A copula goodness-of-fit approach based on the conditional probability integral transformation $**$

Daniel Berg

The Norwegian Computing Center and University of Oslo.
E-mail: daniel@danielberg.no

Henrik Bakken

The Norwegian Computing Center and The Norwegian University of Science and Technology.
E-mail: Bakken.Henrik@bcg.com

[Revised April 17th 2007]

Abstract. Goodness-of-fit testing for copulae recently emerged as a challenging inferential problem and some approaches have been proposed. We investigate such an approach based on the conditional probability integral transformation. This approach implicitly weights observations at corners and edges of the unit hypercube which makes it very powerful at detecting tail heaviness for large sample sizes. However, it is shown to perform rather poor for small sample sizes. We propose a generalization that allows for any weighting, making it more robust and more powerful for small sample sizes. Another weakness is that some deviations from the null hypothesis may be neglected. We show an example and propose an extension. The original approach is shown to be a special case of our generalized and extended approach. Results from extensive Monte Carlo experiments show that the our approach keeps prescribed levels well and that certain weighting schemes produce superior power for three alternative hypotheses and for various combinations of problem dimension and sample size. The margins are treated as unknown nuisance parameters and are replaced by their empirical distribution functions. In addition, since we are testing a parametric null hypothesis requiring parameter estimation, a parametric bootstrap procedure is required to obtain reliable p-value estimates. Applied to daily log-returns of collections of large cap stock portfolios the Gaussian- and one-parameter Clayton- and Gumbel copulae are all strongly rejected, increasingly so for increasing dimension and sample size. The Student's t copula on the other hand, provides a good fit, indicating the presence of tail dependence in the daily log-returns of stock data.

Keywords: Copula, goodness-of-fit, conditional probability integral transformation, order statistic, parametric bootstrap

1. Introduction

Copulae have proved to be a very useful tool in the analysis of dependency structures. The concept of copulae was introduced by Sklar (1959), but was first used for financial applications by Embrechts et al. (1999). Since then we have seen a tremendous increase of copula related research and applications. The limitation of the copula approach is the lack of a recommended way of checking whether the dependency structure of a data set is appropriately modeled by a chosen family of copulae. Information criterions, such as Aikaike's Information Criterion (AIC), are commonly employed for model selection. Such pure model selection criterions do not provide us with any understanding of the size of the decision rule employed, nor its power. This means that one can not say how well the selected family of copulae fits the data. Neither can we say whether one family of copulae fits the data significantly better than another. A goodness-of-fit (gof) approach on the other hand, will provide this information.

Copula gof testing recently emerged as a challenging inferential problem and some approaches have been proposed. Genest et al. (1995) assess the fit of bivariate Archimedean copulae. Shih (1998), Glidden (1999) and Cui and Sun (2004) test the Clayton model (also referred to as the gamma frailty model in survival analysis). Breymann et al. (2003), Chen et al. (2004) and Dobrić and Schmid (2007) apply the conditional probability integral transformation (cpit) and tests for independence. Malevergne and Sornette (2003) compare the empirical distribution of the data with a χ^2-distribution using a bootstrap method, testing the Gaussian copula hypothesis for financial asset dependencies. Fermanian (2005) approximates the underlying probability density function by kernel smoothing of the empirical density.

$**$This research was supported by the Norwegian Research Council, grant number 154079/420.

Scaillet (2005) propose a test statistic based on the integrated squared difference between kernel estimators of the copula density and the parametric copula density. Dobrić and Schmid (2005) propose a chi-squared- and a likelihood ratio test, both based on partitioning the probability space. Panchenko (2005) focuses on positive definite bilinear forms, while Genest and Rémillard (2005) compare the empirical copula function to the parametric, null hypothesis copula function. Savu and Trede (2004) and Genest et al. (2006) assess the cumulative distribution function (cdf) of the copula function. Finally, Genest et al. (2007) propose an approach based on the cpit and the copula function.

We will consider in more detail the approach proposed by Breymann et al. (2003), based on the cpit and henceforth denoted the cpit-approach. The cpit, also known as Rosenblatt's transformation, transforms a set of dependent variables into a set of independent variables, given the multivariate distribution. Breymann et al. (2003) perform a cpit under a parametric null hypothesis copula. Then they employ a dimension reduction technique to the d-variate cpit copula and compute a univariate test statistic on the resulting univariate vector. Their dimension reduction strongly weights data along the boundaries of the cpit copula, i.e. corners and edges of the d-dimensional unit hypercube. This makes it less robust for small sample sizes. We generalize the cpit-approach to allow for any weight function in the dimension reduction. In addition, the dimension reduction is not consistent in the sense that some deviations from the null hypothesis may be neglected. We show an example and propose an extension using an additional cpit, based on order statistics. Our generalized and extended approach is henceforth denoted the cpit2-approach.

The paper is organized as follows. Section 2 presents preliminaries. In Section 3 we introduce copula gof testing, the cpit-approach and our cpit2-approach. In Section 4 we present the results from an extensive Monte Carlo study under several weighting schemes. The study visualizes the size and power of the cpit2-approach in distinguishing the Gaussian copula from the Student's t-, Clayton- and Gumbel copulae under various dimensions and sample sizes. In Section 5 we apply the best performing weighting scheme to analyze the dependence structure of the daily log-returns of some large cap stock portfolios. Finally, Section 6 summarizes our results and concludes.

2. Preliminaries

2.1. Empirical Marginals

Suppose we have n samples of the d-variate vector $\mathbf{X} = (X_1, \ldots, X_d)$. This vector comes from a population with unknown margins and linking copula C. We wish to test the hypotheses that the linking copula belongs to some parametric copula family C_θ:

$$\mathcal{H}_0 : C \in \{C_\theta; \theta \in \Theta\} \quad \text{vs.} \quad \mathcal{H}_a : C \notin \{C_\theta; \theta \in \Theta\}.$$

To extract the copula we transform the vector $\mathbf{X}$ into a pseudo-vector $\mathbf{Z}$, through the empirical marginal distribution functions, $\mathbf{Z}_j = (Z_{j1}, \ldots, Z_{jd}) = (\widehat{F}_1(X_{j1}), \ldots, \widehat{F}_d(X_{jd})),\ j = 1, \ldots, n$, where

$$\widehat{F}_i(x) = \frac{1}{n+1} \sum_{j=1}^{n} I\{X_{ji} \le x\}. \tag{1}$$

Equivalently, the pseudo-vector can be expressed in terms of normalized ranks,

$$\mathbf{Z}_j = (Z_{j1}, \ldots, Z_{jd}) = \left(\frac{R_{j1}}{n+1}, \ldots, \frac{R_{jd}}{n+1}\right),\ j = 1, \ldots, n. \tag{2}$$

Here R_{ji} is the rank of X_{ji} in $X_{1i}, \ldots, X_{ni}$.

2.2. Anderson-Darling test statistic

Suppose we have a random vector $\mathbf{W} = (w_1, \ldots, w_n)$ which is iid $U(0,1)^n$ and that the cdf of $\mathbf{W}$ is $F(w) = w$. The AD statistic is then defined as

$$\mathcal{T} = n \int \frac{\left\{\hat{F}(w) - w\right\}^2}{w(1-w)}\, \mathrm{d}w, \qquad w \in [0,1]. \tag{3}$$

The AD statistic strongly weights deviations near $w = 0$ and $w = 1$. This is justified since the experimental deviations are small here due to the constraints $\{\hat{F}(w) - F(w)\} = 0$ at $w = (0, 1)$ (Aslan and Zech, 2002).

The empirical version of the AD statistic for uniform variables can be shown to be (Marsaglia and Marsaglia, 2004):

$$\widehat{T} = -n - \frac{1}{n}\sum_{j=1}^{n}(2j-1)\left\{\ln\left[\widehat{F}\left(\frac{j}{n+1}\right)\right] + \ln\left[1 - \widehat{F}\left(\frac{n+1-j}{n+1}\right)\right]\right\}, \tag{4}$$

where the empirical cdf, $\widehat{F}$, is given by (1).

2.3. The conditional probability integral transform

There are several probability integral transformations, see e.g. D'Agostino and Stephens (1986) for a discussion. We will consider the transformation proposed by Rosenblatt (1952). This transformation was denoted the conditional probability integral transform (cpit) by D'Agostino and Stephens (1986) and it transforms a set of dependent variables into a new set of independent $U(0, 1)$ variables, given their multivariate distribution. The cpit is a universally applicable way of creating a set of iid $U(0, 1)$ variables from any data set with known distribution. Given a test for multivariate, independent uniformity, this transformation can be used to test the fit of any assumed model.

Definition 1 (CONDITIONAL PROBABILITY INTEGRAL TRANSFORM).
Let $\mathbf{Z} = (Z_1, \ldots, Z_d)$ *denote a random vector with marginal distributions* $F_i(z_i) = P(Z_i \le z_i)$ *and conditional distributions* $F_{i|1\ldots i-1}(Z_i \le z_i|Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})$ *for* $i = 1, \ldots, d$*. The cpit of* $\mathbf{Z}$ *is defined as* $T(\mathbf{Z}) = (T_1(Z_1), \ldots, T_d(Z_d))$ *where*

$$\begin{aligned}
T_1(Z_1) &= P(Z_1 \le z_1) = F_1(z_1), \\
T_2(Z_2) &= P(Z_2 \le z_2|X_1 = z_1) = F_{2|1}(z_2|z_1), \\
&\vdots \\
T_d(Z_d) &= P(Z_d \le x_d|Z_1 = z_1, \ldots, Z_{d-1} = z_{d-1}) = F_{d|1\ldots d-1}(z_d|z_1, \ldots, z_{d-1}).
\end{aligned}$$

The random variables $V_i = T_i(Z_i),\ i = 1, \ldots, d$ *are uniformly and independently distributed on* $[0, 1]^d$.

A recent application of the cpit is to multivariate gof tests. A cpit is applied to a data set, assuming a multivariate null distribution, and then a test of multivariate independence is carried out on the resulting, transformed data set. The null hypothesis in our setting is a parametric copula family. The parameters of this copula family needs to be estimated before applying the cpit. We shortly present parameter estimation in Section 2.4.

An advantage with the cpit in a gof setting is that the null- and alternative hypotheses are the same, regardless of the distribution before the cpit. The cpit also enables weighting in a simple way since the data, after the cpit, is i.i.d. $U(0, 1)$ under the null hypothesis. Hong and Li (2002) report Monte Carlo evidence of multivariate tests using cpit variables outperforming tests using the original random variables. Chen et al. (2004) believe that a similar conclusion also applies to gof tests for copulae.

A disadvantage with the cpit is the invariance with respect to the permutation of the variables since there are $d!$ possible permutations. However, as long as the permutation is decided randomly, the results will not be influenced in any particular direction. D'Agostino and Stephens (1986) discuss this issue and propose solutions for some special cases, e.g. the cpit based on ordered variables, which does not suffer from permutation invariance. We will consider this in more detail when presenting our cpit2-approach in Section 3.2.

2.4. Parameter estimation

There are two main ways of estimating the parameters of a copula, the fully parametric method or a semi-parametric method. The fully parametric method, termed the inference functions for margins (IFM) method (Joe, 1997), relies on the assumption of parametric, univariate margins. First, the parameters of the margins are estimated and then each parametric margin is plugged into the copula likelihood which is then maximized with respect to the copula parameters. Since we treat the margins as nuisance parameters, we rather proceed with the pseudo-vector $\mathbf{Z}$ and the semi-parametric method. This method is denoted the pseudo-likelihood (Demarta and McNeil, 2005) or the canonical maximum likelihood (CML) (Romano, 2002) method and is described in Genest et al. (1995) and in Shih and Louis (1995) in the presence of censorship. Having obtained the pseudo-vector $\mathbf{Z}$, using (2), the copula parameters can be estimated using either maximum likelihood (ML) or using the well-known relations to Kendall's tau (for a survey of copulae and their relationship with measures of association, see Nelsen (1999)).

For the elliptical copulae in higher dimensions, we pairwise invert the sample Kendall's tau. This gives the correlation- and scale matrix for the Gaussian and Student's t copulae, respectively. For the Student's t copula we also need to estimate the degrees of freedom. Genest et al. (2007) estimate the scale matrix by inversion of Kendall's tau but assume the degrees of freedom to be known/fixed. We rather follow the approach used by Mashal and Zeevi (2002) and Demarta and McNeil (2005). This is a two-stage approach in which the scale matrix is first estimated by inversion of Kendall's tau, and then the pseudo-likelihood function is maximized with respect to the degrees of freedom, given the estimate of the scale matrix. For the Archimedean copulae, we consider the so-called exchangeable construction with one dependency parameter. We estimate this parameter by numerically maximizing the pseudo-likelihood.

3. Copula goodness-of-fit testing

For univariate distributions, the gof assessment can be performed by e.g. the well-known Anderson-Darling (Anderson and Darling, 1954) test, or less quantitatively using a QQ-plot. In the multivariate domain there are fewer alternatives. Economic theory sheds little light on the dependence structure between financial assets, and multivariate normality is often assumed a priori. Evidence shows, however, that more appropriate dependence structures are available (Chen et al., 2004; Dobrić and Schmid, 2005).

Several approaches (e.g. Breymann et al. (2003); Genest et al. (2006)) project the multivariate problem to a univariate problem applying some dimension reduction technique and then compute a univariate test statistic. This leads to numerically efficient algorithms even for problems of high dimension. Any univariate statistic may be used, e.g. Kolmogorov-Smirnov, Anderson-Darling, Cramér-von Mises or kernel smoothing based L2 statistics. For a thorough treatment of these and other statistics we refer to D'Agostino and Stephens (1986). In this paper we focus on the Anderson-Darling (AD) statistic.

For copula gof testing we are interested in the fit of the copula alone, hence the margins are commonly treated as nuisance parameters. I.e. we use empirical margins (or equivalently, normalized ranks). The use of empirical margins will alter the asymptotics of any test statistic. In addition, since we are testing a hypothesized, parametric, copula, parameter estimation error will influence the asymptotics. Breymann et al. (2003) fail to recognize these issues. They assume that the limiting distribution of their statistic is the same whether the margins and parameters are estimated or not. As a result, the p-values that they report are not correct. This erroneous assumption is pointed out by Genest and Rémillard (2005). It is also thoroughly investigated by Dobrić and Schmid (2007) who modify the test procedure by Breymann et al. (2003) such that the p-value estimates become reliable. Henceforth, when referring to the approach by Breymann et al. (2003), the cpit-approach, we mean the approach proposed by Breymann et al. (2003) but using the test procedure of Dobrić and Schmid (2007).

3.1. The cpit-approach

The approach proposed by Breymann et al. (2003) coincides with the approach proposed by Malevergne and Sornette (2003) when the latter is based on cpit data. It also coincides with the second approach in Chen et al. (2004). It is a dimension reduction approach and we will denote the test observator by G.

The testing is based on the pseudo-vector $\mathbf{Z}$, see (2). A cpit is applied to $\mathbf{Z}$, assuming a null hypothesis copula C_θ. The d-variate vector $\mathbf{V} = (V_1, \ldots, V_d)$, resulting from the cpit, is i.i.d. $U(0,1)^d$ under the null hypothesis. Due to parameter- and margin estimation errors, this is only close to, but not exactly true. We will consider this issue in Section 3.3. Until then we assume that this holds. The dimension reduction is now performed as

$$W_G = \sum_{i=1}^{d} \Phi^{-1}(V_i)^2. \tag{5}$$

The variable W_G should, under the null hypothesis, be χ^2_d distributed. The test observator G can now be defined.

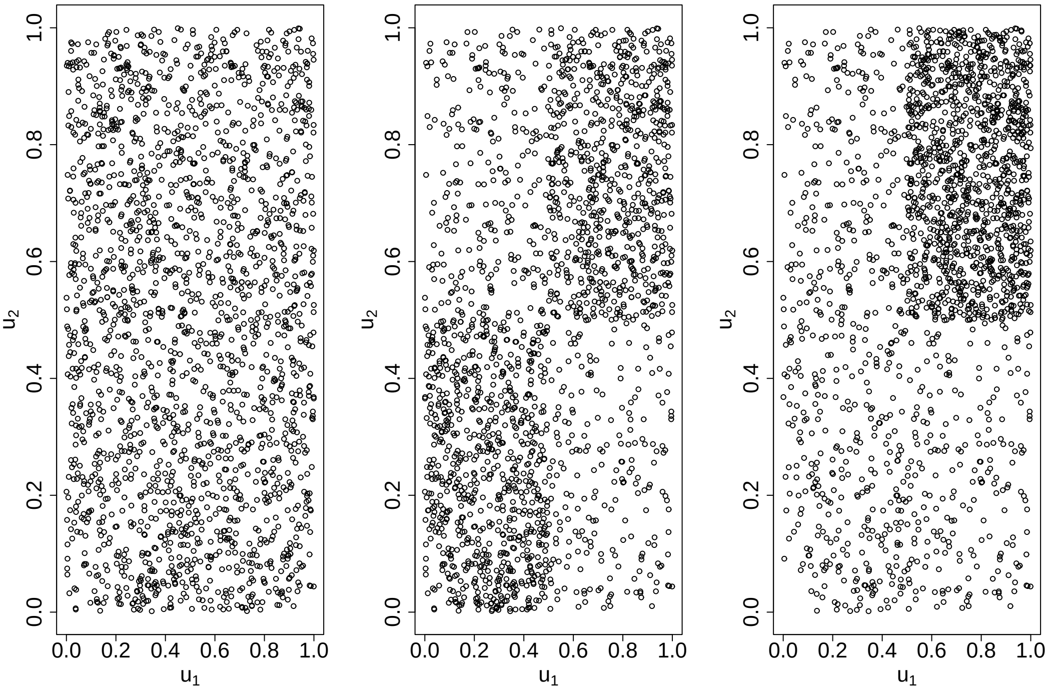

Figure 1. Three $U(0,1)^2$ data sets, one that is independent (left panel) and two that are clearly dependent (center- and right panels). $\widehat{G}(w)$ is equivalent for all three.

Definition 2 (Cpit test observator G).
Let W_G be defined by (5) *and $F_{\chi^2_d}(\cdot)$ be the χ^2_d cdf. $G(w)$ is then defined as the cdf of $F_{\chi^2_d}(W_G)$:*

$$G(w) = P[F_{\chi^2_d}(W_G) \leq w]. \tag{6}$$

Under $\mathcal{H}_0$, all V_i are i.i.d. $U(0,1)$, hence $G(w) = w$ and the density of $G(w)$ is $g(w) = 1$.

Suppose we have n samples of $\mathbf{V}$, $\mathbf{v}_j = (v_{j1}, \dots, v_{jd})$, $j = 1, \dots, n$. After performing the dimension reduction in (5), we have n samples of W_G. The empirical version of the approach then becomes

$$\widehat{G}(w) = \frac{1}{n+1} \sum_{j=1}^{n} I\{F_{\chi^2_d}(W_{G,j}) \leq w\}, \qquad w = \frac{1}{n+1}, \dots, \frac{n}{n+1}. \tag{7}$$

In the cpit-approach $\widehat{G}(w)$ is plugged in for $\widehat{F}(w)$ in the expression for the AD statistic (4).

The cpit-approach is computationally very efficient and conceptually simple. However, it has its weaknesses. First of all, the dimension reduction, through the use of $\Phi^{-1}(\cdot)^2$, strongly weights data along the boundaries of the d-dimensional unit hypercube. This may be appropriate when the sample size is large. However, for small sample sizes, this weighting makes the approach less robust and less powerful since there will be few observations in the boundary regions. We will see the effects of this in Section 4. In addition, some deviations from the null hypothesis may be overlooked by the cpit-approach. Figure 1 shows three constructed bivariate data sets, one that is independent in the left panel and two that are clearly dependent in the center- and right panels. Recall the null hypothesis of independence. We thus wish for the lack of independence in these panels to be detected. However, $G(w)$ will be exactly the same for all three data sets. The explanation is that a value of 0.2 and a value of 0.8 will both contribute with the exact same value to W_G, since $\Phi^{-1}(0.2)^2 = \Phi^{-1}(0.8)^2$. Hence, we suspect the approach to perform poor in cases where the cpit data set is radially asymmetric.

3.2. *The cpit2-approach: a generalization and extension*

With the weaknesses of the cpit-approach in mind, we propose a new approach, denoted the cpit2-approach. This approach generalizes and extends the cpit-approach. First, any weight function can be employed in the dimension reduction (5). In addition, through the use of an additional cpit, based on order statistics, we are able to detect radial asymmetry in the cpit data as illustrated in Figure 1.

We interpret the problem of multivariate gof testing as follows. To perform a gof assessment of a multivariate data set we can essentially perform two univariate gof tests. First, we test the fit in the d-space, e.g. through some dimension reduction technique such as (5). The result of this first, univariate, gof test is n values of a test statistic. If we know the distribution of these n samples, under the null hypothesis, we can perform another, univariate, gof test in the n-space. This will give us the desired test statistic for the multivariate problem. In what follows we are mainly concerned with the first gof test, in the d-space. For the second, in the n-space, we use the Anderson-Darling statistic, as in the cpit-approach.

We first perform the cpit on the original copula data set $\mathbf{Z}$. The resulting data $\mathbf{V}$ should be i.i.d. $U(0,1)^d$ under the null hypothesis. We now propose to test whether this is true, i.e. an additional test in the d-space, testing for independent uniformity of $\mathbf{V}$. This problem is well known and is discussed in great detail in D'Agostino and Stephens (1986). They suggest a cpit, based on ordered variables, that will be permutation invariant. Thus, we perform a regular cpit first, on $\mathbf{Z}$, and then a second cpit on $\mathbf{V}$ which is based on the order statistics of $\mathbf{V}$.

As before, let $\mathbf{V} = (V_1, \ldots, V_d)$ be the i.i.d. $U(0,1)^d$ random vector, obtained from applying the cpit to $\mathbf{Z}$. For $d = 1, 2, \ldots,$ we denote the order statistics of $V_1, \ldots, V_d$ by

$$V_{(1)} \le V_{(2)} \le \cdots \le V_{(d-1)} \le V_{(d)}. \tag{8}$$

If $V_{(1)}, \ldots, V_{(d)}$ are the order statistics of a sample from a $U(0,1)$ parent distribution, then $V_{(i)}$ is a beta distributed variable with parameters $(i, d-(i-1))$ (D'Agostino and Stephens, 1986, ch. 8). To compute the expressions for the order statistic cpit, we resort to David (1981, Theorem 2.7) who shows the Markov nature of the order statistics. Using Deheuvels (1984, Theorem 1) and the fact that $\mathbf{V}$ is an i.i.d. $U(0,1)^d$ random vector under the null hypothesis, we obtain the following expression for the order statistic cpit of V:

$$H_i = F_{V_{(i)}|V_{(i-1)}}(v_{(i)}) = 1 - \left(\frac{1 - v_{(i)}}{1 - v_{(i-1)}}\right)^{d-(i-1)}, \; i = 1, \ldots, d, \; v_{(0)} = 0. \tag{9}$$

Intuitively, poor fit in the d-space is indicated by extreme values of H. Any H too low or too high can indicate a poor fit (Glen et al., 2001). We can now conduct the dimension reduction based on $\mathbf{V}$ and $\mathbf{H}$:

$$W_B = \sum_{i=1}^{d} \Gamma_V(V_{(i)}; \boldsymbol{\alpha}) \cdot \Gamma_H(H_i; \boldsymbol{\alpha}), \tag{10}$$

where Γ_V and Γ_H are weight functions used for weighting the information in V and H, respectively, and $\boldsymbol{\alpha}$ is the set of weight parameters. Any weight function may be used, depending on the use and the region of the copula one wishes to emphasize. Some obvious candidates for both $\Gamma_V(X; \boldsymbol{\alpha})$ and $\Gamma_H(X; \boldsymbol{\alpha})$ are:

(i) $\Phi^{-1}(X)^2$,

(ii) $|X - 0.5|$,

(iii) $(X - 0.5)^{\alpha}, \; \alpha = (2, 4, \ldots)$.

Consider for example the special case $\Gamma_V(X; \boldsymbol{\alpha}) = \Phi^{-1}(X)^2$ and $\Gamma_H(X; \boldsymbol{\alpha}) = 1$. We then obtain (5), the cpit-approach. Since both $\mathbf{V}$ and $\mathbf{H}$ are i.i.d. $U(0,1)^d$ under the null hypothesis we have the following result. By choosing $\Gamma_V(X; \boldsymbol{\alpha}) = 1$ and $\Gamma_H(X; \boldsymbol{\alpha}) = \Phi^{-1}(X)^2$, W_B in (10), as for (5), should follow a χ^2_d distribution under the null hypothesis. However, in general, the distribution of W_B is not known and we must turn to a double bootstrap to approximate the cdf. Suppose we have computed W_B, using some weight functions $\Gamma_V(\cdot; \boldsymbol{\alpha})$ and $\Gamma_H(\cdot; \boldsymbol{\alpha})$. Now we simply draw d i.i.d. $U(0,1)$ variables $\widetilde{V}$, compute $\widetilde{H}$ and $\widetilde{W}_B$ using the same weight functions as for W_B. By repeating this a large number of times (10000 times in this paper), we can approximate the cdf of W_B, F_B, under the null hypothesis. Again, as for W_G, W_B is only close to, but not exactly distributed according to F_B. This discussion is deferred to Section 3.3. Our new test observator B can now be defined.

Definition 3 (Cpit2 test observator B).
Let W_B be defined by (10) *and $F_B(\cdot)$ be the cdf of W_B. $B(w)$ is then defined as the cdf of $F_B(W_B)$:*

$$B(w) = P[F_B(W_B) \le w]. \tag{11}$$

Under $\mathcal{H}_0$, all V_i are i.i.d. $U(0,1)$, hence $B(w) = w$ and the density of $B(w)$, $b(w) = 1$.

Suppose we have n samples of $\mathbf{V}$, $\mathbf{v}_j = (v_{j1}, \dots, v_{jd})$, $j = 1, \dots, n$. The empirical version then becomes

$$\widehat{B}(w) = \frac{1}{n+1} \sum_{j=1}^{n} I\{F_B(W_{B,j}) \leq w\}, \qquad w = \frac{1}{n+1}, \dots, \frac{n}{n+1}, \tag{12}$$

which can be plugged in for $\widehat{F}(w)$ in the expression for the AD statistic (4).

To summarize our cpit2-approach, we have performed two cpit's, the first to $\mathbf{Z}$ and the second to the order statistics of $\mathbf{V}$. By doing this additional order statistic cpit, our dimension reduction approach becomes more robust to phenomena like the one in Figure 1. For the data sets in Figure 1 we obtain $G(w) = (0.91, 0.91, 0.91)$ for the left-, center- and right panels, respectively, using the AD statistic. With the cpit2-approach, using $\Gamma_V(X; \boldsymbol{\alpha}) = 1$ and $\Gamma_H(X; \boldsymbol{\alpha}) = \Phi^{-1}(X)^2$, we obtain $B(w) = (0.36, 66.46, 111.62)$. We clearly see that our extension H detects the asymmetry. Figure 2 shows W_B as a surface, with respect to the cpit data set $\mathbf{V}$ and we see that the cpit-approach heavily emphasizes the boundaries. The use of the cpit2-approach, with weight combination (ii) still emphasizes these regions but less extremely. We also see that the Γ_H term adds weight to the diagonal as well. This is why this extension will help detect radial asymmetry in the cpit data. Finally, the generalization adds flexibility and robustness to small sample sizes. Both weight functions for the dimension reduction, Γ_V and Γ_H, can be decided freely, depending on the specific use.

3.3. Testing Procedure

In Section 3.1 we assume that W_G in (5) follow a χ^2_d distribution. Similarly, in Section 3.2 we assume that the distribution of W_B in (10) can be approximated by a double bootstrap procedure. The estimation of the margins and the parameters of the null copula, introduces dependence in the cpit data. Hence, W_G is only close to, but not exactly χ^2_d distributed. Similarly, W_B is only close to, but not exactly distributed according to F_B.

To cope with this issue and obtain a proper estimate of the p-value of the observed statistic, one should perform a parametric bootstrap procedure, where both margin- and parameter estimation effects are accounted for. We adopt the parametric bootstrap procedure used in Genest et al. (2006), the validity of which is established in Genest and Rémillard (2005). Dobrić and Schmid (2007) propose a very similar procedure in their modification of the original procedure used in Breymann et al. (2003). The asymptotic validity of the bootstrap procedure, applied to our test observator, has not yet been proved. However, our numerical results in Section 4 strongly indicates that the procedure is valid.

Suppose we have a sample $\mathbf{x}$, n observations of the d-variate vector $\mathbf{X}$. The testing procedure for the cpit2-approach is then as given below. Remember that the cpit-approach is a special case, hence the same test procedure can be applied for this approach. Note the use of empirical margins (normalized ranks) in step (1) and (9a), the parametric bootstrap procedure in step (9) and the double bootstrap procedure in steps (6) and (9f).

(1) Extract the pseudo-observations $(\mathbf{z}_1, \dots, \mathbf{z}_n)$ by converting the sample data $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ into normalized ranks according to (2).

(2) Estimate the parameters Θ of the null hypothesis copula, by a consistent estimator $\widehat{\Theta} = \widehat{\mathcal{L}}(\mathbf{z}_1, \dots, \mathbf{z}_n)$.

(3) Compute the cpit sample data $(\mathbf{v}_1, \dots, \mathbf{v}_n)$ by applying the cpit to $(\mathbf{z}_1, \dots, \mathbf{z}_n)$ assuming the parametric null hypothesis copula $C_{\widehat{\Theta}}$.

(4) Compute the cpit2 sample data $(\mathbf{h}_1, \dots, \mathbf{h}_n)$ by applying the cpit (9) to $(\mathbf{v}_1, \dots, \mathbf{v}_n)$.

(5) Compute W_B according to (10), using weight functions Γ_V and Γ_H.

(6) If W_B follows a known distribution under the null hypothesis, compute $F_B(W_B)$ accordingly. If not, approximate F_B as follows. For some large integer m, repeat the following steps for every $l \in \{1, \dots, m\}$:

   (i) Generate a random sample $(v^*_{1,l}, \dots, v^*_{d,l})$ from the null hypothesis copula, namely an i.i.d. $U(0,1)^d$ vector.

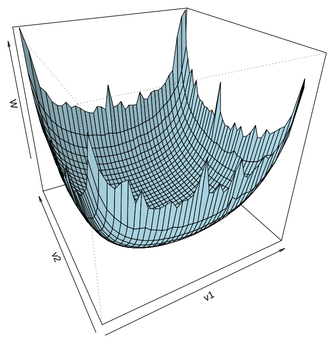

(a) $\Gamma_V(X;\boldsymbol{\alpha}) = \Phi^{-1}(X),\ \Gamma_H(X;\boldsymbol{\alpha}) = 1.$

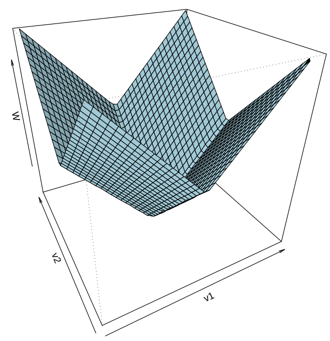

(b) $\Gamma_V(X;\boldsymbol{\alpha}) = |X - 0.5|,\ \Gamma_H(X;\boldsymbol{\alpha}) = 1.$

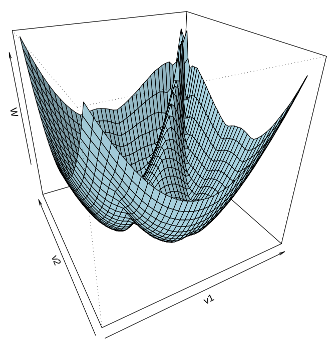

(c) $\Gamma_V(X;\boldsymbol{\alpha}) = |X - 0.5|,\ \Gamma_H(X;\boldsymbol{\alpha}) = |X - 0.5|.$

Figure 2. Weight (W_B) surfaces with respect to cpit data, $\mathbf{V}$. Data generating process is the Gaussian copula ($\rho = 0.71$, corresponding to a Kendall's tau of 0.5) and we perform the cpit assuming the true null hypothesis of a Gaussian copula.

(ii) Compute $(h^*_{1,l}, \ldots, h^*_{d,l})$ by applying the cpit (9) to $(v^*_{1,l}, \ldots, v^*_{d,l})$.

(iii) Compute $\widehat{W}^*_{B,l}$ according to (10) using the same weight functions as in step (5).

(iv) Compute $F_B(W) = \frac{1}{m+1}\sum_{l=1}^{m} I(\widehat{W}'_{B,l} > \widehat{W}_B)$.

(7) Compute $\widehat{B}(w)$ according to (12).

(8) The estimated AD statistic $\widehat{\mathcal{T}}$ is obtained by plugging $\widehat{B}(w)$ into (4).

(9) For some large integer N_b, repeat the following steps for every $k \in \{1, \ldots, N_b\}$:

(a) Generate a random sample $(\mathbf{x}^*_{1,k}, \ldots, \mathbf{x}^*_{n,k})$ from the null hypothesis copula $C_{\widehat{\Theta}}$ and compute the associated pseudo-samples $(\mathbf{z}^*_{1,k}, \ldots, \mathbf{z}^*_{n,k})$ according to (2).

(b) Estimate the parameters Θ, of the null hypothesis copula, with $\widehat{\Theta}^*_k = \widehat{\mathcal{L}}(\mathbf{z}^*_{1,k}, \ldots, \mathbf{z}^*_{n,k})$.

(c) Compute the cpit sample data $(\mathbf{v}^*_{1,k}, \ldots, \mathbf{v}^*_{n,k})$ by applying the cpit to $(\mathbf{z}^*_{1,k}, \ldots, \mathbf{z}^*_{n,k})$, assuming the parametric null hypothesis copula $C_{\widehat{\Theta}^*_k}$.

(d) Compute the cpit2 sample data $(\mathbf{h}^*_{1,k}, \ldots, \mathbf{h}^*_{n,k})$ by applying the cpit (9) to $(\mathbf{v}^*_{1,k}, \ldots, \mathbf{v}^*_{n,k})$.

(e) Compute $W^*_{B,k}$ according to (10), using the same weight functions as in step (5).

(f) If $W^*_{B,k}$ follows a known distribution, compute $F_B(W_B)$ accordingly. If not, approximate F_B as follows. For some large integer m, repeat the following steps for every $l \in \{1, \ldots, m\}$:

(i) Generate a random sample $(v^{**}_{1,l,k}, \ldots, v^{**}_{d,l,k})$ from the null copula, an i.i.d. $U(0,1)^d$ vector.

(ii) Compute $(h^{**}_{1,l,k}, \ldots, h^{**}_{d,l,k})$ by applying the cpit (9) to $(v^{**}_{1,l,k}, \ldots, v^{**}_{d,l,k})$

(iii) Compute $\widehat{W}^{**}_{B,l,k}$ according to (10) using the same weight functions as in step (5).

(iv) Compute $F_B(W^{**}_B) = \frac{1}{m+1}\sum_{k=1}^{m} I(\widehat{W}^{**}_{B,l,k} > \widehat{W}^*_{B,k})$.

(g) Compute $\widehat{B}^*_k(w)$ according to (12).

(h) The estimated AD statistic $\widehat{\mathcal{T}}^*_k$ is obtained by plugging $\widehat{B}^*_k(w)$ into (4).

(10) An approximate p-value for the cpit2 test observator B is then given by

$$\widehat{p} = \frac{1}{N_b+1}\sum_{k=1}^{N_b} I(\widehat{\mathcal{T}}^*_k > \widehat{\mathcal{T}}).$$

Steps $6-9$ may seem abundant and arbitrary. We could have used the W_B's directly and performed some test of its distribution. However, the distribution of W_B is usually not known and numerical- or simulation procedures are needed to approximate F_B.

4. Monte Carlo study

By performing so-called mixing tests we examine the ability of the cpit2-approach to keep nominal sizes and detect tail heaviness and skewness properties. The tests are performed by mixing a Gaussian copula with an alternative copula to construct a mixed copula C^{Mix}:

$$C^{Mix} = (1-\beta)\cdot C^{Ga} + \beta \cdot C^{Alt}, \quad \beta \in [0,1],$$

where β is the mixing parameter, C^{Ga} denotes the Gaussian copula and C^{Alt} denotes the alternative copula. For $\beta = 0$, $C^{Mix} = C^{Ga}$ while for $\beta = 1$, $C^{Mix} = C^{Alt}$. For $0 < \beta < 1$ we sample from the Gaussian copula with probability $(1-\beta)$ and from the alternative copula with probability β.

The alternative copulae considered in this paper are the Student's t-, Clayton- and Gumbel copulae. The ability to distinguish the Gaussian from the Student's t copula indicates the power at detecting lower and upper tail dependency, while the ability to distinguish the Gaussian from the Clayton- and Gumbel copulae indicates the power at detecting lower and upper tail dependency, respectively. For all copulae, the dependency parameter is set to correspond to a Kendall's tau of 0.2, i.e. a weak level of dependence. This should make the various copulae hard to distinguish. For the Student's t copula, the degree of freedom ν, is set to 4, i.e. very heavy tails. For the Gaussian copula, the upper and lower tail dependencies are both 0 while for the Student's t copula the lower and upper tail dependencies, for a Kendall's tau of 0.2, both equal 0.17. For the Clayton copula the lower- and upper tail dependencies equal 4 and 0, respectively. Finally, for the Gumbel copula, the lower- and upper tail dependencies equal 0 and 0.26, respectively. See Nelsen (1999) for the definition of tail dependency.

For the cpit-approach we examined all possible combinations of the weight functions $\Gamma_V(X;\boldsymbol{\alpha})$ and $\Gamma_H(X;\boldsymbol{\alpha})$, listed in Section 3.2, namely $\Phi^{-1}(X)^2$, $|X-0.5|$, $(X-0.5)^{\alpha}$, $\alpha=(2,8,20)$. Note again that the cpit-approach is a special case of the cpit2-approach, with $\Gamma_V(X;\boldsymbol{\alpha})=\Phi^{-1}(X)^2$ and $\Gamma_H(X;\boldsymbol{\alpha})=1$.

Our null hypothesis is that the mixed copula is a Gaussian copula. $\widehat{\mathcal{T}}$ and the corresponding estimate of the p-value is computed according to the test procedure in Section 3.3, using $N_b=500$ for the parametric bootstrap and $m=10000$ for the double bootstrap. The entire procedure is repeated $N_{mix}=2000$ times in order to obtain rejection rates and corresponding power curves. The resulting rejection rates for the best performing weight combinations (at $\beta=1$), are given in Tables 1-3. The weight combination corresponding to the cpit-approach is also included for comparison, although it did not perform very well compared to other combinations.

First, we examine the effect of dimension and sample size. For all combinations of Γ_V and Γ_H, the power increases with dimension and sample size, as expected and visualized in Figure 3. We next examine the nominal levels, i.e. the rejection rates for $\beta=0$, and they all roughly match the prescribed level of 5%. This indicates the validity of our bootstrap procedure. Finally, for all combinations of $d=(2,5)$, $n=(125,250,500)$ and C^{Alt} =(Student's t, Clayton, Gumbel), we examine the power. The best combinations varies with dimension, sample size and whether we consider lower-, upper- or both lower and upper tail dependency. All over, the combinations

(i) $\Gamma_V(X;\boldsymbol{\alpha})=|X-0.5|,\ \Gamma_H(X;\boldsymbol{\alpha})=1$ and

(ii) $\Gamma_V(X;\boldsymbol{\alpha})=(X-0.5)^2,\ \Gamma_H(X;\boldsymbol{\alpha})=1$

stand out as superior. I.e. combinations where we only consider the cpit data $\mathbf{V}$, however with a different weight function than the one in the cpit-approach. When the alternative copula is the Student's t copula, the combinations

(iii) $\Gamma_V(X;\boldsymbol{\alpha})=(X-0.5)^8,\ \Gamma_H(X;\boldsymbol{\alpha})=1$ and

(iv) $\Gamma_V(X;\boldsymbol{\alpha})=(X-0.5)^8,\ \Gamma_H(X;\boldsymbol{\alpha})=|X-0.5|$

perform equally well as combinations (i) and (ii). For this case, the cpit-approach also perform quite well. If the alternative copula is either the Clayton- or the Gumbel copula, combinations (i) and (ii) are, by far, the best. In addition, the combinations

(v) $\Gamma_V(X;\boldsymbol{\alpha})=1,\ \Gamma_H(X;\boldsymbol{\alpha})=\Phi^{-1}(X)^2$,

(vi) $\Gamma_V(X;\boldsymbol{\alpha})=1,\ \Gamma_H(X;\boldsymbol{\alpha})=(X-0.5)^{\alpha},\ \alpha=(2,8,20)$ and

(vii) $\Gamma_V(X;\boldsymbol{\alpha})=|X-0.5|,\ \Gamma_H(X;\boldsymbol{\alpha})=|X-0.5|$

perform quite well.

Consider the particular case where we have very few samples, i.e. $n=125$, and high dimension, i.e. $d=5$ and the alternative copula is the Clayton copula. In this case the Γ_H term adds power compared to combinations only includingthe Γ_V term. However, for large sample sizes (i) and (ii) are superior.

Figure 4 illustrates the difference in power for some combinations. We see that the cpit-approach ($\Gamma_V(X;\boldsymbol{\alpha})=\Phi^{-1}(X)^2,\ \Gamma_H(X;\boldsymbol{\alpha})=1$) has quite low power in some cases, while the cpit2-approach with weight combinations (i), (ii) and several other combinations perform better.

5. Application

The choice of dependency structure can have big impacts in several applications, e.g. capital allocation and the pricing of credit derivatives, such as basket default swaps. We analyze the dependency structure of stock portfolios by looking at their daily log-returns. The historical data consists of 1000 samples of 45 large cap stocks from the New York Stock Exchange, spanning the period from January 13th, 2003 to December 29th, 2006.

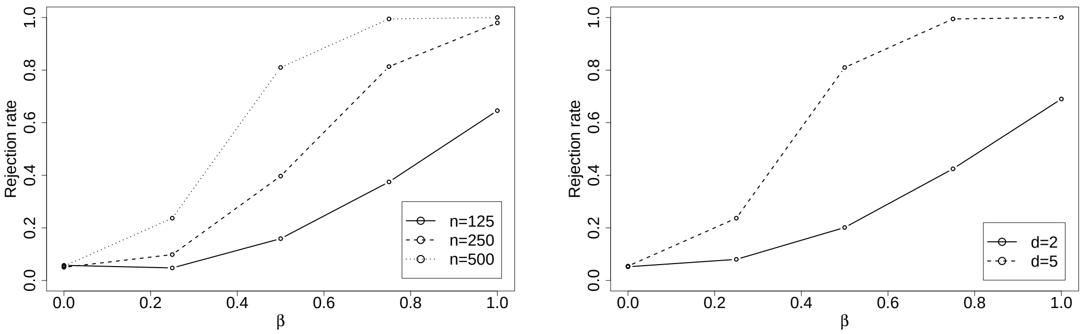

(a) Effect of n, the sample size. $C^{Alt} = C^{t_4}$, $\Gamma_V(x) = |x - 0.5|$, $\Gamma_H(x) = 1$, $d = 5$.

(b) Effect of d, the dimension. $C^{Alt} = C^{t_4}$, $\Gamma_V(x) = |x - 0.5|$, $\Gamma_H(x) = 1$, $n = 500$.

Figure 3. Power curves for the approach B, for varying sample sizes and dimension. On the x-axis we see the mixing parameter β, while on the y-axis we see the portion of times the Gaussian copula (i.e. the null copula) is rejected. 5% significance level.

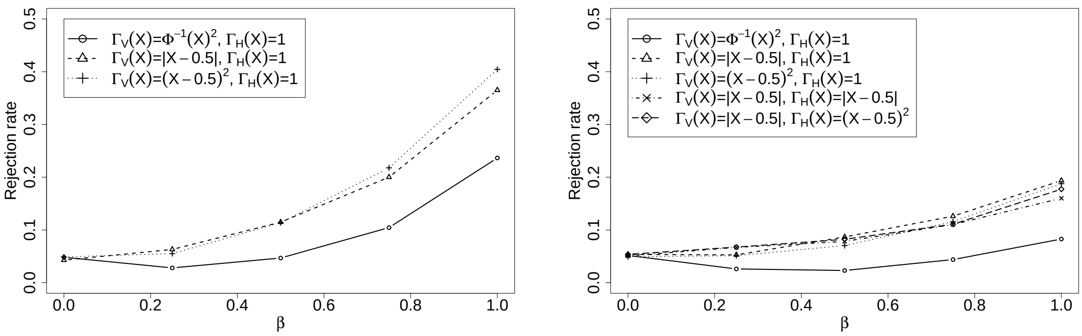

(a) $d = 2$, $n = 250$, C^{Alt} is Student's t.

(b) $d = 5$, $n = 500$, C^{Alt} is Clayton.

Figure 4. Power curves for the approach B, comparing various weight combinations $(\Gamma_V,\ \Gamma_H)$. On the x-axis we see the mixing parameter β, while on the y-axis we see the portion of times the Gaussian copula (i.e. the null copula) is rejected. 5% significance level.

Table 1. Rejection rates, in percent, of the Gaussian copula under various dimensions, sample sizes and weight combinations. The alternative hypothesis is the Student's t copula with 4 degrees of freedom. The dependency parameter corresponds to a Kendall's tau of 0.2. 5% significance level.

d	n	$\Gamma_V(X),\ \Gamma_H(X)$	β				
			0	0.25	0.5	0.75	1
2	125	$\Phi^{-1}(X)^2,\ 1$	4.93	2.71	2.22	3.00	6.06
		$\|X-0.5\|,\ 1$	5.67	4.93	7.24	8.72	16.45
		$(X-0.5)^2,\ 1$	5.67	4.24	6.95	9.51	16.16
		$(X-0.5)^8,\ 1$	4.63	2.86	2.51	3.3	7.34
		$(X-0.5)^8,\ \|X-0.5\|$	5.17	3.00	1.92	2.86	6.31
	250	$\Phi^{-1}(X)^2,\ 1$	4.83	2.81	4.68	10.44	23.65
		$\|X-0.5\|,\ 1$	4.29	6.31	11.48	20	36.50
		$(X-0.5)^2,\ 1$	4.88	5.52	11.38	21.77	40.44
		$(X-0.5)^8,\ 1$	4.53	3.30	4.88	12.81	28.62
		$(X-0.5)^8,\ \|X-0.5\|$	5.02	3.10	3.84	11.97	26.11
	500	$\Phi^{-1}(X)^2,\ 1$	5.42	4.73	15.12	38.33	67.88
		$\|X-0.5\|,\ 1$	5.22	8.03	20.15	42.46	69.01
		$(X-0.5)^2,\ 1$	5.12	7.93	21.72	48.92	77.00
		$(X-0.5)^8,\ 1$	5.22	4.83	15.07	38.33	66.6
		$(X-0.5)^8,\ \|X-0.5\|$	5.57	4.24	13.84	34.19	64.19
5	125	$\Phi^{-1}(X)^2,\ 1$	4.58	1.13	12.71	42.96	79.85
		$\|X-0.5\|,\ 1$	5.76	4.78	15.91	37.49	64.58
		$(X-0.5)^2,\ 1$	5.81	5.91	22.41	52.51	81.92
		$(X-0.5)^8,\ 1$	4.48	3.10	20.49	56.55	87.83
		$(X-0.5)^8,\ \|X-0.5\|$	4.53	1.87	14.53	45.96	82.07
	250	$\Phi^{-1}(X)^2,\ 1$	4.48	6.06	52.41	95.07	100.00
		$\|X-0.5\|,\ 1$	5.02	9.85	39.70	81.33	97.93
		$(X-0.5)^2,\ 1$	4.63	13.00	56.6	93.94	99.85
		$(X-0.5)^8,\ 1$	5.57	10.94	63.15	96.70	100.00
		$(X-0.5)^8,\ \|X-0.5\|$	4.93	8.67	53.55	93.69	99.9
	500	$\Phi^{-1}(X)^2,\ 1$	5.57	26.9	95.86	100.00	100.00
		$\|X-0.5\|,\ 1$	5.47	23.69	80.99	99.46	100.00
		$(X-0.5)^2,\ 1$	5.52	32.32	93.45	100.00	100.00
		$(X-0.5)^8,\ 1$	5.62	36.16	96.75	100.00	100.00
		$(X-0.5)^8,\ \|X-0.5\|$	4.88	29.01	94.19	100.00	100.00

Table 2. Rejection rates, in percent, of the Gaussian copula under various dimensions, sample sizes and weight combinations. The alternative hypothesis is the Clayton copula. The dependency parameter corresponds to a Kendall's tau of 0.2. 5% significance level.

d	n	$\Gamma_V(X),\ \Gamma_H(X)$	β				
			0	0.25	0.5	0.75	1
2	125	$\lvert X-0.5\rvert,\ 1$	5.32	4.93	4.68	4.78	6.40
		$(X-0.5)^2,\ 1$	5.62	5.32	4.58	4.58	5.12
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	5.67	5.02	4.43	4.29	4.48
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	5.32	5.22	4.98	4.38	5.12
		$1,\ \lvert X-0.5\rvert$	5.76	6.06	6.45	5.71	6.40
	250	$\lvert X-0.5\rvert,\ 1$	5.37	5.02	6.11	6.55	6.75
		$(X-0.5)^2,\ 1$	5.62	4.73	5.17	5.17	5.57
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	4.48	6.06	5.27	4.63	5.32
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	4.68	5.32	4.98	4.58	4.43
		$1,\ \lvert X-0.5\rvert$	5.42	6.01	5.27	6.55	7.29
	500	$\lvert X-0.5\rvert,\ 1$	4.04	4.53	6.16	8.03	9.61
		$(X-0.5)^2,\ 1$	4.38	4.53	6.11	6.90	7.83
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	4.58	4.58	6.21	4.68	6.90
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	4.43	4.78	5.12	5.62	5.81
		$1,\ \lvert X-0.5\rvert$	5.27	5.52	6.75	7.14	8.62
5	125	$\lvert X-0.5\rvert,\ 1$	4.29	5.22	5.02	6.45	5.62
		$(X-0.5)^2,\ 1$	5.42	4.24	3.74	4.73	4.48
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	4.78	5.91	5.37	6.21	6.70
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	4.53	6.06	5.57	7.00	6.85
		$1,\ \lvert X-0.5\rvert$	5.07	5.86	5.22	5.71	5.47
	250	$\lvert X-0.5\rvert,\ 1$	4.33	5.37	6.11	8.18	11.63
		$(X-0.5)^2,\ 1$	4.68	4.88	4.19	6.26	8.92
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	4.83	5.12	6.50	7.24	10.20
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	4.53	5.52	6.90	7.83	11.77
		$1,\ \lvert X-0.5\rvert$	5.57	4.93	5.91	6.35	7.54
	500	$\lvert X-0.5\rvert,\ 1$	5.37	5.27	8.67	12.61	19.36
		$(X-0.5)^2,\ 1$	4.93	5.12	7.04	11.58	18.77
		$\lvert X-0.5\rvert,\ \lvert X-0.5\rvert$	5.07	6.70	7.83	11.03	16.01
		$\lvert X-0.5\rvert,\ (X-0.5)^2$	5.37	6.75	8.23	11.08	17.73
		$1,\ \lvert X-0.5\rvert$	5.42	5.57	7.04	7.29	8.62

Table 3. Rejection rates, in percent, of the Gaussian copula under various dimensions, sample sizes and weight combinations. The alternative hypothesis is the Gumbel copula. The dependency parameter corresponds to a Kendall's tau of 0.2. 5% significance level.

d	n	$\Gamma_V(X),\ \Gamma_H(X)$	β				
			0	0.25	0.5	0.75	1
2	125	$\|X-0.5\|,\ 1$	5.91	4.88	5.71	6.06	6.55
		$(X-0.5)^2,\ 1$	6.16	4.88	5.76	5.02	5.37
		$1,\ \Phi^{-1}(X)^2$	5.96	5.22	4.88	4.58	4.73
		$1,\ (X-0.5)^8$	5.91	5.71	4.58	4.98	5.07
	250	$\|X-0.5\|,\ 1$	5.07	5.27	5.91	6.75	8.37
		$(X-0.5)^2,\ 1$	5.67	4.78	4.93	5.52	6.90
		$1,\ \Phi^{-1}(X)^2$	4.93	5.57	5.22	4.19	4.63
		$1,\ (X-0.5)^8$	5.22	5.17	5.91	3.99	4.73
	500	$\|X-0.5\|,\ 1$	5.52	6.40	6.75	9.85	12.51
		$(X-0.5)^2,\ 1$	5.27	5.22	6.75	9.06	12.32
		$1,\ \Phi^{-1}(X)^2$	5.32	5.12	5.07	5.22	6.11
		$1,\ (X-0.5)^8$	5.22	5.12	5.17	5.42	5.67
5	125	$\|X-0.5\|,\ 1$	4.68	4.09	4.29	4.14	3.94
		$(X-0.5)^2,\ 1$	4.88	4.09	3.30	3.15	3.60
		$1,\ \Phi^{-1}(X)^2$	5.32	6.06	4.98	5.22	5.17
		$1,\ (X-0.5)^8$	5.22	5.76	4.93	5.07	4.68
	250	$\|X-0.5\|,\ 1$	5.22	4.04	3.99	5.27	6.60
		$(X-0.5)^2,\ 1$	5.07	4.04	4.29	4.88	7.00
		$1,\ \Phi^{-1}(X)^2$	5.71	4.63	5.47	5.42	5.52
		$1,\ (X-0.5)^8$	5.32	5.37	5.76	5.71	5.47
	500	$\|X-0.5\|,\ 1$	4.63	4.24	4.98	8.33	12.81
		$(X-0.5)^2,\ 1$	4.53	3.65	5.57	9.41	17.09
		$1,\ \Phi^{-1}(X)^2$	4.98	5.57	5.32	6.06	7.44
		$1,\ (X-0.5)^8$	6.06	4.83	4.53	6.11	6.50

Asset collections of dimension 2 and 5 were randomly selected 2000 times from the full data set. As in Chen et al. (2004) and Panchenko (2005) we examine the raw returns and the GARCH$(1,1)$ filtered returns, i.e. each individual assets return is filtered through a standard GARCH$(1,1)$ process. This filtering is done to remove serial dependence in each individual time series. For details of GARCH processes, see e.g. Bollerslev (1986). We fit the Gaussian-, Student's t-, Clayton- and Gumbel copulae to the portfolios and apply the cpit2-approach, with $\Gamma_V(X;\boldsymbol{\alpha}) = |X-0.5|$ and $\Gamma_H(X;\boldsymbol{\alpha}) = 1$, to investigate how often each copula is rejected.

Table 4 shows the rejection rates for the raw and filtered returns. We see that for all but the Student's t copula, the rejection rate is increasing with sample size. For $d = 5$ the rejection rates for the Clayton- and Gumbel copulae are very high. This is not surprising since we are only considering the so-called exchangeable Clayton- and Gumbel copulae, having only one dependency parameter. Fitting a 5-dimensional distribution with only one parameter is usually not sufficient. The Gaussian copula is not that easily rejected for small sample sizes in the bivariate case. However, for higher dimensions and sample sizes, we see that the Gaussian copula is strongly rejected for both raw and filtered returns. The Student's t copula seems to provide a very good fit for all dimensions and sample sizes and for both raw and filtered returns. It is not surprising that the Student's t copula outperforms the other copulae since it has more parameters. Nevertheless, the low rejection rates for the Student's t copula are interesting. Also, note the reduced rejection rates in most cases for the filtered returns. This is also expected since serial dependence is removed.

6. Concluding remarks

We have generalized and extended the copula gof approach proposed by Breymann et al. (2003). The main contribution is the flexibility in the dimension reduction function. The generalization enables the user to apply any weight function combination to the cpit data sets $\mathbf{V}$ and $\mathbf{H}$, depending on the use. The additional cpit step, based on order statistics, should make the dimension reduction more robust to the inconsistency issue illustrated in Figure 1. We have not been able to reconstruct the inconcistency issue in our Monte Carlo study, except when the alternative copula is the Clayton copula, for $d = 5$, $n = 125$ where the Γ_H term seems to add power. However, the added power in this case may also be due to the high dimension and very few samples. Neither have we found a real world data set where this issue manifests itself. However, the danger of this issue coming into play will always be there, justifying our extension.

Monte Carlo results show that our approach keeps the prescribed nominal level for all weight combinations examined. We also see that the cpit-approach has low power in some circumstances, particularly for low sample sizes. The reason is that the dimension reduction strongly weights the boundaries of the d-dimensional unit hypercube. If we have few samples there are few observations in the boundary regions and the cpit-approach becomes less robust and less powerful. An important result is the superior performance, in all our numerical tests, of the two weight combinations (i): $\Gamma_V(X;\boldsymbol{\alpha}) = |X-0.5|$, $\Gamma_H(X;\boldsymbol{\alpha}) = 1$ and (ii): $\Gamma_V(X;\boldsymbol{\alpha}) = (X-0.5)^2$, $\Gamma_H(X;\boldsymbol{\alpha}) = 1$. Hence, based on our numerical experiments, these combinations are recommended. For skewness properties, it may seem like the additional Γ_H term has some effect, in particular for higher dimensions and small sample sizes. With the inconsistency issue from Figure 1 in mind, the additional use of $\Gamma_V(X;\boldsymbol{\alpha}) = |X-0.5|$, $\Gamma_H(X;\boldsymbol{\alpha}) = |X-0.5|$ is recommended.

Application of the cpit2-approach to a collection of large cap stock portfolios show that the Student's t copula provide a fairly good fit to the data while the Gaussian copula is strongly rejected for higher dimensions. A GARCH(1,1) filtering of the original data only marginally reduced the rejection of the Gaussian copula. This is in accordance with the findings of Dobrić and Schmid (2005) and Chen et al. (2004) and indicates that the Student's t copula, in general provides a superior fit to daily log-returns for equity prices.

Further work involve comparison of the cpit2-approach with other approaches, e.g. Panchenko (2005), Genest and Rémillard (2005) and Genest et al. (2006). Further tests of various weight combinations and their relative performance with respect to various degrees of dependence, dimensions, sample sizes and null- and alternative hypotheses, is also of interest. Finally work needs to be done to better understand how the weight combinations relate to the original data set $\mathbf{Z}$, not only the cpit data set $\mathbf{V}$.

Table 4. Rejection rates, in percent, of the Gaussian-, Student's t- and one-parameter Clayton- and Gumbel copulae, applied to the raw and GARCH$(1,1)$ filtered returns. Dimensions $d = (2,5)$, sample sizes $n = (250,500,1000)$, $\Gamma_V(X) = |X - 0.5|$ and $\Gamma_H(X) = 1$. 5% significance level.

Gaussian copula			
Dimension	Sample size	Raw returns	Filtered returns
2	250	6.23	5.15
	500	11.42	5.25
	1000	26.18	11.67
5	250	12.75	7.35
	500	16.52	12.06
	1000	49.71	25.10
Student's t copula			
Dimension	Sample size	Raw returns	Filtered returns
2	250	5.10	7.50
	500	6.18	5.44
	1000	6.42	3.53
5	250	7.30	5.29
	500	9.80	6.37
	1000	10.74	10.29
Clayton copula			
Dimension	Sample size	Raw returns	Filtered returns
2	250	7.30	7.25
	500	9.85	5.64
	1000	21.52	8.43
5	250	32.06	19.61
	500	41.32	35.39
	1000	77.84	60.44
Gumbel copula			
Dimension	Sample size	Raw returns	Filtered returns
2	250	5.88	3.09
	500	5.34	7.89
	1000	14.36	5.44
5	250	44.80	31.62
	500	44.02	33.33
	1000	75.29	50.98

Acknowledgements

The authors acknowledge the support and guidance of colleagues at the Norwegian Computing Center, in particular Assistant Research Director Kjersti Aas and Chief Research Scientist Xeni Kristine Dimakos. We would especially like to express our gratitude to Kjersti Aas for introducing us to the field of copulae and copula gof testing and for her continuous interest and drive. Also, thank is due to an anonymous referee and participants at conferences and workshops, in particular Professor Christian Genest and Professor Dr. Friedrich Schmid, for valuable comments.

References

Anderson, T. W. and D. A. Darling (1954). A test of goodness of fit. *Journal of the American Statistical Association 49*, 765–769.

Aslan, B. and G. Zech (2002). Comparison of different goodness-of-fit tests. In *Advanced Statistical Techniques in Particle Physics*. Grey College, Durham.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*, 307–327.

Breymann, W., A. Dias, and P. Embrechts (2003). Dependence structures for multivariate high-frequency data in finance. *Quantitative Finance 1*, 1–14.

Chen, X., Y. Fan, and A. Patton (2004). Simple tests for models of dependence between multiple financial time series, with applications to U.S. equity returns and exchange rates. Financial Markets Group, London School of Economics, Discussion Paper 483. Revised July 2004.

Cui, S. and Y. Sun (2004). Checking for the gamma frailty distribution under the marginal proportional hazards frailty model. *Statistica Sinica 14*, 249–267.

D'Agostino, R. B. and M. A. Stephens (1986). *Goodness-of-Fit Techniques*. New York: Marcel Dekker Inc.

David, H. A. (1981). *Order Statistics* (2 ed.). John Wiley and Sons.

Deheuvels, P. (1984). The characterization of distributions by order statistics and record values: A unified approach. *Annals of Eugenics 21*(2), 326–334.

Demarta, S. and A. J. McNeil (To appear in 2005). The t copula and related copulas. *International Statistical Review*.

Dobrić, J. and F. Schmid (2005). Testing goodness of fit for parametric families of copulas - application to financial data. *Communications in Statistics: Simulation and Computation 34*(4), 1053–1068.

Dobrić, J. and F. Schmid (2007). A goodness of fit test for copulas based on rosenblatt's transformation. Forthcoming in Computational Statistics and Data Analysis.

Embrechts, P., A. McNeil, and D. Straumann (1999, May). Correlation and dependence in risk management: Properties and pitfalls. *RISK*, 69–71.

Fermanian, J. (2005). Goodness of fit tests for copulas. *Journal of Multivariate Analysis 95*, 119–152.

Genest, C., K. Ghoudi, and L. Rivest (1995). A semi-parametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika 82*, 543–552.

Genest, C., J.-F. Quessy, and B. Rémillard (2006). Goodness-of-fit procedures for copula models based on the probability integral transform. *Scandinavian Journal of Statistics 33*.

Genest, C. and B. Rémillard (2005). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Les Cahiers du GERAD G-2005-51.*

Genest, C., B. Rémillard, and D. Beaudoin (2007). Omnibus goodness-of-fit tests for copulas: A review and a power study. *To appear in Insurance: Mathematics and Economics*.

Glen, A. G., L. M. Leemis, and D. R. Barr (2001). Order statistics in goodness-of-fit testing. *IEEE Transactions on Reliability 50*(2), 209–213.

Glidden, D. V. (1999). Checking the adequacy of the gamma frailty model for multivariate failure times. *Biometrika 86*, 381–393.

Hong, Y. and H. Li (2002). Nonparametric specification testing for continuous-time models with application to spot interest rates. Working Paper, Department of Economics, Cornell University.

Joe, H. (1997). *Multivariate Models and Dependence Concepts.* London: Chapman & Hall.

Malevergne, Y. and D. Sornette (2003). Testing the gaussian copula hypothesis for financial assets dependence. *Quantitative Finance 3*, 231–250.

Marsaglia, G. and J. Marsaglia (2004). Evaluating the Anderson-Darling distribution. *Journal of Statistical Software 9.*

Mashal, R. and A. Zeevi (2002). Beyond correlation: Extreme co-movements between financial assets. Technical report, Columbia University.

Nelsen, R. B. (1999). *An Introduction to Copulas.* New York: Springer Verlag.

Panchenko, V. (2005). Goodness-of-fit test for copulas. *Physica A 355(1)*, 176–182.

Romano, C. (2002). Calibrating and simulating copula functions: An application to the Italian stock market. Working paper n. 12, CIDEM, Universit'a degli Studi di Roma "La Sapienza".

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *The Annals of Mathematical Statistics 23*, 470–472.

Savu, C. and M. Trede (2004). Goodness-of-fit tests for parametric families of archimedean copulas. CAWM discussion paper, No. 6.

Scaillet, O. (2005, May). Kernel based goodness-of-fit tests for copulas with fixed smoothing parameters. FAME Research Paper Series rp145, International Center for Financial Asset Management and Engineering. available at http://ideas.repec.org/p/fam/rpseri/rp145.html.

Shih, J. H. (1998). A goodness-of-fit test for association in a bivariate survival model. *Biometrika 85*, 189–200.

Shih, J. H. and T. A. Louis (1995). Inferences on the association parameter in copula models for bivariate survival data. *Biometrics 51*, 1384–1399.

Sklar, A. (1959). Fonctions dé repartition á n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris 8*, 299–231.