Modeling multivariate operational losses via copula-based distributions with g-and-h marginals

Marco Bee
Department of Economics and Management, University of Trento, Italy

Julien Hambuckers
Department of Finance, HEC Liège, University of Liège, Belgium

September 21, 2021

Abstract. We propose a family of copula-based multivariate distributions with g-and-h marginals. After studying the properties of the distribution, we develop a two-step estimation strategy and analyze via simulation the sampling distribution of the estimators. The methodology is used for the analysis of a 7-dimensional dataset containing 40,871 operational losses. The empirical evidence suggests that a distribution based on a single copula is not flexible enough, thus we model the dependence structure by means of vine copulas. We show that the approach based on regular vines improves the fit. Moreover, even though losses corresponding to different event types are found to be dependent, the assumption of perfect positive dependence is not supported by our analysis. As a result, the Value-at-Risk of the total operational loss distribution obtained from the copula-based technique is substantially smaller at high confidence levels, with respect to the one obtained using the common practice of summing the univariate Value-at-Risks.

JEL codes. C460; C630.

Keywords. Loss model; dependence structure; vine copula; Value-at-Risk.

Key messages.

- We develop a two-step estimation strategy for copula-based distributions with g-and-h marginals.
- We fit a 7-dimennsional loss dataset containing 40,871 operational losses.
- We find that the dependence structure should be modeled by means of vine copulas.
- Losses corresponding to different event types are found to be dependent.

1 Introduction

The empirical analysis of financial data suggests that the normal distribution is often inadequate. In a univariate setup, it cannot be used for asymmetric and/or leptokurtic data, whereas in a multidimensional framework the dependence structure of the multivariate normal distribution implies tail independence, an unrealistic assumption (Donnelly and Embrechts, 2010). Hence, the presence of skewness, excess kurtosis and tail dependence in the data may lead to non-negligible inaccuracies in the estimation of joint tail probabilities, with serious consequences in terms of financial stability. A well-known example is the underestimation of joint default probabilities caused by the uncritical application of the Li model (Li, 2000) for the pricing of Credit Default Obligations during the Great Financial Crisis.

Thus, the need for distributions that incorporate skewed and/or heavy-tailed marginals as well as less standard dependence structures has become more and more evident over the years. In order to construct models that take into account these requirements, the statistical community has explored two directions.

First, new multivariate distributions have been developed; one such example is the family of normal mixture distributions, and especially the subclass of multivariate generalized hyperbolic distributions (MGHD) introduced by Barndorff-Nielsen and Kendall (1977).

Second, copula-based distributions have become increasingly popular, because they allow the investigator to separately model and estimate the marginals and the dependence structure. This feature gives a substantial modeling flexibility as well as a lighter computational burden. In particular, estimating sequentially the parameters of each marginal and of the copula usually simplifies computations with respect to the single large-dimensional optimization required by classical multivariate estimation procedures.

In this paper we propose a family of copula-based multivariate distributions with g-and-h marginals that is well suited for operational risk measurement for at least three reasons. First, on the empirical side, univariate operational losses are often modeled by means of the g-and-h distribution: see Dutta and Babbel (2002); Moscadelli (2004); Peters and Sisson (2006); Fischer et al. (2007); Jiménez and Arunachalam (2011); Cruz et al. (2015); Peters et al. (2016); Bee and Trapin (2016); Bee et al. (2019, 2021). Second, Degen et al. (2007) show that the g-and-h converges very slowly to the Generalized Pareto Distribution (GPD), so that the Peaks-over-Threshold approach may be inaccurate if the data follow a g-and-h distribution. Moreover, fitting the tail is not enough for the computation of Capital-at-Risk (CaR), which is defined as the difference between an high quantile and the expected value. Hence, a GPD-based approach requires a spliced distribution to compute CaR, whereas the g-and-h is a model of the whole distribution. Third, investigating the dependence structure of operational losses is an important task that has not received much attention in the literature, even though some models are available (Böcker and Klüppelberg, 2010; Cruz et al., 2015, Chapter 12). To the best of our knowledge, only a handful of studies investigate this problem: Chapelle et al. (2008) use a splicing approach combined with Fréchet copulas to study the joint distribution of a four-dimensional dataset. Gourier et al. (2009) investigate the dependence of operational losses across four business lines of a single bank using Gaussian, Student and Frank copulas. Aulbach et al. (2012) illustrate the usefulness of a piece-wise copula approach in the bivariate case. Brechmann et al. (2014) estimate first the marginal distributions of event type total losses with simulated compound processes, before fitting seven-dimensional Archimedean and vine copulas. The use of vine copulas for portfolio management and VaR estimation has been considered by Low et al. (2013), Low et al. (2016) and Low (2018).

Until recently, the absence of a closed-form density hindered the use of g-and-h distributions. However, new estimation methods exploiting indirect inference and numerical maximum likelihood techniques were recently proposed in Bee et al. (2019) and Bee et al. (2021), making it easy to rely on g-and-h distributions in applied work. These methods are computationally affordable and were shown to improve the fit with respect to the traditional quantile-based estimation procedure proposed in Hoaglin (1985). The construction of a multivariate version of the g-and-h distribution is still an open topic, though. The only attempt was made by Field and Genton (2006). However, they model the dependence structure via linear correlation, so that nonlinear dependencies in the data are not captured. In addition, their fitting procedure is based on multivariate quantiles. This raises computational and identification challenges since the estimators depend on the definition of multivariate quantiles. The copula-based g-and-h distribution developed in the present paper aims at solving these shortcomings.

In the second part of our article, we conduct an empirical study using the proposed approach: we estimate the multivariate distribution of operational losses recorded at Unicredit, one of Italy's most important banks, between 2005 and 2014. The data, first used in Hambuckers et al. (2018), come from all the seven event types listed in the Basel II Accord. Therefore we are able to investigate the joint distribution of event type-specific operational losses, at a monthly frequency. We find the univariate fit of the marginal g-and-h distributions to be excellent, whereas models for the dependence structure across event types based on elliptical copulas do not pass the goodness-of-fit tests. Hence, we explore a more general approach based on regular vine (R-vine) copulas. The Value-at-Risk (VaR) of the total loss distribution is obtained by means of simulations from the Gaussian copula, the R-vine copula and the perfect dependence approach. The two copula-based methods produce smaller VaRs, especially at high coverage levels.

This paper makes a twofold contribution. On the theoretical side, we propose a new multivariate distribution for dependent loss data, which is both flexible and easy to estimate. On the empirical side, we study the dependence structure of a large database of operational losses and its relation to the VaR. We find that event type-specific operational losses are dependent but not perfectly, and that a precise estimate of the dependence structure has a non-negligible impact on the estimated VaR for the total loss distribution, a result consistent with findings in Brechmann et al. (2014).

Finally, our work has also several implications for practitioners and regulators. For banks willing to quantify their operational risk level and to derive appropriate capital reserves, our results emphasize the need to account for both the marginal tail risks and the correlation structure across event types to obtain economically meaningful measures. In particular, a flexible parametric density like the g-and-h distribution must be favored for the marginal distributions, to reflect properly the probability of (marginal) extreme events and to not underestimate capital reserves. Second, the use of a flexible copula structure should be favored over a perfect dependence (resp. independence) approach which could lead to an overestimation (resp. underestimation) of the required capital for high confidence levels. Similarly, for regulators, our results demonstrate the potential of the g-and-h distribution to be a valid loss distribution candidate for CaR calculations (BCBS, 2016). Consequently, the g-and-h distribution could be considered in future versions of the revised Standardized Approach.

The rest of this work is organized as follows. In Section 2, after a brief review the g-and-h distribution, we introduce the family of copula-based multivariate g-and-h distributions, study its properties and develop estimation and simulation procedures. In Section 3 we show the results of some simulation experiments. In Section 4 we analyze a large operational loss dataset. Finally, Section 5 concludes and outlines some directions for future research. Further results are reported in the online supplementary material.

2 A multivariate g-and-h distribution

The g-and-h distribution (Tukey, 1977) is obtained by modifying a standard normal random variable via the following non-linear transformation:

$$X = a + b\frac{e^{gZ} - 1}{g} e^{\frac{hZ^2}{2}}, \quad Z \sim N(0, 1), \tag{1}$$

where $a \in \mathbb{R}$ is a location parameter, $b \in \mathbb{R}^+$ is a scale parameter, $g \in \mathbb{R}$ and $h \in \mathbb{R}^+$ are shape parameters. From the perspective of our application, X refers here to the sum of all operational losses related to a given event type, over a particular time period. Thanks to the two shape parameters, its skewness-kurtosis region is larger with respect to other size distributions (Dutta and Perry, 2006, Figure 3). The support of X is the entire real line, so that it is important to consider parameter settings that minimize the probability of negative values (Cruz et al., 2015, p. 311).

The joint distribution studied in this paper has g-and-h marginals and a dependence structure defined by a copula. Without loss of generality, in the following we assume that the d marginals are standardized, i.e. $(a_j, b_j) =$

$(0,1)$, $j = 1,\dots,d$. In practical applications one has n d-dimensional observations $\tilde{\boldsymbol{x}}_1,\dots,\tilde{\boldsymbol{x}}_n$ stored in a data matrix $\tilde{\boldsymbol{X}} = \{\tilde{x}\}_{ij}$, $i = 1,\dots,n$, $j = 1,\dots,d$, where $\tilde{x}_{ij}$ is the i-th observation from the j-th marginal $\tilde{X}_j \sim \mathrm{gh}(a_j, b_j, g_j, h_j)'$. When $(a_j, b_j) \neq (0,1)$, the observations $\tilde{x}_{ij}$ must be standardized in advance by means of the transformation $(\tilde{x}_{ij} - \tilde{a}_j)/\tilde{b}_j$, where $\tilde{a}_j$ and $\tilde{b}_j$ are the Hoaglin (1985) quantile estimators of a_j and b_j.

The general representation of the joint cumulative distribution function (cdf) $F_{\boldsymbol{X}}$ of the proposed distribution is

$$F_{\boldsymbol{X}}(x_1,\dots,x_d;\boldsymbol{\eta},\boldsymbol{\theta}) = C_{\boldsymbol{\eta}}(F_1(x_1;\boldsymbol{\theta}_1),\dots,F_d(x_d;\boldsymbol{\theta}_d)), \tag{2}$$

where $C_{\boldsymbol{\eta}}$ is the cdf of a d-dimensional copula with parameter $\boldsymbol{\eta}$, $\boldsymbol{\theta}_j = (0,1,g_j,h_j)'$ and $F_j(\cdot;\boldsymbol{\theta}_j)$ is the cdf of the g-and-h random variable $X_j \sim \mathrm{gh}(0,1,g_j,h_j)$, $j = 1,\dots,d$. The parameter $\boldsymbol{\eta}$ is specific to the chosen copula: for example, in the Gaussian and Student-t case, it is a correlation matrix, denoted $\boldsymbol{R}$ later on.

The key issue is that the g-and-h cdf is not known in closed form. Bee et al. (2021) have proposed an approximation exploiting the quantile function, explicitly given by

$$Q(p;\boldsymbol{\theta}) = a + b\frac{e^{gz_p} - 1}{g}e^{\frac{hz_p^2}{2}}, \tag{3}$$

where $p \in (0,1)$ and z_p is the p-quantile of the standard normal distribution.

Assuming that the random variable V has density f_V and $h(v)$ is a differentiable 1-1 transformation, the density of $W = h(V)$ is equal to

$$f_W(w) = \frac{f_V(v)}{h'(v)}, \quad \text{where } v = h^{-1}(w). \tag{4}$$

From (1) and (4), and setting $h(\cdot)$ equal to the quantile function (3), the approximation of the density is

$$f(x) = \frac{\phi(p)}{Q'(p;\boldsymbol{\theta})}, \quad p = F(x;\boldsymbol{\theta}), \tag{5}$$

where $\phi(\cdot)$ is the standard normal density, $F(x;\boldsymbol{\theta})$ is the g-and-h distribution function and $Q'(p;\boldsymbol{\theta})$ is the derivative of (3), explicitly given by (Cruz et al., 2015, Eq. 9.33):

$$Q'(p;\boldsymbol{\theta}) = e^{gp + \frac{hp^2}{2}} + \frac{h}{g}pe^{\frac{hp^2}{2}}(e^{gp} - 1).$$

Since $F(x;\boldsymbol{\theta})$ in (5) is not known explicitly, we evaluate it pointwise by inverting numerically (3) via a root-finding technique. Here is where the approximation comes in, but this is a rather easy (yet time-consuming) numerical problem that can be solved by means of standard optimization routines; see Bee et al. (2021) for details.

In principle, any copula can be used in (2). For example, Chapelle et al. (2008) use Fréchet copulas, Gourier et al. (2009) consider Frank copulas whereas Brechmann et al. (2014) rely on Archimedean and vine copulas. The decision should be motivated by the dependence structure of the data at hand, as well as by the computational burden and analytical tractability of the chosen copula, especially in high dimensions. We will discuss this issue in Section 4[1].

2.1 Properties

Basic properties of copulas imply that the marginal distributions of (2) are $\mathrm{gh}(0, 1, g_i, h_i)$, regardless of the copula C. Hence, the expected value and variance of the i-th marginal are given by the formulas derived by Martinez and Iglewicz (1984); see also Cruz et al. (2015, Proposition 9.3). From (2) we see that, when using the Gaussian or Student-t copula, the correlation matrix $\boldsymbol{R}$ incorporates all the information about the dependence between the variables. However, in general it is preferable to describe the dependence structure by means of rank correlations, commonly measured by the Kendall's τ (ρ_τ) and Spearman's ρ (ρ_S) correlation coefficients; see McNeil et al. (2015, Sect. 7.2.3). Both measures are direct functionals of the copula and in various important cases there is a closed form-relationship between ρ_τ and the copula parameter. For example, for the Gaussian(ρ) and Student-$t(\rho)$ copulas $\rho_\tau = (2/\pi)\arcsin\rho$ (Demarta and McNeil, 2005), for the Gumbel(α) copula $\rho_\tau = 1 - 1/\alpha$ (McNeil et al., 2015, Table 7.5).

Another relevant concept is extremal dependence. The coefficients of upper and lower tail dependence measure asymptotic dependence in the upper (lower) tail of the bivariate distribution of (X_1, X_2). Assuming that X_1 and X_2 are continuous random variables with cdf F_1 and F_2 respectively, the coefficients are defined as

$$\lambda_u = \lim_{q\to 1^-} P(X_2 > F_2^{\leftarrow}(q)|X_1 > F_1^{\leftarrow}(q)) = \lim_{q\to 1^-} \frac{\bar{C}(q,q)}{1-q},$$
$$\lambda_\ell = \lim_{q\to 0^+} P(X_2 \leq F_2^{\leftarrow}(q)|X_1 \leq F_1^{\leftarrow}(q)) = \lim_{q\to 0^+} \frac{C(q,q)}{1-q},$$

[1]A natural competitor to the generic equation (2) is the multivariate version of the g-and-h distribution introduced in Field and Genton (2006), who also propose an estimation method exploiting multivariate quantiles, defined in terms of a norm minimization problem. The estimators depend on the chosen norm, leading to non-trivial computational challenges (Field and Genton, 2006, p. 105), and dependence is only modeled via linear correlation (Field and Genton, 2006, p. 107). For comparison purposes, we have fitted our model to one of the datasets used in Field and Genton (2006). Full results are reported in Section **??** of the online supplementary material.

where $F^{\leftarrow}$ is the generalized inverse of the cdf, C is the copula defining their joint distribution and $\bar{C}(u,u) \overset{\text{def}}{=} 1 - 2u + C(u,u)$ is the so-called survivor function of C. Whereas the Gaussian copula has no tail dependence, i.e. $\lambda_u = \lambda_\ell = 0$ (McNeil et al., 2015, Example 7.38), the Student-$t(\rho)$ copula has the same asymptotic tail dependence in both tails (Demarta and McNeil, 2005):

$$\lambda_u = \lambda_\ell = 2t_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\rho)}{1+\rho}}\right),$$

where t_ν is the cdf of the Student-t distribution with ν degrees of freedom. The Gumbel copula has upper tail dependence $\lambda_u = 2 - 2^{1/\alpha}$ (McNeil et al., 2015, Example 7.37).

2.2 Copula estimation

Given the multivariate distribution (2), the *inference functions for margins* copula estimation method (Joe, 1997, 2005) splits the optimization of the likelihood in $d+1$ separate maximizations. For simplicity, we explain the method for the distribution defined in (2).

First, the estimators $\hat{g}_j$ and $\hat{h}_j$ of g_j and h_j, $j = 1, \ldots, d$, are computed separately, through solving d independent maximization problems. Then, in a second step, we build the pseudo-observations

$$\boldsymbol{u}_i = (F_1(x_i; 0, 1, \hat{g}_1, \hat{h}_1), \ldots, F_p(x_i; 0, 1, \hat{g}_p, \hat{h}_p))', \quad i = 1, \ldots, n, \tag{6}$$

and use them in a last optimization where we estimate the copula parameters. Given a random sample $\tilde{\boldsymbol{x}}_1, \ldots, \tilde{\boldsymbol{x}}_n$ from (2), the details are as follows.

1. For $j = 1, \ldots, d$:
   (a) estimate a_j and b_j by means of the quantile estimation method; let the estimators be denoted by $\tilde{a}_j$ and $\tilde{b}_j$;
   (b) compute the standardized observations $x_{ij} = (\tilde{x}_{ij} - \tilde{a}_j)/\tilde{b}_j$;
   (c) estimate g_j and h_j via numerical maximum likelihood estimation (MLE) using the standardized observations x_{ij} $(i = 1, \ldots, n)$;
2. For $i = 1, \ldots, n$, compute the pseudo-observations $\boldsymbol{u}_i$ defined in (6);
3. Estimate the parameters of the copula via MLE using the pseudo-observations $\boldsymbol{u}_i$ $(i = 1, \ldots, n)$.

The numerical MLE at Step 2 uses the approximation of the g-and-h density described in Section 2. When the copula is Gaussian or Student-t, the output of step 3 is $\hat{\boldsymbol{R}}$, i.e. the MLE of the correlation matrix. The goodness-of-fit of the chosen copula can be measured via the Cramér-von Mises functional defined in Genest et al. (2009). Both the estimation of the copula parameters at Step 3 and the computation of the goodness-of-fit test are performed via the `copula` `R` package (Hofert et al., 2018).

2.3 Copula Simulation

Once the parameters have been estimated, we might be interested in simulating from this joint distribution to obtain, e.g. the distribution of the total loss across event types by means of Monte Carlo simulations. Sampling one observation from (2) with the Gaussian copula proceeds as follows:

Algorithm 1

1. *Sample a random vector* $\boldsymbol{x} = (x_1, \ldots, x_d)'$ *from the* $N_d(\mathbf{0}, \hat{\boldsymbol{R}})$ *distribution;*

2. *evaluate* $u_j = \Phi(x_j)$, $j = 1, \ldots, d$, *where* Φ *is the standard normal cdf;*

3. *compute* $y_j = Q(u_j; \hat{\boldsymbol{\theta}}_j)$, *where* Q *is the g-and-h quantile function (3) and* $\hat{\boldsymbol{\theta}}_j = (0, 1, \hat{g}_j, \hat{h}_j)'$, $j = 1, \ldots, d$.

If one needs to simulate an unstandardized observation, Step 3 above is modified as follows:

3a. *compute* $y_j = Q(u_j; \hat{\hat{\boldsymbol{\theta}}}_j)$, *where* $\hat{\hat{\boldsymbol{\theta}}}_j = (\tilde{a}_j, \tilde{b}_j, \hat{g}_j, \hat{h}_j)'$, $j = 1, \ldots, d$, *and* $\tilde{a}$ *and* $\tilde{b}$ *are the Hoaglin (1985) quantile estimators.*

If the Student-t copula is used in place of the Gaussian, it is enough to replace the multivariate normal with the multivariate Student-t distribution at Step 1 and the univariate standard normal with the Student-t cdf at Step 2. The implementation of Algorithm 1 is straightforward in both cases, since the quantile function of the g-and-h distribution is explicit. Hence, random number generation is very fast.

Figure 1 shows 5000 observations simulated from the bivariate distribution with the Gaussian copula (panel A), Student-t copula with $\nu = 2$ degrees of freedom (panel B) and Gumbel copula with parameter $\alpha = 3$ (panel C). The marginals are $X_1 \sim \text{gh}(0, 1, 0.1, 0.02)$ and $X_2 \sim \text{gh}(0, 1, 0.08, 0.15)$ in all cases. The remaining parameters are $\rho \approx 0.7071$ for the Gaussian and Student-t copulas, and $\alpha = 2$ for the Gumbel copula, so that in all graphs the observations are simulated from a joint distribution with $\rho_\tau = 0.5$.

2.4 Vine copulas

When the dimension of the problem is larger than two, classical copulas such as those mentioned above have a dependence structure with uniform pairwise dependence. For practical purposes, this is often a limitation. A major step towards a more general model has been the introduction of pair-copula constructions (Aas et al., 2009), also called *vine copulas* because they can be efficiently described by means of a graphical model called vine.

Vine copulas are flexible and computationally efficient, and have proved to be quite successful in high-dimensional applications to financial data.

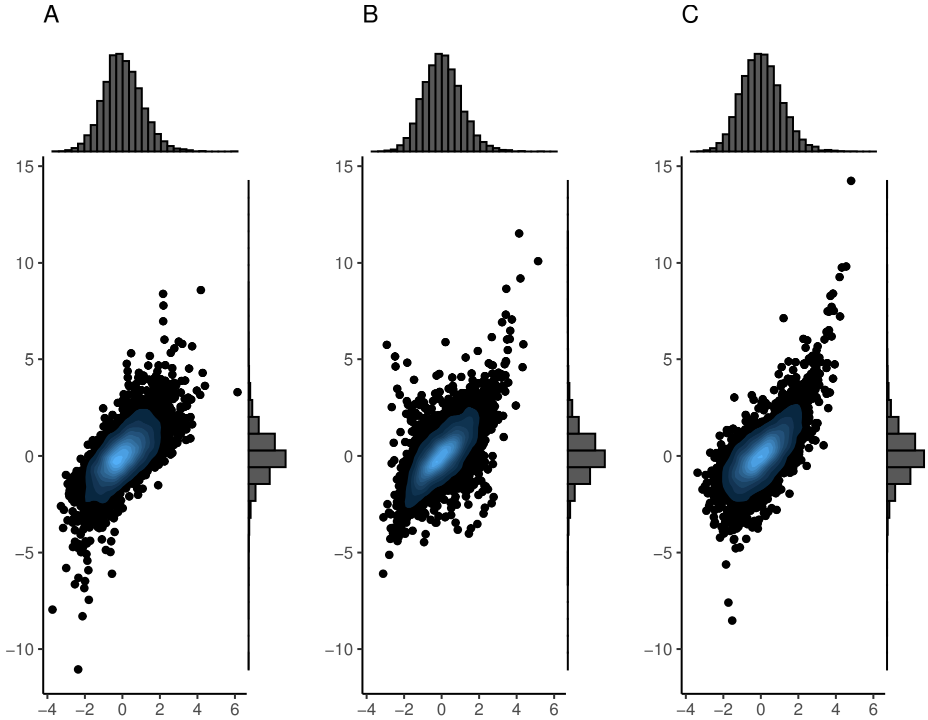

Figure 1: 5000 observations simulated from the Gaussian (panel A), Student-t (panel B) and Gumbel (panel C) copulas with marginals $X_1 \sim \text{gh}(0, 1, 0.1, 0.02)$ and $X_2 \sim \text{gh}(0, 1, 0.08, 0.15)$. The Gaussian and Student-t copulas have $\rho \approx .7071$, the t copula has $\nu = 2$ degrees of freedom, the Gumbel copula has dependence parameter $\alpha = 2$.

Their main advantage is that a joint density can be decomposed as a product of pair copula densities and marginal densities, represented by a nested set of trees. This implies that one can use different pair copula densities for different pairs of marginals, with an increased modeling flexibility. Furthermore, after determining the dependence structure of the data via graph theory, the parameters can be estimated using sequential maximum likelihood of bivariate copulas, thus avoiding large-dimensional optimizations. The quality of fit can be assessed via the test based on White's information matrix equality (White, 1982).

The analysis carried out in this paper is mostly based on Mai and Scherer (2017) and Czado (2019, Chapter 5), to which the reader is referred for further details. All the computations are performed using the `vinecopula` `R` package (Nagler et al., 2019).

3 Simulation experiments

The performance of the estimation procedures introduced in Section 2 is assessed via simulation experiments. We sample 500 observations from the 5-dimensional distribution (2) with the Gaussian copula and parameters

$$\boldsymbol{g} = \begin{pmatrix} 0.05 \\ 0.10 \\ 0.15 \\ 0.20 \\ 0.25 \end{pmatrix}, \ \boldsymbol{h} = \begin{pmatrix} 0.15 \\ 0.25 \\ 0.05 \\ 0.20 \\ 0.10 \end{pmatrix}, \ \boldsymbol{R} = \begin{pmatrix} 1.00 & 0.25 & 0.20 & 0.30 & 0.80 \\ 0.25 & 1.00 & 0.10 & 0.65 & 0.40 \\ 0.20 & 0.10 & 1.00 & 0.75 & 0.45 \\ 0.30 & 0.65 & 0.75 & 1.00 & 0.50 \\ 0.80 & 0.40 & 0.45 & 0.50 & 1.00 \end{pmatrix}. \tag{7}$$

The bias, standard deviation and relative-RMSE of $\hat{\boldsymbol{g}}$ and $\hat{\boldsymbol{h}}$ are shown in Figure 2. The relative RMSE of an estimator $\hat{\theta}$ is defined as $\mathrm{RMSE}(\hat{\theta})/\hat{\theta}$. Figure 3 displays the simulated distributions of the estimators of g and h for the first two marginals; the outcomes for the other marginals are reported in figures **??** and **??** of the online supplementary material.

Figure 4 shows bias, standard deviation and relative-RMSE of the estimators of the elements of the correlation matrix $\boldsymbol{R}$ of the Gaussian copula. Finally, Figure 5 reports the simulated distributions of the estimators of the first four entries of the correlation matrix $\boldsymbol{R}$. The histograms of the remaining correlation coefficients are shown in figures **??** and **??** in the online supplementary material. The estimators seem to follow an approximately normal distribution, with the possible exception of $\hat{\rho}_{3,4}$.

Figures 2 and 4 suggest that all the parameters, and especially the correlations, are estimated accurately.

The same experiment has been performed by simulating from a t copula with 5 degrees of freedom instead of a Gaussian copula. The outcomes shown in plots (**??**)-(**??**) in Appendix **??**, are quite similar to the Gaussian case.

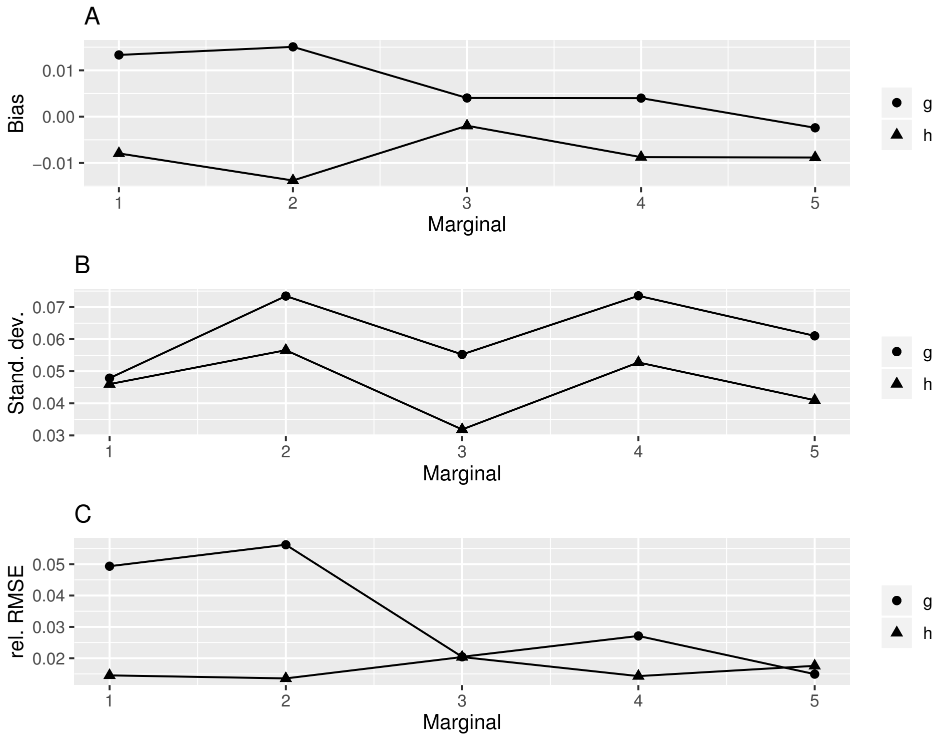

Figure 2: Bias (panel A), standard deviation (panel B) and relative-RMSE (panel C) of the estimators of the parameters g and h of the marginal distributions. Estimators are computed from 500 observations sampled from the 5-dimensional distribution based on the Gaussian copula. The true values of the parameters are shown in (7).

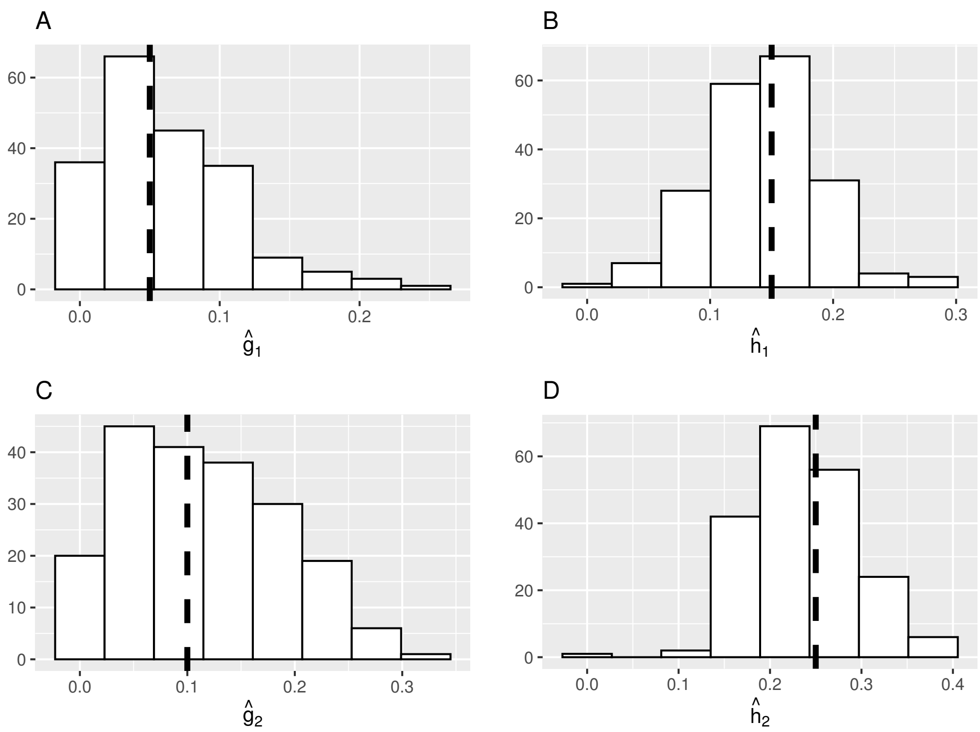

Figure 3: Simulated distributions of the estimators of g and h for the first two marginals. Estimators are the averages of 200 estimates obtained from samples of size 500 from the 5-dimensional distribution with numerical values of the parameters shown in (7). The dashed vertical lines represent the true values of the parameters g and h.

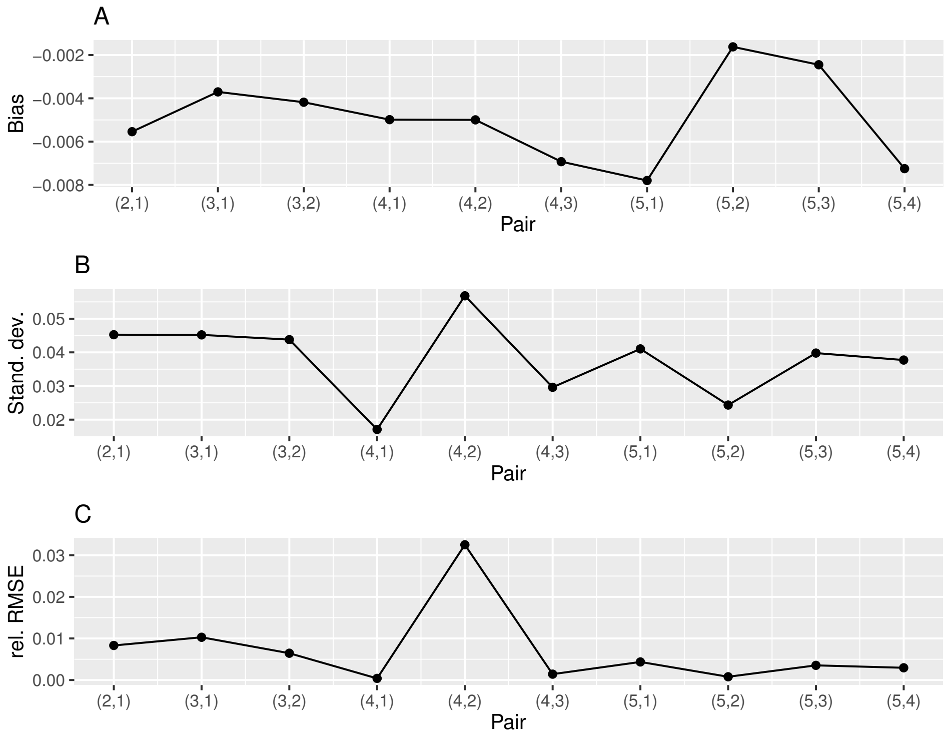

Figure 4: Bias (panel A), standard deviation (panel B) and relative-RMSE (panel C) of the estimators of the distinct elements of the correlation matrix $\boldsymbol{R}$ of the Gaussian copula. Estimators are the averages of 200 estimates obtained from samples of size 500 from the 5-dimensional distribution with numerical values of the parameters shown in (7).

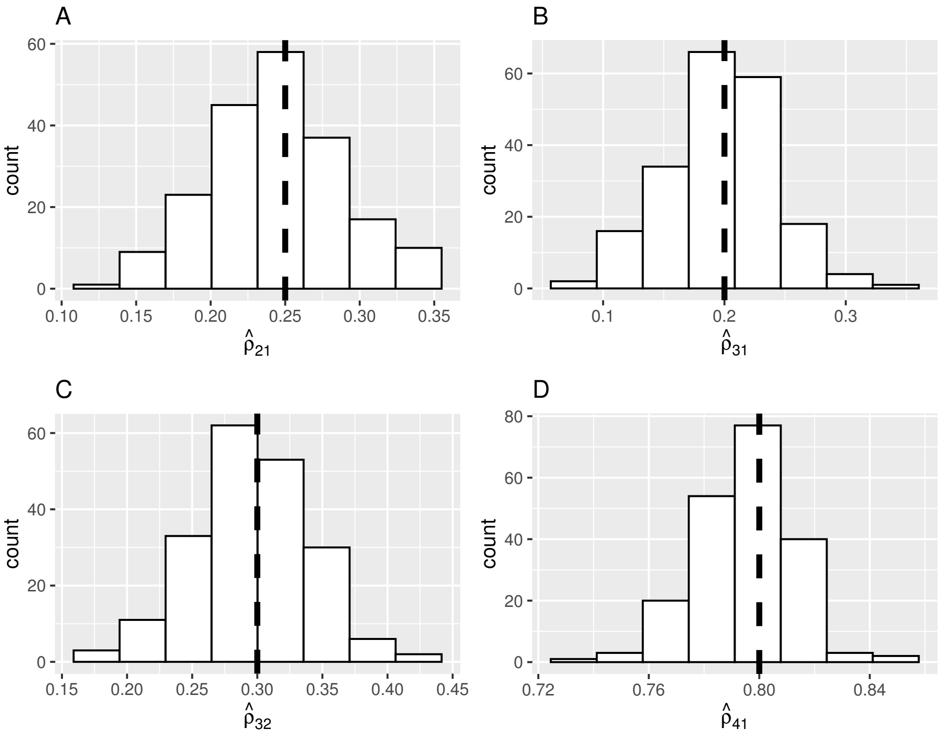

Figure 5: Simulated distributions of the estimators of the first four elements of the correlation matrix $\boldsymbol{R}$. Estimators are the averages of 200 estimates obtained from samples of size 500 from the 5-dimensional distribution with numerical values of the parameters shown in (7). The dashed vertical lines represent the true values of the parameters.

3.1 Estimation under a misspecified model

To assess the degree of flexibility of vine copulas, we carry out another simulation experiment.

1. Sample $n = 500$ observations from the five-variate Gaussian copula-based g-and-h distribution with parameters (7);
2. Estimate the marginal distributions (in a univariate setup) and the parameters of the Gaussian and R-vine copulas;
3. Simulate the loss distribution in the three cases using the estimated parameters and assuming an equally-weighted portfolio;
4. Compute the VaR in the three setups as well as the VaR of the loss simulated at Step 1.

The steps above are repeated 200 times, and the simulated distributions of the four 99.9% VaRs are shown in Fig. 6. Panel A displays the estimated VaR distribution using the R-Vine approach to approximate the dependence structure, whereas Panel B displays the estimated distribution obtained with the true Gaussian model (although with estimated paramaters). Panel C displays the results for a perfect dependence structure, whereas Panel D shows the true (finite sample) VaR distribution. The histograms of the VaRs at different levels are displayed in figures **??**-**??** in the online supplementary material. The average of the estimated VaRs obtained in the 200 replications are shown in Table 1.

Table 1: Average of the estimated VaRs (quantiles of the simulated total loss distribution) at levels 90, 95, 99, 99.5 and 99.9% in the R-vine, Gaussian and perfect positive dependence cases.

α	90%	95%	99%	99.5%	99.9%
R-vine	1.27 (0.10)	1.79 (0.16)	2.99 (0.43)	3.58 (0.66)	5.23 (2.28)
Gaussian	1.29 (0.13)	1.78 (0.18)	2.99 (0.45)	3.56 (0.63)	4.87 (1.37)
Perfect dependence	1.61 (0.08)	2.31 (0.15)	4.31 (0.58)	5.40 (0.87)	8.73 (2.06)
True	1.29 (0.09)	1.77 (0.13)	3.02 (0.42)	3.60 (0.56)	5.05 (1.28)

The R-vine VaR is almost identical to both the true and the Gaussian VaR in terms of point estimation; the standard errors are also practically the same, with the exception of the 99.9% case. These results suggest that the pair copula construction approach works very well even in cases where the true model is specified by a single copula.

Table **??** in Appendix **??** of the online supplementary material shows the results obtained when sampling the t copula. Because of the tail properties of the t copula, all the values are higher, especially for large α. However, also in this case the VaR computed via the vine copula approach is approximately

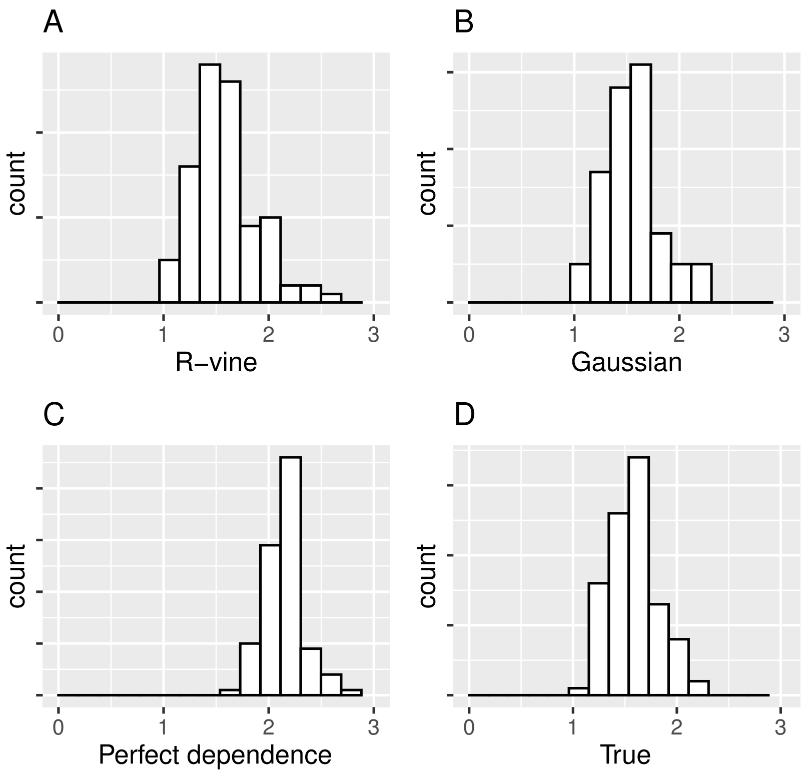

Figure 6: Simulated distributions of the 99.9% VaRs. The histograms show 200 estimates obtained from samples of size 500 from the 5-dimensional distribution with numerical values of the parameters shown in (7).

as precise as the VaR obtained from the true data-generating process, i.e. the t copula.

4 Empirical analysis: operational risk

Univariate distributions of operational risk data are typically very skewed and leptokurtic. Hence, the g-and-h model is a good candidate, and the extensive empirical investigation of Dutta and Perry (2006) suggests that it is the best-fitting distribution for operational risk data.

We estimate the joint loss distribution of operational losses recorded in different event types by means of the model developed in Section 2. To accomplish this goal, we exploit a large database containing 40,871 operational losses recorded at the Italian bank Unicredit between January 2005 and June 2014 for the seven Basel II event types: Internal fraud (IFRAUD), External fraud (EFRAUD), Employment Practices and Workplace Safety (EPWS), Clients, Products, and Business Practice (CPBP), Damage to Physical Assets (DPA), Business Disruption and Systems Failures (BDSF), Execution, Delivery, and Process Management (EDPM). See Hambuckers et al. (2018) for more details about the data.

For estimation purposes we compute monthly total losses for each event type and use them to fit both the g-and-h marginals and the copula. That is, for a given month and event type, we sum up all registered losses exhibiting these characteristics. The final estimated joint distribution is therefore the joint distribution of the event type-specific total losses. Similar approaches are proposed by Aulbach et al. (2012) and Eling and Jung (2018); see also Chapelle et al. (2008) and Brechmann et al. (2014). The resulting sample size is $n = 114$. The univariate estimates of the marginal distributions are displayed in Table 2, along with bootstrap standard errors based on 500 replications. The last column reports the p-value of the Kolmogorov-Smirnov goodness-of-fit test statistics, which shows that the g-and-h distributional assumption is never rejected. Notice also that the computational burden is acceptable: the estimation procedure takes approximately 150 seconds using `R` 3.6.1 on an i7-6700 3.40GHz processor.

From an economic standpoint, these estimated distributions indicate a significant likelihood of extremely high losses for all event types, especially for IFRAUD and CPBP. For these two event types, the kurtosis levels are of magnitude 10^{20} and 10^{10}, respectively. On the contrary, an event type like EFRAUD exhibits a much smaller kurtosis (around 10). These results are consistent with historical rogue trading events (e.g., J. Kerviel, N. Leeson and H. Hubler) or legal fines related to business practices (e.g. during the Great Financial Crisis) belonging to these event types, and found to be the main contributors to total operational losses (see, e.g. Kley et al., 2020, for additional examples).

Table 2: Estimates of the parameters for the marginal distributions, with associated bootstrap standard errors (in parentheses) and p-value of the KS goodness-of-fit test statistics.

	a	b	g	h	KS p-value
IFRAUD	342 070 (75 646.7)	733 287.7 (208 093.3)	2.181 (0.587)	0 (0.179)	0.612
EFRAUD	665 850 (42 771.2)	320 046.2 (63 303.9)	0.552 (0.204)	0 (0.257)	0.108
EPWS	265 580 (42 783.2)	367 644 (83 243.1)	1.428 (0.289)	0.013 (0.137)	0.348
CPBP	3 572 250 (521 981)	4 414 807 (984 099.9)	1.940 (0.444)	0.058 (0.14)	0.265
DPA	40 220.5 (4029.7)	33 550.82 (7426.7)	1.098 (0.233)	0.125 (0.168)	0.821
BDSF	43 772.5 (12 009)	80 324.58 (13 941.5)	1.809 (0.519)	0 (0.081)	0.111
EDPM	2 262 750 (328 025.7)	2 119 520 (385 399.5)	1.498 (0.499)	0.092 (0.141)	0.318

The bootstrap distribution of the estimated parameters for the EDPM event type is shown in Fig. 7 along with the normal density with parameters equal to the mean and standard deviation of the bootstrap distribution. No normal density is shown in panel D, because the distribution of $\hat{h}$ has a lower bound at 0 and therefore cannot be normal. The histograms of the remaining marginals are in figures **??**-**??** of the online supplementary material. Estimates of b and g are approximately normal, whereas $\hat{a}$ is not. This may be justified by the fact that $\hat{a}$ is a quantile estimator, not a MLE.

The empirical evidence provided so far suggests that the g-and-h is an appropriate model for the marginals. We now turn to the analysis of the dependence structure via two different copula-based approaches.

4.1 A single-copula approach

We first assume that the joint distribution is defined by (2) with C equal to the Gaussian or Student-t copula. This implies that the dependence structure is assumed identical for each pair of variables, although with different strengths, as the correlation parameter ρ_{ij} can be different across pairs.

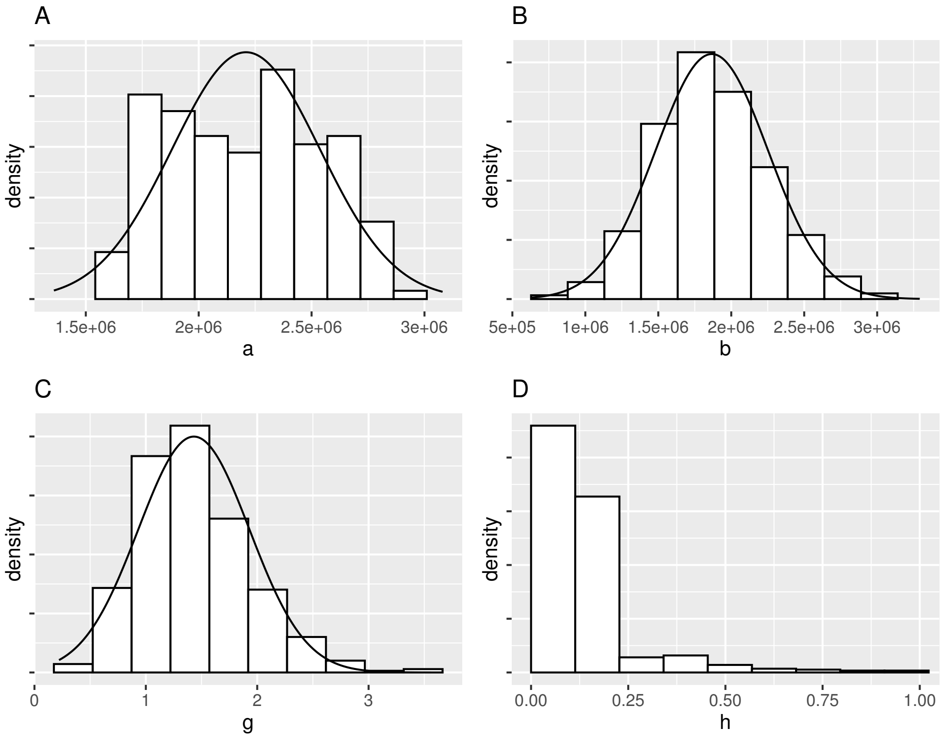

Figure 7: Bootstrap distributions of the estimators of the parameters for the EDPM event type and estimated normal density.

The estimated correlation matrix of the Gaussian copula is given by

$$\hat{\boldsymbol{R}}^G = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.241 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0.409 & 0.397 & 1 & 0 & 0 & 0 & 0 \\ 0.048 & 0.042 & 0.389 & 1 & 0 & 0 & 0 \\ 0.168 & 0.233 & 0.114 & 0.305 & 1 & 0 & 0 \\ 0.322 & 0.293 & 0.258 & 0.025 & 0.394 & 1 & 0 \\ 0.255 & 0.072 & 0.596 & -0.005 & 0.273 & 0.241 & 1 \end{pmatrix},$$

and the corresponding standard errors are

$$\mathrm{se}(\hat{\boldsymbol{R}}^G) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.084 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.098 & 0.104 & 0 & 0 & 0 & 0 & 0 \\ 0.090 & 0.150 & 0.104 & 0 & 0 & 0 & 0 \\ 0.087 & 0.079 & 0.074 & 0.092 & 0 & 0 & 0 \\ 0.095 & 0.118 & 0.089 & 0.110 & 0.109 & 0 & 0 \\ 0.093 & 0.117 & 0.097 & 0.098 & 0.097 & 0.123 & 0 \end{pmatrix}.$$

The results for the Student-t copula are almost identical and thus not reported here. In particular, the estimated number of degrees of freedom is 2,584, so that there is essentially no difference between the Gaussian and the Student-t copula.

The correlation coefficients of the various pairs are substantially larger than zero, but the estimated dependence structure is very different from perfect positive dependence. The goodness of fit test statistics is 0.0436, with p-value equal to 0.0105, for both copulas. Since neither copula is accepted at the 5% test level, we prefer to carry out a more detailed study of the dependence structure based on vine copulas.

4.2 A vine copula approach

To get some preliminary information about the empirical features of the bivariate distribution of each pair, we display in Figure 8 contour plots of the bivariate copula densities estimated non-parametrically via the `kdecop` command of the `kdecopula R` package (Nagler and Wen, 2018). The plots suggest that the dependence structure is unlikely to be the same for each pair. This provides further support for an approach where different pairs are possibly modeled by different copulas.

In the first step of the analysis we look for the most appropriate dependence structure by selecting R-vine trees via maximum spanning trees with respect to edge weights given by the absolute value of the Kendall's τ coefficient. We consider two different sets of copula families:

1. All copula families considered in the `VineCopula R` package;

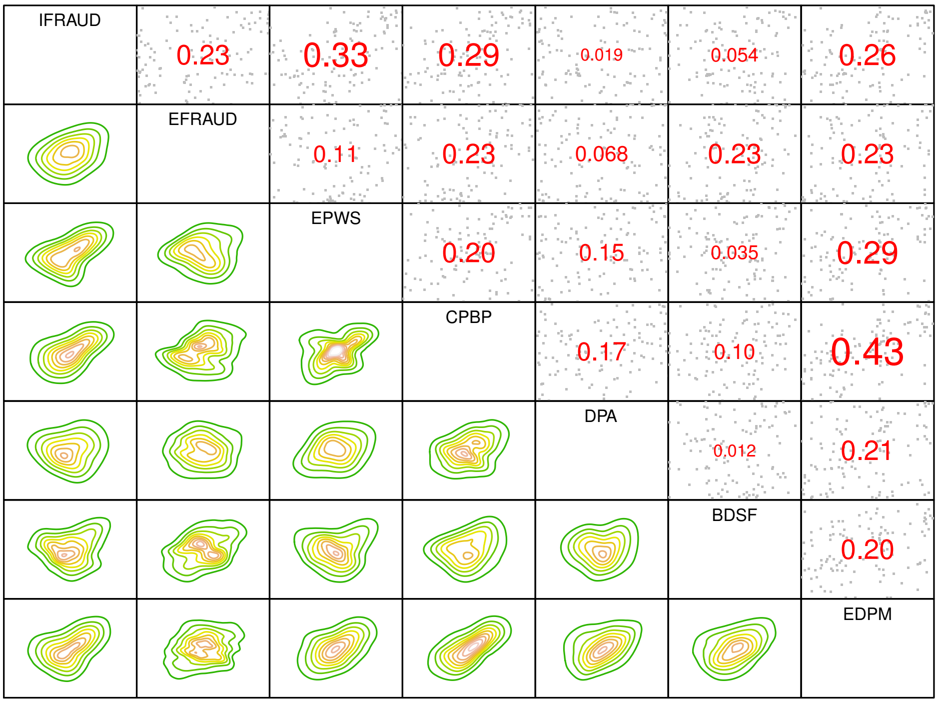

Figure 8: Contour plots of the bivariate non-parametric estimates of the copula densities of each pair of variables (below the diagonal) and Kendall's τ estimates (above the diagonal) for the operational risk data.

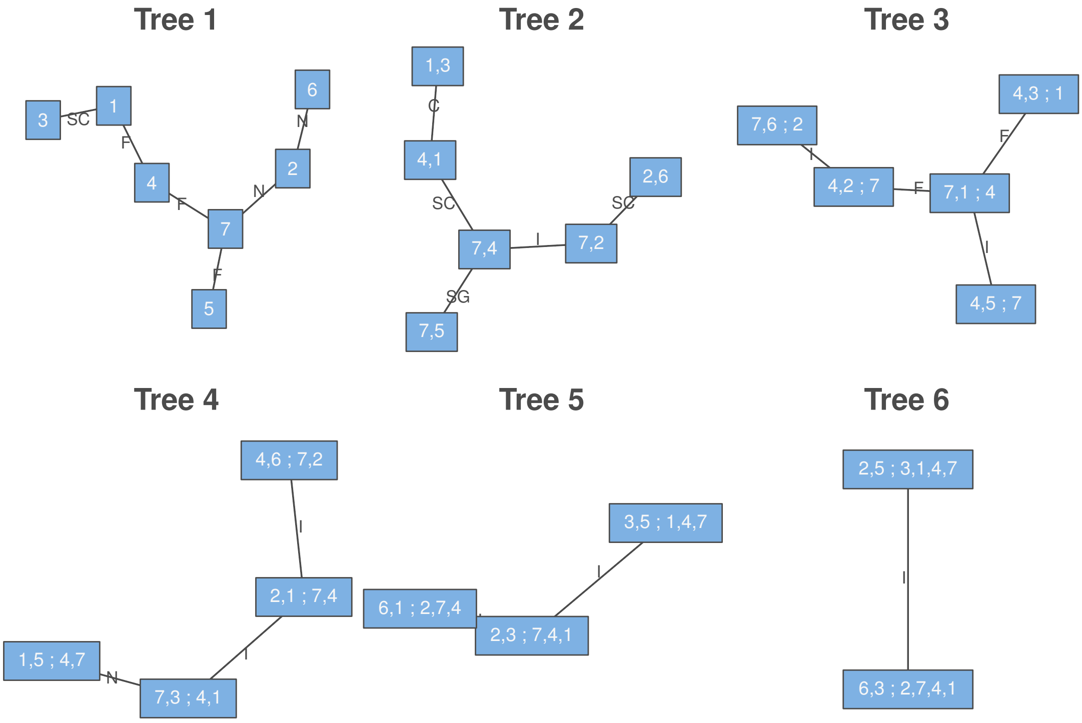

Figure 9: Tree plots of the R-vines based on all copula families.

2. the Independence (I), Gaussian (N) and Student-t (t) copulas.

In each case, once the structure has been selected, we estimate the parameters via MLE.

The trees describing the conditional distributions are displayed in figures 9 (when all copulas are used) and 10 (when only the Gaussian and Student-t are considered). The letter near each edge identifies the pair copula family linking the two nodes. It is worth noting that the structure is simplified by allowing for the independence copula, since in both figures almost all trees use at least once the independence copula for specific pairs.

The fitted contour plots are displayed in figures 11 and 12, whereas detailed estimation results are given in tables 3 and 4.

Although the two trees in figures 9 and 10 are essentially identical, the fourth column of tables 3 and 4 reveals that the copulas used to describe the dependence structure are different. Some tail dependence is implied only in a few cases: notably, two cases with tail dependence involve the EDPM event type, in agreement with the contour plots in the last line of Figure 8. Moreover, the first tree suggests that EDPM and CPBP are the most important nodes in the dependence network described by the estimated R-

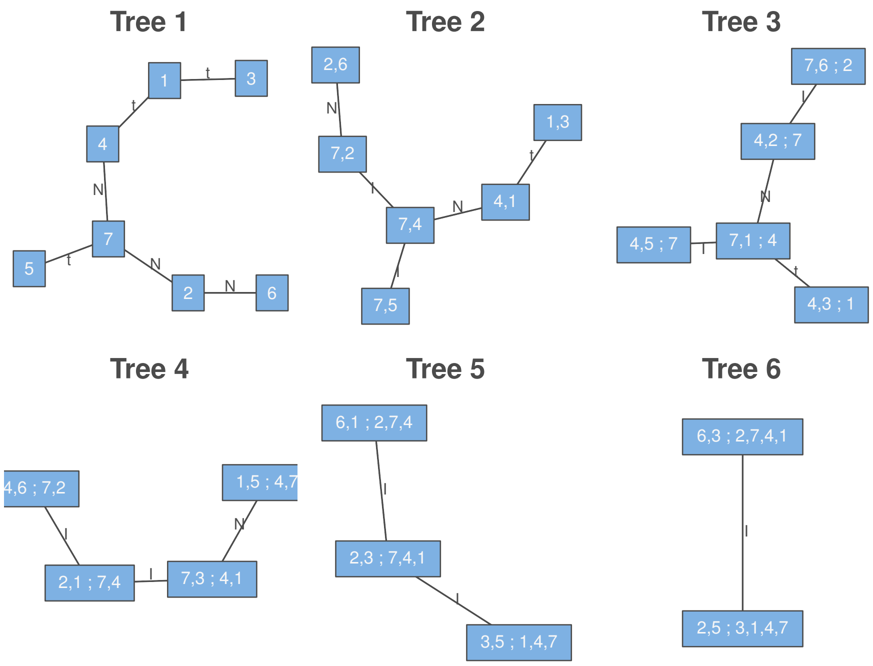

Figure 10: Tree plots of the R-vines based on the independence, Gaussian and Student-t copulas.

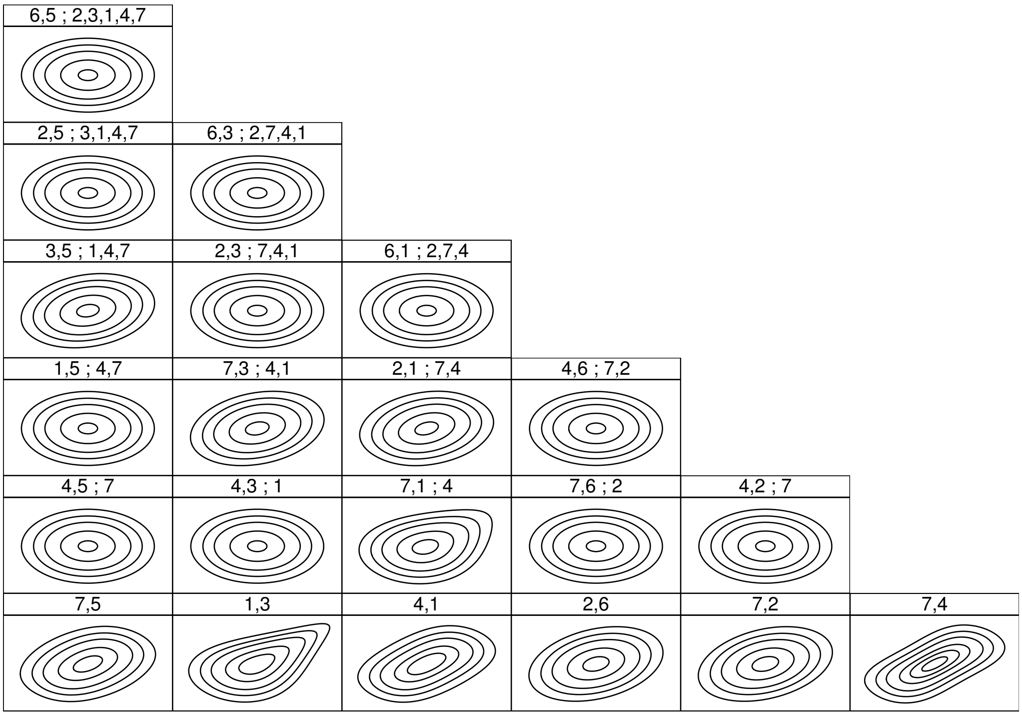

Figure 11: Fitted contour plots of the R-vines based on all copula families.

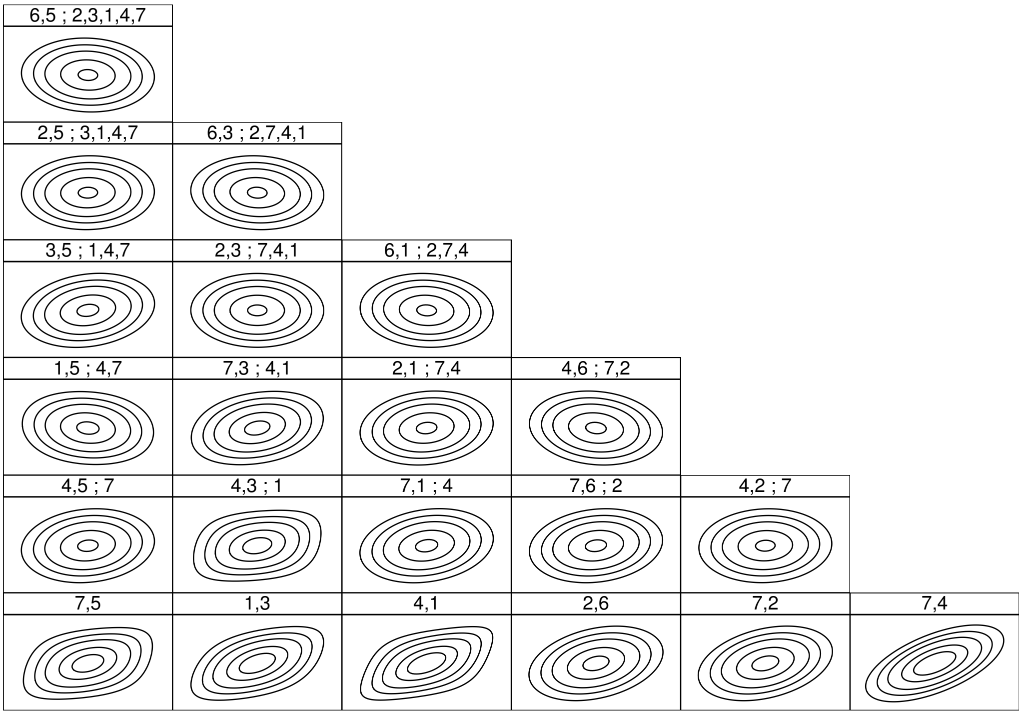

Figure 12: Fitted contour plots of the R-vines based on the independence, Gaussian and Student-t copulas.

Table 3: Sequential estimates of the parameters for the model based on all copulas and corresponding Kendall's τ, upper (utd) and lower (ltd) tail dependence coefficient estimates.

tree	edge	family	cop	par	par_2	τ	utd	ltd
1	7,5	5	F	1.87	0.00	0.20	-	-
	1,3	13	SC	0.84	0.00	0.30	0.44	-
	4,1	5	F	3.42	0.00	0.34	-	-
	2,6	1	N	0.31	0.00	0.20	-	-
	7,2	1	N	0.31	0.00	0.20	-	-
	7,4	5	F	4.84	0.00	0.45	-	-
2	4,5;7	14	SG	1.11	0.00	0.10	-	0.13
	4,3;1	3	C	0.24	0.00	0.11	-	0.05
	7,1;4	13	SC	0.31	0.00	0.14	0.11	-
	7,6;2	13	SC	0.44	0.00	0.18	0.21	-
	4,2;7	0	I	-	-	0.00	-	-
3	1,5;4,7	0	I	-	-	0.00	-	-
	7,3;4,1	5	F	1.59	0.00	0.17	-	-
	2,1;7,4	5	F	1.05	0.00	0.12	-	-
	4,6;7,2	0	I	-	-	0.00	-	-
4	3,5;1,4,7	1	N	0.18	0.00	0.12	-	-
	2,3;7,4,1	0	I	-	-	0.00	-	-
	6,1;2,7,4	0	I	-	-	0.00	-	-
5	2,5;3,1,4,7	0	I	-	-	0.00	-	-
	6,3;2,7,4,1	0	I	-	-	0.00	-	-
6	6,5;2,3,1,4,7	0	I	-	-	0.00	-	-
	logLik: 88.4		AIC: -151		BIC: -115			

vine copula.

From the economic point of view, the central place of CPBP and EDPM in the dependence structure makes sense since operational events can be traced back to weak internal control environments (Chernobai et al., 2011). In particular, CPBP and EDPM event types are mostly driven by "a lack of controls, and failures to comply with procedures" (Chernobai et al., 2011, Table 3). Although different internal control weaknesses intervene for the other event types, lack of controls and failure to comply still play a significant role in the determination of each loss severity. Consequently, a stronger (bivariate) dependence should be expected between CPBP and EDPM, and the other event types.

Table 4: Sequential estimates of the parameters for the model based on the Gaussian and Student-t copulas and corresponding Kendall's τ, upper (utd) and lower (ltd) tail dependence coefficient estimates.

tree	edge	family	cop	par	par_2	τ	utd	ltd
1	7,5	2	t	0.30	6.24	0.19	0.09	0.09
	1,3	2	t	0.45	10.69	0.29	0.06	0.06
	4,1	2	t	0.47	5.87	0.31	0.16	0.16
	2,6	1	N	0.31	0.00	0.20	-	-
	7,2	1	N	0.31	0.00	0.20	-	-
	7,4	1	N	0.61	0.00	0.41	-	-
2	4,5;7	0	I	-	-	0.00	-	-
	4,3;1	2	t	0.22	7.07	0.14	0.05	0.05
	7,1;4	1	N	0.21	0.00	0.13	-	-
	7,6;2	1	N	0.19	0.00	0.12	-	-
	4,2;7	0	I	-	-	0.00	-	-
3	1,5;4,7	0	I	-	-	0.00	0.00	0.00
	7,3;4,1	2	t	0.24	30.00	0.15	0.00	0.00
	2,1;7,4	1	N	0.13	0.00	0.08	-	-
	4,6;7,2	0	I	-	-	0.00	-	-
4	3,5;1,4,7	1	N	0.18	0.00	0.11	-	-
	2,3;7,4,1	0	I	-	0.00	0.00	-	-
	6,1;2,7,4	0	I	-	0.00	0.00	-	-
5	2,5;3,1,4,7	0	I	-	-	0.00	-	-
	6,3;2,7,4,1	0	I	-	-	0.00	-	-
6	6,5;2,3,1,4,7	0	I	-	-	0.00	-	-
	logLik: 76.6		AIC: -119		BIC: -72.6			

Table 4 identifies the same network, but the values of the log-likelihood and of the information criteria shown at the bottom of tables 3 and 4 allow us to unambiguously conclude that the model based on all copulas yields a better fit. Since the Student-t copula is never selected in Table 3, this is not surprising.

4.3 Simulation and VaR computation

After finding the joint loss distribution $\boldsymbol{L} = (L_1, \ldots, L_7)'$ and estimating its parameters, we compute the VaR, at different levels, of the total loss over all event types (i.e. aggregating across event types). The variability of our estimation is estimated by means of a bootstrap analysis. To take into account the loss frequency, we define the total loss as

$$L_T = \sum_{i=1}^{7} w_i L_i,$$

where w_i is equal to the total number of losses for the i-th event type divided by the total number of losses for all event types. At each replication and for each of the three models (R-vine copula, Gaussian copula and perfect dependence), we draw a bootstrap sample from our historical sample, re-estimate the parameters and simulate B seven-dimensional vectors of observations

$$\boldsymbol{L}_i^s = (L_{i,1}^s, \ldots, L_{i,7}^s)'.$$

Then, we compute the corresponding value of

$$L_{T,i}^s = \sum_{j=1}^{7} w_i \boldsymbol{L}_{i,j}^s,$$

for $i = 1, \ldots, B$, with $B = 10^5$. The VaR at level α is simply the α-quantile of $L_T^s = (L_{T,1}^s, \ldots, L_{T,B}^s)'$. We repeat the previous steps 500 times, i.e. $s = 1, \ldots, 500$ to obtain the bootstrap distribution of the estimated VaR.

Figure 13 shows, on a logarithmic scale, the simulated distribution of the 99.9% VaR for the total distribution[2], obtained from the R-vine approach (panel A), the Gaussian copula (panel B) and under the assumption of perfect positive dependence (panel C). We see that the distribution under the hypothesis of perfect positive dependence is the most skewed one, and is therefore characterized by the largest VaR at high coverage levels. The average VaRs are displayed in Table 5 for the three approaches.

Table 5: Estimated VaRs (quantiles of the simulated total loss distribution) at the 90, 95, 99, 99.5 and 99.9% level in the R-vine, Gaussian and perfect positive dependence cases.

α	90%	95%	99%	99.5%	99.9%
R-vine	4.43	8.29	31.07	58.86	508.95
Gaussian	4.26	7.95	30.75	58.47	506.76
Perfect dependence	4.06	8.30	35.96	69.48	548.49

[2]The histograms of the simulated distributions at level 90, 95, 99 and 99.5% are reported in figures ??-?? of the online supplementary material.

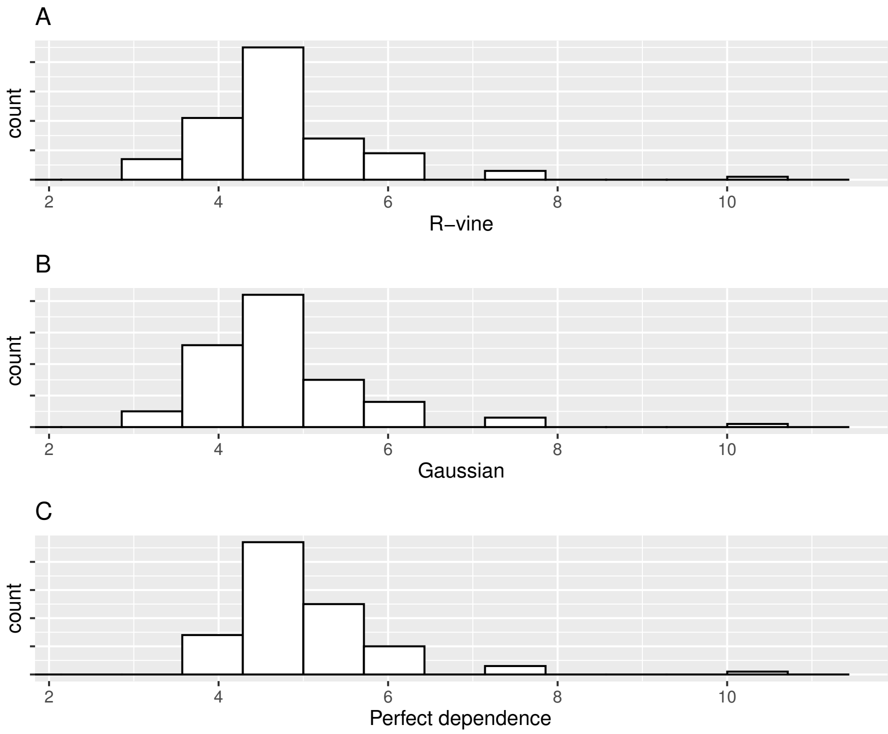

Figure 13: Simulated distribution of the 99.9% VaR (on a logarithmic scale) using the estimated R-vine copula (panel A), the estimated Gaussian copula (panel B), and the perfect positive dependence approach (panel C).

Table 6: Average and 95% confidence interval of the difference between the estimated Value-at-Risk obtained with the perfect dependence and the R-vine copula approach at the 90, 95, 99, 99.5 and 99.9% level.

α	90%	95%	99%	99.5%	99.9%
Average	-0.37	0.02	4.89	10.63	39.54
95% CI	$(-1.28, 0.17)$	$(-1.67, 0.76)$	$(2.23, 7.72)$	$(4.30, 17.83)$	$(7.84, 87.25)$

Despite the rejection of the Gaussian copula by the goodness-of-fit test in Section 4.1, the VaRs resulting from the two copula-based approaches are very similar to each other. On the other hand, under the assumption of perfect positive dependence, the results are quite different. To assess the significance of the difference, we have computed the average difference across the 500 replications between the perfect dependence VaR and the R-vine VaR, as well as the 95% confidence intervals. The results in Table 6 confirm that the former is significantly larger when $\alpha \geq 0.95$. Finally, since the dependence found by both the Gaussian and the R-vine copula approaches is far from perfect positive dependence, and that the R-vine copula found in Section 4.2 gives the best fit, we can conclude that the R-vine copula is the most appropriate model for the data at hand. Therefore, the corresponding VaR is also the most reliable one. In an application to market risk data, Low et al. (2013) also find that vine copulas outperform standard copulas when the dimension of the problem is large.

Three main messages arise from this analysis. First, the VaRs obtained with the two copula-based approaches are almost identical; second, the estimated dependence is much weaker than perfect positive dependence; third, accounting for a realistic dependence structure implies a lower total risk when $\alpha \geq 95\%$. The last outcome is in line with the discussion of the relationship between dependence and total capital charge put forth by Chapelle et al. (2008) and Brechmann et al. (2014).

4.4 Non-parametric estimation of the marginal distributions

Even though the results of the Kolmogorov-Smirnov goodness-of-fit test in Table 2 strongly support the choice of the g-and-h, we double check this assumption by means of a model that avoids any parametric assumption about the univariate marginal distributions. In particular, we repeat the steps performed in sections 4.1 to 4.3 by estimating the marginal distributions via the modified empirical cdf proposed by McNeil et al. (2015, p. 270).

Since the results are similar to the previous sections, we only show the estimated VaRs. Accordingly, Table 7 is identical to Table 5, except for the fact that the marginals are now estimated non-parametrically.

Table 7: Estimated VaRs (quantiles of the simulated total loss distribution) at the 90, 95, 99, 99.5 and 99.9% level in the R-vine, Gaussian and perfect positive dependence cases, using the modified empirical cdf to estimate the marginal distributions.

α	90%	95%	99%	99.5%	99.9%
R-vine	4.01	7.77	27.19	48.06	421.32
Gaussian	3.81	7.17	23.22	46.97	366.50
Perfect dependence	3.94	7.98	33.61	60.12	478.03

The outcomes suggest that, even though the VaRs are somewhat smaller, the three approaches behave similarly to the case where the marginals are modelled via the g-and-h distribution.

5 Conclusion

In this paper we develop a family of copula-based multivariate distributions with g-and-h marginals. The distribution has a high degree of flexibility, both because the g-and-h is characterized by a very large skewness-kurtosis range and because copulas can model a wide range of dependence structures. Since we have included vine copulas in the analysis, the flexibility is even bigger.

Our large database of operational losses suggests that the fit of the model based on regular vine copulas is good, and that the VaR at a high confidence level, obtained from the best-fitting multivariate distribution, is smaller than the VaR under the hypothesis of perfect positive dependence. Thus, for banks willing to quantify their operational risk level and to derive corresponding capital reserves, our results emphasize the need to account for both the marginal tail risks and the correlation structure across event types to obtain economically meaningful measures. Similarly, for regulators, our results demonstrate that the g-and-h distribution is a valid candidate as loss distribution for CaR calculations (BCBS, 2016) and could be included in future versions of the revised Standardized Approach.

Overall, our results hint also at a wider applicability of the proposed methodology for, e.g., portfolio optimization under conditional VaR constraints (Low et al., 2013) or risk assessment of credit derivatives (Aas, 2016), when the focus is on extreme quantiles. Under this scenario, the g-and-h copula-based approach would provide accurate quantification of both the tail of the marginal distribution, and the multivariate structure of a portfolio of financial assets. A third promising field of application for our work consists in the trading of securities related to extreme tail-risk of operational losses (for hedging purposes), as suggested in Kley et al. (2020) and already in use in some financial institutions.

Finally, notice that these results are based on aggregated monthly losses, treated as if they were observed simultaneously in each event type. A more sophisticated approach would model each marginal distribution as a univariate compound Poisson process and the dependence structure via Lévy copulas; see Cruz et al. (2015, Chapter 12). This way of proceeding needs to be studied in future research.

Declarations of Interest

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

Acknowledgements

We thank Fabio Piacenza (UniCredit SpA) for providing us with the operational risk data.

References

Aas, K. (2016). Pair-copula constructions for financial applications: A review. *Econometrics*, 4(4).

Aas, K., Czado, C., Frigessi, A., and Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44:182–198.

Aulbach, S., Bayer, V., and Falk, M. (2012). A multivariate piecing-together approach with an application to operational loss data. *Bernoulli*, 18(2):455–475.

Barndorff-Nielsen, O. and Kendall, D. G. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 353(1674):401–419.

BCBS (2016). Standardised Measurement Approach for operational risk. Technical report, Bank of International Settlements, Basel, Switzerland.

Bee, M., Hambuckers, J., and Trapin, L. (2019). Estimating Value-at-Risk for the g-and-h distribution: an indirect inference approach. *Quantitative Finance*, 19(8):1255–1266.

Bee, M., Hambuckers, J., and Trapin, L. (2021). Estimating large losses in insurance analytics and operational risk using the g-and-h distribution. *Quantitative Finance*, 21(7):1207–1221.

Bee, M. and Trapin, L. (2016). A simple approach to the estimation of Tukey's gh distribution. *Journal of Statistical Computation and Simulation*, 86(16):3287–3302.

Böcker, K. and Klüppelberg, C. (2010). Multivariate models for operational risk. *Quantitative Finance*, 10(8):855–869.

Brechmann, E., Czado, C., and Paterlini, S. (2014). Flexible dependence modeling of operational risk losses and its impact on total capital requirements. *Journal of Banking & Finance*, 40:271 – 285.

Chapelle, A., Crama, Y., Hübner, G., and Peters, J.-P. (2008). Practical methods for measuring and managing operational risk in the financial sector: A clinical study. *Journal of Banking and Finance*, 32:1049–1061.

Chernobai, A., Jorion, P., and Yu, F. (2011). The determinants of operational risk in u.s. financial institutions. *The Journal of Financial and Quantitative Analysis*, 46(6):1683–1725.

Cruz, M., Peters, G., and Shevchenko, P. (2015). *Fundamental Aspects of Operational Risk and Insurance Analytics: A Handbook of Operational Risk*. Wiley.

Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas: a Practical Guide with R*. Springer.

Degen, M., Embrechts, P., and Lambrigger, D. D. (2007). The quantitative modeling of operational risk: between g-and-h and EVT. *Astin Bulletin*, 37(2):265–291.

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

Donnelly, C. and Embrechts, P. (2010). The devil is in the tails: Actuarial mathematics and the subprime mortgage crisis. *ASTIN Bulletin*, 40(1):1–33.

Dutta, K. K. and Babbel, D. F. (2002). On measuring skewness and kurtosis in short rate distributions: The case of the US dollar London inter bank offer rates. *Wharton School Center for Financial Institutions, University of Pennsylvania*, 02-25.

Dutta, K. K. and Perry, J. (2006). A tale of tails: an empirical analysis of loss distribution models for estimating operational risk capital. Technical Report 06-13, Federal Reserve Bank of Boston.

Eling, M. and Jung, K. (2018). Copula approaches for modeling cross-sectional dependence of data breach losses. *Insurance: Mathematics and Economics*, 82:167–180.

Field, C. and Genton, M. G. (2006). The multivariate g-and-h distribution. *Technometrics*, 48(1):104–111.

Fischer, M., Horn, A., and Klein, I. (2007). Tukey-type distributions in the context of financial data. *Communications in Statistics Theory and Methods*, 36(1):23–35.

Genest, C., Rémillard, B., and Beaudoin, D. (2009). Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, 44:199–214.

Gourier, E., Farkas, W., and Abbate, D. (2009). Operational risk quantification using extreme value theory and copulas: From theory to practice. *Journal of Operational Risk*, 4(3):3–24.

Hambuckers, J., Groll, A., and Kneib, T. (2018). Understanding the economic determinants of the severity of operational losses: A regularized Generalized Pareto regression approach. *Journal of Applied Econometrics*, 33(6):898–935.

Hoaglin, D. C. (1985). *Summarizing Shape Numerically: The g-and-h Distributions*, chapter 11, pages 461–513. Wiley.

Hofert, M., Kojadinovic, I., Maechler, M., and Yan, J. (2018). *copula: Multivariate Dependence with Copulas.* R package version 0.999-19.1.

Jiménez, J. A. and Arunachalam, V. (2011). Using Tukey's g and h family of distributions to calculate Value-at-Risk and conditional Value-at-Risk. *Journal of Risk*, 13(4):95–116.

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

Joe, H. (2005). Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, 94:401–419.

Kley, O., Klüppelberg, C., and Paterlini, S. (2020). Modelling extremal dependence for operational risk by a bipartite graph. *Journal of Banking & Finance*, 117:105855.

Li, D. X. (2000). On default correlation: A copula function approach. *Journal of Fixed Income*, 9:43–54.

Low, R. K. Y. (2018). Vine copulas: modelling systemic risk and enhancing higher-moment portfolio optimisation. *Accounting & Finance*, 58(S1):423–463.

Low, R. K. Y., Alcock, J., Faff, R., and Brailsford, T. (2013). Canonical vine copulas in the context of modern portfolio management: Are they worth it? *Journal of Banking & Finance*, 37(8):3085–3099.

Low, R. K. Y., Faff, R., and Aas, K. (2016). Enhancing mean-variance portfolio selection by modeling distributional asymmetries. *Journal of Economics and Business*, 85(C):49–72.

Mai, J.-F. and Scherer, M. (2017). *Simulating Copulas: Stochastic Models, Sampling Algorithms, and Applications.* World Scientific.

Martinez, J. and Iglewicz, B. (1984). Some properties of the Tukey g and h family of distributions. *Communications in Statistics - Theory and Methods*, 13(3):353–369.

McNeil, A., Frey, R., and Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques, Tools*. Princeton University Press, second edition.

Moscadelli, M. (2004). *The modelling of operational risk: experience with the analysis of the data collected by the Basel Committee*, volume 517 of *Temi di discussione*. Banca d'Italia.

Nagler, T., Schepsmeier, U., Stoeber, J., Brechmann, E. C., Graeler, B., Erhardt, T., Almeida, C., Min, A., Czado, C., Hofmann, M., Killiches, M., Joe, H., and Vatter, T. (2019). *VineCopula: Statistical Inference of Vine Copulas*. R package version 2.2.0.

Nagler, T. and Wen, K. (2018). *kdecopula: Kernel Smoothing for Bivariate Copula Densities*. R package version 0.9.2.

Peters, G. W., Chen, W. Y., and Gerlach, R. H. (2016). Estimating quantile families of loss distributions for non-life insurance modelling via L-moments. *Risks*, 4(2):14.

Peters, G. W. and Sisson, S. (2006). Bayesian inference, Monte Carlo sampling and operational risk. *Journal of Operational Risk*, 1(3):27–50.

Tukey, J. W. (1977). Modern techniques in data analysis. In *NSF-Sponsored Regional Research Conference at Southern Massachusetts University, North Dartmouth*.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50:1–26.