Hierarchical Archimedean Copulae: The HAC Package

Ostap Okhrin
Humboldt-Universität zu Berlin

Alexander Ristig
Humboldt-Universität zu Berlin

Abstract

This paper aims at explanation of the R-package **HAC**, which provides user friendly methods for dealing with hierarchical Archimedean copulae (HAC). Computationally efficient estimation procedures allow to recover the structure and the parameters of HAC from data. In addition, arbitrary HAC can be constructed to sample random vectors and to compute the values of the corresponding cumulative distribution plus density functions. Accurate graphics of the HAC structure can be produced by the generic `plot` function.

Keywords: copula, R, hierarchical Archimedean copula, HAC.

1. Introduction

The use of copulae in applied statistics began in the end of the 90th, when Embrechts, McNeil, and Straumann (1999) introduced copula to empirical finance in the context of risk management. Nowadays, quantitative orientated sciences like biostatistics and hydrology use copulae to attempt measuring the dependence of random variables, e.g., Lakhal-Chaieb (2010); Acar, Craiu, and Yao (2011); Bárdossy (2006); Genest and Favre (2007); Bárdossy and Li (2008). In finance, copulae became a standard tool, explicitly on VaR measurement and in valuation of structured credit portfolios, see Mendes and Souza (2004); Junker and May (2005) and Li (2000). This paper aims at providing the necessary tools for academics and practitioners for simple and effective use of hierarchical Archimedean copulae (HAC) in their statistical analysis.

Copula is the function splitting a multivariate distribution into its margins and a pure dependency component. Formally, copulae are introduced in Sklar (1959) stating that if F is an arbitrary d-dimensional continuous distribution function of the random vector $X = (X_1, \ldots, X_d)^\top$, then the associated copula is unique and defined as the continuous mapping $C : [0,1]^d \to [0,1]$ which satisfies the equality

$$C(u_1, \ldots, u_d) \quad = \quad F\{F_1^{-1}(u_1), \ldots, F_d^{-1}(u_d)\}, \quad u_1, \ldots, u_d \in [0,1],$$

where $F_1^{-1}(\cdot), \ldots, F_d^{-1}(\cdot)$ are the quantile functions of the corresponding continuous marginal distribution functions $F_1(x_1), \ldots, F_d(x_d)$. Accordingly, a d-dimensional density $f(\cdot)$ can be split in the copula density $c(\cdot)$ and the product of the marginal densities. For an overview and recent developments of copulae we refer to Nelsen (2006), Cherubini, Luciano, and Vecchiato (2004), Joe (1997) and Jaworski, Durante, and Härdle (2013). If $F(\cdot)$ belongs to the class of elliptical distributions, then $C(\cdot)$ is an elliptical copula, which in most cases cannot be

given explicitly because the distribution function $F(\cdot)$ and the inverse marginal distributions $F_j(\cdot)$ usually have integral representations. One of the classes that overcomes this drawback of elliptical copulae is the class of Archimedean copulae, which, however, is very restrictive yet for moderate dimensions. Among other R packages (R Core Team (2013)) dealing with Archimedean copula, we would like to mention the **copula** and the **fCopulae** package, c.f. Yan (2007); Kojadinovic and Yan (2010); Hofert and Mächler (2011); Hofert, Kojadinovic, Mächler, and Yan (2013) and Wuertz, many others, and see the SOURCE file (2013).

HAC generalize the concept of simple Archimedean copulae by substituting (a) marginal distribution(s) by a further HAC. This class is thoroughly analyzed in Embrechts, Lindskog, and McNeil (2003); Whelan (2004); Savu and Trede (2010); Hofert (2011); Okhrin, Okhrin, and Schmid (2013b). The first sampling algorithms for special HAC structures were provided by the **QRMlib** package of McNeil and Ulman (2011), which is not updated anymore, but several functions were ported to the **QRM** package, see Pfaff and McNeil (2013). Hofert and Mächler (2012) presented the comprehensive **nacopula** package which, among other features, allows sampling from arbitrary HAC and was integrated into the package **copula** from version 0.8-1. The central contribution of the **HAC** package is the estimation of the parameter and the structure for this class of copulae, as discussed in Okhrin, Okhrin, and Schmid (2013a), including a simple and intuitive representation of HAC as R-objects of the class `hac`. The main estimation procedure relies on a recursive multi-stage Maximum Likelihood (ML) procedure, which determines the parameter and the structure simultaneously. This elegant procedure endows the estimator with the usual asymptotic properties but avoids the computationally intensive one-step ML estimation, which is also implemented for a predetermined structure. Besides, the package offers functions for producing graphics of the copula's structure, for sampling random vectors from a given copula and for computing values of the corresponding distribution and density.

The paper is organized as follows. Section 2 describes shortly the theoretical aspects of HAC and its estimation. Section 3 presents the functions of the **HAC** package and Section 4 a simulation study. Section 5 concludes.

2. Hierarchical Archimedean copulae

As mentioned above, the large class of copulae, which can describe tail dependency, non-ellipticity, and, most importantly, has close form representation

$$C(u_1, \ldots, u_d; \theta) \;\; = \;\; \phi_\theta \left\{ \phi_\theta^{-1}(u_1) + \cdots + \phi_\theta^{-1}(u_d) \right\}, \quad u_1, \ldots, u_d \in [0,1], \tag{1}$$

where $\phi_\theta(\cdot) \in \mathfrak{L} = \{ \phi_\theta : [0;\infty) \to [0,1] \,|\, \phi_\theta(0) = 1,\, \phi_\theta(\infty) = 0;\, (-1)^j \phi_\theta^{(j)} \geq 0;\, j \in \mathbb{N} \}$ and $(-1)^j \phi_\theta^{(j)}(x)$ being non-decreasing and convex on $[0,\infty)$, for $x > 0$, is the class of Archimedean copulae. The function $\phi(\cdot)$ is called the generator of the copula and commonly depends on a single parameter θ. For example, the Gumbel generator is given by $\phi_\theta(x) = \exp(-x^{1/\theta})$ for $0 \leq x < \infty$, $1 \leq \theta < \infty$. Detailed reviews of the properties of Archimedean copulae can be found in McNeil and Nešlehová (2009) and in Joe (1996).

A disadvantage of Archimedean copulae is the fact that the multivariate dependency structure is very restricted, since it typically depends on a single parameter of the generator function $\phi(\cdot)$. Moreover, the rendered dependency is symmetric with respect to the permutation of variables, i.e., the distribution is exchangeable. HAC (also called nested Archimedean copulae)

overcome this problem by considering the compositions of simple Archimedean copulae. For example, the special case of four-dimensional fully nested HAC can be given by

$$
\begin{aligned}
C(u_1, u_2, u_3, u_4) &= C_3\{C_2(u_1, u_2, u_3), u_4\} \\
&= \phi_3\{\phi_3^{-1} \circ C_2(u_1, u_2, u_3) + \phi_3^{-1}(u_4)\},
\end{aligned}
\tag{2}
$$

where $C_j(u_1, \ldots, u_{j+1}) = \phi_j[\phi_j^{-1}\{C_{j-1}(u_1, \ldots, u_j)\} + \phi_j^{-1}(u_{j+1})]$, $j = 2, \ldots, d-1$, and $C_1 = \phi_1\{\phi_1^{-1}(u_1) + \phi_1^{-1}(u_2)\}$. The functional form of $C_j(\cdot)$ indicates that the composition can be applied recursively. A different segmentation of the variables leads naturally to more complex HAC. In the following, let d-dimensional HAC be denoted by $C(u_1, \ldots, u_d; s, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ denotes the vector of feasible dependency parameters and $s = (\ldots (i_g i_k) i_\ell \ldots)$ the structure of the entire HAC, where $i_m \in \{1, \ldots, d\}$ is a reordering of the indices of the variables with $m = 1 \ldots, d$, and $g, k, \ell \in \{1, \ldots, d : g \neq k \neq \ell\}$. Structures of subcopulae are denoted by s_j with $s = s_{d-1}$. For instance, the structure according to Equation 2 is $s = (s_2)4$ with $s_j = (s_{j-1}(j+1))$, $j = 2, 3$, for the sucopulae and $s_1 = (12)$. A clear definition of the structure is essential, as s is in fact a parameter to estimate. Thus, Equation 2 can be rewritten as

$$
\begin{aligned}
C(u_1, u_2, u_3, u_4; s = (((12)3)4), \boldsymbol{\theta}) &= C\{u_1, u_2, u_3, u_4; (s_2 4), (\theta_1, \theta_2, \theta_3)^\top\} \\
&= \phi_{\theta_3}(\phi_{\theta_3}^{-1} \circ C_2\{u_1, u_2, u_3; (s_1(3)), (\theta_1, \theta_2)^\top\} + \phi_{\theta_3}^{-1}(u_4)).
\end{aligned}
$$

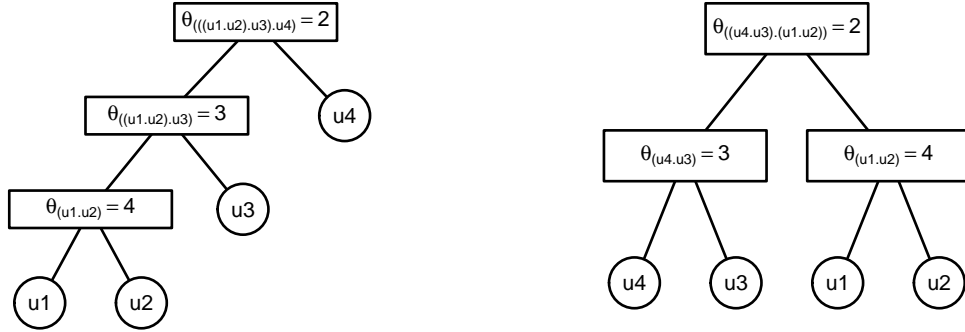Figure 1 presents the four-dimensional fully and partially nested Archimedean copula.

Figure 1: Fully and partially nested Archimedean copulae of dimension $d = 4$ with structures $s = (((12)3)4)$ on the left and $s = ((43)(12))$ on the right.

HAC can adopt arbitrarily complex structures s. This makes it a very flexible and simultaneously parsimonious distribution model. The generators $\phi_{\theta_j}(\cdot)$ within a single nested Archimedean copula can come either from a single generator family or from different generator families. If the $\phi_{\theta_j}(\cdot)$'s belong to the same family, then the required complete monotonicity of $\phi_{\theta_{i+j}}^{-1}(\cdot) \circ \phi_{\theta_j}(\cdot)$ usually imposes some constraints on the parameters $\theta_1, \ldots, \theta_{d-1}$. Theorem 4.4 of McNeil (2008) provides sufficient conditions on the generator functions to guarantee that C is a copula. It holds that if $\phi_{\theta_j}(\cdot) \in \mathfrak{L}$, for $j = 1, \ldots, d-1$, and $\phi_{\theta_{j+1}}^{-1}(\cdot) \circ \phi_{\theta_j}(\cdot)$ have completely monotone derivatives, then $C(\cdot)$ is a copula for $d \geq 2$. For the majority of generators feasible HAC require decreasing parameters from the highest to the lowest hierarchical level. However, in the case of different families within a single HAC, the condition of complete monotonicity is not always fulfilled, see Hofert (2011). In our study, we consider HAC with generators from the same family only. If we use the same single-parameter generator function

on each level, but with a different value of θ, we may specify the whole distribution with at most $d-1$ parameters. From this point of view, the HAC approach can be seen as an alternative to covariance driven models. Nevertheless, for HAC not only the parameters are unknown, but also the structure has to be determined. One possible procedure for estimating both the parameters and the structure is to enumerate all possible structures and to estimate at first the parameters only. Next, the optimal structure can be determined by a suitable goodness-of-fit test. This approach is, however, unrealistic in practice because the variety of different structures is enormously large even in moderate dimensions. Okhrin *et al.* (2013a) suggest computationally efficient procedures, which allow to estimate HAC recursively. The **HAC** package provides these methods for estimating the parameters and structure in a user-friendly way.

2.1. Estimation of HAC

In most cases the discussion is constrained to binary copulae, i.e., at each node of structure only two variables are joined together. We impose this restriction because the numerical values of the estimated parameters are compared during the procedure. Certainly, these comparisons cannot be justified if one parameter corresponds to a tertiary or higher dimensional copula. The entire procedure can be described in a recursive way where at the first iteration step we fit a bivariate copula to every couple of the variables. The couple of variables with the strongest dependency is selected. We denote the respective estimator of the parameter at the first level by $\hat{\theta}_1$ and the set of indices of the variables by I_1. The selected couple is joined together to define the pseudo-variable $Z_{I_1} \stackrel{\text{def}}{=} C\{(I_1); \hat{\theta}_1, \phi_1\}$. At the next step, we proceed in the same way by considering the remaining variables and the new pseudo-variable as the new set of variables. This procedure allows us to determine the estimated structure of the copula. If the restrictions on the parameters are always fulfilled, it leads to a feasible copula function with $d-1$ parameters. Nevertheless, if the true copula is not binary, the procedure might return a slightly misspecified structure. Despite a difference in the structures, the difference in the distribution functions is in general minor. To allow more sophisticated structures, we aggregate the variables of the estimated copula afterwards. This is possible if the absolute value of the difference of two successive nodes is smaller than a fixed small threshold, i.e., $\theta_1 - \theta_2 < \epsilon$, with $\theta_1 > \theta_2$, as suggested by Okhrin *et al.* (2013a).

For better understanding, let us consider a three-dimensional example with u_j, $j = 1, 2, 3$, being uniformly distributed on $[0, 1]$. All possible pairs $C_{(12)}(u_1, u_2, \hat{\theta}_{(12)})$, $C_{(13)}(u_1, u_3, \hat{\theta}_{(13)})$ and $C_{(23)}(u_2, u_3, \hat{\theta}_{(23)})$ are estimated by regular ML, see Franke, Härdle, and Hafner (2011). To compare the strengths of the fit one can use computationally complicated goodness-of-fit tests, which do not necessarily lead to a function which will be a copula on the final level of aggregation due to the restrictions on $\boldsymbol{\theta}$. For that reason we compare simply the estimates $\hat{\theta}_{(12)}$, $\hat{\theta}_{(13)}$ and $\hat{\theta}_{(23)}$. This is due to the fact that for most Archimedean copulae, the larger the parameter the stronger is the dependency (the larger the parameter the larger is Kendall's correlation coefficient). Let the strongest dependence be in the first pair $\hat{\theta}_1 \stackrel{\text{def}}{=} \hat{\theta}_{(12)} = \max\{\hat{\theta}_{(12)}, \hat{\theta}_{(13)}, \hat{\theta}_{(23)}\}$, then $I_1 = \{1, 2\}$ and we introduce the pseudo-variable $Z_1 \stackrel{\text{def}}{=} C_1(I_1; \hat{\theta}_1) = C_1(u_1, u_2; \hat{\theta}_{(12)})$. At the next and final step for this example we join together u_3 and Z_1. The theoretical validation is also reported by Proposition 1 of Okhrin *et al.* (2013b) stating that HAC can be uniquely recovered from the marginal distribution functions and all bivariate copula functions. Crucially for the superior recursive ML estimation procedure,

pseudo-variables are regarded as functions of the underlying random variables $X_1, \ldots, X_d$ and are not explicitly computed.

In practice, the marginal distributions F_j, $j = 1, \ldots, d$, are either parametrically $\widehat{F}_j(\cdot) = F_j(\cdot, \hat{\boldsymbol{\alpha}}_j)$, where $\boldsymbol{\alpha}_j$ denotes the vector of parameters of the j-th margin, or non-parametrically

$$\widehat{F}(x) \;\; = (n+1)^{-1} \sum_{i=1}^{n} \mathbf{I}\left(X_i \leq x\right) \tag{3}$$

estimated in advance. Accordingly, the marginal densities $\hat{f}_j(\cdot)$, $j, \ldots, d$, are estimated by an appropriate kernel density estimator or using a parametric density.

Following Okhrin *et al.* (2013a), the estimation of the copula parameters at each step of the iteration can be sketched as follows: at first stage, we estimate the parameter of the copula at the first hierarchical level assuming that the marginal distributions are known. At further stages, the next level copula parameter is estimated assuming that the margins as well as the copula parameters at lower levels are known. Let $\mathbf{X} = \{x_{ij}\}^{\top}$ be the respective sample, for $i = 1, \ldots, n$, $j = 1, \ldots, d$, and $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_{d-1})^{\top}$ be the parameters of the copula starting with the lowest up to the highest level. The recursive multi-stage ML estimator $\hat{\boldsymbol{\theta}}$ solves the system

$$\left(\frac{\partial \mathcal{L}_1}{\partial \theta_1}, \ldots, \frac{\partial \mathcal{L}_{d-1}}{\partial \theta_{d-1}} \right)^{\top} = \mathbf{0}, \tag{4}$$

$$\text{where} \quad \mathcal{L}_j \;\; = \;\; \sum_{i=1}^{n} l_j(\mathbf{X}_i), \text{ for } j = 1, \ldots, d-1,$$

$$l_j(\mathbf{X}_i) \;\; = \;\; \log \left\{ c_j \big[\{\widehat{F}_m(x_{im})\}_{m \in s_j}; s_j, \theta_j \big] \prod_{m \in s_j} \hat{f}_m(x_{im}) \right\}$$

$$\text{for } j = 1, \ldots, d-1, \; i = 1, \ldots, n,$$

where s_j is referred to the (pseudo)-variables considered at the j-th estimation stage. Chen and Fan (2006) and Okhrin *et al.* (2013a) provide asymptotic behaviour of the estimates. At the moment, there are three different ways to estimate HAC:

(i) Ordinary (full) ML estimation, also denoted by `FML`, which is based on the complete log-likelihood and hence on a predetermined structure.

(ii) The `ML` setup is based on realized pseudo-variables, i.e., the values of the pseudo-variables for the given sample are explicitly computed, so that the bivariate density is maximized with respect to the copula parameter at each step of the procedure. The benefits of the ML method hold in particular for binary and non-complex structures. If the structure of the HAC is, however, complex, the estimates around the initial node seem to be slightly biased. Note that this procedure is not supported by asymptotic theory.

(iii) More precise results can be obtained by the recursive ML (`RML`) procedure discussed in Okhrin *et al.* (2013a). The difference between the ML method and the recursive ML procedure results from the maximized log-likelihood. While the bivariate log-likelihood is considered at each estimation step of the ML method, the log-likelihood of the recursive ML procedure corresponds at each estimation step to the full log-likelihood for the marginal HAC regarded at that step. Compared to the full ML approach, the log-likelihood is only optimized with respect to the parameter at the root node taken the estimated parameter(s) at lower

hierarchical levels as given, so that the final HAC being a copula is ensured by shortening the feasible parameter interval from above. From this point of view, the computational challenge is to build the log-likelihood for the full ML estimation, which is almost solved by constructing the d-dimensional density, see Section 3.4.

3. Applications of HAC

Core of the **HAC** package is the function `estimate.copula` estimating the parameter and determining the structure for given data. Let us consider the dataset `finData` included in the **HAC** package. It contains the residuals of the filtered daily log-returns of four oil corporations: Chevron Corporation (`CVX`), Exxon Mobil Corporation (`XOM`), Royal Dutch Shell (`RDSA`) and Total (`FP`), covering $n = 283$ observations from 20110202 to 20120319. Intertemporal dependence is removed by usual ARMA-GARCH models, whose standardized residuals are plotted in Figure 2 and used in the subsequent analysis:

```
R> library("HAC")
R> data("finData")
R> system.time(result <- estimate.copula(finData, margins = "edf"))

   user   system  elapsed
  0.095    0.008    0.175


R> result


Class: hac
Generator: Gumbel
((XOM.CVX)_{2.83}.(FP.RDSA)_{2.09})_{1.55}
```

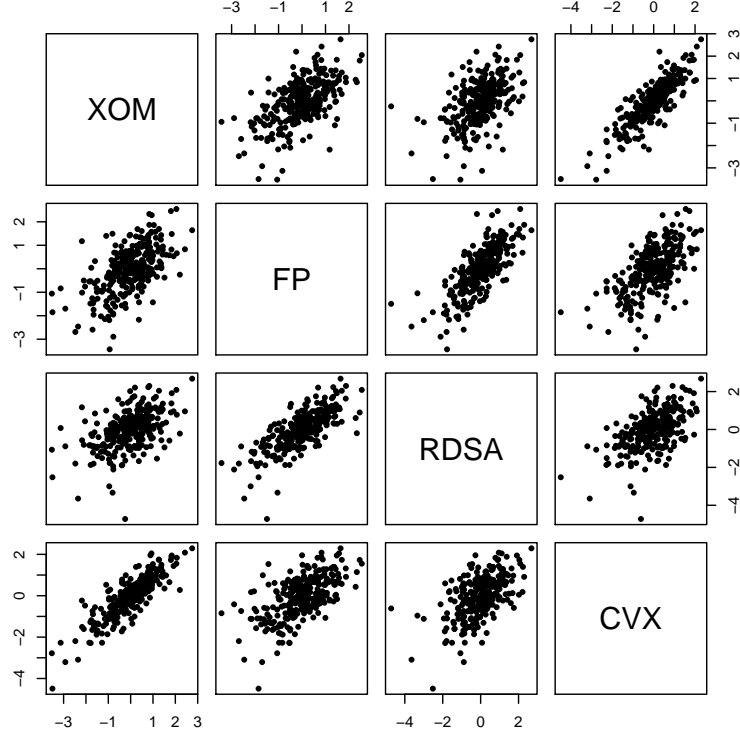The returned object `result` is of class `hac`, whose properties are explored below.

The multi-stage estimation procedure is illustrated in Table 1 for the four-dimensional example from above. At the lowest hierarchical level, the parameter of all bivariate copulae are estimated. The couple $(X_{\text{CVX}}, X_{\text{XOM}})$ produces the strongest dependency, hence the best fit. Then, the pseudo-variable

$$Z_{(\text{CVX.XOM})} \quad \overset{\text{def}}{=} \quad \phi_{\hat{\theta}_{(\text{CVX.XOM})}} \left[\phi_{\hat{\theta}_{(\text{CVX.XOM})}}^{-1} \left\{ \widehat{F}_{\text{XOM}} \left(X_{\text{XOM}} \right) \right\} + \phi_{(\hat{\theta}_{\text{CVX.XOM}})}^{-1} \left\{ \widehat{F}_{\text{CVX}} \left(X_{\text{CVX}} \right) \right\} \right] \tag{5}$$

is defined. As the `RML` method is used by default, the involved variables X_{XOM} and X_{CVX} are not substituted by the in the dataset. At the next nesting level the parameters of all bivariate subsets are estimated and the variables X_{FP} and X_{RDSA} exhibit the best fit. Finally, the realizations of the remaining random variables $Z_{(\text{CVX.XOM})}$ and $Z_{(\text{FP.RDSA})}$ are grouped at the highest level of the hierarchy, where $Z_{(\text{FP.RDSA})}$ is defined analogously to $Z_{(\text{CVX.XOM})}$.

In general, `estimate.copula` includes the following arguments:

```
R> names(formals(estimate.copula))
```

Figure 2: Scatterplot of the sample `finData`.

$$z_{i,(\text{CVX.XOM})} \stackrel{\text{def}}{=} \widehat{C}\{\widehat{F}_{\text{CVX}}(x_{i,\text{CVX}}), \widehat{F}_{\text{XOM}}(x_{i,\text{XOM}})\} \qquad z_{i,(\text{FP.RDSA})} \stackrel{\text{def}}{=} \widehat{C}\{\widehat{F}_{\text{FP}}(x_{i,\text{FP}}), \widehat{F}_{\text{RDSA}}(x_{i,\text{RDSA}})\}$$

$$
\begin{array}{l}
(\text{CVX.FP}) \rightsquigarrow \hat{\theta}_{(\text{CVX.FP})} \\
(\text{CVX.XOM}) \rightsquigarrow \hat{\theta}_{(\text{CVX.XOM})} \\
(\text{FP.RDSA}) \rightsquigarrow \hat{\theta}_{(\text{FP.RDSA})} \\
(\text{FP.XOM}) \rightsquigarrow \hat{\theta}_{(\text{FP.XOM})} \\
(\text{RDSA.XOM}) \rightsquigarrow \hat{\theta}_{(\text{RDSA.XOM})}
\end{array}
\quad
\underset{(\text{CVX.XOM})}{\text{best fit}} \Rightarrow
\begin{array}{l}
(\text{CVX.XOM})\text{FP} \rightsquigarrow \hat{\theta}_{(\text{CVX.XOM})\text{FP}} \\
(\text{CVX.XOM})\text{RDSA} \rightsquigarrow \hat{\theta}_{(\text{CVX.XOM})\text{RDSA}} \\
(\text{FP.RDSA}) \rightsquigarrow \hat{\theta}_{(\text{FP.RDSA})}
\end{array}
\quad
\underset{(\text{FP.RDSA})}{\text{best fit}} \Rightarrow
\begin{array}{l}
((\text{CVX.XOM})(\text{FP.RDSA})) \\
\rightsquigarrow \hat{\theta}_{((\text{CVX.XOM})(\text{FP.RDSA}))}
\end{array}
$$

Table 1: The estimation procedure in practice.

```
[1] "X"          "type"      "method"    "hac"       "epsilon"
[6] "agg.method" "margins"   "na.rm"     "max.min"   "..."
```

The whole procedure is divided in three (optional) computational blocks. First, the margins are specified. Secondly, the copula parameter, $\boldsymbol{\theta}$, is estimated and finally the HAC is checked for aggregation possibilities. The `margins` of the $(n \times d)$ data matrix, `X`, are assumed to follow the standard Uniform distribution by default, i.e., `margins = NULL`, but the function also permits non-uniformly distributed data as input if the argument `margins` is specified. The marginal distributions can be determined non-parametrically, `margins = "edf"`, or in a parametric way, e.g., `margins = "norm"`. Following the latter approach, the log-likelihood of the marginal `Distributions` is optimized with respect to the first (and second) parameter(s) of the density `dxxx`. Based on these estimates, the values of the univariate margins are computed. If the argument is defined as scalar, all margins are computed according to this specification. Otherwise, different margins can be defined, e.g., `margins = c("norm", "t", "edf")` for a three-dimensional sample. Except the Uniform distribution, all continuous `Distributions` of the **stats** package provided by the R Core Team (2013) are available: `"beta"`, `"cauchy"`, `"chisq"`, `"exp"`, `"f"`, `"gamma"`, `"lnorm"`, `"norm"`, `"t"` and `"weibull"`. The values of non-parametrically estimated distributions are computed according to Equation 3.

Inappropriate usage of this argument might lead to misspecified margins, e.g., `margins = "exp"` although the sample contains negative values. Even though the margins might be assumed to follow parametric distributions if `margins != NULL`, no joint log-likelihood is maximized, but the margins are estimated in advance. As the asymptotic theory works well for parametric and nonparametric estimation of margins, for the univariate analysis we refer to other built-in packages. In practice, the column names of `X` should be specified, as the default names `X1, X2, ...` are given otherwise.

A further optional argument of `estimate.copula` determines the estimation `method`. As discussed above, we present three procedures: `ML`, which is based on the bivariate density, full ML (`FML`) and recursive ML (`RML`) respectively. The routines of the **copula** package are imported if a simple Archimedean copula is fitted to the data, see Yan (2007); Kojadinovic and Yan (2010); Hofert and Mächler (2011).

At the final computational step of the procedure the binary HAC is checked for aggregation possibilities, if `epsilon > 0`. The new dependency parameter is computed according to the specification `agg.method`, i.e., the `"min"`, `"max"` or `"mean"` of the original parameters. To emphasize this point, recall the four-dimensional binary HAC

$$C(u_1, \ldots, u_4; (((12)3)4), \boldsymbol{\theta}) \quad = \phi_{\theta_3}\left\{\phi_{\theta_3}^{-1} \circ C\{u_1, \ldots, u_3; ((12)3), (\theta_1, \theta_2)^\top\} + \phi_{\theta_3}^{-1}(u_4)\right\}, \tag{6}$$

from Section 2. If we assume additionally $\theta_1 \approx \theta_2$, such that $\theta_1 - \theta_2 < \varepsilon$, the copula $C(\cdot)$ can be approximated by

$$C^*(u_1, \ldots, u_4; ((123)4), \boldsymbol{\theta}) \quad = \phi_{\theta_3}\left\{\phi_{\theta_3}^{-1} \circ C\{u_1, \ldots, u_3; (123), \theta^*\} + \phi_{\theta_3}^{-1}(u_4)\right\}, \tag{7}$$

where $\theta^* = (\theta_1 + \theta_2)/2$ for instance. This is referred to as the associativity property of Archimedean copulae, see Theorem 4.1.5 of Nelsen (2006). If the variables of two nodes are aggregated, the new copula is checked for aggregation possibilities as well. Beside this

Generator	Parameter	$\phi(u; \theta)$	$\tau(\theta)$
AMH	$\theta \in [0, 1)$	$(1 - \theta)/\{\exp(u) - \theta\}$	$1 - 2/(3\theta^2)\{\theta + (1 - \theta)^2 \log(1 - \theta)\}$
CLAYTON	$\theta \in (0, \infty)$	$(u + 1)^{-1/\theta}$	$\theta/(\theta + 2)$
FRANK	$\theta \in (0, \infty)$	$-\log[1 - \{1 - \exp(-\theta)\}\exp(-u)]/\theta$	$1 + 4/\theta\{D(\theta) - 1\}$
GUMBEL	$\theta \in [1, \infty)$	$\exp(-u^{1/\theta})$	$1 - 1/\theta$
JOE	$\theta \in [1, \infty)$	$1 - \{1 - \exp(-u)\}^{1/\theta}$	$1 - 4\sum_{\ell=1}^{\infty}[\ell(\theta\ell + 2)\{2 + \theta(\ell - 1)\}]^{-1}$

Table 2: Generator functions and the relations between the copula parameter and Kendall's $\tau(\cdot)$. The Debye function $D(\cdot)$ involved in $\tau(\cdot)$ for the family Frank is given by $D(\theta) = 1/\theta \int_0^\theta u/\{\exp(u) - 1\} du$.

threshold approach, the realized estimates $\hat{\theta}_1$ and $\hat{\theta}_2$ can obviously be used to test H_0 : $\theta_1 - \theta_2 = 0$, since the asymptotic distribution is known. On the other hand, this approach is extremely expensive computationally. The estimation results for the non-aggregated and the aggregated cases are presented in the following:

```
R> result.agg = estimate.copula(sample, margins = "edf", epsilon = 0.3)
R> plot(result, circles = 0.3, index = TRUE, l = 1.7)
R> plot(result.agg, circles = 0.3, index = TRUE, l = 1.7)
```


Figure 3: Plot of `result` on the left and `result.agg` on the right hand side.

3.1. The `hac` object

`hac` objects can be constructed by the general function `hac`, with the same name as the object it creates, and its simplified version `hac.full` for building fully nested HAC. For instance, consider the construction of a four-dimensional fully nested HAC with Gumbel generator, i.e.,

```
R> G.cop = hac.full(type  = HAC_GUMBEL,
+                   y     = c("X4", "X3", "X2", "X1"),
+                   theta = c(1.1, 1.8, 2.5))
R> G.cop

Class: hac
Generator: Gumbel
(((X1.X2)_{2.5}.X3)_{1.8}.X4)_{1.1}
```

where y denotes the vector of variables of class `character` and `theta` denotes the vector of dependency parameters. The parameters should be ascending ordered, so that the first parameter, `1.1`, is referred to the initial node of the HAC and the last parameter, `2.5`, corresponds to the first hierarchical level with variables `"X1"` and `"X2"`. The vector y has to contain one element more than the vector `theta`.

The returned output of `hac` objects is structured by three lines: (i) the object's `Class`, (ii) the `Generator` family and (iii) the HAC structure s. The structure can also be produced by the supplementary function `tree2str`. Variables, grouped at the same node are separated by a dot "." and the dependency parameters are printed within the curly parentheses.

Partially nested Archimedean copulae are constructed by `hac` with the main argument `tree`. For a better understanding let us first consider a four-dimensional simple Archimedean copula with dependency parameter $\theta = 2$:

```
R> hac(tree = list("X1", "X2", "X3", "X4", 2))
```

```
Class: hac
Generator: Gumbel
(X1.X2.X3.X4)_{2}
```

The copula `tree` is constructed by a `list` consisting of four `character` objects, i.e., `"X1"`, `"X2"`, `"X3"`, `"X4"`, and a number, which denotes the dependency parameter of the Archimedean copula. According to the theoretical construction of HAC in Section 2, we can induce structure by substituting margins through a subcopula. The four variables `"X1"`, `"X2"`, `"X3"`, `"X4"` can, for example, be structured by

```
R> hac(tree = list(list("X1", "X2", 2.5), "X3", "X4", 1.5))
```

```
Class: hac
Generator: Gumbel
((X1.X2)_{2.5}.X3.X4)_{1.5}
```

where the nested component, `list("X1", "X2", 2.5)`, is the subcopula at the lowest hierarchical level. Note that the nested component is of the same general form `list(...,  numeric(1))` as the simple Archimedean copula, where `numeric(1)` denotes the dependency parameter and "..." refers to arbitrary variables and subcopulae, which may contain subcopulae as well, like presented in the following:

```
R> HAC = hac(tree = list(list("Y1", list("Z3", "Z4", 3), "Y2", 2.5),
+                        list("Z1", "Z2", 2), list("X1", "X2", 2.4),
+                        "X3", "X4", 1.5))
R> HAC
```

```
Class: hac
Generator: Gumbel
((Y1.(Z3.Z4)_{3}.Y2)_{2.5}.(Z1.Z2)_{2}.(X1.X2)_{2.4}.X3.X4)_{1.5}
```
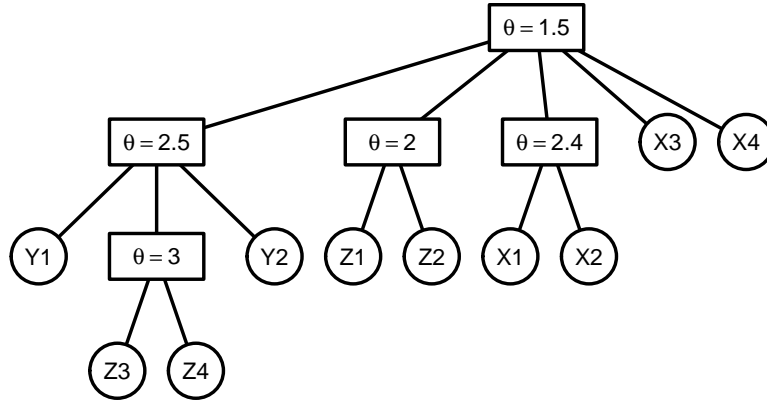
Figure 4: Plot of the object `HAC`.

We cannot avoid the notation becoming more cumbersome for higher dimensions, but the principle stays the same for arbitrary dimensions, i.e., variables are substituted by `list`s of the general form `list(..., numeric(1))`. The function `hac` provides a further argument for specifying the `type` of the HAC.

3.2. Graphics

As the string representation of the structure becomes more unclear as dimension increases, the package allows to produce graphics of `hac` objects by the generic `plot` function. Figure 4 illustrates for example the dependence structure of the lastly defined object `HAC`.

```
R> plot(HAC, cex = 0.8, circles = 0.35)
```

The explanatory power of these plots can be enhanced by several of the usual `plot` parameters:

```
R> names(formals(plot.hac))
```

```
 [1] "x"         "xlim"       "ylim"      "xlab"      "ylab"
 [6] "col"       "fg"         "bg"        "col.t"     "lwd"
[11] "index"     "numbering"  "theta"     "h"         "l"
[16] "circles"   "digits"     "..."
```

The optional, boolean argument `theta` determines whether the dependency parameter of the copula θ or Kendall's τ is printed, whereby Kendall's τ cannot be easily interpreted in the usual way for more than two dimensions. The supplementary function `theta2tau` computes Kendall's rank correlation coefficient based on the value of the dependency parameter, whereas `tau2theta` corresponds to the inverse function, see Table 2. If `index = TRUE`, strings illustrating the subcopulae of the nodes are used as subscripts of the dependency parameters. If, additionally, `numbering = TRUE`, the parameters are numbered, such that the subscripts correspond to the estimation stages if the non-aggregated output of `estimate.copula` is plotted. The radius of the `circles`, the width `l` and the height `h` of the rectangles and the specific

colors of the lines and the text can be adjusted. Further arguments "`...`" can, for example, be used to modify the font size `cex` or to include a subtitle `sub`.

3.3. Random sampling

To be in line with other R-packages providing tools for different univariate and multivariate distributions we provide: (i) `dHAC` for computing the values of the copula density, (ii) `pHAC` for the cumulative distribution function and (iii) `rHAC` for simulations. Sampling methods are imported from the **copula** package and rely on the algorithm suggested in Hofert and Mächler (2011), who summarize the procedure for the former **nacopula** package as follows:

Algorithm 1. *Let $C(\cdot)$ be a nested Archimedean copula with root copula $C_0(\cdot)$ generated by ϕ_0. Let U be a vector of the same dimension as $C_0(\cdot)$.*

  1. *sample from inverse Laplace transform $\mathcal{LS}^{-1}$ of ϕ_0, i.e., $V_0 \sim F_0(\cdot) \stackrel{\text{def}}{=} \mathcal{LS}^{-1}(\phi_0)$*

  2. *for all components u of $C_0(\cdot)$ that are nested Archimedean copulae do:*

      (a) *set $C_1(\cdot)$ with generator $\phi_1(\cdot)$ to the nested Archimedean copula u*

      (b) *sample $V_{01} \sim F_{01}(\cdot) \stackrel{\text{def}}{=} \mathcal{LS}^{-1}\{\phi_{01}(\cdot; V_0)\}$*

      (c) *set $C_0(\cdot) \stackrel{\text{def}}{=} C_1(\cdot), \phi_0(\cdot) \stackrel{\text{def}}{=} \phi_1(\cdot)$, and $V_0 \stackrel{\text{def}}{=} V_{01}$ and continue with 2.*

  3. *for all other components u of $C_0(\cdot)$ do*

      (a) *sample $R \sim Exp(1)$*

      (b) *set the component of U corresponding to u to $\phi_0(R/V_0)$*

  4. *return U*

The function `rHAC` requires only two arguments: (i) the sample size `n` and (ii) an object of the class `hac` specifying the characteristics of the underlying HAC, e.g.,

```
R> sim.data = rHAC(500, G.cop)
R> pairs(sim.data, pch = 20)
```

In particular, the contributions of McNeil (2008), Hofert (2008) and Hofert (2011) provide the theoretical foundations to sample computationally efficient random vectors from HAC. Algorithm 1 exploits the recursively determined structure of HAC and samples from F_0 and F_{01}, which are comprehensively discussed in Hofert (2011) and Hofert and Mächler (2011).

3.4. The cdf and density

The arguments for `pHAC` are a `hac` object and a sample `X`, whose column names should be identical to the variables' names of the `hac` object, e.g.,

```
R> probs = pHAC(X = sim.data, hac = G.cop)
```

Figure 5: Scatterplot of the sample `sim.data`, which is simulated from `G.cop` associated with a four dimensional HAC-based Gumbel copula.

As the copula density is defined as d-th derivative of the copula $C(\cdot)$ with respect to the arguments u_j, $j = 1, \ldots, d$, c.f. Savu and Trede (2010), the explicit form of the density varies with the structure of the underlying HAC. Hence, including the explicit form of all possible d-dimensional copula densities is absolutely unrealistic. Our function `dHAC` derives an analytical expression of the density for a given `hac` object, which can be instantaneously evaluated if `eval = TRUE`. The analytical expression of the density is found by subsequently using the `D` function to differentiate the algebraic form of the copula "symbolically" with respect to the variables of the inserted `hac` object. Although the derivation and evaluation of the density is computationally and numerically demanding, `dHAC` provides a flexible way to work with HAC densities in practice, because they do not need to be manually derived or numerically approximated. Since the densities of the two-dimensional Archimedean copulae are frequently called during the pseudo multi-stage estimation procedure (`ML`), their closed form expressions are given explicitly.

3.5. Empirical copula

As long as our package does not cover goodness-of-fit tests, which are difficult to implement in general and involve computational intensive techniques via bootstrapping, see Genest, Rémillard, and Beaudoin (2009), it might be difficult to justify the choice of a parametric assumption. However, the values of `probs` can be compared to those of the empirical copula,

Figure 6: The values of `probs` on the x-axis against the values of `probs.emp`.

i.e.,

$$\widehat{C}(u_1, \ldots, u_d) \;\; = \;\; n^{-1} \sum_{i=1}^{n} \prod_{j=1}^{d} \mathbf{I}\left\{\widehat{F}_j\left(X_{ij}\right) \leq u_j\right\}, \tag{8}$$

where $\widehat{F}_j(\cdot)$ denotes the estimated marginal distribution function of variable X_j. Figure 6 suggests a proper fit of the empirical copula computed by

```
R> probs.emp = emp.copula.self(sim.data, proc = "M")
```

There are two functions which can be used for computing the empirical copula:

```
R> emp.copula(u, x, proc = "M", sort = "none", margins = NULL,
+             na.rm = FALSE, ...)
R> emp.copula.self(x, proc = "M", sort = "none", margins = NULL,
+                  na.rm = FALSE, ...)
```

The difference between the arguments of these functions is that `emp.copula` requires a matrix `u`, at which the estimated function is evaluated. This can, in particular, be helpful, when the sample is decomposed to evaluate the out-of-sample performance as the empirical copula can be regarded as natural benchmark. In contrast, `emp.copula.self` evaluates $\widehat{C}(\cdot)$ at the sample `x` used for the estimation and thus, the returned values can be considered as in-sample fit. The argument `proc` enables the user to choose between two computational methods. We

Figure 7: The figure shows the computational times for an increasing sample size but a fixed dimension $d = 5$ on a log-log scale. The solid line is referred to `proc = "M"` and the dashed line to `proc = "A"`.

recommend to use the default method, `proc = "M"`, which is based on `matrix` manipulations, because its computational time is just a small fraction of the taken time of method `"A"`, which is based on `apply`, see Figure 7. However, method `"M"` is sensitive with respect to the size of the working memory and therefore inapplicable for very large datasets. Note that standard applications, e.g., measuring the VaR of a portfolio, are based on 250 or 500 observations. Figure 7 illustrates rapidly increasing computational times of the `matrix` based method for more than 5000 observations until the method collapses. In contrast, the runtimes of the alternative method `proc = "A"` are more robust against an increasing sample size. The computational times are less sensitive with respect to the dimension and we recommend using the default method up to $d = 100$ for non-large sample sizes. Another possibility to deal with large datasets is specifying the matrix `u` manually in order to reduce the number of vectors which are to be evaluated.

4. Simulation study

To ensure the accuracy of the proposed methods, we generate random data from six copula models of different dimension $C_i^j \stackrel{\text{def}}{=} C^j(\cdot; s_i, \boldsymbol{\theta}_i)$, for $i = 1, 2, 3$ and $j = \mathrm{C}, \mathrm{G}$, and show that the estimates almost coincide with the true model specification. Here, j denotes the copula family (Clayton or Gumbel) and the structures are given by $s_1 = ((12)3)$, $s_2 = ((((12)3)4)5)$ and $s_3 = ((12)(34)5)$. The values of $\boldsymbol{\theta}_i$ are presented in Table 3 and Table 4. They are chosen such that a similar strength of dependence is produced by the Clayton and Gumbel based models.

The summary statistics of Table 3 and Table 4 rely on $n = 1000$ estimates, whereby only estimates with the same structure can be compared. For this reason the procedure was m times replicated till $n = 1000$ estimates were available. As `estimate.copula` approximates

Model	θ	Statistics				
		min	median	mean	max	sd
C_1^{G}	$\theta_2 = 1.500$	1.31	1.49	1.50	1.77	0.07
	$\theta_1 = 3.000$	2.56	3.00	3.00	3.51	0.15
C_1^{C}	$\theta_2 = 1.000$	0.61	0.99	1.00	1.47	0.13
	$\theta_1 = 4.000$	3.13	4.02	4.03	5.09	0.28
C_2^{G}	$\theta_4 = 1.125$	1.00	1.10	1.11	1.24	0.04
	$\theta_3 = 1.500$	1.28	1.45	1.45	1.71	0.07
	$\theta_2 = 2.250$	1.92	2.24	2.24	2.81	0.12
	$\theta_1 = 4.500$	3.86	4.49	4.50	5.20	0.23
C_2^{C}	$\theta_4 = 0.250$	0.00	0.17	0.17	0.39	0.07
	$\theta_3 = 1.000$	0.54	0.92	0.92	1.28	0.12
	$\theta_2 = 2.500$	1.88	2.48	2.50	3.21	0.20
	$\theta_1 = 7.000$	5.77	7.02	7.04	8.91	0.45

Table 3: The models for the Gumbel family C_1^{G}, C_2^{G} and for the Clayton family C_1^{C}, C_2^{C}, where θ denotes the true copula parameters.

the true structure, we set `epsilon = 0.15` for C_3^{G} and `epsilon = 0.20` for C_3^{C}, which are not based on a binary structure and employ the `RML` procedure. Note that the `RML` procedure attempts at aggregating the copula tree after each estimation step. The simulated samples for the copula estimation consist of 250 observations for the copula types in order to illustrate the finite sample properties of the procedures. Table 3 and Table 4 indicate, that the estimation procedure works properly for the suggested models, as the estimates are on average consistent with the true parameters. Nevertheless, a few points deserve being mentioned: (i) The multi-stage procedure detects the true structure for the binary HAC in $n/m = 100\%$ and the recursive ML procedure for the non-binary HAC in at least $n/m = 99\%$ of the cases as long as the parameters exhibit the imposed distance and the permutation symmetry of the variables at the same node is taken into consideration. (ii) The estimates at lower hierarchical levels show a higher volatility than the estimates close to the initial node and the estimates for the Clayton models are more volatile than the estimates of the Gumbel based HAC. (iii) All estimated models indicate more imprecise estimates for higher nesting levels, but the gains from full ML estimation regarding the precision are only observable for the estimates at the root node of C_3^{G}, see Table 4. However, this minor improvement is costly since the results are based on a preestimated structure. (iv) These observations justify choosing different values of `epsilon` for C_3^{G} and C_3^{C}, as the tuning parameter should reflect the variability of the parameters. Theoretically, `epsilon` can be different for each aggregation of the structure so that the parameter variability is correctly represented. This, however, becomes infeasible in practice, because the number of nodes contained in the true structure is generally unknown. If the parameters are closer and/or the value of `epsilon` is chosen smaller, the amount of correctly classified structures declines. On the other hand, larger sample sizes permit smaller values of `epsilon` as the parameters are more precisely estimated.

5. Conclusion

Model	θ	Statistics for recursive ML						
		$\bar{s}$	min	median	mean	max	sd
C_3^{G}	$\theta_3 = 1.125$	$((12)(534)) = 00.20\%$	1.00	1.10	1.11	1.24	0.03
	$\theta_2 = 1.500$	$((12)((34)5)) = 00.40\%$	1.31	1.50	1.50	1.77	0.07
	$\theta_1 = 3.000$	$((12)534) = 00.10\%$	2.59	3.01	3.01	3.59	0.15
C_3^{C}	$\theta_3 = 0.250$	$((12)((34)5)) = 00.30\%$	0.05	0.25	0.25	0.47	0.06
	$\theta_2 = 1.000$	$(((34)(12))5) = 00.39\%$	0.63	1.00	1.01	1.53	0.12
	$\theta_1 = 4.000$		3.19	4.00	4.02	4.94	0.29
		Statistics for full ML						
C_3^{G}	$\theta_3 = 1.125$		1.03	1.13	1.13	1.23	0.03
	$\theta_2 = 1.500$	$-$	1.31	1.50	1.50	1.77	0.07
	$\theta_1 = 3.000$		2.59	3.01	3.01	3.60	0.15
C_3^{C}	$\theta_3 = 0.250$		0.09	0.25	0.25	0.47	0.05
	$\theta_2 = 1.000$	$-$	0.64	1.00	1.01	1.53	0.12
	$\theta_1 = 4.000$		3.20	3.99	4.02	4.94	0.29

Table 4: The model for the Gumbel family C_3^{G} and for the Clayton family C_3^{C}, where θ denotes the true copula parameters and the column $\bar{s}$ refers to the percentage of incorrectly classified structures based on $n = 1000$ replications.

The **HAC** package focuses on the computationally efficient estimation of hierarchical Archimedean copula, which is based on grouping binary structures within a recursive multi-stage ML procedure. Its theoretical and practical advantages are (i) avoiding the demanding asymptotic theory, which arises due to constrained one-step ML estimation and (ii) the consecutive optimization of the log-likelihood instead of the singular optimization of the d-dimensional one with respect to several parameters. Since HAC permit modeling large-dimensional random vectors, the package provides a function for producing graphics of the related `hac` objects. According to the usual naming of distributions in R, we provide `dHAC`, `pHAC` and `rHAC` to compute the values of density- and distribution functions or to sample from arbitrary HAC. Finally, the accuracy of the methods has been shown in a small simulation study.

Acknowledgments

References

Acar EF, Craiu RV, Yao F (2011). "Dependence Calibration in Conditional Copulas: A Nonparametric Approach." *Biometrics*, **67**(2), 445–453.

Bárdossy A (2006). "Copula-Based Geostatistical Models for Groundwater Quality Parameters." *Water Resources Research*, **42**(11), 1–12.

Bárdossy A, Li J (2008). "Geostatistical Interpolation Using Copulas." *Water Resources Research*, **44**(7), 1–15.

Chen X, Fan Y (2006). "Estimation and Model Selection of Semiparametric Copula-Based Multivariate Dynamic Models under Copula Misspecification." *Journal of Econometrics*, **135**, 125–154.

Cherubini U, Luciano E, Vecchiato W (2004). *Copula Methods in Finance*. John Wiley & Sons, New York.

Embrechts P, Lindskog F, McNeil AJ (2003). "Modeling Dependence with Copulas and Applications to Risk Management." In ST Rachev (ed.), *Handbook of Heavy Tailed Distributions in Finance*. Elsevier, North-Holland.

Embrechts P, McNeil AJ, Straumann D (1999). "Correlation and Dependence in Risk Management: Properties and Pitfalls." In *Risk Management: Value at Risk and Beyond*, pp. 176–223. Cambridge University Press.

Franke J, Härdle WK, Hafner C (2011). *Statistics of Financial Markets: An Introduction*. Universitext, 3 edition. Springer-Verlag.

Genest C, Favre AC (2007). "Everything You Always Wanted to Know about Copula Modeling but Were Afraid to Ask." *Journal of Hydrologic Engineering*, **12**(4), 347–368.

Genest C, Rémillard B, Beaudoin D (2009). "Goodness-of-Fit Tests for Copulas: A Review and a Power Study." *Insurance: Mathematics and Economics*, **44**(2), 199–213.

Hofert M (2008). "Sampling Archimedean Copulas." *Computational Statistics & Data Analysis*, **52**, 5163–5174.

Hofert M (2011). "Efficiently Sampling Nested Archimedean Copulas." *Computational Statistics & Data Analysis*, **55**, 57–70.

Hofert M, Kojadinovic I, Mächler M, Yan J (2013). *copula: Multivariate Dependence with Copulas*. R package version 0.999-7, URL http://CRAN.R-project.org/package=nacopula.

Hofert M, Mächler M (2011). "Nested Archimedean Copulas Meet R: The **nacopula** Package." *Journal of Statistical Software*, **39**(9), 1–20. URL http://www.jstatsoft.org/v39/i09/.

Hofert M, Mächler M (2012). *nacopula: Nested Archimedean Copulas*. R package version 0.8-1, URL http://CRAN.R-project.org/package=nacopula.

Jaworski P, Durante F, Härdle WK (2013). *Copulae in Mathematical and Quantitative Finance*, volume 213 of *Lecture Notes in Statistics*. Springer.

Joe H (1996). "Families of m-Variate Distributions with Given Margins and $m(m-1)/2$ Bivariate Dependence Parameters." In L Rüschendorf, B Schweizer, M Taylor (eds.), *Distribution with Fixed Marginals and Related Topics*, IMS Lecture Notes – Monograph Series. Institute of Mathematical Statistics.

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

Junker M, May A (2005). "Measurement of Aggregate Risk with Copulas." *Econometrics Journal*, **8**(3), 428–454.

Kojadinovic I, Yan J (2010). "Modeling Multivariate Distributions with Continuous Margins Using the **copula** R Package." *Journal of Statistical Software*, **34**(9), 1–20. URL http://www.jstatsoft.org/v34/i09/.

Lakhal-Chaieb ML (2010). "Copula Inference under Censoring." *Biometrika*, **97**(2), 505–512.

Li DX (2000). "On Default Correlation: A Copula Function Approach." *Journal of Fixed Income*, **9**(4), 43–54.

McNeil AJ (2008). "Sampling Nested Archimedean Copulas." *Journal of Statistical Computation and Simulation*, **78**, 567–581.

McNeil AJ, Nešlehová J (2009). "Multivariate Archimedean Copulas, d-Monotone Functions and l_1 Norm Symmetric Distributions." *The Annals of Statistics*, **37**(5b), 3059–3097.

McNeil AJ, Ulman S (2011). ***QRMlib**: Provides R-Language Code to Examine Quantitative Risk Management Concepts.* R package version 1.4.5.1, URL http://cran.r-project.org/src/contrib/Archive/QRMlib/.

Mendes BVM, Souza RM (2004). "Measuring Financial Risks with Copulas." *International Review Of Financial Analysis*, **13**, 53–77.

Nelsen RB (2006). *An Introduction to Copulas.* 2 edition. Springer Verlag, New York.

Okhrin O, Okhrin Y, Schmid W (2013a). "On the Structure and Estimation of Hierarchical Archimedean Copulas." *Journal of Econometrics*, **173**, 189–204.

Okhrin O, Okhrin Y, Schmid W (2013b). "Properties of Hierarchical Archimedean Copulas." *Statistics & Risk Modeling*, **30**(1), 21–54.

Pfaff B, McNeil A (2013). ***QRM**: Provides R-language Code to Examine Quantitative Risk Management Concepts.* R package version 0.4-9, URL http://CRAN.R-project.org/package=QRM.

R Core Team (2013). *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Savu C, Trede M (2010). "Hierarchies of Archimedean Copulas." *Quantitative Finance*, **10**(3), 295–304.

Sklar A (1959). "Fonctions de Répartition à n Dimension et Leurs Marges." *Publications de l'Institut de Statistique de l'Université de Paris*, **8**, 299–231.

Whelan N (2004). "Sampling from Archimedean Copulas." *Quantitative Finance*, **4**, 339–352.

Wuertz D, many others, see the SOURCE file (2013). ***fCopulae**: Rmetrics - Dependence Structures with Copulas.* R package version 3000.79, URL http://CRAN.R-project.org/package=fCopulae.

Yan J (2007). "Enjoy the Joy of Copulas: With a Package **copula**." *Journal of Statistical Software*, **21**(4), 1–21. URL http://www.jstatsoft.org/v21/i04/.

Affiliation:

Ostap Okhrin
C.A.S.E - Center for Applied Statistics and Economics
Ladislaus von Bortkiewicz Chair of Statistics
Humboldt-Universität zu Berlin
10099 Berlin, Germany
E-mail: ostap.okhrin@wiwi.hu-berlin.de
URL: http://lehre.wiwi.hu-berlin.de/Professuren/quantitativ/statistik/members/personalpages/o2

Alexander Ristig
C.A.S.E - Center for Applied Statistics and Economics
Ladislaus von Bortkiewicz Chair of Statistics
Humboldt-Universität zu Berlin
10099 Berlin, Germany
E-mail: alexander.ristig@hu-berlin.de
URL: http://lehre.wiwi.hu-berlin.de/Professuren/quantitativ/statistik/members/personalpages/ar