Flexible dependence modeling of operational risk losses and its impact on total capital requirements

Eike Christian Brechmann[*]   Claudia Czado[*]   Sandra Paterlini[†]

March 14, 2013

Abstract

Operational risk data, when available, are usually scarce, heavy-tailed and possibly dependent. In this work, we introduce a model that captures such real-world characteristics and explicitly deals with heterogeneous pairwise and tail dependence of losses. By considering flexible families of copulas, we can easily move beyond modeling bivariate dependence among losses and estimate the total risk capital for the seven- and eight-dimensional distributions of event types and business lines. Using real-world data, we then evaluate the impact of realistic dependence modeling on estimating the total regulatory capital, which turns out to be, as often expected, up to 38% smaller than what the standard Basel approach would prescribe.

Keywords: operational risk, risk capital, dependence modeling, zero inflation, Student's t copula, vine copula,

1 Introduction

The magnitude of operational losses observed in recent years and their potential systemic effects has pointed out the need of development of realistic and therefore often more sophisticated quantitative risk management models (see Basel Committee on Banking Supervision [2009b]). Among the main challenges in operational risk modeling, we have the presence of very heterogeneous losses, usually classified in a matrix of 56 risk classes (seven event types (ETs) × eight business lines (BLs)[1]; see Basel Committee on Banking Supervision [2006]), scarcity of data and large numbers of zero losses for some classes, short time series with extreme tails and the need to estimate quantiles at very high confidence levels. In fact, banks are required to calculate the minimum capital requirement as the 99.9%-Value-at-Risk of the loss distribution such that

$$\text{MCR} = \text{VaR}_{99.9\%}\left(\sum_{j=1}^{56} S_j\right), \tag{1.1}$$

where S_j is the aggregate loss of one of the 56 BL-ET combinations. It is clear that this quantity is influenced by the dependencies among the different risk classes. The standard approach of the Basel Committee on Banking Supervision [2006] requires banks to marginally calculate the risk

---

[*]Center for Mathematical Sciences, Technische Universität München, Germany

[†]European Business School, Wiesbaden, Germany

[1]Business lines: 1. Corporate Finance, 2. Trading and Sales, 3. Retail Banking, 4. Commercial Banking, 5. Payment and Settlement, 6. Agency and Custody, 7. Asset Management, 8. Retail Brokerage. Event types: 1. Internal Fraud, 2. External Fraud, 3. Employment Practices & Workplace Safety, 4. Clients, Products & Business Practices, 5. Damage to Physical Assets, 6. Business Disruption & System Failures, 7. Execution, Delivery & Process Management.

capital of all 56 BL-ET combinations and then determine the risk capital as the sum of these 56 figures, that is

$$\text{MCR}^{\text{Basel}} = \sum_{j=1}^{56} \text{VaR}_{99.9\%}(S_j).$$ (1.2)

This corresponds to the assumption of comonotonicity (perfect dependence) among all 56 BL-ET combinations, which is often perceived by banks as a worst-case scenario for the MCR, assuming that $\text{MCR}^{\text{Basel}} \geq \text{MCR}$. However, due to the lack of subadditivity of the VaR measure for non-elliptical distributions (see Artzner et al. [1999]), it may also happen that $\text{MCR} \geq \text{MCR}^{\text{Basel}}$. The question if the standard Basel approach is too conservative has therefore been raised many times by practitioneers and researchers.

Regulators allow then, with due diligence, explicit dependence modeling under the Advanced Measurement Approach (AMA)[2]. The supervisory guidelines for the AMA banks [Basel Committee on Banking Supervision, 2011] underline that dependence modeling for operational risk is an evolving area, where many approaches are currently used by banks with potential impact on the capital requirements. Results from the 2008 Loss Data Collection Exercises [Basel Committee on Banking Supervision, 2009a] suggest that, among the AMA banks only 17% use correlation coefficients, while most rely on copulas (43%), with a preference for Gaussian copulas, and 31% AMA banks use other methods.

Research is high then on the regulators agenda to avoid spurious differences in exposure estimates (see Basel Committee on Banking Supervision [2011, §224]) and to provide sound guidelines for dependence modeling in operational risk, explicitly suggesting to move beyond Gaussian copulas and correlations coefficients. Theoretical and empirical evidence so far has mostly supported the idea that the assumption of perfect positive dependence is unduly strong and using internally determined correlations, as Basel II allows, could lead to lower the risk capital requirements while still providing adequate coverage for future losses (see Artzner et al. [1999], Chapelle et al. [2008] and Frachot et al. [2004]). However, recently Mittnik et al. [2011] have also shown that, despite only for a small number of risk classes, modeling bivariate dependence could also lead to increase the risk capital and tail dependence, which cannot be captured by Pearson's correlation coefficient, should not be ignored. What is the total impact on risk capital of explicit dependence modeling among all BLs and ETs is still, to our knowledge, a question with no answer, as estimating realistic multivariate operational risk management models with more than two cells can be computationally challenging and data are often scarce for model validation.

In this work, we aim at analyzing how much the risk capital estimate may change when modeling dependencies in multivariate settings. That is, we consider the total impact of explicit dependence modeling within the eight- or seven-dimensional BL and ET distributions by introducing a statistical model, which allows to explicitly consider the presence of extreme tails, heterogeneous pairwise dependence and large numbers of zero observations. In particular, we propose a flexible approach that, inspired by the work of Deb et al. [2009] on drug expenditures and Erhardt and Czado [2012] on dependent health insurance claims, directly models the dependence between the aggregate losses in BL-ET combinations using copulas. Since the non-occurrence of losses (zero events) also conveys information about the dependence characteristics, we explicitly incorporated it to allow a more accurate assessment of dependence. Finally,

---

[2]§669d in Basel Committee on Banking Supervision [2006]: *"Risk measures for different operational risk estimates must be added for purposes of calculating the regulatory minimum capital requirement. However, the bank may be permitted to use internally determined correlations in operational risk losses across individual operational risk estimates, provided it can demonstrate to the satisfaction of the national supervisor that its systems for determining correlations are sound, implemented with integrity, and take into account the uncertainty surrounding any such correlation estimates (particularly in periods of stress). The bank must validate its correlation assumptions using appropriate quantitative and qualitative techniques."*

given that no excessive data aggregation is required, parameter estimation can be based on the maximum amount of available observations.

By using real-world data from the Italian Database of Operational Losses (DIPO) in the period from January 2003 to June 2011, we can provide new and much needed insights on the impact of different dependence modeling strategies on total capital requirements and their validation on real-world data. In fact, our results suggest that explictly modeling dependence and zero inflation can lead to a reduction, as often expected, of the total regulatory capital, which might turn out to be up to 38% smaller than what the Basel comonotonicity approach would prescribe.

The paper is organized as follows. Section 2 describes the general modeling framework by discussing the key components of our modeling strategy. Dependence modeling of positive losses is treated in Section 3, after introducing the key properties we consider as relevant to compare the four major copula classes that we investigate for multivariate operational risk modeling. Marginal modeling and a detailed description of the copulas families considered are reported in Appendix A and B, respectively. Modeling dependence among zero events as additional model component is then described in Section 4, and the computation of risk measures is subsequently discussed in Section 5. Section 6 finally provides the results of the empirical investigation on real-world data, while Section 7 concludes.

2 Zero-inflated dependence model

Common characteristics of operational risk data can be summarized as follows. First, if losses are modeled on a weekly or monthly basis, it may frequently occur that there are no losses observed for a particular BL or ET. An excessive number of zero losses is called *zero inflation*. Second, distributions of operational losses per BL or ET may be *heavy-tailed*, i.e. there is a significant probability of extreme losses that has to be taken into account. Third, different BLs and ETs are not independent. Most importantly, the type of dependence, especially the so-called *tail dependence*, may have a huge impact on risk capital estimates.

While the heavy tails in marginal distributions have already been extensively discussed in the literature (see for example Chavez-Demoulin et al. [2006] and Gourier et al. [2009]), appropriate zero-inflated dependence models for aggregate operational losses have, to our knowledge, not yet been proposed.

Copulas have been established as the fundamental concept of statistical dependence modeling. A d-dimensional copula is a multivariate distribution function on $[0,1]^d$ with uniformly distributed marginals. Its central role in dependence modeling is due to the theorem by Sklar [1959], which states that any multivariate distribution can be decomposed into its margins and a copula. More precisely, let $\boldsymbol{X} = (X_1, ..., X_d)' \sim F$ with marginal distributions $F_1, ..., F_d$, then

$$F(x_1, ..., x_d) = C_{1,...,d}(F_1(x_1), ..., F_d(x_d)), \quad x_1, ..., x_d \in (\mathbb{R} \cup \{-\infty, \infty\})^d, \qquad (2.1)$$

where $C_{1,...,d}$ is a d-dimensional copula. If $\boldsymbol{X}$ is a continuous random vector, then the copula $C_{1,...,d}$ is unique and the multivariate density f of $\boldsymbol{X}$ can be decomposed as

$$f(x_1, ..., x_d) = c_{1,...,d}(F_1(x_1), ..., F_d(x_d))f_1(x_1)...f_d(x_d), \qquad (2.2)$$

where $c_{1,...,d}$ is the copula density and $f_1, ..., f_d$ are the marginal densities of f. Comprehensive reference books on copulas are Joe [1997] and Nelsen [2006].

Now, suppose that we want to model the multivariate distribution of d BLs, ETs or cells of the 7×8 BL-ET matrix. For brevity, we henceforth always speak of d cells with $d \in \{7, 8, 56\}$. Let $S_j \geq 0$, $j = 1, ..., d$, denote the aggregate loss of the jth cell. Then, we explicitly model the

presence of zero inflation in the aggregate loss by defining the following binary random variable $W_j \sim P_{W_j}$ for each cell $j \in \{1, ..., d\}$ as

$$W_j := \begin{cases} 1, & \text{zero loss in cell } j \\ 0, & \text{otherwise} \end{cases}.$$

That is, W_j is the zero inflation component of the otherwise positive continuous part of S_j, which we denote by $S_j^+ > 0$. We obtain the following decomposition:

$$S_j = W_j \times 0 + (1 - W_j)S_j^+ = (1 - W_j)S_j^+ \geq 0. \tag{2.3}$$

If $P_{W_j}(0) = P(W_j = 0) = 1$, then there is no zero inflation.

Inspired by the work of Deb et al. [2009] on drug expenditures and Erhardt and Czado [2012], who build a three-dimensional model for dependent health insurance claims, we introduce the multivariate zero-inflated density of $\boldsymbol{S} = (S_1, ..., S_d)'$ and $\boldsymbol{W} := (W_1, ..., W_d)'$ as

$$f_{\boldsymbol{W},\boldsymbol{S}}(\boldsymbol{w}, \boldsymbol{s}) = p_{\boldsymbol{W}}(\boldsymbol{w})f_{\{S_j^+, \; j \in D(\boldsymbol{w})\}}(s_j, \; j \in D(\boldsymbol{w})), \tag{2.4}$$

where $\boldsymbol{w} := (w_1, ..., w_d)'$, $\boldsymbol{s} := (s_1, ..., s_d)'$ and $D(\boldsymbol{w}) = \{i \in \{1, ..., d\} : \; w_i = 0\}$. That is, $D(\boldsymbol{w})$ contains all indices, for which the respective component of $\boldsymbol{w}$ is equal to 0. In other words, these are the cells with non-zero events. The $|D(\boldsymbol{w})|$-dimensional density of $S_j^+, \; j \in D(\boldsymbol{w})$, is denoted by $f_{\{S_j^+, \; j \in D(\boldsymbol{w})\}}$; if $D(\boldsymbol{w}) = \emptyset$, then $f_\emptyset := 1$. The multivariate binary probability mass function of $\boldsymbol{W}$ is denoted by $p_{\boldsymbol{W}}$. More explicitly, we can write Expression (2.4) also as

$$\begin{aligned} f_{\boldsymbol{W},\boldsymbol{S}}(\boldsymbol{w}, \boldsymbol{s}) = p_{\boldsymbol{W}}(\boldsymbol{w}) \Big[&\mathbf{1}_{\{\boldsymbol{w}=(1,...,1)'\}} + \mathbf{1}_{\{\boldsymbol{w}=(0,1,...,1)'\}} f_{S_1^+}(s_1) + \mathbf{1}_{\{\boldsymbol{w}=(1,0,1,...,1)'\}} f_{S_2^+}(s_2) \\ &+ ... + \mathbf{1}_{\{\boldsymbol{w}=(1,...,1,0)'\}} f_{S_d^+}(s_d) + \mathbf{1}_{\{\boldsymbol{w}=(0,0,1,...,1)'\}} f_{S_1^+,S_2^+}(s_1, s_2) \\ &+ \mathbf{1}_{\{\boldsymbol{w}=(0,1,0,1,...,1)'\}} f_{S_1^+,S_3^+}(s_1, s_3) + ... + \mathbf{1}_{\{\boldsymbol{w}=(1,...,1,0,0)'\}} f_{S_{d-1}^+,S_d^+}(s_{d-1}, s_d) \\ &+ \mathbf{1}_{\{\boldsymbol{w}=(0,0,0,1,...,1)'\}} f_{S_1^+,S_2^+,S_3^+}(s_1, s_2, s_3) + ... + \mathbf{1}_{\{\boldsymbol{w}=(0,...,0)'\}} f_{S_1^+,...,S_d^+}(\boldsymbol{s}) \Big], \end{aligned}$$

where $\mathbf{1}$ is the indicator function. Clearly, only exactly one of the indicator functions is different from zero.

In this way, we separate the dependence of the zero inflation component from the dependence of the positive losses. By applying Sklar's theorem (2.2) to $f_{\{S_j^+, \; j \in D(\boldsymbol{w})\}}$, we further obtain

$$f_{\{S_j^+, \; j \in D(\boldsymbol{w})\}}(s_j, \; j \in D(\boldsymbol{w})) = c_{\{j \in D(\boldsymbol{w})\}}(F_{S_j^+}(s_j), \; j \in D(\boldsymbol{w})) \prod_{j \in D(\boldsymbol{w})} f_{S_j^+}(s_j),$$

where $c_{\{j \in D(\boldsymbol{w})\}}$ is the copula density of $C_{\{j \in D(\boldsymbol{w})\}}$, the $|D(\boldsymbol{w})|$-dimensional margin of the d-dimensional copula $C_{1,...,d}$ for all cells, which is assumed to be absolutely continuous such that the density exists. More explicitly, it holds that

$$C_{\{j \in D(\boldsymbol{w})\}}(u_j, \; j \in D(\boldsymbol{w})) = C_{1,...,d}(\boldsymbol{v}), \quad u_j \in [0, 1] \; \forall j \in D(\boldsymbol{w}),$$

where $\boldsymbol{v} := (v_1, ..., v_d)'$ with

$$v_j = \begin{cases} u_j, & j \in D(\boldsymbol{w}) \\ 1, & \text{otherwise} \end{cases} \quad \forall j = 1, ..., d.$$

Then, it follows that we can state Expression (2.4) in terms of the copula $C_{1,...,d}$ and its margins as

$$f_{\boldsymbol{W},\boldsymbol{S}}(\boldsymbol{w}, \boldsymbol{s}) = p_{\boldsymbol{W}}(\boldsymbol{w})c_{\{j \in D(\boldsymbol{w})\}}(F_{S_j^+}(s_j), \; j \in D(\boldsymbol{w})) \prod_{j \in D(\boldsymbol{w})} f_{S_j^+}(s_j). \tag{2.5}$$

4

This means that, as in the classical copula approach, we can separate the dependence modeling in terms of the copula $C_{1,...,d}$ from the modeling of the marginal distributions $F_{S_j^+}$, $j = 1, ..., d$, in (2.5). In addition, the multivariate binary distribution $p_{\boldsymbol{W}}$ has to be modeled to account for the zero inflation component. Appropriate choices for these three model components are discussed in the following. We begin with the copula of positive losses (Section 3) and then turn to the zero loss distribution $p_{\boldsymbol{W}}$ (Section 4). Marginal distributions are briefly treated in Appendix A, while a detailed description of pros and cons of different copulas families is reported in Appendix B.

3   Copula modeling of positive losses

The between-cell copula is central to the model in order to appropriately capture dependencies in operational risk data. In the simplest setting, the cells are simply assumed to be either independent or perfectly dependent. As this is not necessarily the case, the question is what properties a reasonable copula for operational losses should exhibit.

(i) *Pairwise dependence:* Clearly, dependence among different pairs of cells may be heterogeneous. Therefore the between-cell copula should be flexible enough to allow for different strengths of dependence for different pairs.

(ii) *Tail dependence:* The between-cell copula should allow for the modeling of tail dependence. The presence of upper tail dependence means that very large losses tend to occur jointly rather than independently.

From a statistical point of view, the upper tail dependence coefficient is defined as a bivariate property: the probability of an extreme event of S_j^+ given an extreme event of S_k^+, $k \neq j$, that is

$$\lambda_{j,k}^U = \lim_{t \to 1^-} P(S_j^+ > F_{S_j^+}^{-1}(t) | S_k^+ > F_{S_k^+}^{-1}(t)) = \lim_{t \to 1^-} \frac{1 - 2t + C_{j,k}(t,t)}{1 - t},$$

which only depends on the copula of the pair (S_j^+, S_k^+). Obviously, $\lambda_{j,k}^U \in [0, 1]$ and, if $\lambda_{j,k}^U = 0$, (S_j^+, S_k^+) don't exhibit upper tail dependence.

(iii) *Interpretability:* Copulas may be specified in terms of many parameters with many different meanings. For reasons of internal and external communication, for example to the senior management or regulators, it is however desirable to work with a model that has parameters whose meaning can actually be interpreted. As we will see below, often parameters are directly linked to tail dependence or to measures of general dependence, such as the Kendall's τ rank correlation.

(iv) *Computational tractability:* To work with model (2.5) requires the availability of all multivariate copula margins. This is particularly important for statistical inference, where the density expression is required for likelihood-based techniques.

While most of the literature on dependence modeling of operational losses (see Dalla Valle et al. [2008] and Giacometti et al. [2008]) and also practitioners (see Basel Committee on Banking Supervision [2009a]) focus on elliptical copulas, we here more generally evaluate a range of different state-of-art multivariate copula models in light of the above four characteristics and discuss how useful they are for modeling operational risk data. In particular, we focus on Archimedean, elliptical, individual Student's t and vine copulas. While Archimedian and elliptical copulas have already found applications in operational risk modeling, individual Student's t and vine copulas
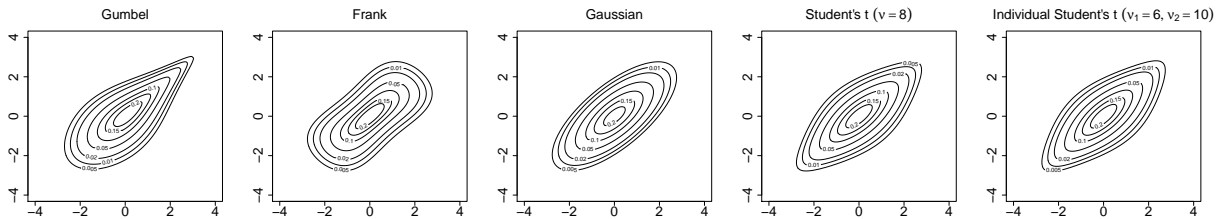
Figure 1: Contour plots of bivariate Archimedian Gumbel, Archimedian Frank, Gaussian, Student's t with 8 degrees of freedom as well as individual Student's t (6 and 10 degrees of freedom) copulas. Kendall's τ is 0.5 for all copulas. Margins are standard normal.

	Archim.	Gaussian	Student's t	Indiv. t	Vine
Pairwise dependence	−	+	+	+	+
Tail dependence	○	−	○	+	+
Interpretability	+	+	+	+	○
Comput. tractability	+	+	+	○	−

Table 1: Overview of copula characteristics: positive (+), neutral (○), and negative (−).

are here proposed, implemented and empirically tested for the first time. Moreover, Table 1 summarizes how the different copulas models behave (that is, either positively (+), neutrally (○), or negatively (−)) with respect to the desired characteristics of being capable of modeling pairwise and tail dependence while still being easily interpretable and computationally tractable.

While Archimedean copulas fail to model heterogeneous pairwise dependence, this is straightforward using elliptical copulas, which also are easily interpretable in terms of their underlying correlation matrices. However, standard elliptical copulas are somewhat restrictive when it comes to the tails. A flexible modeling of the tails is instead facilitated by the individual Student's t copula and vine copulas, at the expense of an increased computational complexity and, to some extent, a less obvious interpretation of parameters, especially for vine copulas. Contour plots of the different copula families are shown in Figure 1, especially to illustrate the different behavior in the upper tail, as indicated by sharper contour lines in the upper right corner of the plots: the tail-asymmetric Gumbel copula with only upper tail dependence is very sharp in the upper right corner, while Frank and Gaussian copulas do not show such behavior. The standard Student's t and also the individual Student's t copula are also both sharper than Gaussian copulas in the tails, corresponding to their non-zero tail dependence. We refer the reader to Appendix B for a detailed description of the pros and cons of the different copulas model, with respect to the four key-characteristics we have proposed above.

4   Copula modeling of zero losses

The random vector $\boldsymbol{W}$ is multivariate binary, for which distributions are however rather nonstandard and often require an excessive number of parameters (see for example Johnson et al. [1997]). We propose to use a copula approach here, which for clarity we first illustrate in a bivariate example. For this, let $\boldsymbol{W} = (W_1, W_2)'$ and recall that $P(W_i = w_i) = P(W_i \leq w_i) - P(W_i \leq w_i - 1)$, $i = 1, 2$. Similarly, it holds for the bivariate probability mass function

$p_{\boldsymbol{W}}$ that

$$
\begin{aligned}
p_{\boldsymbol{W}}(\boldsymbol{w}) &= P(W_1 = w_1, W_2 = w_2) \\
&= P(W_1 \leq w_1, W_2 \leq w_2) - P(W_1 \leq w_1 - 1, W_2 \leq w_2) \\
&\quad - P(W_1 \leq w_1, W_2 \leq w_2 - 1) + P(W_1 \leq w_1 - 1, W_2 \leq w_2 - 1) \\
&= C_{\boldsymbol{W}}(P_{W_1}(w_1), P_{W_2}(w_2)) - C_{\boldsymbol{W}}(P_{W_1}(w_1 - 1), P_{W_2}(w_2)) \\
&\quad - C_{\boldsymbol{W}}(P_{W_1}(w_1), P_{W_2}(w_2 - 1)) + C_{\boldsymbol{W}}(P_{W_1}(w_1 - 1), P_{W_2}(w_2 - 1)),
\end{aligned}
$$

where we used Sklar's theorem (2.1) with an appropriate bivariate copula $C_{\boldsymbol{W}}$. In general, $p_{\boldsymbol{W}}$ can be represented as

$$
p_{\boldsymbol{W}}(\boldsymbol{w}) = \sum_{k_1=1}^{2} ... \sum_{k_d=1}^{2} (-1)^{k_1 + ... + k_d} C_{\boldsymbol{W}}\left(u_1^{(k_1)}, ..., u_d^{(k_d)}\right), \tag{4.1}
$$

where $u_j^{(1)} = P_{W_j}(w_j)$ and $u_j^{(2)} = P_{W_j}(w_j - 1)$ for $j = 1, ..., d$ (see Song [1997]). For binary margins P_{W_j}, it is either $w_j = 0$ or $w_j = 1$. If $w_j = 0$, then $u_j^{(1)} = P_{W_j}(0)$ is the probability of a non-zero loss and $u_j^{(2)} = P_{W_j}(-1) = 0$. Conversely, if $w_j = 1$, then $u_j^{(1)} = P_{W_j}(1) = 1$ and $u_j^{(2)} = P_{W_j}(0)$.

The copula $C_{\boldsymbol{W}}$ can be any d-dimensional copula. As before, we recommend to use a copula that allows for heterogeneous pairwise dependence, which makes Archimedean copulas of no interest here. As vine copulas, on the other hand, do not have a closed-form copula expression, they should also be ruled out (see Panagiotelis et al. [2012] for an alternative, Pair Copula Construction (PCC)-based approach to modeling multivariate discrete data). Finally, to evaluate Expression (4.1), 2^d evaluations of the copula are needed, resulting in being very time-consuming and such that also the individual Student's t copula, which is of more complicated form than standard elliptical copulas, typically is not a good choice. Gaussian or Student's t copulas may, on the other hand, be used here, since fast algorithms for the evaluation of $C_{\boldsymbol{W}}$ are available [Genz and Bretz, 2009]. Estimation is carried out by maximum likelihood techniques.

After defining the copulas to be used for modeling positive losses and zero inflation, we can then estimate the model in (2.5) and evaluate the impact of realistic dependence modeling on total capital requirement estimates, as described in the next sections.

5 Risk measures

A major purpose of a multivariate model for operational losses is of course an accurate assessment of the regulatory risk capital to be held to cover future losses. The standard risk measure for computing the operational risk capital under Basel II is the Value-at-Risk (VaR) at the 99.9% level, as reported in (1.1). In general, for a level $\alpha \in [0, 1]$ the α-VaR is defined as

$$
\text{VaR}_\alpha(L) = F_L^{-1}(\alpha), \tag{5.1}
$$

where $L = \sum_{j=1}^{d} S_j$ is the total operational loss over d cells and F_L its distribution function (see for example McNeil et al. [2005]). Since F_L is however not known in closed form, it has to be obtained by simulation. To simulate N losses s_{ij}, $i = 1, ..., N$, for each cell $j \in \{1, ..., d\}$ using our multivariate model (2.5), we proceed as follows.

(i) Obtain N samples u_{ij}, $i = 1, ..., N$, $j = 1, ..., d$, from the copula for positive losses. For Archimedean, elliptical and vine copulas see Mai and Scherer [2012], for the individual Student's t copula see Luo and Shevchenko [2010].

7

(ii) Set $s_{ij}^+ := F_{S_j^+}^{-1}(u_{ij}) > 0$ for $i = 1, ..., N$ and $j = 1, ..., d$ to invert (B.9).

(iii) Obtain N samples w_{ij}, $i = 1, ..., N$, $j = 1, ..., d$, from the copula for zero losses

(iv) Set $s_{ij} := (1 - w_{ij})s_{ij}^+ \geq 0$ for $i = 1, ..., N$ and $j = 1, ..., d$ according to (2.3).

In this way, we obtain N samples of the total operational loss $l_i := \sum_{j=1}^d s_{ij}$, $i = 1, ..., N$, and the α-quantile $F_L^{-1}(\alpha)$, that is the α-VaR, can be approximated by the corresponding empirical quantile of $l_1, ..., l_N$.

6 Empirical Analysis

After introducing our multivariate model for operational losses, we now evaluate its practical implications by using real-world data from the Italian Database of Operational Losses (DIPO). So far, the number of studies using real-world data is still small in operational risk due to the scarcity of available dataset and the reluctance of the financial industry to share their private and often very sensitive information. Using data from a pool of banks and financial institutions, we can provide very much needed insights on the effects of using more sophisticated risk management modeling techniques on risk capital estimation and accuracy.

6.1 Data

Our dataset comprises operational losses reported from 33 Italian banking groups with about 180 entities to the Italian Database of Operational Losses (DIPO). The time period we consider is from January 2003 to June 2011 for a total of 102 months or 451 weeks. The reporting threshold is 5000 Euro, below which no loss is reported. Further, it is known for each loss which ET and BL are affected. Figure 2 displays a pie chart for the proportion of aggregate losses (left) and of numbers of losses (right) for each BL-ET combination with share larger than 1%. The figure shows that the losses per cell are rather heterogeneous. That is, operational losses in certain BLs and ETs occur more often and are of different magnitude. Especially Retail Banking (BL3) and Retail Brokerage (BL8) are often subject to operational losses, while the most frequent event types are External Fraud (ET2), Execution, Delivery & Process Management (ET7) and, in particular, Clients, Products & Business Practices (ET4).

We choose here to model on a weekly basis to balance the trade-off between having sufficiently many observations to avoid a large number of zero losses in each cell and also to being able of an accurate marginal and dependence modeling. In our investigation, we will both model the multivariate dependence between the eight BLs and between the seven ETs, to evaluate also the impact of such choice in the total risk capital estimate. Both approaches are followed in practice (see Basel Committee on Banking Supervision [2009a]), but so far no investigation showed the implications behind them. A 56-dimensional model for all BL-ET combinations is subject to future research, as more data will become available. Figure 3 shows pairwise scatter plots of the aggregate losses of the seven ETs on the log scale. The general level of dependence appears to be rather weak, as the observations are rather dispersed, but there are also cases with seemingly stronger dependence such as between ETs 4 and 7. The dependence between positive losses will be investigated in more detail in Section 6.3. Figure 3 also indicates that there is a significant number of zero events that needs to be taken into account explicitly (see Section 6.4).

6.2 Marginal modeling

For the estimation of weekly operational losses per BL and ET we consider the following frequency and severity distributions: zero-truncated Poisson, generalized Poisson [Consul and Jain,
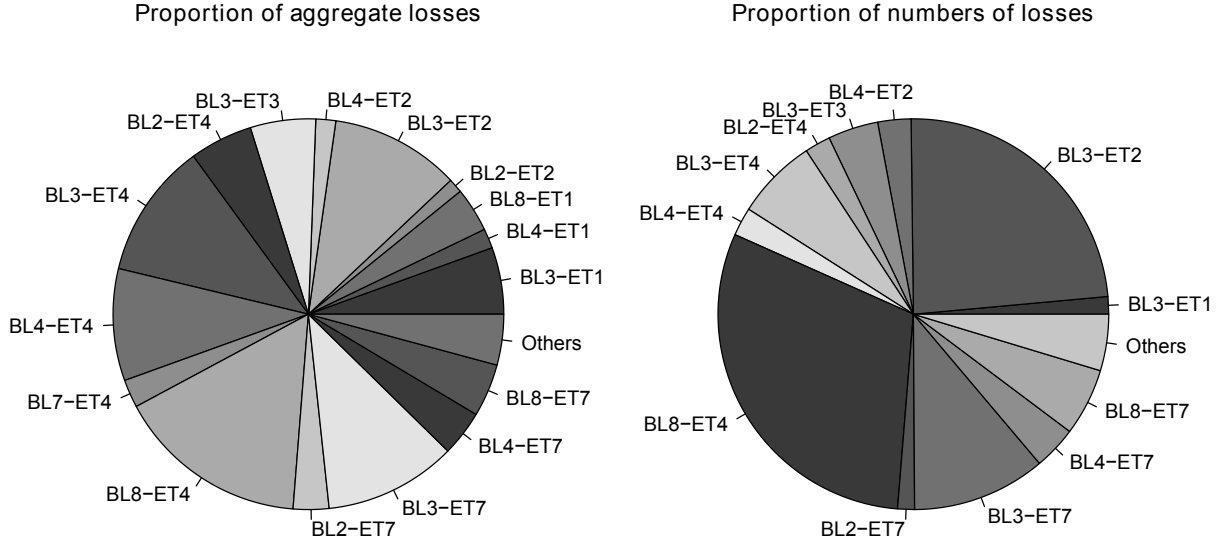
Figure 2: Proportion (in %) of aggregate losses S_j (upper table) and numbers of losses N_j (lower table) for each BL-ET combination.

1973] and negative binomial for loss frequency and gamma, Weibull, log normal, Pareto and generalized Pareto for loss severity, where each severity distribution is translated to the reporting threshold of 5000 Euro (see Expression (A.1)).

Using QQ-plots and goodness-of-fit tests, it turns out that negative binomial fits for the loss frequency are quite good and always superior to the respective Poisson fits, which postulate equidispersion that cannot be observed in the available data. Generalized Poisson fits, which also allow for non-equidispersion, gave similar results as the negative binomial, so that we decided to use the latter model. In an exploratory pre-analysis, we also determined that significantly fewer losses are observed in three weeks in August. We account for this holiday season effect by fitting appropriate mean regressions with indicator variable for these three weeks[3]. Severities were not observed to be different in summer compared to the rest of the year.

For the loss severities, we also used QQ-plots and goodness-of-fit tests to determine the best fitting distributions per BL and ET. From the above list of distributions the log normal gave the best fit for the individual positive losses. Figure 4 shows that the fitted log normal distribution functions in fact very closely follow the empirical distribution functions of the different BLs. A

---

[3]More precisely, let f_{N+} denote the probability mass function of the zero-truncated negative binomial distribution for the number of losses of a given cell. Then,

$$f_{N+}(n) = \frac{\Gamma(n+k)}{\Gamma(k)\,n!} \left(\frac{k}{k+\mu_i}\right)^k \left(\frac{\mu_i}{k+\mu_i}\right)^n \Big/ \left(1 - \left(\frac{k}{k+\mu_i}\right)^k\right), \quad n \in \mathbb{N},$$

with season-dependent mean parameter $\mu_i > 0$ and size $k > 0$. We model μ_i as

$$\mu_i = \exp\left(\beta_0 + \beta_1 I_{\text{summer}}(i)\right),$$

where $\beta_0, \beta_1 \in \mathbb{R}$ denote regression parameters, and the indicator variable I_{summer} is defined as

$$I_{\text{summer}}(i) := \begin{cases} 1, & \text{observation } i \text{ occurs in one of the three summer weeks} \\ 0, & \text{otherwise} \end{cases}.$$

That is, we have two different frequency distributions for each BL and ET, depending on whether a loss occurs in summer or not.
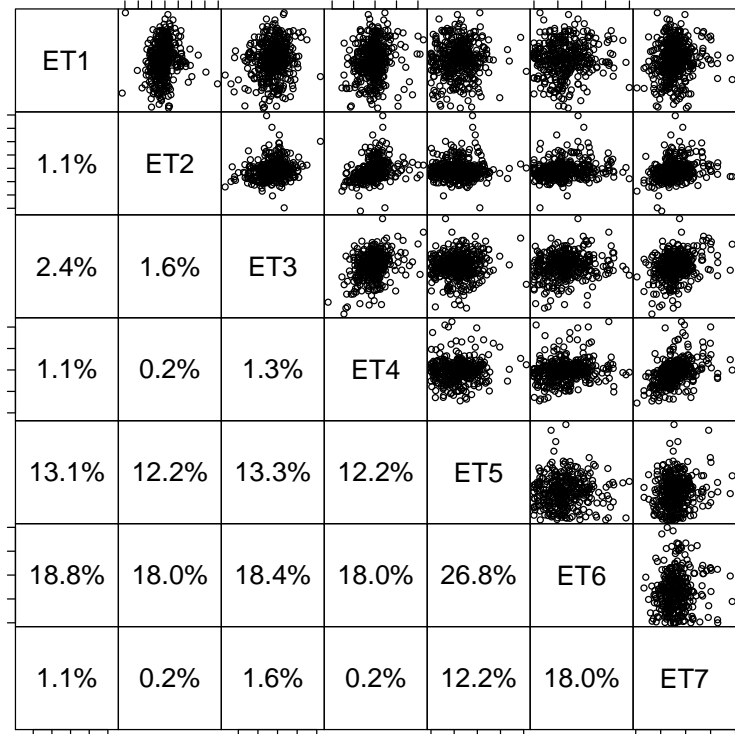
Figure 3: Pairwise scatter plots of aggregate losses S_j^+ per ET on the log scale. The lower triangle shows the percentage of joint zero events of the pairs (S_j^+, S_k^+). Labels are omitted from the axes for confidentiality reasons.

look at the very tail of the distributions (see Figure 5) underlines this. For ETs similar results hold.

In the next step, we then compute the convolution of the chosen severity and frequency distributions for each BL and ET by Monte Carlo simulation with 100'000 samples (see Klugman et al. [2008]). For each BL and ET we obtain two convolutions: one for the three summer weeks with significantly fewer losses and one for the rest of the year. This provides us with estimates $\widehat{F}_{S_j^+}$ and $\widehat{F}_{S_j^+}^s$ of the marginal aggregate loss distribution functions $F_{S_j^+}$ and $F_{S_j^+}^s$, where the latter indicates the one for the three summer weeks. (Note that we have $j = 1, ..., 8$ for BLs and $j = 1, ..., 7$ for ETs.) These estimates are then used to transform the losses s_{ij}, $i = 1, ..., 451$, to approximately uniform data following the IFM approach described in Appendix B.5. That is, we set

$$\widehat{u}_{ij} := \begin{cases} \widehat{F}_{S_j^+}^s(s_{ij}), & \text{observation } i \text{ occurs in one of the three summer weeks} \\ \widehat{F}_{S_j^+}(s_{ij}), & \text{otherwise} \end{cases}. \tag{6.1}$$

These pseudo observations are then used in the following dependence analysis.

6.3 Copula modeling of positive losses

Based on the marginal transformations, we then fitted the copula models for positive losses of Section 3 to the $\widehat{u}_{ij}$s (6.1). As Archimedean copulas we chose the Frank, which is symmetric
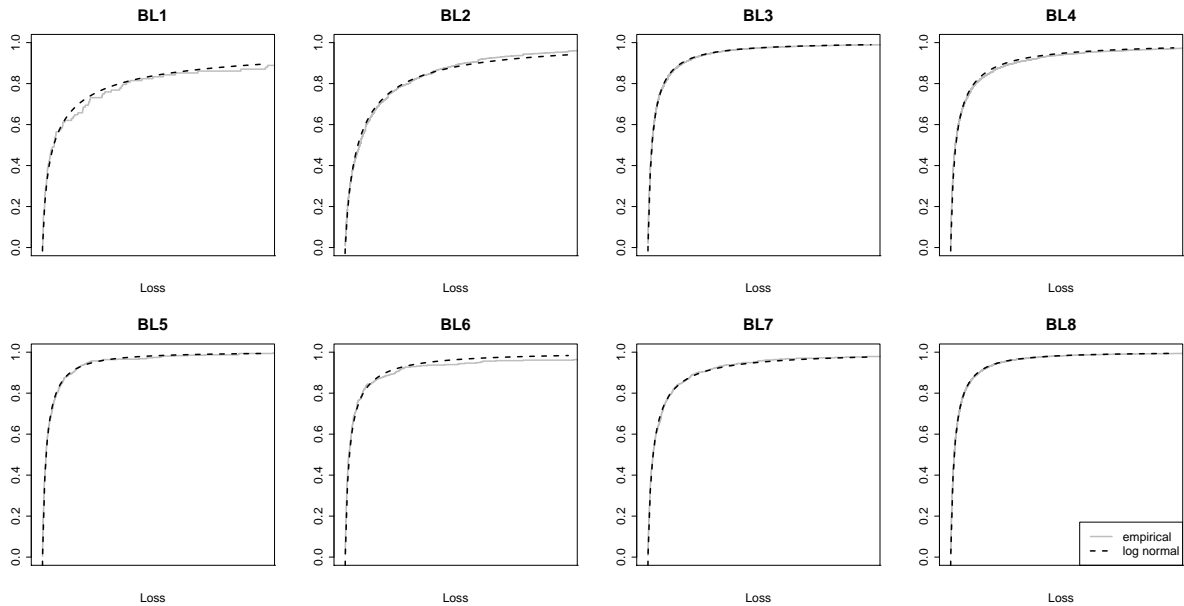
Figure 4: Comparison of empirical and fitted log normal distribution functions for all eight BLs and the main range of the individual losses X_{ij}, $i = 1, ..., N_j$. Labels are omitted from the horizontal axes for confidentiality reasons.

BL/ET	1	2	3	4	5	6	7	8
1		-0.14	0.02	-0.07	0.08	0.18	0.05	-0.10
2	0.12		0.12	0.16	0.00	0.06	0.10	0.20
3	0.08	0.15		0.23	0.01	0.12	0.10	0.23
4	0.10	0.30	0.16		0.01	0.04	0.06	0.22
5	0.04	-0.01	0.06	0.04		-0.06	0.12	0.07
6	0.05	0.13	0.10	0.12	0.05		0.18	0.06
7	0.07	0.20	0.18	0.29	0.07	0.08		0.10

Table 2: Empirical Kendall's τ values of jointly observed pairs of positive aggregate losses (S_j^+, S_k^+) per BL (upper triangle) and per ET (lower triangle).

and does not exhibit any tail dependence, and the Gumbel, which has upper but no lower tail dependence. Upper tail dependence is of particular interest here, because it describes the joint probability of very large losses and therefore needs to be accounted for in risk capital calculations (see Section 6.5). General dependence is however rather weak between BLs and ETs, respectively: Pairwise empirical Kendall's τ values between BLs range from -0.14 to 0.23 and between ETs from -0.01 to 0.30 (see Table 2, and also Figure 3). This is in line with previous studies (see Dalla Valle et al. [2008], Cope and Antonini [2008], and Giacometti et al. [2008]) and clearly different from comonotonicity as postulated by the standard approach (see Expression (1.2)).

As described in Appendix B.4, the use of vine copulas in our multivariate model for operational losses defined in (2.5) is quite challenging, since multivariate margins are not available in closed form but involve possibly high-dimensional integration. We therefore checked both for BLs and ETs which combinations of non-zero events are actually observed, that is which margins need to evaluated. For a vine copula to be tractable for statistical inference, it is important to construct it such that only low-dimensional integrations are required for the margins. Such a
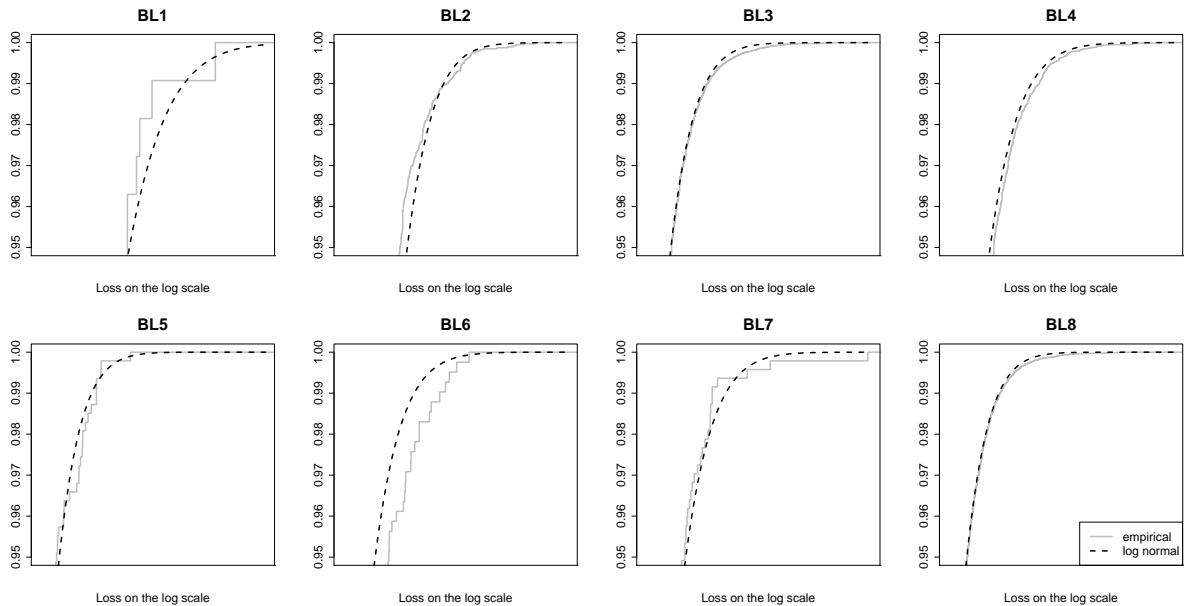
Figure 5: Comparison of empirical and fitted log normal distribution functions on the log scale for all eight BLs in the upper tail of the individual losses X_{ij}, $i = 1, ..., N_j$. Labels are omitted from the horizontal axes to maintain confidentiality.

vine copula can be determined as described in Appendix C.

We found a PCC for ETs such that only one-dimensional integration is needed to evaluate the multivariate operational loss model (2.5). This is similar to the individual Student's t copula. Unfortunately, for BLs the best possible vine copula still required up to three-dimensional integrations, which are numerically very demanding. We therefore do not consider a vine copula for BLs. For ETs, we fit two different vine copulas: one with only Student's t copulas and one with Student's t copulas for all unconditional variable pairs and Gaussian copulas for all pairs that are specified conditionally.

Log likelihoods, numbers of parameters as well as AIC values (to penalize the log likelihood for the number of parameters) of the considered models for BLs and ETs are shown in Table 3. While both Archimedean copulas are clearly not appropriate for the operational risk data per BL or ET, elliptical copulas provide a good fit. The individual Student's t copula however provides an even better fit than the standard Student's t copula. This indicates that tail dependencies are not only present but also quite heterogeneous, implying that the standard Student's t copula with only one parameter for the degrees of freedom is too restrictive. In fact, the estimated numbers of degrees of freedom of the standard Student's t copula are 42.04 for BLs and 54.78 for ETs, while the estimated individual degrees of freedom of the individual Student's t copula vary between 6.30 and more than 300 for BLs (of which three out of eight are below 50) and between 7.62 and also more than 300 for ETs (of which three out of seven are below 50). The standard Student's t copulas therefore create the false impression that dependence is almost Gaussian. The large estimated numbers of degrees of freedom can be seen as averages of the respective estimated individual degrees of freedom, which are rather heterogeneous. Entries of the correlation matrices range between -0.23 and 0.31 for BLs and between 0 and 0.33 for ETs, corresponding to a rather weak to medium level of dependence. Computing tail dependence coefficients for these parameters however results in values very close to 0. This is a very interesting result, since the use of the two Student's t copulas allows for an exact quantification of the tail dependence ($\lambda_{j,k}^U \in [0,1]$; see (B.4)), which the non-tail dependent Gaussian copula does not allows for

12

	BLs			ETs		
	log lik.	#par.	AIC	log lik.	#par.	AIC
Gumbel	29.14	1	-56.27	51.99	1	-101.99
Frank	36.62	1	-71.23	70.96	1	-139.92
Gaussian	92.92	28	-129.83	107.14	21	-172.28
Student's t	96.01	29	-134.02	111.32	22	-178.65
Individual Student's t	105.35	36	-138.70	119.50	28	-183.01
Vine (uncond.: t, rest: Gaussian)	-	-	-	112.80	27	-171.61
Vine (Student's t)	-	-	-	121.39	42	-158.78

Table 3: Log likelihoods, numbers of parameters and AIC values of copulas for dependence among BLs.

$(\lambda^U_{j,k} = 0)$.

The two vine copulas for dependence between ETs as alternative extension of the standard Student's t copula also improve the log likelihood, as larger values imply better model fitting. This also stresses the need for a flexible tail dependence modeling, but as before estimated tail dependence coefficients are very small and close to 0. In contrast to the standard and the individual Student's t copulas, the vine copulas however strongly suffer from the large number of model parameters, such that AIC values are better for the elliptical copulas. The number of parameters could however be reduced by setting Student's t copulas with, for example, more than 100 degrees of freedom to Gaussian copulas. This would hardly impact the log likelihood value, but significantly reduce the number of parameters. If a correlation parameter is very small and close to independence, one could even think of using independence copulas to reduce the number of parameters even further. Also the individual Student's t copula could be simplified to some extent, either by grouping variables with similar numbers of degrees of freedom (the resulting model being the grouped Student's t copula by Daul et al. [2003]) or by removing the χ^2 mixing variables in the definition of the individual Student's t copula (B.5) for components with a large number of degrees of freedom.

6.4 Copula modeling of zero losses

As discussed in Section 4, essentially any copula could be used to construct a flexible multivariate binary distribution for the zero losses, but due to computational and other limitations we focus here on a Gaussian copula. Empirical marginal probabilities of zero losses per BL range between 0 and 0.81 in the first BL, while for ETs the largest number of zero losses is observed for the sixth ET with about 0.20.

In a pre-analysis, it turns out that many of the empirical pairwise probabilities of zero losses are very small (below 0.01). We therefore restrict our copula analysis to those BLs and ETs with significant pairwise probabilities. These are the BLs 1, 5, 6 and 7 and the ETs 3, 5 and 6. That is, we fit a four-dimensional Gaussian copula for zero losses of BLs and a three-dimensional one for the ETs. Estimated entries of the correlation matrix range between 0.02 and 0.28 for BLs and between 0.05 and 0.51 for ETs.

6.5 Risk measures

After carefully modeling the dependence between positive and zero losses for BLs and ETs, we evaluate the different models in terms of their risk capital estimates and assess the diversification benefit compared to comonotonicity assumption (see also Embrechts et al. [2009]). In particular, we focus on the Gaussian, the Student's t, the individual Student's t and the vine copula with
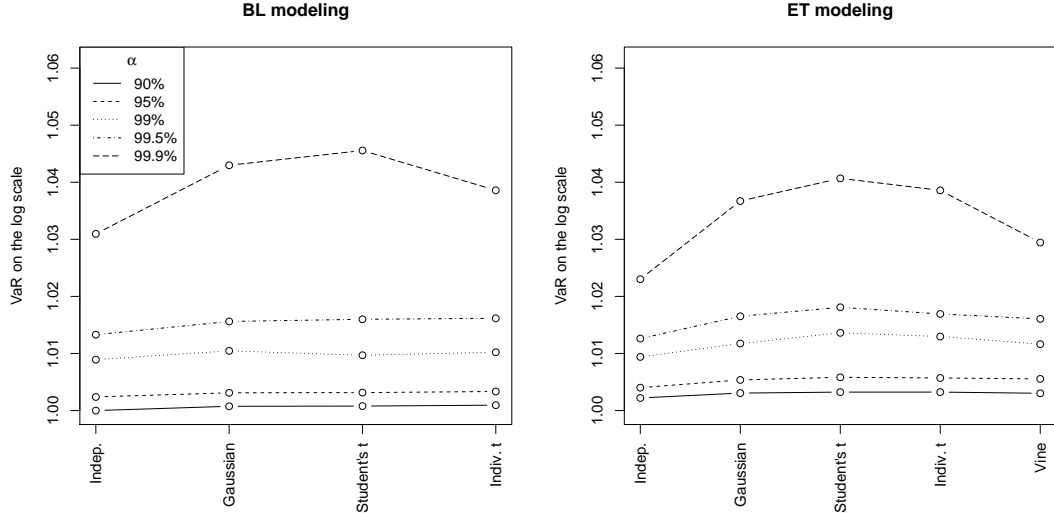
Figure 6: Annual α-VaRs on the log scale for $\alpha = 90\%, 95\%, 99\%, 99.5\%, 99.9\%$ and for the different copula models for positive losses. The left panel shows the results for BL modeling, the right for ET modeling. Vertical axes have been scaled to maintain confidentiality.

only Student's t copulas, since they provided the best fit. They differ mainly in if and how tail dependencies are modeled. For comparison, a model with independence between positive and zero losses is also used.

We obtained 100'000 yearly samples from the different models, that is we simulated each week of the year respecting the different marginal distributions in summer and then summed over the weeks. Based on these samples we computed annual VaRs (5.1) for different levels (i.e. $\alpha = 90\%, 95\%, 99\%, 99.5\%, 99.9\%$) as shown in Figure 6. Obviously, risk measures are lowest when no dependence among BLs and ETs is assumed. Using a Student's t copula as underlying dependence model however results in the highest VaR estimates, while its extensions, the individual Student's t and the vine copula with only Student's t copulas, lead to smaller estimates. This is very interesting, since a standard Student's t copula is often chosen instead of a Gaussian copula to respect tail dependence, but it apparently leads to overestimation of quantiles. This is certainly due to the restrictive modeling approach of allowing only one parameter to govern the overall level of tail dependence. Although tail dependence here is very small, its impact is obvious and a more accurate modeling of it, using extensions of the Student's t copula, yields refined estimates of risk measures. Also, notice that when $\alpha = 99.9\%$ as under Basel II, differences between the Gaussian, Student's t and individual Student's t copula are more pronounced compared to the independence assumption. Comparing the BL and the ET modeling shows that the picture is essentially the same for both. This is reassuring given that the modeling is based on the same data, which only have been grouped differently. nual VaR estimates are not entirely unrealistic, we obtained twelve months of additional data, that is one new annual observation.

The second question of interest is the diversification effect of considering the VaR of the total loss rather than the sum of the individual VaRs (see Expressions (1.1) and (1.2)). That is, we are interested in the fraction

$$\text{Div}_\alpha := \frac{\text{VaR}_\alpha(\sum_{j=1}^d S_j) - \sum_{j=1}^d \text{VaR}_\alpha(S_j)}{\sum_{j=1}^d \text{VaR}_\alpha(S_j)}, \tag{6.2}$$

which gives the relative reduction in the VaR compared to the Basel comonotonicity assumption. If $\text{Div}_\alpha < 0$, there is a diversification benefit.
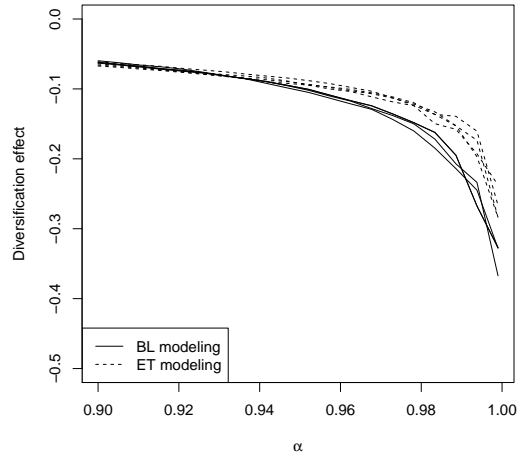
Figure 7: VaR diversification effect Div_α (6.2) for $\alpha \in [90\%, 99.9\%]$ and for the Gaussian, Student's t and individual Student's t copula models for positive losses of BLs and ETs. Since differences between lines are very small, no distinction between the different models is made.

The diversification effect for the copula models considered here is displayed in Figure (6.2). There is a clear diversification benefit of up to 38% for BLs and 35% for ETs when explicitly modeling the dependence. This mirrors the fact that dependence between BLs and ETs, both in general and in the tails, is observed to be rather weak and therefore clearly differs from the Basel comonotonicity standard approach, where $\text{Div}_\alpha = 0$. Interestingly, there is no clear-cut difference between the different models that we choose as the most accurate ones among a larger pool of models (see Section 6.3). This indicates that the diversification effect mainly is driven by the general level of dependence modeled, which is about similar for the different models considered here.

7   Conclusion

In this work, we introduce a rather general and flexible multivariate modeling approach for operational risk losses, which explicitly takes into account the multivariate dependence among losses and the presence of scarce data. Our aim is to introduce a more accurate model and then evaluate its implications on the estimation of the total risk capital compared to the Basel II comonotonicity assumption for the entire set of BLs and ETs using real-world data. Explicit dependence modeling is discussed critically by considering different copula classes and introducing, from both statistical and business perspectives, four key characteristics the ideal model should allow to take into account: heterogeneous pairwise dependence, tail dependence, interpretable parameters and computational tractability. It turns out that from a theoretical perspective, the individual Student's t copulas is probably the most appealing for operational losses as it only poses a moderate computational challenge, while still being easily interpretable in its parameters and structure and allowing for the presence of different pairwise and tail dependencies. Compared to the Student's t copula with only one degrees of freedom parameter, tail dependence is determined by the degrees of freedom of the individual variables and therefore flexible in considering potentially heterogeneous behaviors among losses. Vine copulas add even more flexibility, but at the expense of a more complex parameter interpretability and computational tractability, which make it less appealing in our zero-inflated model framework. Gaussian copulas are an

15

option when it is reasonable to assume no tail dependence, while Archimedean copulas, despite used in practice, are of less interest for operational risk modeling as they cannot account for heterogeneous pairwise dependence.

Empirical results on real-world data suggest that Gaussian and even better Student's t copulas can provide a good fit to positive losses grouped by BLs or ETs, despite the Gaussian implies zero tail dependence and the Student's t copula results in a potential overestimation of tail dependence due to its inflexible degrees of freedom parameter. However, as expected, the individual Student's t copula shows the best fit in terms of AIC and log likelihood to the data at hand and allows to build an easily interpretable model for operational losses, from which to run Monte Carlo estimation to determine the overall effect on risk capital. In fact, while the Student's t copula results in the highest VaR estimates for both BL and ET modeling, with an increase with respect to the independence assumption of up to 35% and 43% for BLs and ETs, respectively, individual Student's t copula estimates result only in an increase of 17% and 37%. This reflects the more accurate assessment of tail dependence by the individual Student's t copula, which can account for heterogeneous tail dependence through the individual degrees of freedom per variable, of which the standard Student's t copula's degrees of freedom can be regarded as an average.

Finally, when considering the diversification ratio to evaluate the effect of a potential reduction of risk capital estimates compared to the standard Basel assumption of comonotonic losses, results suggest that a more realistic modeling of the multivariate distribution of BLs and ETs leads to a reduction in capital of up to 38% for BLs and 35% for ETs for high quantiles ($\alpha = 99.9\%$, with few differences between models). This confirms, as also previously reported in the literature for much simpler bivariate settings, that the comonotonicity assumption of the standard Basel framework is unduly strong when evaluating the overall effect. Banks might therefore have an incentive to move towards more sophisticated but then realistic and accurate risk management models. However, further investigations on different real-world data and larger sample sizes for out-of-sample evaluation are still required to draw irrefutable conclusions.

Acknowledgment

A   Marginal modeling

The topic of marginal modeling of operational losses has already been extensively discussed in the literature (see, amongst others, Aue and Kalkbrener [2006]). We therefore only briefly mention the most relevant points. Since here we only deal with one particular cell, we suppress indices relating to the different cells and simply denote the aggregate loss by S.

The aggregate loss S is the sum of individual operational losses within a given cell, that is

$$S = \sum_{i=1}^{N} X_i,$$

with the number of losses $N \sim F_N$ and the individual losses $X_i \sim F_X$, $i = 1, ..., N$, independent and identically distributed, where the severity distribution F_X is positive and continuous and the frequency distribution F_N is a count distribution with support on $\mathbb{N}_0 = \{0, 1, 2, ...\}$. The case $N = 0$ means that no losses are observed, that is $S = 0$ (the empty sum is taken to be zero). As this case is explicitly taken into account by the zero loss variable W, we denote the positive number of losses as $N^+ \in \mathbb{N} = \{1, 2, ...\}$ with zero-truncated distribution F_{N^+}. Therefore, we have

$$S^+ = \sum_{i=1}^{N^+} X_i.$$

Typical choices for F_{N^+} are the zero-truncated Poisson and negative binomial distributions (see, e.g., Grogger and Carson [1991]). As is commonly known, the Poisson distribution assumes equidispersion, that is it fixes mean and variance to be equal. Since this is a very strict assumption, the negative binomial distribution may often be a more realistic choice.

It is typically the case that operational losses below some threshold are not reported. That is, only losses above a certain threshold are observed. In this case, the severity distribution has to be translated to this threshold. Let F_0 be an arbitrary positive continuous distribution, then translation to the threshold $m > 0$ means that we define the severity distribution for $x \in (m, \infty)$ as

$$F_X(x) = \frac{F_0(x) - F_0(m)}{1 - F_0(m)}. \tag{A.1}$$

An important issue of the distribution of individual operational losses is that it may be heavy-tailed as noted above. Examples of heavy-tailed (or subexponential) distributions are in particular the log normal, Pareto, Weibull (with shape parameter smaller than 1) and generalized Pareto, while the also commonly used gamma distribution is light-tailed (see Embrechts et al. [1997]). Furthermore, according to a classification based on moments, Pareto tails can be shown to be heavier than log normal tails. An alternative proposed in extreme value theory is to use the generalized Pareto as explicit tail distribution above a certain threshold, while a different distribution is used for the body of the distribution (see Chavez-Demoulin et al. [2006] for more details).

B   Copulas Families

B.1   Archimedean copulas

A d-dimensional Archimedean copula is defined in terms of its generator function $\varphi : [0, 1] \to [0, \infty)$, whose inverse φ^{-1} is d-monotone on $[0, \infty)$ (see McNeil and Nešlehová [2009]):

$$C_{1,...,d}(u_1, ..., u_d) = \varphi^{-1} \left(\varphi(u_1) + ... + \varphi(u_d) \right), \quad u_1, ..., u_d \in [0, 1]. \tag{B.1}$$

Popular Archimedean copulas are Clayton, Gumbel and Frank. According to our proposed key-characteristics, this class exhibits:

(i) *Pairwise dependence:* Each margin of an Archimedean copula is again of the same copula type with the same parameter. That is, pairwise dependence is fixed to be the same for all pairs. This quite restrictive property is the reason why Archimedean copulas are mainly used in bivariate applications, but are not quite appropriate for multivariate applications with heterogeneous pairwise dependence.

(ii) *Tail dependence:* Depending on the chosen generator function, Archimedean copulas can exhibit different tail behavior. For instance, the Gumbel copula is upper tail dependent, while the Frank does not have any tail dependence.

(iii) *Interpretability:* For most Archimedean copulas there is a close correspondence between copula parameter(s) and the Kendall's τ rank correlation, which measures the general level of dependence. The Gumbel copula, which has just one parameter θ, for example implies the following Kendall's τ value: $\tau_{j,k} := \tau_{j,k}(\theta) = 1 - 1/\theta$ for each pair of cells j and k. The corresponding upper tail dependence coefficient is $\lambda_{j,k}^{U} := \lambda_{j,k}^{U}(\theta) = 2 - 2^{1/\theta}$. For other such relationships see, amongst others, Brechmann and Schepsmeier [2013].

(iv) *Computational tractability:* The major advantage of Archimedean copulas is that all margins are readily available and the copula often only depends on one or two parameters. This strongly simplifies statistical inference [Hofert et al., 2012].

Clearly, because of the insufficient flexibility in pairwise dependence, Archimedean copulas should ruled out as reasonable models for multivariate operational losses when moving beyond the bivariate case.

B.2 Elliptical copulas

By inverting Sklar's theorem (2.1), copulas can be constructed for arbitrary multivariate distribution functions F and marginal distribution functions $F_1, ..., F_d$ as

$$C_{1,...,d}(u_1, ..., u_d) = F(F_1^{-1}(u_1), ..., F_d^{-1}(u_d)), \quad u_1, ..., u_d \in [0, 1]. \tag{B.2}$$

Elliptical copulas are obtained by letting F be an elliptical distribution function (see Fang et al. [1990] and McNeil et al. [2005]) and $F_1, ..., F_d$ the corresponding margins. The most popular examples are the Gaussian and the Student's t copulas. Their properties are:

(i) *Pairwise dependence:* Pairwise dependence of Gaussian and Student's t copulas can be different for each pair. In the $d \times d$-correlation matrix each pair has a corresponding entry.

(ii) *Tail dependence:* While the Gaussian copula does not exhibit any tail dependence, the Student's t copula has symmetric upper and lower tail dependence. The degree of tail dependence for each pair is controlled by the degrees of freedom: the smaller the degrees of freedom, the stronger the tail dependence. Conversely, the Student's t copula converges to the Gaussian copula when the degrees of freedom go to infinity.

(iii) *Interpretability:* As noted above, pairwise dependence is essentially determined by the correlation matrix, whose entries can be directly related to the general level of dependence of each pair. The pairwise Kendall's τ of cells j and k is

$$\tau_{j,k} := \tau_{j,k}(\rho_{j,k}) = \frac{2}{\pi} \arcsin(\rho_{j,k}), \tag{B.3}$$

where $\rho_{j,k} \in (-1, 1)$ is the $(j, k)$the entry of the correlation matrix [McNeil et al., 2005]. In case of the Student's t copula with $\nu > 2$ degrees of freedom the upper (and lower) tail dependence of each pair can be obtained as

$$\lambda_{j,k}^{U} := \lambda_{j,k}^{U}(\rho_{j,k}, \nu) = 2t_{\nu+1}\left(-\sqrt{\nu+1}\frac{\sqrt{1-\rho_{j,k}}}{\sqrt{1+\rho_{j,k}}}\right), \tag{B.4}$$

where t_ν is the univariate Student's t distribution function with ν degrees of freedom.

(iv) *Computational tractability:* The margins of elliptical copulas are again elliptical of the same class. The correlation matrix is a sub-matrix of the full correlation matrix. Although easily tractable density expressions are available, statistical inference is complicated by the fact that the correlation matrix has to be positive definite. Good starting values for numerical optimization are often provided by computing pairwise empirical Kendall's τ estimates and then inverting Expression (B.3).

The above properties present elliptical copulas as an appealing model for multivariate operational losses. Their major disadvantage certainly is that the Gaussian copula does not exhibit any tail dependence, while the tail dependence of the Student's t copula is symmetric in both tails and governed by only one parameter for all pairs.

To illustrate the model defined in (2.5) in terms of elliptical copulas, we consider a trivariate example. If $d = 3$, then Expression (2.5) can be written as

$$
\begin{aligned}
f_{\boldsymbol{W},\boldsymbol{S}}(\boldsymbol{w},\boldsymbol{s}) = p_{\boldsymbol{W}}(\boldsymbol{w}) \Big[& \mathbf{1}_{\{\boldsymbol{w}=(1,1,1)'\}} + \mathbf{1}_{\{\boldsymbol{w}=(0,1,1)'\}} f_{S_1^+}(s_1) \\
& + \mathbf{1}_{\{\boldsymbol{w}=(1,0,1)'\}} f_{S_2^+}(s_2) + \mathbf{1}_{\{\boldsymbol{w}=(1,1,0)'\}} f_{S_3^+}(s_3) \\
& + \mathbf{1}_{\{\boldsymbol{w}=(0,0,1)'\}} c_{1,2}\big(F_{S_1^+}(s_1), F_{S_2^+}(s_2)\big) f_{S_1^+}(s_1) f_{S_2^+}(s_2) \\
& + \mathbf{1}_{\{\boldsymbol{w}=(0,1,0)'\}} c_{1,3}\big(F_{S_1^+}(s_1), F_{S_3^+}(s_3)\big) f_{S_1^+}(s_1) f_{S_3^+}(s_3) \\
& + \mathbf{1}_{\{\boldsymbol{w}=(1,0,0)'\}} c_{2,3}\big(F_{S_2^+}(s_2), F_{S_3^+}(s_3)\big) f_{S_2^+}(s_2) f_{S_3^+}(s_3) \\
& + \mathbf{1}_{\{\boldsymbol{w}=(0,0,0)'\}} c_{1,2,3}\big(F_{S_1^+}(s_1), F_{S_2^+}(s_2), F_{S_3^+}(s_3)\big) f_{S_1^+}(s_1) f_{S_2^+}(s_2) f_{S_3^+}(s_3) \Big] .
\end{aligned}
$$

That is, dependence is modeled by the trivariate copula $C_{1,2,3}$ with bivariate margins $C_{1,2}$, $C_{1,3}$ and $C_{2,3}$. If $C_{1,2,3}$ is an elliptical copula with correlation matrix

$$
R(\rho_{1,2}, \rho_{1,3}, \rho_{2,3}) = \begin{pmatrix} 1 & \rho_{1,2} & \rho_{1,3} \\ \rho_{1,2} & 1 & \rho_{2,3} \\ \rho_{1,3} & \rho_{2,3} & 1 \end{pmatrix},
$$

then each bivariate margin $C_{j,k}$, $1 \le j < k \le 3$, is also elliptical of the same class and has a 2×2-correlation matrix with parameter $\rho_{j,k}$. In case of the Student's t copula, each margin also has the same degrees of freedom as $C_{1,2,3}$.

B.3 Individual Student's t copula

Recently, Luo and Shevchenko [2010] defined the Student's t copula with multiple degrees of freedom or individual Student's t copula as follows. Let $\boldsymbol{Z} = (Z_1, ..., Z_d)'$ follow a multivariate normal distribution with zero mean vector, unit variances and positive definite correlation matrix R. Further, let Q be uniformly distributed on $[0, 1]$ and independent of $\boldsymbol{Z}$. Then define $V_j = \sqrt{\nu_j / \chi_{\nu_j}^{-1}(Q)}$ for constants $\nu_j > 2$ and $j = 1, ..., d$, where χ_ν^{-1} is the inverse χ^2 distribution function with ν degrees of freedom. This means that $V_1, ..., V_d$ are perfectly dependent.

The individual Student's t copula is then defined as the copula obtained by inverting Sklar's theorem as in (B.2) for the distribution of

$$
\boldsymbol{X} := (V_1 Z_1, ..., V_d Z_d)'. \tag{B.5}
$$

The univariate margins of $\boldsymbol{X}$ are univariate Student's t distributions with ν_j degrees of freedom, $j \in \{1, ..., d\}$. In other words, each component of an individual Student's t copula has an individual degrees of freedom parameter. The standard Student's t copula is obtained when $\nu_1 = \nu_2 = ... = \nu_d$. Another special case is the grouped Student's t copula with fixed degrees of freedom for groups of variables by Daul et al. [2003]. Luo and Shevchenko [2010] show that the individual Student's t copula has the following properties.

(i) *Pairwise dependence:* As an extension of the standard Student's t copula, the pairwise dependence of the individual Student's t copula can also vary between different pairs. The dependence of a pair of variables is determined in terms of the corresponding entry in the correlation matrix and by the two parameters for the degrees of freedom of the variables.

(ii) *Tail dependence:* The tail dependence of the individual Student's t copula is again symmetric in both tails. In contrast to the standard Student's t copula, it is however decisively determined by the degrees of freedom of the individual variables (see below) and therefore much more flexible.

(iii) *Interpretability:* In terms of general dependence the pairwise Kendall's τ of cells j and k is approximately given by $\tau_{j,k} := \tau_{j,k}(\rho_{j,k}) = \frac{2}{\pi} \arcsin(\rho_{j,k})$ as for elliptical copulas (see (B.3)). According to Daul et al. [2003] and Luo and Shevchenko [2010] the approximation error typically is very small.

The determination of tail dependence coefficients is however less simple. The upper tail dependence of cells j and k can be shown to be

$$\lambda_{j,k}^U := \lambda_{j,k}^U(\rho_{j,k}, \nu_k, \nu_j) = A(\rho_{j,k}, \nu_j, \nu_k) + A(\rho_{j,k}, \nu_k, \nu_j),$$

with

$$A(\rho, \nu_1, \nu_2) = \int_0^\infty f_{\chi_{\nu_1+1}^2}(t) \Phi\left(-\frac{B(\nu_1, \nu_2) t^{\nu_1/(2\nu_2)} - \rho t^{1/2}}{\sqrt{1-\rho^2}} \right) dt,$$

$$B(\nu_1, \nu_2) = \left(\frac{2^{\nu_2/2}\Gamma((1+\nu_2)/2)}{2^{\nu_1/2}\Gamma((1+\nu_1)/2)} \right)^{1/\nu_2},$$

where $f_{\chi_\nu^2}$ denotes the χ^2 density function with ν degrees of freedom and Φ the standard normal distribution function. If $\nu = \nu_j = \nu_k$, this reduces to Expression (B.4) for the standard Student's t copula with ν degrees of freedom. As before, small degrees of freedom indicate higher tail dependence. It holds that $\lambda_{j,k}^U(\rho_{j,k}, \nu_k^{(1)}, \nu_j) > \lambda_{j,k}^U(\rho_{j,k}, \nu_k^{(2)}, \nu_j)$ if $\nu_k^{(2)} > \nu_k^{(1)} > \nu_j$.

(iv) *Computational tractability:* Despite its more sophisticated structure, the individual Student's t copula is still computationally tractable. The margins are simply individual Student's t copulas with corresponding degrees of freedom parameters and sub-matrix of the full correlation matrix. Furthermore the multivariate density only involves a one-dimensional integration, as

$$c_{1,...,d}(u_1, ..., u_d) = \int_0^1 \phi_R(z_1(u_1, q), ..., z_d(u_d, q)) \prod_{j=1}^d (v_j(q))^{-1} dq \left(\prod_{j=1}^d f_{t_{\nu_j}}(t_{\nu_j}^{-1}(u_j)) \right)^{-1},$$

where $v_j(q) = \sqrt{\nu_j/\chi_{\nu_j}^{-1}(q)}$ and $z_j(u_j, q) = t_{\nu_j}^{-1}(u_j)/v_j(q)$ for $j = 1, ..., d$. Here, f_{t_ν} denotes the univariate Student's t density function with ν degrees of freedom and ϕ_R the multivariate normal density function with zero mean vector, unit variances and correlation matrix R.

Statistical inference of the individual Student's t copula is therefore feasible also in higher dimensions. To obtain good starting values for numerical optimization, it is convenient to use inverted pairwise empirical Kendall's τ values as for elliptical copulas. For the individual degrees of freedom parameters, we propose to preliminarily fit bivariate individual Student's t copulas for each pair and then take the average estimated degrees of freedom of each variable as starting value. In our numerical examples, this typically provides reasonably good starting values, which speeds up the numerical optimization.

The individual Student's t copula hence extends the standard Student's t copula at the critical point: Individual degrees of freedom parameters for each variable allow for a more flexible range of tail dependence of the different pairs, while the assumption of one common parameter for the degrees of freedom can be very restrictive if larger numbers of variables are considered. On the other hand, the individual Student's t copula is harder to interpret in terms of its parameters, as its tail dependence coefficients are of more complicated form. Statistical inference is also more difficult due to a one-dimensional integration in the density expression.

B.4 Vine copulas

Vine copulas are a completely different approach to construct multivariate copulas and can also be used to generalize the Student's t copula. They are based on bivariate copulas as building blocks, which made them to be called pair copula constructions (PCCs) as introduced by Aas et al. [2009]. A special case of vine copulas was already discussed by Joe [1996], however Bedford and Cooke [2001, 2002] independently constructed a general class of multivariate distributions, called regular vine distributions, which depend only on bivariate and univariate distributions. They are best illustrated in a trivariate example.

Let $\boldsymbol{X} = (X_1, X_2, X_3)' \sim F$. The multivariate density f of $\boldsymbol{X}$ can then be decomposed as

$$f(x_1, x_2, x_3) = f_1(x_1) f_{2|1}(x_2|x_1) f_{3|1,2}(x_3|x_1, x_2). \tag{B.6}$$

With Sklar's theorem (2.2) it follows that

$$
\begin{aligned}
f_{2|1}(x_2|x_1) &= \frac{f_{1,2}(x_1, x_2)}{f_1(x_1)} = \frac{c_{1,2}(F_1(x_1), F_2(x_2)) f_1(x_1) f_2(x_2)}{f_1(x_1)} \\
&= c_{1,2}(F_1(x_1), F_2(x_2)) f_2(x_2),
\end{aligned} \tag{B.7}
$$

where $C_{1,2}$ is the bivariate copula of the pair $(1, 2)$. In the same way, it holds that

$$
\begin{aligned}
f_{3|1,2}(x_3|x_1, x_2) &= \frac{f_{2,3|1}(x_2, x_3|x_1)}{f_{2|1}(x_2|x_1)} = \frac{c_{2,3|1}(F_{2|1}(x_2|x_1), F_{3|1}(x_3|x_1)) f_{2|1}(x_2|x_1) f_{3|1}(x_3|x_1)}{f_{2|1}(x_2|x_1)} \\
&= c_{2,3|1}(F_{2|1}(x_2|x_1), F_{3|1}(x_3|x_1)) f_{3|1}(x_3|x_1) \\
&\stackrel{(B.7)}{=} c_{2,3|1}(F_{2|1}(x_2|x_1), F_{3|1}(x_3|x_1)) c_{1,3}(F_1(x_1), F_3(x_3)) f_3(x_3),
\end{aligned}
$$

where

$$F_{2|1}(x_2|x_1) = C_{2|1}(F_2(x_2)|F_1(x_1)) := \frac{\partial C_{1,2}(F_1(x_1), F_2(x_2))}{\partial F_1(x_1)} \tag{B.8}$$

and similarly for $F_{3|1}$.

Thus, we have decomposed the density f of $\boldsymbol{X}$ in terms of the marginal densities and the three bivariate copulas $C_{1,2}$, $C_{1,3}$ and $C_{2,3|1}$ with densities $c_{1,2}$, $c_{1,3}$ and $c_{2,3|1}$, respectively. In particular, we have decomposed the three-dimensional copula of $\boldsymbol{X}$ into these bivariate copulas, of which two are unconditional and one is conditional on another variable. The copulas can be chosen independently of each other from different copula classes. For example, $C_{1,2}$ could be a Gaussian, $C_{1,3}$ a Student's t and $C_{2,3|1}$ a Gumbel copula. This flexible way of constructing multivariate copulas therefore allows for a wide range of different dependence structures to be captured.

Clearly, the above decomposition can be extended to the general multivariate case. A d-dimensional PCC of $d(d-1)/2$ bivariate copulas then requires different choices regarding the order of variables in the decomposition as in (B.6). To organize the different decompositions, Bedford and Cooke [2001, 2002] therefore introduced vines as a graphical representation of PCCs. Rather than stating the general definition of vines and vine copulas, we concentrate here

on the properties and refer to Kurowicka and Cooke [2006] and Kurowicka and Joe [2011] for more details. Statistical inference and model selection are considered in Brechmann et al. [2012] and Dißmann et al. [2013].

(i) *Pairwise dependence:* Due to the flexibility in the choice of copulas in the decomposition, different pairs of cells can have very different dependence structures such as asymmetry or tail dependence. However, vine copulas only allow for $d-1$ unconditional bivariate copulas, which are straightforward to interpret (in the example above: for the pairs $(1,2)$ and $(1,3)$), while the dependence of the pairs, for which copulas are specified conditionally, is harder to work with (see the discussion below; in the example above: the pair $(2,3)$ given cell 1).

(ii) *Tail dependence:* Joe et al. [2010] showed that for each pair of cells to have tail dependence, it is sufficient for the unconditional bivariate copulas to have tail dependence. That is, if for instance all $d-1$ unconditional copulas are specified as Student's t, then each pair of cells is tail dependent.

(iii) *Interpretability:* As noted above, interpretation of dependence patterns of vine copulas is complicated by the fact that most pairs are specified conditionally. Only for the dependence of the $d-1$ pairs, for which the copula is given unconditionally, the interpretation is as simple as for the copulas discussed previously. However, simulation from vine copulas is very simple so that properties can be assessed empirically based on sufficiently large simulated data sets.

(iv) *Computational tractability:* Statistical inference of vine copulas is in principle rather straightforward, since the density of a vine copula is conveniently given in terms of a product of bivariate copulas. In the example above we have

$$c_{1,2,3}(u_1, u_2, u_3) = c_{1,2}(u_1, u_2)c_{1,3}(u_1, u_3)c_{2,3|1}(C_{2|1}(u_2|u_1), C_{3|1}(u_3|u_1)),$$

where $C_{2|1}$ and $C_{3|1}$ depend on the copulas $C_{1,2}$ and $C_{1,3}$, respectively (see Expression (B.8)). For log likelihood calculations, this nicely transforms into a sum of log copula densities.

To be useful for our multivariate operational loss model (2.5), we however also require the margins of vine copulas. The availability of the multivariate margins is unfortunately a major issue of vine copulas. Clearly, some margins are available in closed form, in particular those of the $d-1$ pairs with unconditional copula. However, in most cases integration is needed to compute the margins. For example in the above illustration, the margin of the pair $(2,3)$ is given by

$$c_{2,3}(u_2, u_3) = \int_0^1 c_{1,2,3}(u_1, u_2, u_3)du_1.$$

In the worst case, where a pair of cells (j, k) is specified conditionally on all other cells $\{1, ..., d\} \setminus \{j, k\}$, $(d-2)$-dimensional integration may be needed to compute a bivariate margin. This renders the use of vine copulas in model (2.5) hardly feasible even if d is only as large as 5 or 6. Only if there are few zero events, margins may stay sufficiently well tractable. In our application in Section 6, we will show how to calibrate a seven-dimensional vine copula for operational losses per ET.

Although building blocks of vine copulas can be of arbitrary type, we focus here on elliptical copulas. On the one hand, this narrows the wide range of possible constructions, from which it is difficult to choose (see Czado et al. [2013] for an overview on this issue). On the other

hand, this allows us to define an extension of the Student's t copula, which stays interpretable for the following reason: For example, Stöber et al. [2012] show that a Student's t copula can be represented as a vine copula, where the parameters of the bivariate copulas are obtained as partial correlations and degrees of freedom are increased by one for each additional variable that is conditioned on. The Student's t copula can therefore be generalized by a vine copula with Student's t building blocks, where each bivariate Student's t copula is allowed to have different numbers of degrees of freedom. Since this results in a model with a large number of parameters (two parameters per building block), a simpler model can be constructed by only choosing Student's t copulas for the $d-1$ unconditional pairs and Gaussian copulas for all conditioned pairs. According to Joe et al. [2010], this construction also has tail dependence for all pairs. A further model simplification could be achieved by even choosing independence copulas for pairs that are conditioned on a certain number of other variables as discussed by Brechmann et al. [2012]. Note that if a vine copula is specified completely in terms of Gaussian copulas, it is in fact equivalent to a Gaussian copula.

B.5 Estimation of copulas

The number of copula parameters to be estimated depends on the family: While Archimedean copulas typically have just one or at most two parameters, elliptical copulas require the specification of the full correlation matrix with $d(d-1)/2$ entries. Student's t copulas have an additional parameter for the degrees of freedom, while the individual Student's t copula has individual degrees of freedom parameters per variable, that is $d(d-1)/2 + d$ parameters in total. If a vine copula has only Student's t building blocks, it has twice $d(d-1)/2$ parameters, that is $d(d-1)$ parameters. This can be reduced by using Gaussian copula for the conditional pairs. Then, the number of parameters is $d(d-1)/2 + d - 1$.

These parameters are typically estimated by maximum likelihood estimation. Often the method of inference functions for margins (IFM; Joe and Xu [1996]) is used for this purpose. To avoid estimation of too many parameters at the same time, marginal parameters are estimated first and then hold fixed in copula parameter estimation. That means, if we have n (i.i.d.) observations s_{ij} of the total loss in cell j for $j = 1, ..., d$ and $i = 1, ..., n$, then we estimate the marginal distribution $F_{S_j^+}$ by $\widehat{F}_{S_j^+}$ in a first step (see Appendix A) and set

$$\widehat{u}_{ij} := \widehat{F}_{S_j^+}(s_{ij}). \tag{B.9}$$

Using these pseudo observations, the copula log likelihood ℓ_C is then maximized in terms of the copula parameters $\boldsymbol{\theta}$ (see Expression (2.5)):

$$\ell_C(\boldsymbol{\theta}|\widehat{\boldsymbol{u}}_1, ..., \widehat{\boldsymbol{u}}_d, \boldsymbol{w}) = \sum_{i=1}^{n} \log c_{\{j \in D(\boldsymbol{w})\}}(\widehat{u}_{ij}, \ j \in D(\boldsymbol{w}_i)|\boldsymbol{\theta}), \tag{B.10}$$

where $\widehat{\boldsymbol{u}}_j = (\widehat{u}_{1j}, ..., \widehat{u}_{nj})'$ and $\boldsymbol{w} = (\boldsymbol{w}_1, ..., \boldsymbol{w}_n)'$ indicates whether zero losses were observed or not. Note that the zero event distribution $p_{\boldsymbol{W}}$ is independent of the copula parameters and therefore does not need to considered in the estimation here.

For the trivariate Gaussian copula example above, it is $\boldsymbol{\theta} = (\rho_{1,2}, \rho_{1,3}, \rho_{2,3})'$ and (B.10) can be written as

$$\ell_C(\rho_{1,2}, \rho_{1,3}, \rho_{2,3}|\widehat{\boldsymbol{u}}_1, \widehat{\boldsymbol{u}}_2, \widehat{\boldsymbol{u}}_3, \boldsymbol{w})$$
$$= \sum_{i:\ \boldsymbol{w}_i=(0,0,1)'} \log c_{1,2}(\widehat{u}_{i1}, \widehat{u}_{i2}|\rho_{1,2}) + \sum_{i:\ \boldsymbol{w}_i=(0,1,0)'} \log c_{1,3}(\widehat{u}_{i1}, \widehat{u}_{i3}|\rho_{1,3})$$
$$+ \sum_{i:\ \boldsymbol{w}_i=(1,0,0)'} \log c_{2,3}(\widehat{u}_{i2}, \widehat{u}_{i3}|\rho_{2,3}) + \sum_{i:\ \boldsymbol{w}_i=(0,0,0)'} \log c_{1,2,3}(\widehat{u}_{i1}, \widehat{u}_{i2}, \widehat{u}_{i3}|R(\rho_{1,2}, \rho_{1,3}, \rho_{2,3})).$$
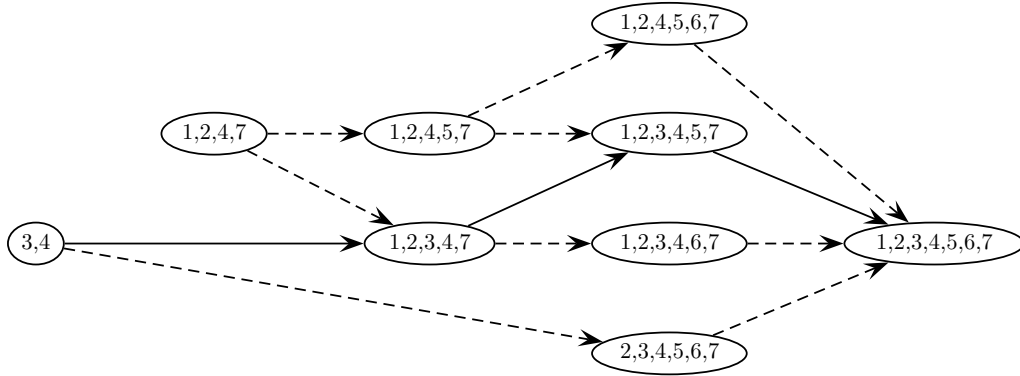
Figure 8: Combinations of observed non-zero events in ETs, where arrows indicate that a set is the subset of another.

For instance, the parameter ρ_{12} is estimated based on the information when $\boldsymbol{w}_i = (0, 0, 1)'$ and when $\boldsymbol{w}_i = (0, 0, 0)'$, that is every time when the pair $(1, 2)$ is observed.

C Vine copula selection for zero-inflated data

From the conditional density decomposition (B.6) it is clear that a d-dimensional PCC can always be extended to a $(d + 1)$-dimensional one by adding the term $f(x_{d+1}|x_1, ..., x_d)$ to the existing PCC. That is, we can construct the vine copula by starting with a bivariate one and then iteratively extending it to the seven-dimensional one, which is needed to describe the dependence between all ETs. For this, we check which set of jointly observed ETs (these are {3,4}, {1,2,4,7}, {1,2,3,4,7}, {1,2,4,5,7}, {1,2,3,4,5,7}, {1,2,3,4,6,7}, {2,3,4,5,6,7}, {1,2,4,5,6,7}, and {1,2,3,4,5,6,7}) is a subset of another as displayed in Figure 8. One option is to start with a vine copula for {3,4}. It is reasonable to extend it to {1,2,3,4,7}, which is the smallest set of which {3,4} is a subset. Then, we choose {1,2,3,4,5,7} and finally {1,2,3,4,5,6,7} as indicated by solid lines in Figure 8. The reason why we choose {1,2,3,4,5,7} instead of {1,2,3,4,6,7}, of which {1,2,3,4,7} also is a subset, is that {1,2,4,5,7}, which is not explicitly included in the vine copula, is a subset of {1,2,3,4,5,7}. In this way, every multivariate margin, that is not an explicit sub-model of the PCC, is a subset of a set that is included and has only one element more. For the vine copula, this means that only one-dimensional integration is needed to integrate out this one additional element. For instance, the margin of {1,2,4,7} is obtained by integrating out variable 3 from the sub-model for {1,2,3,4,7}, since $\{1, 2, 4, 7\} \subset \{1, 2, 3, 4, 7\}$.

References

K. Aas, C. Czado, A. Frigessi, and H. Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2):182–198, 2009.

P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.

F. Aue and M. Kalkbrener. LDA at work: Deutsche Bank's approach to quantifying operational risk. *Journal of Operational Risk*, 1(4):49–93, 2006.

Basel Committee on Banking Supervision. International convergence of capital measurement and capital standards: A revised framework. Technical report, Bank for International Settlements, 2006.

Basel Committee on Banking Supervision. Observed range of practice in key elements of Advanced Measurement Approaches. Technical report, Bank for International Settlements, 2009a.

Basel Committee on Banking Supervision. Results from the 2008 Loss Data Collection Exercise. Technical report, Bank for International Settlements, 2009b.

Basel Committee on Banking Supervision. Operational risk – Supervisory guidelines for the advanced measurement approaches. Technical report, Bank for International Settlements, 2011.

T. Bedford and R. M. Cooke. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial intelligence*, 32: 245–268, 2001.

T. Bedford and R. M. Cooke. Vines - a new graphical model for dependent random variables. *Annals of Statistics*, 30:1031–1068, 2002.

E. C. Brechmann and U. Schepsmeier. Modeling dependence with C- and D-vine copulas: The R-package CDVine. *Journal of Statistical Software*, 52(3):1–27, 2013.

E. C. Brechmann, C. Czado, and K. Aas. Truncated regular vines in high dimensions with applications to financial data. *Canadian Journal of Statistics*, 40(1):68–85, 2012.

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

V. Chavez-Demoulin, P. Embrechts, and J. Nešlehová. Quantitative models for operational risk: Extremes, dependence and aggregation. *Journal of Banking and Finance*, 30(10):2635–2658, 2006.

P. C. Consul and G. C. Jain. A generalization of the Poisson distribution. *Technometrics*, 15 (4):791–799, 1973.

E. Cope and G. Antonini. Observed correlations and dependencies among operational losses in the ORX consortium database. *Journal of Operational Risk*, 3(4):47–74, 2008.

C. Czado, E. C. Brechmann, and L. Gruber. Selection of vine copulas. In P. Jaworski, F. Durante, and W. Härdle, editors, *Copulae in Mathematical and Quantitative Finance*, Lecture Notes in Statistics - Proceedings. Springer, Berlin, 2013.

L. Dalla Valle, D. Fantazzini, and P. Giudici. Copulae and operational risks. *International Journal of Risk Assessment and Management*, 9(3):238–257, 2008.

S. Daul, E. De Giorgi, F. Lindskog, and A. McNeil. The grouped t-copula with an application to credit risk. *RISK*, 16:73–76, 2003.

P. Deb, P. K. Trivedi, and D. M. Zimmer. Dynamic cost-offsets of prescription drug expenditures: Panel data analysis using a copula-based hurdle model. NBER Working paper No. 15191, 2009.

J. Dißmann, E. C. Brechmann, C. Czado, and D. Kurowicka. Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, 59(1):52–69, 2013.

P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, New York, 1997.

P. Embrechts, D. D. Lambrigger, and M. V. Wüthrich. Multivariate extremes and the aggregation of dependent risks: examples and counter–examples. *Extremes*, 12:107–127, 2009.

V. Erhardt and C. Czado. Modeling dependent claim totals including zero claims in private health care insurance. *Scandinavian Actuarial Journal*, 2012(2):106–299, 2012.

K. Fang, S. Kotz, and K. Ng. *Symmetric Multivariate and Related Distributions*. Chapman & Hall, London, 1990.

A. Frachot, T. Roncalli, and E. Salomon. The correlation problem in operational risk. Working paper, Crédit Lyonnais, 2004.

A. Genz and F. Bretz. *Computation of Multivariate Normal and t Probabilities*. Springer, New York, 2009.

R. Giacometti, S. Rachev, A. Chernobai, and M. Bertocchi. Aggregation issues in operational risk. *Journal of Operational Risk*, 3(3):3–23, 2008.

E. Gourier, W. Farkas, and D. Abbate. Operational risk quantification using extreme value theory and copulas: from theory to practice. *Journal of Operational Risk*, 4(3):3–26, 2009.

J. T. Grogger and R. T. Carson. Models for truncated counts. *Journal of Applied Econometrics*, 6(3):225–238, 1991.

M. Hofert, M. Mächler, and A. J. McNeil. Likelihood inference for Archimedean copulas in high dimensions under known margins. *Journal of Multivariate Analysis*, 110:133–150, 2012.

H. Joe. Families of m-variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. In L. Rüschendorf, B. Schweizer, and M. D. Taylor, editors, *Distributions with fixed marginals and related topics*, pages 120–141. Institute of Mathematical Statistics, Hayward, 1996.

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

H. Joe and J. J. Xu. The estimation method of inference functions for margins for multivariate models. Technical Report 166, Department of Statistics, University of British Columbia, 1996.

H. Joe, H. Li, and A. K. Nikoloulopoulos. Tail dependence functions and vine copulas. *Journal of Multivariate Analysis*, 101(1):252–270, 2010.

N. Johnson, S. Kotz, and N. Balakrishnan. *Discrete Multivariate Distributions*. John Wiley, New York, 1997.

S. A. Klugman, H. H. Panjer, and G. E. Willmot. *Loss Models: From Data to Decisions*. John Wiley, New York, 3rd edition, 2008.

D. Kurowicka and R. M. Cooke. *Uncertainty Analysis with High Dimensional Dependence Modelling*. John Wiley, Chichester, 2006.

D. Kurowicka and H. Joe. *Dependence Modeling: Vine Copula Handbook*. World Scientific Publishing Co., Singapore, 2011.

X. Luo and P. V. Shevchenko. The t copula with multiple parameters of degrees of freedom: bivariate characteristics and application to risk management. *Quantitative Finance*, 10(9): 1039–1054, 2010.

J.-F. Mai and M. Scherer. *Simulating Copulas: Stochastic Models, Sampling Algorithms, and Applications*. World Scientific Publishing Co., Singapore, 2012.

A. J. McNeil and J. Nešlehová. Multivariate Archimedean copulas, d-monotone functions and ℓ_1-norm symmetric distributions. *Annals of Statistics*, 37(5B):3059–3097, 2009.

A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts Techniques and Tools*. Princeton University Press, Princeton, 2005.

S. Mittnik, S. Paterlini, and T. Yener. Operational-risk dependencies and the determination of risk capital. Preprint, `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1905600`, 2011.

R. B. Nelsen. *An Introduction to Copulas*. Springer, Berlin, 2nd edition, 2006.

A. Panagiotelis, C. Czado, and H. Joe. Pair copula constructions for multivariate discrete data. *Journal of the American Statistical Association*, 107(499):1063–1072, 2012.

A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de L'Université de Paris*, 8:229–231, 1959.

P. X.-K. Song. *Correlated Data Analysis: Modeling, Analytics, and Applications*. Springer, New York, 1997.

J. Stöber, H. Joe, and C. Czado. Simplified Pair Copula Constructions - Limits and Extensions. Preprint, `http://arxiv.org/abs/1205.4844`, 2012.