 0$.

Lemma 1. *By setting the null hypothesis as* $\text{H}_0\colon \{\theta = 0\}$*, we obtain the fixed alternative* $\text{H}_A(\eta)\colon \{\theta = \eta\}$ *and the local alternative* $\bar{\text{H}}_A(\eta)\colon \{\theta = \eta/\sqrt{b}\}$.

(i) *Let the sample size* $b \geq 1$ *be fixed and* $\eta \in \mathbb{R}$*. Then, for* $\gamma_b(\eta) = \mathbb{P}_{\text{H}_A(\eta)}(T_b \in K(\eta))$*, we have*

$$\mathbb{P}_{\text{H}_A(\eta)}(T_b \in K(\eta)) = (1 + \text{o}(1)) \begin{cases} 1 - \Phi(q_\alpha - \sqrt{b}\eta), & \eta > 0, \\ \Phi(-q_\alpha - \sqrt{b}\eta), & \eta < 0, \\ 1 - \Phi(q_{\alpha/2} - \sqrt{b}\eta) + \Phi(-q_{\alpha/2} - \sqrt{b}\eta), & \eta \neq 0. \end{cases}$$

Note that the asymptotic equality of the left- and right-hand sides is with respect to $b \to \infty$.

(ii) *For* $\eta \in \mathbb{R}$*, we asymptotically have*

$$\bar{\gamma}(\eta) = \lim_{b \to \infty} \mathbb{P}_{\bar{\text{H}}_A(\eta)}(T_b \in K(\eta)) = \begin{cases} 1 - \Phi(q_\alpha - \eta), & \eta > 0, \\ \Phi(-q_\alpha - \eta), & \eta < 0, \\ 1 - \Phi(q_{\alpha/2} - \eta) + \Phi(-q_{\alpha/2} - \eta), & \eta \neq 0, \end{cases}$$

where Φ is the distribution function of the standard normal law and $\mathbb{P}_{\mathrm{H}_A}$ denotes probability under the (alternative) hypothesis H_A.

The proof is obvious by plugging the (local) alternatives into $\gamma_b(\eta)$ and $\bar{\gamma}(\eta)$ respectively (see Lemma 4.2.2 of [30]). Note that, for the fixed alternative, we consider γ_b, i.e. the power of the test depending on the sample size b. For $b \to \infty$, we have $\gamma_b(\eta) \to 1$ for any η in either the one-sided or two-sided testing case. The advantage of local alternatives becomes apparent here as the power $\bar{\gamma}$ is independent of the sample size and the limit for $b \to \infty$ can be established right away.

Furthermore, we see that $\gamma(0) = \alpha$ and $\gamma(\eta) > \alpha$ if $\eta \neq 0$. Consequently, the following theorem formally shows that the test is unbiased, which in the case of the local alternative can be used for a result concerning the local power and its order.

Theorem 2. *Let $\gamma_b(\eta)$ and $\bar{\gamma}(\eta)$ be given as in Lemma 1(i) and (ii), i.e. the fixed and local power of the test to the level $\alpha \in (0,1)$ of $\mathrm{H}_0\colon \{\theta = 0\}$ against $\mathrm{H}_A(\eta)\colon \{\theta = \eta\}$ and $\bar{\mathrm{H}}_A(\eta)\colon \{\theta = \eta/\sqrt{b}\}$, respectively. Then we have, for $\eta \neq 0$,*

$$\gamma_b(\eta) > \alpha \quad \textit{and} \quad \bar{\gamma}(\eta) > \alpha,$$

i.e. the test is asymptotically unbiased against both alternatives H_A and $\bar{\mathrm{H}}_A$. Furthermore, the test against $\bar{\mathrm{H}}_A$ has local power of order $\sqrt{b}$.

We have shown that the asymptotically normal test statistic T_b based on an estimator $\hat{\theta}$ for the copula parameter θ on a subsample of size b is suitable to test $\mathrm{H}_0\colon \{\theta = \theta_0\}$. The test against the local alternative $\bar{\mathrm{H}}_A(\eta)\colon \{\theta = \theta_b(\eta) = \theta_0 + \eta/\sqrt{b}\}$ has been shown to have local power of order $\sqrt{b}$. We now take advantage of the fact that the estimators $\hat{\theta}_1, \dots, \hat{\theta}_n$ as in Definition 3(i) are established on subsamples that do not overlap, i.e. are independent, to construct a multiple test.

3.1.2. *Multiple test.* Our aim is to simultaneously test the following pairs of hypotheses:

$$\mathrm{H}_{0k}\colon \{\theta_{t_k} = \theta_0\} \quad \text{versus} \quad \bar{\mathrm{H}}_{Ak}(\eta)\colon \{\theta_{t_k} = \theta_b(\eta)\}$$

for all $t_k = T-(n-k)b$, $n = \lfloor (T-b)/b \rfloor$. This notation also implies that we limit the following considerations to the local alternatives $\theta_b(t) = \theta_0 + \eta/\sqrt{b}$, $\eta \neq 0$. For every $k = 1, \dots, n$, construct the partial test statistics

$$T_{bk} := \frac{\sqrt{b}(\hat{\theta}_k - \theta_0)}{\sqrt{\Sigma}}.$$

Consider the random variables $Y_k = \mathbf{1}_{\{T_{bk} \in K(\eta)\}}$ that indicate whether the kth partial test rejects the null hypothesis H_{0k} to the level $\alpha \in (0,1)$. Summing them up, we obtain

$$S_n = \sum_{k=1}^{n} Y_k \sim \mathrm{Bin}(n, \alpha) \tag{10}$$

under the null hypothesis of parameter constancy. Note that, since the partial tests are asymptotical tests, the size of the above binomial distribution is only approximately α. The binomial test statistic S_n is now suitable to test the simultaneous null hypothesis $\mathbb{H}_0 = \bigcap_{k=1}^{n} \mathrm{H}_{0k}$ of all the partial null hypotheses being true. The alternatives might range from one false hypothesis

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

up to all hypotheses being false. Formally, we specify the alternative hypothesis for $1 \le m \le n$ by $\mathbb{H}_{A,m}(\eta) = \bigcap_{k=1}^{m} \bar{\mathrm{H}}_{Ak}(\eta) \cap \bigcap_{k=m+1}^{n} \mathrm{H}_{0k}$, that is,

$$\mathbb{H}_{A,m}(\eta)\colon \{\theta_{t_k} = \theta_b(\eta) \text{ for } k \le m,\ \theta_{t_k} = \theta_0 \text{ for } k > m\}.$$

Note that in this setup all the partial alternatives have to be of the same absolute value with respect to η.

The null hypothesis $\mathbb{H}_0$ is rejected to the significance level $\alpha^* \in (0,1)$ if $S_n > c_{\alpha^*}$, where c_{α^*} is the $(1-\alpha^*)$-quantile of the binomial distribution with size n and probability α. The critical region has the form $K = \{c_{\alpha^*}+1, c_{\alpha^*}+2, \dots, n\}$. Note that, with fixed $\alpha, \alpha^* \in (0,1)$, the rejection probability under the null hypothesis $\mathbb{H}_0$ is in general

$$\mathbb{P}_{\mathbb{H}_0}(S_n > c_{\alpha^*}) = \sum_{j=c_{\alpha^*}+1}^{n} \binom{n}{j} \alpha^j (1-\alpha)^{n-j} \le \alpha^*, \tag{11}$$

i.e. the size of the test is less than the desired significance level α^* most of the time. This is usually the case when the distribution of the test statistic is discrete. In order to actually attain the proper size, binomial tests of such kind have to be *randomized*. The existence of p in the following modification of the sample function φ is due to the *fundamental lemma of Neyman and Pearson* (see Theorem 3.2.1 of [26]). The lemma basically states that, for every given significance level $\alpha \in (0,1)$, a test that exactly attains this level can be constructed.

Definition 4. Let the pair of hypotheses H_0 versus H_A and the significance level $\alpha \in (0,1)$ be given. Denote by ∂K the boundary of the respective critical region K and by $\bar{K}$ its closure, and set

$$\tilde{\varphi}(x_1, \dots, x_b) = \begin{cases} 1, & (x_1, \dots, x_b) \in K, \\ R, & (x_1, \dots, x_b) \in \partial K, \\ 0, & (x_1, \dots, x_b) \notin \bar{K}, \end{cases}$$

where $R \sim \mathrm{Bin}(1, p)$, and the value $p \in (0,1)$ can be chosen such that $\mathbb{P}_{\mathrm{H}_0}(\tilde{\varphi} = 1) = \alpha$ and $\mathbb{P}_{\mathrm{H}_A}(\tilde{\varphi} = 0) \ge \alpha$. This test is called a *randomized test* to the level α.

In the case of the binomial test statistic S_n in (10) we set

$$\tilde{\varphi}(s) = \begin{cases} 1, & c_{\alpha^*} \le s-1, \\ R, & s-1 < c_{\alpha^*} \le s, \\ 0, & s < c_{\alpha^*}, \end{cases}$$

for a realization $0 \le s \le n$ of $S \sim \mathrm{Bin}(n, \alpha)$, and choose $p \in (0,1)$ such that

$$1 \cdot \mathbb{P}(S \ge c_{\alpha^*} + 1) + p \cdot \mathbb{P}(S = c_{\alpha^*}) = \alpha^*, \tag{12}$$

i.e. $p = (\alpha^* - \mathbb{P}(S \ge c_{\alpha^*} + 1))/\mathbb{P}(S = c_{\alpha^*})$. Practically, we sample r from $R \sim \mathrm{Bin}(1, p)$ whenever the specific realization s of S_n is equal to c_{α^*}, and reject $\mathbb{H}_0$ if $s + r = c_{\alpha^*} + 1$.

It is highly recommended to avoid randomized tests, as the outcome for a specific sample with a fixed significance level is not deterministic, and hence not reproducible (see [26, p. 75]). Since we sum over integer-valued j in (11), it is possible to fix $\kappa := c_{\alpha^*} + 1 \in \mathbb{N}$ and then solve

$$\sum_{j=\kappa}^{n} \binom{n}{j} \alpha^j (1-\alpha)^{n-j} = \alpha^* \tag{13}$$

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

for the partial level $\alpha \in (0,1)$. With this α, it holds that $\kappa - 1 = c_{\alpha^*}$ and that

$$\mathbb{P}_{\mathbb{H}_0}(S_n > c_{\alpha^*}) = \alpha^*.$$

No clear advice can be given on how to choose κ. Its choice depends very much on the number n of partial hypotheses and on the alternative. We discuss this further at the end of this section. Formula (13) is also given in [35] (note that there is an error in the lower bound of the second sum due to a misprint of 1 for $l = n - m$). The authors compared the binomial multiple test to other testing procedures for multiple comparisons, such as the [33] or the [18] test.

Now we turn to the local power $\mathbb{P}_{\mathbb{H}_{A,m}(\eta)}(S_n > c_{\alpha^*})$ of the multiple test and show that the local power of order $\sqrt{b}$ transfers to the multiple test regardless of the number m of false partial alternatives.

Theorem 3. *The binomial test $S_n > c_{\alpha^*}$ against the alternative $\mathbb{H}_{A,m}(\eta)$ is asymptotically unbiased for $\eta \neq 0$, and, thus, has local power of order $\sqrt{b}$ for all $m = 1, \dots, n$.*

For the proof, see [30].

Summarizing, we have constructed a multiple binomial test with proper size and local power of order $\sqrt{b}$. The multiple null hypothesis is composed of n independent partial null hypotheses and is rejected when there are at least κ partial rejections. We are free to choose $1 \leq \kappa \leq \lfloor n/2 \rfloor$, and arrive at a test that attains the significance level α^* under the null hypothesis. In order to maximize the local power, we should take into account the number of false partial hypotheses in our choice of κ.

In the next subsection we turn to another approximation of the global null hypothesis given in (6) that allows us to test against a more general alternative. The test we construct makes use of a greater portion of the estimator series $\{\hat{\theta}_t\}$ by allowing the underlying subsamples to overlap.

3.2. Extremal test

The null hypothesis $\mathbb{H}_0^*\colon \{\theta_t \equiv \theta_0 \text{ for all } t\}$ means that the copula parameter is constant for all points in time t. A consequence of this is that in particular $\max_t \theta_t = \theta_0$. Thus, we subsequently test the following hypothesis and its one-sided alternative:

$$\mathbb{H}_0'\colon \Big\{\max_t \theta_t = \theta_0\Big\} \quad \text{versus} \quad \mathbb{H}_A'\colon \Big\{\max_t \theta_t > \theta_0\Big\} \tag{14}$$

for an arbitrary $\theta_0 \in \Theta$. As the basis for our test statistic, consider a fixed constant c with $1/b < c < 1$ and the thinned estimator series $\{\hat{\theta}_{t_k}\}$ for $t_k = T - (n^{(c)} - k)cb$, $k = 1, \dots, n^{(c)} = \lfloor (T-b)/cb \rfloor$ (cf. Definition 3(ii)).

For notational convenience, we use a centered and scaled version of $\hat{\theta}_{t_k}$ for all $k = 1, \dots, n^{(c)}$ such that, under $\mathbb{H}_0^*$, we have

$$\xi_k := \sqrt{b}(\hat{\theta}_{t_k} - \theta_0) \xrightarrow{\mathrm{D}} \mathcal{N}(0, \Sigma) \quad \text{as } b \to \infty$$

due to Proposition 1(ii).

In order to test the pair of hypotheses in (14), a test statistic based on the vector $\boldsymbol{\xi} = (\xi_1, \dots, \xi_{n^{(c)}})$ is suitable. For assessing the statistic's distribution, we first need the joint distribution of the vector $\boldsymbol{\xi}$, and secondly, extreme value theory for asymptotically normal variables.

3.2.1. *Multivariate central limit theorem.* We trace back the structure of $\boldsymbol{\xi} = (\xi_1, \dots, \xi_{n^{(c)}})$ to the summands in the likelihood estimation, and then establish its asymptotic distribution. To this end, we work on a random quantity $\boldsymbol{S}_b$ exhibiting a special structure that we need to obtain our results. Although the subsample size b grows with the overall observation horizon T, the overlap proportion remains fixed. A reordering of the original summation enables us to apply a recent result [3] on convergence rates in the multivariate central limit theorem that depends on the involved dimension. Together with bounds on the elements in the square-root inverse of the occurring covariance matrices (a result given in [12]), this finally leads to conditions on b and T that guarantee distributional convergence of $\boldsymbol{S}_b$ to a Gaussian distribution (the relevant covariance structure is discussed in [30]).

Now we can proceed to the distributional convergence of $\boldsymbol{S}_b$ as the sample size increases to ∞. In [3], Lyapunov-type Berry–Esseen bounds in the multivariate central limit theorem are given, inheriting a dimensional dependence. A version of this result suitable for our purposes is stated below.

Theorem 4. *Let $\boldsymbol{Y}^{(1)}, \dots, \boldsymbol{Y}^{(n)} \in \mathbb{R}^d$ be independent random vectors with $\mathbb{E}\boldsymbol{Y}^{(k)} = \boldsymbol{0}$ for all k. Let $\boldsymbol{S} = \boldsymbol{Y}^{(1)} + \dots + \boldsymbol{Y}^{(n)}$, and let $\boldsymbol{C}^2 = \operatorname{cov} \boldsymbol{S}$. Denote by $\Phi_{d,\boldsymbol{C}^2}$ the distribution function of the d-variate centered Gaussian distribution with covariance matrix $\boldsymbol{C}^2$.*

Then there exists a positive constant C_2 such that

$$\Delta := \sup_{\boldsymbol{x} \in \mathbb{R}^d} |F_{\boldsymbol{S}}(\boldsymbol{x}) - \Phi_{d,\boldsymbol{C}^2}(\boldsymbol{x})| \le C_2 d^{1/4} \beta, \tag{15}$$

where $F_{\boldsymbol{S}}$ is the d-variate distribution function of $\boldsymbol{S}$ and $\beta = \beta_1 + \dots + \beta_n$ with $\beta_k = \mathbb{E}|\boldsymbol{C}^{-1}\boldsymbol{Y}^{(k)}|^3$.

Proof. See Theorem 1.1 of [3]. □

Besides this convergence result, the special structure of the covariance matrix will now be exploited to derive the limiting distribution of the target variable $\boldsymbol{S}_b$ as well as a convergence rate depending on the sample sizes T and b.

The vector $\boldsymbol{S}_b$ can now be seen as the sum of the i.i.d. vectors $\boldsymbol{Y}^{(k)}$ (with a special structure as defined in Section 4.3.1 of [30]), i.e.

$$\boldsymbol{S}_b = \frac{1}{\sqrt{b}} (\boldsymbol{Y}^{(1)} + \dots + \boldsymbol{Y}^{(cb)}),$$

with covariance matrix

$$\boldsymbol{C}^2 = \operatorname{cov} \boldsymbol{S}_b = c\Sigma M, \tag{16}$$

where c is the thinning constant. We now have the following result (see Theorem 4.3.5 of [30]).

Theorem 5. *For $b = T^\delta$, $\delta \in (0, 1)$, and $n^{(c)} = \mathrm{O}(T/b)$, let Δ_b be defined in analogy to (15), namely,*

$$\Delta_b = \sup_{\boldsymbol{x} \in \mathbb{R}^{n^{(c)}}} |F_{\boldsymbol{S}_b}(\boldsymbol{x}) - \Phi_{n^{(c)},\boldsymbol{C}^2}(\boldsymbol{x})|.$$

If $\delta > \frac{7}{9}$ then $\Delta_b \to 0$ as $T \to \infty$.

To see that Δ_b actually converges to 0 as $T \to \infty$, we apply Theorem 4 to Δ_b and show that β is bounded from above by (const) $n^{(c)3/2} b^{-1/2}$ by applying Hölder's inequality two times and the result in [12] to the Cholesky square root of M.

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

Note that this specific ratio of the subsample size b to the overall sample size T ensures that the vector of normalized subsample sums is asymptotically jointly normal, even if the samples overlap by a certain fixed amount of b.

3.2.2. *Asymptotic extreme value theory.* We show that the maximal component of the asymptotically Gaussian target vector $\boldsymbol{S}_b$ converges in law. To this end, recall that the kth component of $\boldsymbol{S}_b$ is

$$S_{bk} = \frac{1}{\sqrt{b}} S_k = \frac{1}{\sqrt{b}} \sum_{\tau=t_k-b+1}^{t_k} X_\tau.$$

Furthermore, the covariance matrix is denoted by $\boldsymbol{C}^2 = \operatorname{cov}(S_{b1}, \dots, S_{bN}) = \operatorname{cov} \boldsymbol{S}_b = c\Sigma M$, as in (16).

Theorem 6. *If* $T^{7/9}/b = \mathrm{o}(1)$ *and* $n^{(c)} = (T-b)/(cb)$, *then*

$$\mathbb{P}\Big(\max_{k=1,\dots,n^{(c)}} S_{bk} \le z_{n^{(c)}}\Big) \to \mathrm{e}^{-\mathrm{e}^{-z}} \quad \text{as } T \to \infty$$

with $z_{n^{(c)}} = a_{n^{(c)}} z + d_{n^{(c)}}$ *for suitable* $z \in \mathbb{R}$, *where*

$$a_{n^{(c)}} = \sqrt{\frac{\Sigma}{2\log n^{(c)}}}, \qquad d_{n^{(c)}} = \sqrt{\Sigma}\bigg(\sqrt{2\log n^{(c)}} - \frac{\log\log n^{(c)} + \log 4\pi}{\sqrt{2\log n^{(c)}}}\bigg). \tag{17}$$

Proof. Let $n^{(c)} = n^{(c)}(T) = (T-b)/(cb) = \mathrm{O}(T/b)$, and recall that $c = 1/\mathcal{K}$, $\mathcal{K} \in \mathbb{N}$. For $\boldsymbol{C}^2 = (c_{ij}^2)$, it holds that

$$c_{ij}^2 = \gamma_{|i-j|} = \begin{cases} \Sigma\Big(1 - \dfrac{|i-j|}{\mathcal{K}}\Big), & 0 \le |i-j| \le \mathcal{K}, \\ 0, & |i-j| > \mathcal{K}. \end{cases}$$

This means that $\{S_{bk}\}_{k\ge1}$ is a $\mathcal{K}$-dependent sequence. In particular, we have

$$\lim_{h\to\infty} \gamma_h \log h = 0.$$

Moreover, it is stationary, i.e. $S_{bk} \sim F_b$ for all k, since the summation extends over i.i.d. variables X_i. Define $\boldsymbol{z}_{n^{(c)}} = (z_{n^{(c)}}, \dots, z_{n^{(c)}})$ as the $n^{(c)}$-dimensional vector with all components being $z_{n^{(c)}}$, as above. According to Theorem 4, it holds that

$$\begin{aligned} \mathbb{P}\Big(\max_{k=1,\dots,n^{(c)}} S_{bk} \le z_{n^{(c)}}\Big) &= \mathbb{P}(S_{b1} \le z_{n^{(c)}}, \dots, S_{bn^{(c)}} \le z_{n^{(c)}}) \\ &= F_{\boldsymbol{S}_b}(\boldsymbol{z}_{n^{(c)}}) \\ &\le \Phi_{n^{(c)},\boldsymbol{C}^2}(\boldsymbol{z}_{n^{(c)}}) + \Delta_b \\ &= \mathbb{P}(Z_1 \le z_{n^{(c)}}, \dots, Z_{n^{(c)}} \le z_{n^{(c)}}) + \Delta_b, \end{aligned}$$

where the vector $(Z_1, \dots, Z_{n^{(c)}})$ is $n^{(c)}$-variate Gaussian with mean vector $\mathbf{0}$ and covariance matrix $\boldsymbol{C}^2$.

Theorem 5 and the assumption that $T^{7/9}/b = \mathrm{o}(1)$ ensure that the second summand Δ_b converges to 0 as $T \to \infty$. For the first summand, a standard argument on the maximum of random variables can be applied. Since $n^{(c)}$ tends to ∞ as $T \to \infty$ and, owing to the appropriate covariance structure (see (16)), it holds that

$$\mathbb{P}(Z_1 \le z_{n^{(c)}}, \dots, Z_{n^{(c)}} \le z_{n^{(c)}}) = \mathbb{P}(M_{n^{(c)}} \le a_{n^{(c)}} z + d_{n^{(c)}}) \to \mathrm{e}^{-\mathrm{e}^{-z}} \quad \text{as } T \to \infty.$$

This concludes the proof. □

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

3.2.3. *Local power.* Now we are ready to state an asymptotic distributional result for the extremes of ξ_k as an application of Theorem 6.

Theorem 7. *Let $a_{n^{(c)}}$ and $d_{n^{(c)}}$ be as in (17). If $T^{7/9}/b = \mathrm{o}(1)$, we have*

$$\mathbb{P}\Big(\max_{k=1,\dots,n^{(c)}} \xi_k \le a_{n^{(c)}} z + d_{n^{(c)}}\Big) \to \mathrm{e}^{-\mathrm{e}^{-z}} \quad \text{as } T \to \infty \tag{18}$$

for all $z \in \mathbb{R}$.

Proof. Proposition 2 can be extended beyond the bivariate case to show that the $n^{(c)}$-variate vector $(\xi_1, \dots, \xi_{n^{(c)}})$ is asymptotically normally distributed with mean $\mathbf{0}$ and covariance matrix $c\Sigma M$. To be precise, we know that

$$(\xi_1, \dots, \xi_{n^{(c)}}) \xrightarrow{\mathrm{D}} (Z_1, \dots, Z_{n^{(c)}}) \sim \mathcal{N}(\mathbf{0}, c\,\Sigma M).$$

By assumption, $\delta > \frac{7}{9}$, so, by setting $S_{bk} = \xi_k$, Theorem 6 is directly applicable to establish (18). □

This result can now be exploited to construct a global test of parameter constancy over the observation horizon T. Define $M_{n^{(c)}} := \max_{k=1,\dots,n^{(c)}} \xi_k$, and consider the test statistic

$$T'_{n^{(c)}} = \frac{M_{n^{(c)}} - d_{n^{(c)}}}{a_{n^{(c)}}} \xrightarrow{\mathrm{D}} \Lambda \quad \text{as } n^{(c)} \to \infty, \tag{19}$$

where Λ denotes the Gumbel extremal distribution. The extremal test rejects the null hypothesis $\mathbb{H}'_0$ against the alternative that the maximal value of the thinned parameter series exceeds a fixed value $\theta_0 \in \Theta$, to the confidence level $\alpha \in (0,1)$, if $T_{n^{(c)}} > \lambda_\alpha$, where λ_α denotes the $(1-\alpha)$-quantile of the Gumbel extremal distribution. The critical region of the test is then $K' = (\lambda_\alpha, \infty)$. We are again interested in the local power of the test, and, thus, consider the following local hypothesis for $\eta > 0$:

$$\mathbb{H}'_A(\eta)\colon \left\{\theta_{t_k} = \theta'(\eta) = \theta_0 + \eta\sqrt{\frac{\Sigma}{b \log n^{(c)}}}\right\}.$$

The extremal test's properties with respect to its statistical power are summarized in the following theorem.

Theorem 8. *Let $T'_{n^{(c)}}$ be given as in (19), i.e. the centered and normalized maximum of the thinned and re-scaled estimator series $\{\xi_k\}_{k=1,\dots,n^{(c)}}$. Under the conditions of Theorem 7, especially under the assumption that $T^{7/9}/b = \mathrm{o}(1)$, and under the local alternative $\mathbb{H}'_A(\eta)$ for $\eta \ge 0$, consider the extremal test to the level $\alpha \in (0,1)$ as $T'_{n^{(c)}} > \lambda_\alpha$.*

(i) *As $n^{(c)} \to \infty$ with $b \to \infty$, the asymptotic local power of the extremal test is*

$$\mathbb{P}_{\mathbb{H}'_A(\eta)}(T'_{n^{(c)}} > \lambda_\alpha) \to \Gamma'(\eta) = 1 - \Lambda(\lambda_\alpha - \eta\sqrt{2}),$$

where $\Lambda(\cdot)$ is the distribution function of the Gumbel extremal distribution.

(ii) *For $\eta = 0$, we have $\Gamma'(\eta) = \alpha$, and, for $\eta > 0$, we have $\Gamma'(\eta) > \alpha$. Thus, the extremal test has local power of order $\sqrt{b \log n^{(c)}}$.*

Proof. For notational convenience, we set $\mathbb{P} = \mathbb{P}_{\mathbb{H}'_A}$, $n = n^{(c)}$, and $\theta_0 = 0$, and consider the rejection probability under the alternative, i.e.

$$\begin{aligned}
\mathbb{P}(T'_n > \lambda_\alpha) &= \mathbb{P}(\xi_1 > a_n\lambda_\alpha + d_n, \dots, \xi_n > a_n\lambda_\alpha + d_n) \\
&= \mathbb{P}(\sqrt{b}(\hat{\theta}_{t_1} - \theta'(\eta)) > a_n\lambda_\alpha + d_n - \sqrt{b}\,\theta'(\eta), \dots, \\
&\qquad \sqrt{b}(\hat{\theta}_{t_n} - \theta'(\eta)) > a_n\lambda_\alpha + d_n - \sqrt{b}\,\theta'(\eta)) \\
&= \mathbb{P}(\tilde{\xi}_1 > a_n(\lambda_\alpha - \eta\sqrt{2}) + d_n, \dots, \tilde{\xi}_n > a_n(\lambda_\alpha - \eta\sqrt{2}) + d_n) \\
&= \mathbb{P}(\tilde{M}_n > a_n(\lambda_\alpha - \eta\sqrt{2}) + d_n) \to 1 - \Lambda(\lambda_\alpha - \eta\sqrt{2}),
\end{aligned}$$

where $a_n = \sqrt{\Sigma/(2\log n)}$ in the case of centered Gaussian variables with variance Σ, and Theorem 7 was applied. □

Note that, for any sequence $\{z_n\}_{n\in\mathbb{N}}$, we have $\min\{z_n\} = -\max\{-z_n\}$, and, furthermore, we know that, for $X \sim \mathcal{N}(0,1)$, it holds that $X \overset{\mathrm{D}}{=} -X$. Therefore, a similar result for the minimal limiting law of the estimator series $\{\xi_k\}$ holds, namely, for all $z \in \mathbb{R}$, we have

$$\mathbb{P}\Big(\min_{k=1,\dots,n^{(c)}} \xi_k > -a_{n^{(c)}} z - d_{n^{(c)}}\Big) \to \mathrm{e}^{-\mathrm{e}^{-z}}$$

with $a_{n^{(c)}}$ and $d_{n^{(c)}}$ as in Theorem 7. This can be used to test $\mathbb{H}'_0$ against the alternative $\mathbb{H}''_A$: $\{\min_t \theta_t < \theta_0\}$ with the test statistic $T''_{n^{(c)}} = -(m_{n^{(c)}} + d_{n^{(c)}})/a_{n^{(c)}}$, where $m_{n^{(c)}} = \min_{k=1,\dots,n^{(c)}} \xi_k$. Rejecting the null hypothesis upon $T''_{n^{(c)}} < \lambda_{1-\alpha}$ yields an asymptotically unbiased test with local power of the same order as in the maximum case.

Furthermore, as the sequence $\{\xi_k\}_{k=1,\dots,n^{(c)}}$ is $\mathcal{K}$-dependent, we have asymptotic independence of $M_{n^{(c)}}$ and $m_{n^{(c)}}$, so we can also test against the alternative $\mathbb{H}'''_A$: $\{\min_t \theta_t < \theta_0 \text{ or } \max_t \theta_t > \theta_0\}$, i.e. that the extrema of the parameter series exceed some confidence band. The null hypothesis is then rejected if either $T'_{n^{(c)}} > \lambda_{\alpha/2}$ or $T''_{n^{(c)}} < \lambda_{1-\alpha/2}$. This testing procedure results in a symmetric simultaneous test with proper asymptotic size and local power of the same order as in the preceding cases.

3.3. Discussion

To conclude, we discuss the two types of test from the previous sections with respect to their ability to detect the presence and type of deviations from the null hypothesis of parameter homogeneity.

In Section 3.1 we presented a binomial test with local power of order $\sqrt{b}$, where $b = \mathrm{o}(T)$ as $T \to \infty$. The facts that the test does not achieve the usual local power of order $\sqrt{T}$, as might be expected when a sample of size T is available, and that the local estimates are restricted to be independent may seem to be a disadvantage at first sight. At the same time, take into account the fact that in using moving-window estimation we implicitly presume a locally homogeneous structure in the copula parameter anyway. Under this perspective, a rejection of one (or more) of the partial hypotheses hints at one (or more) regime changes in the parameter series. In order to avoid multiple testing, the multiple approach can be used. We again refer to Section 4 for an empirical application of the binomial test on real-world data.

The extremal test presented in Section 3.2 has local power of order $\sqrt{b\log n^{(c)}}$, where $n^{(c)} = \lfloor (T-b)/(cb) \rfloor$ and $T^{7/9}/b = \mathrm{o}(1)$ with respect to the overall observation horizon T. This local power is not least among the considered tests, but the speed of the approximation in the limit law is low. Furthermore, the applicability of the procedure is limited in finite samples due to the restriction $T^{7/9}/b = \mathrm{o}(1)$. This bandwidth condition is practically only fulfilled in

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

samples over a long observation horizon, say for T of the order of decades if we consider daily price data. Nevertheless, the test is suitable to test against the most general alternative in our setup, namely whether there is an exceedance from a long-term constant parameter value at all, and if the deviation is in a sense extreme, i.e. the maximal (or minimal) parameter value peaks high above (or below) an assumed average level. Motivated by the empirical observations in real-world data where the estimated parameter series sometimes exhibit high peaks (mostly above the average parameter level), this test was developed as an omnibus test that allows us to judge if peaks in the parameter value lie within a predefined confidence region. In the course of the data description and empirical analyses in Section 4, we also present what the rejection of the extremal null hypothesis $\mathbb{H}'_A$ could be used for after that.

Summarizing, each of the presented tests has certain pros and cons, and on its own is able to detect only the special alternatives that were outlined in the respective sections. We can though combine the testing of the respective null hypotheses in the same order as they were presented in this discussion: first, the extremal test lends itself to decide whether there is an extremal deviation of an average copula parameter; second, the binomial test provides evidence for moderate deviation and/or the location of possible change points. Finally, a suspected fixed regime change can be verified through a so-called forecast regression based test; see, e.g. [30] and [36].

4. Empirical results

4.1. Empirical size of the binomial test

Under the null hypothesis of parameter constancy, we simulate copula data of different dimensions and with different parameter values. Then the binomial test is run with various types of base data in order to investigate the test's empirical size. Within this scope, the proper determination of the involved estimation variance proves vital for a correct testing procedure. The main outcome is that bootstrapping the variance cannot be avoided in our finite-sample setup notwithstanding the theoretical results. We refer the reader to [30] for details.

4.2. Commodity contracts

In recent years the so-called financialisation of commodities, which implies a growing interaction of commodity markets with classical equity and fixed income markets, has been investigated in a number of studies. We contribute to this discussion by investigating the dependency structure of typical commodity contracts.

4.2.1. *Data description.* We have chosen to work on so-called second front-month forward prices for different types of commodities, such as coal, natural gas, crude oil, and electricity. We have decided to consider the second front month as these are the contracts that are traded most liquidly. Further information on these contracts can be found in [5].

The historical time series data for our analysis were retrieved from Bloomberg's business database. We consider the following contracts identified by their Bloomberg ticker.

`API22MON index`: coal, price per metric ton (based on all published indices (API)), delivered to the Amsterdam, Rotterdam and Antwerp region (ARA) in Northwest Europe.

`API42MON index`: coal, same as above, but the delivery takes place to Richards Bay, South Africa. The pricing source was not Bloomberg, but Credit Suisse.

`TTFG2MON index`: natural gas, price per megawatt-hour, based on the virtual gas hub covering all entry and exit points in the Netherlands.

`NBPG2MON index`: natural gas, price per megawatt-hour, based on the market movements at the virtual gas hub covering mainland Britain.

`CO2 comdty`: Brent crude oil, price per barrel, supplied at Sullom Voe on Shetland, Scotland.

`BRSWMO2 index`: Brent crude oil, price per barrel, calculated by Bloomberg based on the Intercontinental Exchange (ICE) futures prices.

`EL{U,G,B}B2MON index`: price per megawatt-hour of base-load electricity in the United Kingdom (U), Germany (G) and Belgium (B), respectively.

`EL{U,G}P2MON index`: price per megawatt-hour of peak-load electricity in the United Kingdom and Germany, respectively.

After cleaning the data (described in Section 5.2.1 of [30]) the transformed series of log-returns

$$R_t^{(j)} = \log\biggl(\frac{S_t^{(j)}}{S_{t-1}^{(j)}}\biggr)$$

for all $j = 1, \dots, 11,\ t \geq 1$, forms the basis of all further considerations.

In the subsequent subsection we describe the univariate GARCH models for the log-return series, how we proceeded to decide on goodness-of-fit of the different copula classes, and the test results of the binomial and extremal tests.

4.2.2. *SCOMDY modeling and testing.* After data retrieval and preparation, the focus is on fitting appropriate time series models and finding a proper copula family to capture the dependence structure.

In a first analysis we use one year of historical data in the period from 18-03-2008 to 18-03-2009 and consider all of the possible pairwise combinations, resulting in 55 bivariate series vectors. Subsequently, we fit several of the copula families presented in Appendix A of [30] to the rank-transformed log-returns, i.e. we do not deGARCH the series first, but directly investigate the present pairwise dependence structure. According to the maximum-likelihood logic, we choose the minimal value of the negative log-likelihood function of the whole series for $t = 1, \dots, T$ as the goodness-of-fit criterion, i.e. we select a copula family as 'best fit' for a log-return pair when the respective negative log-likelihood function is minimal compared to the other families. The result is that the Clayton copula family with a positive parameter (see Section A.5 of [30]) emerges most frequently as 'best fit'.

The second step is taken with regard to suitable univariate GARCH models. As argued in Section 3.1 of [30] with reference to [24], it suffices in most cases to set up the parsimonious GARCH(1, 1) model. Therefore, we assume that, for the log-returns, we have

$$R_t^{(j)} = \mu_j + Y_{jt}, \qquad Y_{jt} = \sigma_{jt}\varepsilon_{jt}, \qquad \sigma_{jt}^2 = \omega_j + \alpha_j Y_{j,t-1}^2 + \beta_j \sigma_{j,t-1}^2,$$

with parameters $\mu_j, \omega_j, \alpha_j$, and β_j, and univariate white-noise residuals $\varepsilon_{jt} \sim \mathrm{WN}(0,1)$ with respect to t for all $j = 1, \dots, 11$. This time, we take as basis for the estimation two years' time series data from 19-11-2007 to 19-11-2009 and obtain the following parameter estimates in the univariate models. In Table 1 the respective estimates are given together with the sum $\hat{\alpha}_j + \hat{\beta}_j$. A value greater than or equal to 1 indicates that the respective model is not covariance stationary; note, however, that this does not affect the consistency of the quasi-maximum-likelihood estimates.

TABLE 1: Estimated parameters for univariate commodity GARCH(1, 1) models.

j		$\hat{\mu}_j$	$\hat{\omega}_j$	$\hat{\alpha}_j$	$\hat{\beta}_j$	$\hat{\alpha}_j + \hat{\beta}_j$
1	API2	8.7409e−04	6.4357e−06	0.146 726	0.860 399 60	1.007 13
2	API4	1.8548e−03	1.3157e−05	0.350 214	0.703 033 20	1.053 25
3	TTFG	5.7206e−04	1.8513e−05	0.135 967	0.876 727 40	1.012 69
4	NBPG	7.1524e−04	2.9631e−05	0.077 660	0.913 061 80	0.990 72
5	CO2	9.5303e−04	8.7841e−06	0.071 987	0.914 766 80	0.986 75
6	BRSW	9.2096e−04	8.2616e−06	0.063 581	0.922 852 10	0.986 43
7	ELUB	2.1332e−04	8.9599e−05	0.321 522	0.621 478 20	0.943 00
8	ELGB	3.5195e−05	8.0567e−04	0.211 854	0.000 000 01	0.211 85
9	ELBB	2.4977e−04	1.0554e−03	0.160 923	0.036 532 76	0.197 46
10	ELUP	−8.6258e−05	1.3132e−04	0.231 266	0.630 506 10	0.861 77
11	ELGP	2.1356e−04	8.9710e−04	0.358 092	0.000 000 01	0.358 09

Note that we abbreviated the tickers. After deGARCHing the original log-return series, we obtain 11 univariate series of empirical residuals which are rank transformed before the further investigation. All bivariate combinations result in $\binom{11}{2} = 55$ possible pairs. The pairwise dependence is captured by a bivariate Clayton copula. In Table 2 the global parameter estimates $\hat{\theta}$ are given in the lower-left triangle, and the respective variance estimates $\hat{\Sigma}$ obtained by the bootstrap procedure described in Section 5.1.2 of [30] are given in the upper-right triangle.

The two considered trading years sum up to the observation horizon $T = 539$, the bandwidth is chosen as one tenth of T, namely $b = 54$ (which accounts for merely $n = 9$ independent estimators), and the thinning constant is $c = \frac{1}{4}$ (which accounts for $N = 37$ weakly dependent estimators in the thinned series). With a significance level of $\alpha = \alpha^* = 5\%$ and the number of local exceedances less than $q_{\alpha^*} = 2$, the null hypothesis of parameter constancy cannot be rejected. On the other hand, if the number of exceedances is greater than or equal to 3 then the null hypothesis is rejected. For exactly two local exceedances, a randomization has to be conducted, i.e. by generating r as a realization of a Bernoulli random variable with size p as in (12), the null is rejected if $r = 1$ and cannot be rejected if $r = 0$. Furthermore, the extremal test is conducted in its min-max version, i.e. we test for both deviations of the maximum above

TABLE 2: Estimated global parameter (*lower left*) and bootstrapped variance estimate (*upper right*) for bivariate SCOMDY models.

	3.4427	1.5737	1.5563	2.2915	2.3264	1.4272	1.6389	1.4685	1.3076	1.5001	API2
1.2703		1.5303	1.4775	2.0483	2.1947	1.6920	1.7098	1.7401	1.5123	1.5528	API4
0.1407	0.1635		8.2404	1.6686	1.6423	2.2104	2.1753	1.6393	2.0779	2.4668	TTFG
0.1881	0.1702	1.6905		1.8589	1.7635	2.4607	2.8032	1.8740	1.9904	2.5159	NBPG
0.3345	0.3159	0.1743	0.1895		78.7038	1.8754	1.8699	1.5593	1.6921	1.5389	CO2
0.3095	0.3081	0.1151	0.1514	6.0296		1.7605	1.8124	1.5718	1.5494	1.5377	BRSW
0.2659	0.1791	0.6536	0.8289	0.2813	0.2536		3.0064	2.0581	2.4142	2.4250	ELUB
0.3125	0.2288	0.4261	0.4527	0.1774	0.1587	0.6175		4.4794	3.9574	7.2225	ELGB
0.2395	0.1776	0.4166	0.3651	0.2070	0.1873	0.5471	1.2602		2.4444	2.5953	ELBB
0.1943	0.1685	0.6918	0.7331	0.2387	0.2068	1.6195	0.7405	0.7813		4.6296	ELUP
0.1739	0.1172	0.4409	0.4247	0.1221	0.1088	0.6480	1.6319	1.1894	0.9103		ELGP
API2	API4	TTFG	NBPG	CO2	BRSW	ELUB	ELGB	ELBB	ELUP	ELGP	

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

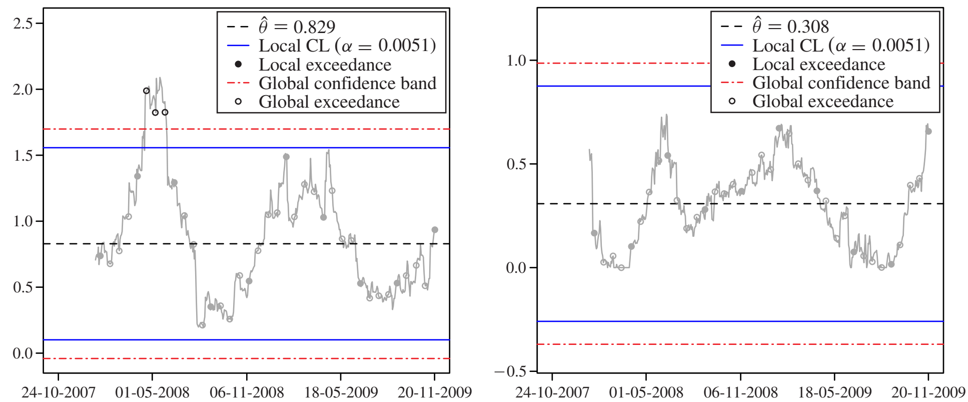

(a) NBPG2MON versus ELUB2MON (b) API42MON versus BRSWMO2

FIGURE 1: Global exceedance and absence of exceedances.

the upper global confidence level (CL), as well as deviations of the minimum below the lower global CL. The test results can be visualized in a very intuitive way (see Figure 1); by giving two examples in the case of the two-year data we demonstrate on the one hand so-called global exceedances which are deviations of the extremal values of the estimator series, and that they possibly may come along with extreme real-world events. On the other hand, we will explain the different occurring elements of the graphics.

The horizontal dashed lines depict the global parameter estimate $\hat{\theta}$; the whole time-varying parameter series $\{\hat{\theta}_t\}$ is displayed in grey. The solid circles mark the independent estimators $\hat{\theta}_k$, whereas the open circles mark the thinned series $\hat{\theta}_{t_k}$. The marked estimators are grey whenever they lie within the local bands (depicted as solid lines) or global bands (depicted as dash–dot lines); otherwise, they are black. The observation horizon is given on the time axis, the parameter values on the y-axis. In particular, we see the following.

`NBPG2MON` versus `ELUB2MON`. The dependence structure between the natural gas and the United Kingdom base-load electricity contract is rather unstable—nevertheless, the binomial test on independent samples does not detect the time inhomogeneity, whereas the extremal test leads to the detection of three exceedances beyond the upper global confidence band. Let us summarize the special features of this pair which might help us to understand this behavior. The global dependence which stems from a parameter value of $\hat{\theta} = 0.83$ is rather high and peaks beyond a value of $\theta = 2$ in its extremes. Lastly, the so-called subprime mortgage crisis in September 2008 falls into the observation horizon, and precisely in this period the dependence exhibits a peak which exceeds the global bands.

`API42MON` versus `BRSWMO2`. The coal contract versus the oil contract depicted here is an example of a parameter series which fluctuates around the global parameter in a very quiet way. Adding the local and global bands confirms the hypothesis of parameter constancy since there are no exceedances at all.

The numerical results that have been described so far represent the first experiments and a preliminary data survey that we conducted to get a sense for the specific features of the data. Now we report on a comprehensive numerical analysis that was conducted in consequence.

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

TABLE 3: Summary statistics for Clayton copula estimation with five-year commodity data.

	$\hat{\theta}$	$\hat{\Sigma}$	$\hat{\Sigma}^{(BS)}$	$\hat{\tau} = \hat{\theta}/(\hat{\theta}+2)$	$\hat{\lambda}_{\mathrm{L}} = 2^{-1/\hat{\theta}}$
∅	0.414 03	0.003 36	0.070 40	0.171 51	0.187 47
min	0.009 98	0.000 67	0.004 05	0.004 97	0.000 00
max	7.309 08	0.101 48	2.711 64	0.785 16	0.909 52

We use five years of historical data from 19-12-2006 to 19-12-2011. After preliminary data preparation, we have chosen to apply more sophisticated univariate GARCH models for the log-return series. We proceed along the lines of [11] and use an exponential GARCH(1, 1) model due to [31]. We add an autoregressive part and arrive at

$$R_t^{(j)} = \mu_j + \rho_j R_{t-1}^{(j)} + Y_{jt}, \tag{20a}$$

$$Y_{jt} = \sigma_{jt}\varepsilon_{jt}, \tag{20b}$$

$$\log(\sigma_{jt}^2) = \omega_j + \alpha_j \varepsilon_{j,t-1} + \gamma_j(|\varepsilon_{j,t-1}| - \mathbb{E}(\varepsilon_{j,t-1})) + \beta_j \log(\sigma_{j,t-1}^2), \tag{20c}$$

with real parameters μ_j, ρ_j, ω_j, α_j, γ_j, and β_j, and univariate white-noise residuals $\varepsilon_{jt} \sim \mathrm{WN}(0,1)$ with respect to t for all $j = 1, \ldots, 11$. Furthermore, in accordance with the ideas presented in [10], univariate structural-break analyses are performed using the fluctuation tests due to [2]. Furthermore, the asymptotic properties of estimators of the break dates are derived. The empirical residuals of the AR(1)–eGARCH(1, 1) model given in (20) are subsequently the basis of a pairwise dependence analysis, i.e. the bivariate vectors of empirical residuals are assumed to be distributed according to a copula-based multivariate distribution. We condense the results we obtained for the parameter estimate of bivariate Clayton copulas in Table 3.

In Table 3 we only give the average (∅), the minimal (min), and the maximal (max) value of the global parameter estimate $\hat{\theta}$ based on the whole sample of about $T = 1300$ observations, as well as of the variance estimate $\hat{\Sigma}$ stemming from the canonical maximum-likelihood estimation and the bootstrapped variance estimate $\hat{\Sigma}^{(BS)}$. The bootstrapped version of the variance estimator was obtained nonparametrically as based on $B = 200$ estimations on a bootstrap sample of the same length as the bandwidth of the moving-window estimates, i.e. $b = T/10$. Furthermore, we give the respective range of the estimator of Kendall's tau and the lower tail dependence. The estimation of $\hat{\theta}$ does actually work for 47 of the 55 possible pairs, so we can conduct the binomial and extremal tests for these commodity pairs. This results in 25 local rejections (53%), i.e. rejections of the null hypothesis $\mathbb{H}_0$ as in (8) based on the binomial test, and 27 global rejections (60%), i.e. rejections of $\mathbb{H}_0'$ as in (14) based on the extremal test for the maximum of the parameter series. After re-estimation for the potential breakpoints indicated by the local exceedances, the ratio of rejections shrinks to 22 (42%) in the binomial and 21 (44%) in the extremal test, respectively.

In Figure 2 we present a graphical display of an example test result on the five-year historical data similar to Figure 1. In Figure 2(a) we depict the estimator series $\hat{\theta}_t$ together with the local confidence regions due to the binomial test and the global confidence bands due to the extremal test. The empirical residuals that form the basis for the copula estimation stem from a SCOMDY model over the whole observation horizon $t = 1, \ldots, T = 1296$.

The binomial test leads to the rejection of the null hypothesis of parameter constancy as there is one exceedance beyond the local confidence band at time $\tau = 1037$, which corresponds to the date 22-12-2010. With this knowledge, we re-estimate the univariate GARCH models with two different regimes: first for $t = 1, \ldots, \tau$ and then for $t = \tau + 1, \ldots, T$. Then the

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

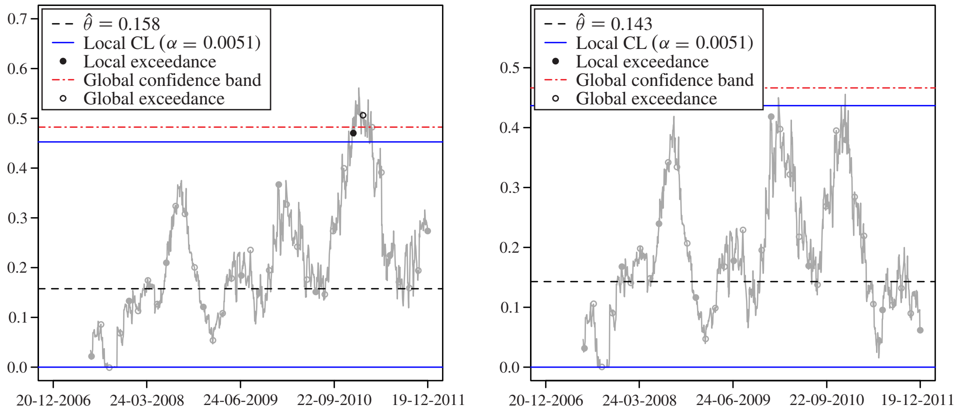

(a) Before change point handling (b) After re-estimation of univariate models

FIGURE 2: Binomial test with proper size for `API22MON` versus `ELGP2MON`.

empirical residuals are concatenated and again used as pseudo-realizations of the residual copula-based distribution. Canonical maximum-likelihood estimation of the global copula parameter $\hat{\theta}$ and the estimator series $\hat{\theta}_t$, as well as bootstrapping a variance estimate $\hat{\Sigma}^{(BS)}$ based on this concatenated series, then results in a binomial test that does not reject the null hypothesis, as Figure 2(b) shows.

This approach is advantageous compared to another attempt to remedy time variation in the copula parameter: the respective regimes are detected separately in the univariate models in order to re-estimate the GARCH parameters on different regimes.

We conclude this section with an additional discussion of the possible economic interpretation of our results.

4.2.3. *Discussion.* Concerning the empirical data, let us note the following: in shorter time series we observed fewer exceedances in sum, but also some extreme exceedances that finally gave the motivation for the extremal test. This happened especially with pairs which have the same underlying commodity. In the five-year time series data, the binomial test gave more overall rejections, but we had fewer rejections when applying the extremal test.

In our subsequent analyses, we found that there are cases when exceedances hint at some regime changes in the univariate models. Taking times of local (or global) exceedances as change points in the GARCH parameters, re-estimating the univariate models, and using the concatenated empirical residuals as basis for a repeated canonical maximum-likelihood estimation of the global and local copula parameters, we found that the dependency represented by the copula parameter series exhibited a time-homogeneous structure at least in some of the investigated cases. However, there were also cases where regime-changed GARCH models did not have any effect on the heterogeneity in the copula parameter. With a merely graphical valuation of the results, the estimator series seemed to have either different regimes themselves, or exhibited an autoregressive structure.

From the economic point of view, our observations and test results are in line with similar econometric findings. For example, Gaisser *et al.* [19] measured a portfolio's dependence structure through Spearman's ρ and found that significant deviations from the hypothesis of constancy come along with unusual economic events (their procedure actually detects the volatility peaks that the financial markets exhibited in 2002 and the beginning of 2004).

https://doi.org/10.1017/apr.2016.48 Published online by Cambridge University Press

Also, Dias and Embrechts [13] modeled the conditional dependence structure using copulas and so were able to detect changes in the dependence structure of bivariate time series, such as changes in the tail of the joint distribution. A time-varying copula approach has already been used in [14], which detected changes in the dependence structure of exchange rates.

A very recent example regarding the assessment of time-dynamic dependencies with the help of copulas is [34]. There a pair copula construction, also called vine copula, was used to model the empirical residuals of high-dimensional financial data. By applying Bayesian estimation techniques to determine Markov-switching parameters, the authors found that different regimes in the dependency parameters are often implied by exceptional economic events (such as the 2008 subprime crisis).

Moreover, Almeida *et al.* [1] investigated the dependence between the DAX 30 constituents with vine copulas and incorporated time dynamics by the stochastic autoregression due to [23]. They found that, according to the BIC model selection criterion, nondynamic models outperform the time-varying models in most of the levels of the vine construction.

Acknowledgements

The authors thank an anonymous referee for many thoughtful suggestions which improved the paper considerably.

Access to data was provided by the LBBW Trading Room operated by the Institute of Financial Mathematics at Ulm University.

References

[1] Almeida, C., Czado, C. and Manner, H. (2012). Modeling high dimensional time-varying dependence using D-vine SCAR models. Preprint. Available at http://arxiv.org/pdf/1202.2008v1.

[2] Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica* **66,** 47–78.

[3] Bentkus, V. (2005). A Lyapunov-type bound in $\mathbb{R}^d$. *Theory Prob. Appl.* **49,** 311–323.

[4] Bücher, A., Kojadinovic, I., Rohmer, T. and Segers, J. (2014). Detecting changes in cross-sectional dependence in multivariate time series. *J. Multivariate Anal.* **132,** 111–128.

[5] Burger, M., Graeber, B. and Schindlmayer, G. (2014). *Managing Energy Risk: An Integrated View on Power and Other Energy Markets* (Wiley Finance Ser. **425**), 2nd edn. John Wiley, Chichester.

[6] Chan, N.-H., Chen, J., Chen, X., Fan, Y. and Peng, L. (2009). Statistical inference for multivariate residual copula of GARCH models. *Statistica Sinica* **19,** 53–70.

[7] Chen, X. and Fan, Y. (2006). Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification. *J. Econometrics* **135,** 125–154.

[8] Cherubini, U., Gobbi, F., Mulinacci, S. and Romagnoli, S. (2010). A copula-based model for spatial and temporal dependence of equity markets. In *Copula Theory and Its Applications* (Warsaw 2009; Lecture Notes Statist. Proc. **198**), Springer, Heidelberg, pp. 257–265.

[9] Cherubini, U., Luciano, E. and Vecchiato, W. (2004). *Copula Methods in Finance*. John Wiley, Chichester.

[10] Chevallier, J. (2011). Detecting instability in the volatility of carbon prices. *Energy Economics* **33,** 99–110.

[11] Chevallier, J. (2012). *Econometric Analysis of Carbon Markets: The European Union Emissions Trading Scheme and the Clean Development Mechanism*. Springer, Dordrecht.

[12] Demko, S., Moss, W. F. and Smith, P. W. (1984). Decay rates for inverses of band matrices. *Math. Comput.* **43,** 491–499.

[13] Dias, A. and Embrechts, P. (2009). Testing for structural changes in exchange rates' dependence beyond linear correlation. *Europ. J. Finance* **15,** 619–637.

[14] Dias, A. and Embrechts, P. (2010). Modeling exchange rate dependence dynamics at different time horizons. *J. Internat. Money Finance* **29,** 1687–1705.

[15] Dobrić, J. and Schmid, F. (2007). A goodness of fit test for copulas based on Rosenblatt's transformation. *Comput. Statist. Data Anal.* **51,** 4633–4642.

[16] Embrechts, P. (2009). Copulas: A personal view. *J. Risk Insurance* **76,** 639–650.

[17] Fernández, B. and Muriel, N. (2009). Regular variation and related results for the multivariate GARCH(p, q) model with constant conditional correlations. *J. Multivariate Anal.* **100,** 1538–1550.

[18] Fisher, R. A. (1970). *Statistical Methods for Research Workers*, 14th edn. Oliver and Boyd, Edinburgh.
[19] Gaisser, S., Memmel, C., Schmidt, R. and Wehn, C. (2009). Time dynamic and hierarchical dependence modelling of an aggregated portfolio of trading books: a multivariate nonparametric approach. Discussion Paper Series **2**: Banking and Financial Studies 2009,07. Deutsche Bundesbank, Frankfurt am Main.
[20] Genest, C., Ghoudi, K. and Rivest, L.-P. (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika* **82,** 543–552.
[21] Genest, C., Rémillard, B. and Beaudoin, D. (2009). Goodness-of-fit tests for copulas: A review and a power study. *Insurance Math. Econom.* **44,** 199–213.
[22] Giacomini, E., Härdle, W. and Spokoiny, V. (2009). Inhomogeneous dependence modelling with time-varying copulae. *J. Business Econom. Statist.* **27,** 224–234.
[23] Hafner, C. M. and Manner, H. (2012). Dynamic stochastic copula models: estimation, inference and applications. *J. Appl. Econometrics* **27,** 269–295.
[24] Hansen, P. R. and Lunde, A. (2005). A comparison of volatility models: does anything beat a GARCH(1, 1)? *J. Appl. Econometrics* **20,** 873–889.
[25] Härdle, W. K., Okhrin, O. and Okhrin, Y. (2008). Modeling dependencies with copulae. In *Applied Quantitative Finance*, 2nd edn., eds W. K. Härdle, N. Hautsch and L. Overbeck, Springer, Berlin, pp. 3–36.
[26] Lehmann, E. L. and Romano, J. P. (2005). *Testing Statistical Hypotheses*, 3rd edn. Springer, New York.
[27] Mai, J.-F. and Scherer, M. (2012). *Simulating Copulas: Stochastic Models, Sampling Algorithms, and Applications* (Ser. Quant. Finance **4**). Imperial College Press, London.
[28] Mai, J.-F. and Scherer, M. (2014). *Financial Engineering with Copulas Explained*. Palgrave Macmillan, Basingstoke.
[29] Mikosch, T. and Starica, C. (2000). Limit theory for the sample autocorrelations and extremes of a GARCH(1, 1) process. *Ann. Statist.* **28,** 1427–1451.
[30] Mroz, M. (2012). Time-varying copula models for financial time series. Doctoral Thesis, Ulm University.
[31] Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* **59,** 347–370.
[32] Patton, A. J. (2012). A review of copula models for economic time series. *J. Multivariate Anal.* **110,** 4–18.
[33] Simes, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* **73,** 751–754.
[34] Stöber, J. and Czado, C. (2012). Detecting regime switches in the dependence structure of high dimensional financial data. Preprint. Available at http://arxiv.org/abs/1202.2009v1.
[35] Teriokhin, A. T., de Meeûs, T. and Guégan, J.-F. (2007). On the power of some binomial modifications of the Bonferroni multiple test. *J. General Biology* **68,** 332–340.
[36] Yaskov, P. (2010). Testing for predictive ability in the presence of structural breaks. *Quantile* **8,** 127–135 (in Russian).

RÜDIGER KIESEL, *University of Duisburg-Essen*

Chair for Energy Trading and Finance, University of Duisburg-Essen, Campus Essen, Universitätsstraße 12, 45141 Essen, Germany. Email address: ruediger.kiesel@uni-due.de

MAGDA MROZ, *Ulm University*

Institut für Zahlentheorie und Wahrscheinlichkeitstheorie, Ulm University, 89069 Ulm, Germany.
Email address: magda.mroz@web.de

ULRICH STADTMÜLLER, *Ulm University*

Institut für Zahlentheorie und Wahrscheinlichkeitstheorie, Ulm University, 89069 Ulm, Germany.
Email address: ulrich.stadtmueller@uni-ulm.de